



UNIVERSITÉ DE LIÈGE

FACULTÉ DE PSYCHOLOGIE ET DES SCIENCES DE L'ÉDUCATION

INFERNET: A NEUROCOMPUTATIONAL MODEL OF BINDING AND INFERENCE

JACQUES P. SOUGNÉ

A thesis submitted for the degree of Docteur es Sciences de l'Éducation
Augustus 26, 1999

Doctoral committee: Professor Véronique **DE KEYSER**, Université de Liège (Supervisor)
Professor Serge **BRÉDART**, Université de Liège
Doctor Axel **CLEEREMANS**, Université Libre de Bruxelles
Professor Robert M. **FRENCH**, Université de Liège
Professor Dieudonné **LECLERCQ**, Université de Liège
Doctor Denis **MARESCHAL**, Birbeck College, University of London

To my cousin André, who would have been happy to see this done

Contents

Contents.....	iii
Abstract.....	vii
Résumé.....	xi
Acknowledgments	xv
1 Introduction.....	1
1.1 Cognitive modeling	1
1.2 Neurons and neural networks.....	2
1.3 How does the brain represents the world?.....	8
1.3.1 What are the units of representation	8
1.3.2 How does the brain communicate	9
Rate codes.....	9
Pulse codes	10
1.4 The 3 generations of artificial neural networks.....	12
1.4.1 McCulloch & Pitts.....	12
1.4.2 Networks with analogical output	12
1.4.3 Pulsed Networks	12
1.5.1 The binding problem.....	14
1.5.2 Connectionist solutions to the binding problem.....	16
Grandmother cells.....	16
Coarse coding	16
Tensor product.....	17
Value associated to role and filler.....	17
Activation similarity & spatial contiguity.....	18
Binding by temporal frequency.....	18
Binding by synchrony	19
SHRUTI.....	19
LISA	21
Neurobiological justifications for synchrony as a binding mechanism.....	22
1.5.3 Other problems for connectionism.....	23
Multiple instantiation	23
Recursive structure	23
1.6 Overview of the thesis and contribution	23
2 INFERNET.....	25
2.1 Connectivity and spikes.....	25
2.2 Synchrony.....	30
2.3 Phases.....	31
2.4 Symbols.....	32
2.5 Binding, discrimination and predicates.....	33
2.6 Temporal Gates and coincidence detection.....	37
2.7 Multiple instantiation.....	39
2.8 Treatment of Simulator Data.....	41
2.8.1 Proportion of response	42
2.8.2 Reaction times.....	44
3 Short term memory as the activated part of long term memory.....	45
3.1 The capacity of short term memory	46

3.2 Forgetting.....	46
3.3 Arguments for STM LTM distinction.....	47
3.3.1 Serial effect.....	47
3.3.2 Coding systems	48
3.3.3 Double dissociation	48
3.4 INFERNET as a model of STM	48
3.4.1 INFERNET STM capacity.....	48
3.4.2 Forgetting.....	51
3.4.3 Serial effects.....	51
3.4.4 Similarity	53
3.4.5 Double dissociation between LTM and STM	53
3.5 Discussion	54
Overview of empirical studies.....	59
4 Constraints on the length of the reasoning chain	61
4.1 Introduction.....	61
4.2 Conditional reasoning.....	62
4.2.1 Models of conditional reasoning.....	63
4.2.2 Content effect on Conditional reasoning	65
Content effect	65
Deontic effect	65
4.2.3 Effect of negated constituents.....	69
4.2.4 Connectionist models of conditional reasoning.....	70
4.3 Effect of the length of reasoning chain: the case of abstract content with negations.....	70
4.3.1 Conditional reasoning with negated constituents in INFERNET.....	71
INFERNET LTM.....	71
Hypotheses	72
Material.....	73
Procedure and parameters.....	73
INFERNET simulation results.....	75
4.3.2 Experiment 1: Conditional reasoning with negated constituents in humans.....	80
Participants and Design	80
Material.....	81
Procedure.....	81
Results.....	81
4.3.3 Conditional reasoning with negated constituents comparison between INFERNET and humans	84
4.4 Discussion	86
5 The number of predicate's arguments constraints processing	89
5.1 The effect of the number of predicate arguments on cognition.....	89
5.2 Effect of the number of Predicate's arguments	94
5.2.1 Effect of the number of predicate's arguments on INFERNET	94
INFERNET LTM.....	94
Hypothesis.....	95
Design	95
Material.....	95
Procedure and parameters.....	96
Results.....	96
5.2.2 Experiment 2: Effect of the number of Predicate's arguments.....	99
Participants and design	99
Material.....	99
Procedure.....	100
Results.....	100
5.2.3 Comparison of simulator and humans data	101
5.2.4 Discussion	103
5.3 Experiment 3: Reducing the number of a predicate's arguments	104
Participants and design	104
Material.....	105
Procedure.....	106
Hypothesis.....	106
Results.....	106
5.4 General discussion	108
6 Constraints on Multiple Instantiation.....	109
6.1 Introduction.....	109
6.1.1 Connectionism and Multiple Instantiation.....	111

6.1.2	Relevance for Cognitive Science	112
	Reasoning.....	112
	Similarity and Working Memory.....	113
	Similarity in perception.....	113
	Repetition blindness	114
	Animal studies.....	114
6.1.3	Connectionist Solutions.....	114
	Multiple copies loaded in a working area.....	115
	Multiple copies of concepts inside the LTKB.....	115
	Period doubling.....	117
6.2	Double instantiations	118
6.2.1	INFERNET's ability to handle double instantiations	118
	INFERNET LTM.....	118
	Hypothesis.....	121
	Design	121
	Material.....	121
	Procedure and parameters.....	122
	Results.....	122
6.2.2	Experiment 4: Human abilities to treat double instantiation.....	125
	Participants and design	125
	Material.....	125
	Procedure.....	125
	Results.....	125
6.2.3	Comparison between simulator and human data.....	127
	Discussion.....	128
6.3	More than 2 instantiations.....	128
6.3.1	INFERNET ability to handle more than two instantiations.....	128
	INFERNET LTM.....	129
	Hypothesis.....	134
	Design	134
	Material.....	134
	Procedure and parameters.....	134
	Results.....	135
6.3.2	Experiment 5: Human ability to handle more than two instantiations	137
	Participants and design	138
	Material.....	138
	Procedure.....	139
	Results.....	139
6.3.3	Many instantiations: Comparison between simulator and humans.....	141
	Discussion.....	142
6.4	Similarity and multiple instantiation	143
6.4.1	Similarity and Multiple Instantiation INFERNET	143
	INFERNET LTM.....	143
	Hypothesis.....	144
	Design	144
	Material.....	144
	Procedure and parameters.....	144
	Results.....	146
6.4.2	Experiment 6: Similarity and Multiple Instantiation for humans	147
	Participants and design	147
	Material.....	147
	Procedure.....	148
	Hypothesis.....	149
	Results.....	149
6.4.3	Similarity and Multiple Instantiation: Comparison between simulator and humans data	150
6.4.4	Discussion.....	152
6.5	General Discussion	152
7	Exploring the INFERNET Simulator	155
7.1	Noise on connection weights	157
7.1.1	Noise and Threshold.....	157
7.1.2	Increasing vs. decreasing connection strength with noise.....	158
7.2	Does stochastic resonance depend on the problem?	159
7.2.1	Noise on connection and Threshold.....	159

7.2.2 Increasing vs. decreasing connection strength with noise	160
7.3 Noise on spike generation	162
7.3.1 Noise and Threshold.....	162
7.3.2 Comparison between forcing and preventing a spike.....	163
7.4 Noise on Delay	164
7.5 Noise and more realistic PSP curve	165
7.6 Discussion	166
8 Conclusions.....	169
8.1 Comparisons with other work.....	170
8.2 Problems and their potential solutions.....	171
8.2.1 Learning.....	171
8.2.2 Embedding of predicates.....	172
8.2.3 Representation.....	174
8.2.4 The control of attention and consciousness.....	175
8.3 Tractability of the solution.....	175
8.4 Epilogue.....	176
Appendices.....	177
A. INFERNET algorithm description with a pseudo-code.....	177
B. Statistical analysis.....	179
B.1 Frequency tables.....	179
B.2 Reaction times.....	181
C. INFERNET code and simulator	183
Glossary of symbols.....	185
References.....	189
Author index	201
Subject index.....	205

Abstract

We know that color and form are processed in distinct areas of the human brain. This information must somehow be brought together. How the brain might achieve these color-form associations, as well as all other associations of this type, is one of the central themes of this dissertation. When looking at a field of poppies on a sunny day, how can we correctly associate the color red with the poppies, green with the grass, and blue with the sky, while avoiding associating the color red with the grass and the color blue with poppies? How can we associate the perception of red poppies with the name “red poppy,” and with its superordinate category “flower?” A red poppy is composed of several features, like its shape, color, texture, etc. How might a cognitive system bind these features to build a coherent whole? If we see Louise picking a red poppy, how can we correctly associate Louise with the picker and the red poppy with the picked object, without making the opposite and incorrect association? These associations may seem easy to us, but how does the brain achieve them? How a cognitive system binds a set of features together, associates a filler with a role, a value with a variable, an attribute with a concept, ... is what we mean by “the binding problem.”

This thesis focuses on the neurobiological processes that enable connectionist cognitive systems to display binding abilities, on the constraints that affect the binding process, and on the cognitive consequences of these constraints. To study these processes, we developed a computer model of them. This method forces a detailed and unequivocal description of processes used by the simulation. This method is also a powerful means of generating new hypotheses. In this study we attempt to link psychological processes with the neuronal constraints that act on brain functioning. The brain is composed of approximately 10 billion highly interconnected neurons. To achieve binding it is necessary for neurons to communicate with each other because it has been shown that different aspects of a perceived object are not processed in the same cortical areas. Therefore, there must be a means for binding neurons responding to each of these different aspects. The neurons responding to

the color red, to the object's shape, and to its name must be linked to produce a coherent whole representing the red poppy.

Neurons are connected by synapses. The functioning of these connections is constrained by the architecture of the brain and by the process of signal transmission. A particular neuron is connected to a relatively small set of other neurons. Therefore, communication between any two neurons generally requires a chain of transmission through intermediate neurons. A pre-synaptic neuron has an effect on another neuron (called the post-synaptic neuron) only if the pre-synaptic neuron emits an action potential (i.e., if it fires). As a consequence, this brief polarization, which lasts a few milliseconds, results in a modification of other neurons' firing potential. Transmission efficiency depends on the strength of the connecting synapses and the state of the post-synaptic neuron. When a neuron emits an action potential, it is completely insensitive to incoming signals for a short period, then its sensitivity slowly increases. A single pre-synaptic cortical neuron cannot alone provoke the post-synaptic neuron firing. This post-synaptic neuron must receive convergent and more or less synchronized signals from many synapses in order to fire.

These neurobiological properties of neurons and neuronal firing constrain the way in which the brain can achieve binding. Among the various hypotheses of how this could be done, we chose synchronization of action potentials for our model. In the red poppy example, neurons responding to the color red will fire in synchrony with those responding to the shape of the flower and to the name "red poppy." This particular synchronized cluster corresponding to "red poppy" must be temporally distinguished from the cluster responding to "green grass." Numerous neurobiological studies seem to confirm this action-potential synchrony hypothesis. They show that synchronization involves a particular timing precision and occurs at a particular oscillation frequency. This oscillation requires participating neurons to fire repeatedly and rhythmically for a particular period of time. These properties of firing timing and duration have been implemented in a computer simulation called INFERNET. This artificial neural network uses integrate-and-fire nodes (artificial neuron-like elements). These nodes fire at a precise moment and transmit their activation, with a particular strength and delay, to nodes connected to them. When the potential of the node reaches a particular threshold, it emits a spike. Thereafter, the potential is reset to a resting value. As with real neurons, this node will then be completely insensitive to incoming signals for a short period, after which its sensitivity will slowly increase.

INFERNET solves the binding problem by means of oscillation synchrony. Symbols are represented by clusters of nodes firing in synchrony. Fillers are also bound to their roles by synchrony. This synchronous activity defines a window of synchrony i.e., an interval during which the required nodes fire. This time interval takes neurally plausible values. Object discrimination is achieved by a succession of windows of synchrony. Bindings are maintained in memory by the use of particular oscillations. The rhythmic activity and the

synchrony precision constrain the number of distinct entities that the system is able to maintain in memory. This represents the short term memory span of INFERNET. We show that this span is comparable with human short term memory span.

The limited number of windows of synchrony also constrains predicate representations. This prediction is tested on human participants. If there are too many windows of synchrony, these will interfere with each other. In addition, binding strength decreases with time. These two properties explain why the short-term memory of INFERNET displays primacy and recency effects similar to those observed in humans. Bindings in INFERNET are also constrained by the number of intermediate steps required for particular role nodes to enter into synchrony with the filler nodes. This constraint is shown to provided a plausible explanation of various differences human reasoning. The last INFERNET constraint concerns multiple instantiation. This problem arises in connectionist networks as soon as a symbol has to be simultaneously used twice in different ways. Since INFERNET's short term memory is the transient activation of parts of long term memory, it cannot make multiple copies of a symbol, in the same way, for example, that a symbolic system does. The INFERNET solution to the multiple instantiation problem involves superposition of different node oscillations. This process is constrained by the refractory period of the nodes. A number of simulations with INFERNET and experiments on humans show that this solution is psychologically plausible. Multiple instantiation is also shown to be a plausible explanation of certain similarity effects in short term memory. INFERNET is also shown to be capable of symbolic processing with using neurologically and psychologically plausible mechanisms that have the advantages of generalization and noise tolerance found in connectionist networks. Finally, under certain circumstances, noise is shown to enhance INFERNET's processing capabilities.

Résumé

Comment parvient-on, lorsque l'on regarde un paysage, à associer correctement la couleur rouge aux coquelicots, la couleur verte à l'herbe, la couleur bleue au ciel,..., tout en évitant d'associer, le rouge à l'herbe et le bleu au coquelicot ? Comment peut-on associer le percept de coquelicot au nom "coquelicot", à sa catégorie de "fleur"? Le coquelicot est composé de plusieurs traits comme sa forme, sa couleur, sa texture, etc. Comment relier ces traits pour en faire un tout cohérent ? Comment peut-on, voyant Louise cueillir le coquelicot, associer correctement Louise à la "cueilleuse" et le coquelicot à l'objet cueilli en évitant d'associer coquelicot au cueilleur et Louise à l'objet cueilli? Les associations de ce genre nous semblent aisées, tant nous les effectuons avec facilité, mais comment le cerveau parvient-il à former ces associations? L'étude de la capacité des systèmes cognitifs à associer un ensemble de traits, à associer un objet à un rôle, une valeur à une variable, un attribut à un concept, réfère au problème de l'affectation. Cette thèse étudie les processus permettant à un système cognitif d'effectuer des affectations ainsi que les contraintes pesant sur ces processus.

Afin d'étudier ces processus, nous avons eu recours à la modélisation informatique. Cette méthode permet une description détaillée et non équivoque des processus utilisés par la simulation. Cette méthode est aussi un moyen de générer de nouvelles hypothèses.

Nous avons voulu dans cette étude relier les aspects psychologiques de l'affectation aux contraintes neurobiologiques pesant sur le fonctionnement neuronal. Le cerveau est composé de 10 à 100 milliards de neurones interconnectés. Pour parvenir à effectuer des affectations, il est nécessaire que les neurones puissent communiquer entre eux. En effet, les différents aspects d'un objet ne sont pas traités par les mêmes aires corticales. Il doit donc exister un moyen de relier les neurones répondant à ces différents aspects. Les neurones répondant à la couleur rouge du coquelicot, ceux répondant à la forme du coquelicot ainsi que ceux représentant le nom "coquelicot" doivent former un tout cohérent.

Les neurones sont reliés entre eux par des synapses. Le fonctionnement de ces connexions est contraint par l'architecture et par le processus de communication. Un neurone

quelconque n'est connecté qu'à un sous ensemble restreint d'autres. Il en résulte que la communication entre la plupart des neurones nécessite un cheminement à travers des neurones relais. Un neurone ne peut avoir un effet sur un autre neurone que lorsqu'il émet un potentiel d'action (brusque polarisation). Celui-ci est très bref et aura comme conséquence de modifier le potentiel d'autres neurones. Cette transmission n'est pas immédiate, elle prend quelques millisecondes. Le signal sur le neurone receveur sera décroissant. L'efficacité de l'effet d'un potentiel d'action d'un neurone sur une cellule recevant le signal, dépend de la force de la synapse les connectant, et de l'état du neurone receveur. En effet, quand un neurone vient d'émettre un potentiel d'action, il reste insensible aux signaux afférents pendant un moment, ensuite, sa sensibilité va doucement croître. Un neurone ne peut à lui seul provoquer le déclenchement d'un autre neurone. Ce déclenchement ne peut intervenir que si plusieurs signaux convergent de manière plus ou moins synchronisée. Ces propriétés ont des contraintes temporelles qui devraient intervenir dans la façon dont le cerveau code les objets de l'univers et relie ces objets entre eux. Il y a plusieurs hypothèses sur la manière dont le cerveau résout le problème de l'affectation. L'hypothèse qui motive notre travail est la synchronisation des potentiels d'action. Pour reprendre l'exemple précédent, les potentiels d'action des neurones détectant la couleur rouge devraient se synchroniser avec les potentiels d'action des neurones détectant la forme du coquelicot ainsi qu'avec ceux répondant au nom "coquelicot". Cet ensemble se distinguant temporellement de l'ensemble représentant l'herbe verte. De nombreuses études neurobiologiques confortent cette hypothèse, et montrent que la synchronisation a une précision temporelle et qu'elle est provoquée par une fréquence spécifique d'oscillations. Une oscillation consiste en une répétition rythmique des potentiels d'action de neurones.

L'ensemble de ces propriétés a été implanté dans un simulateur appelé INFERNET. Ce réseau de neurones artificiels utilise des noeuds (neurones artificiels) à intégration et déclenchement. Ces noeuds déclenchent à un moment précis puis transmettent avec un certain délai leur activation vers les noeuds auxquels ils sont connectés. Quand le potentiel d'un noeud atteint un certain seuil, il émet un potentiel d'action. Après quoi, ce potentiel est ramené à une valeur de repos. Il sera insensible à toute nouvelle stimulation pendant un moment (période réfractaire absolue) avant de devenir de plus en plus sensible (période réfractaire relative). INFERNET résout le problème de l'affectation par la synchronie. Les symboles sont représentés par un ensemble de noeuds déclenchant de manière synchronique. Les objets sont affectés à leur rôle grâce à cette même synchronie. L'activité synchronique définit une fenêtre de synchronie, c'est à dire un délai à l'intérieur duquel les noeuds requis déclenchent. Ce délai est contraint par des valeurs mesurées par des études neurobiologiques. La discrimination des symboles est obtenue par la succession de différentes fenêtres de synchronies. Les affectations sont maintenues en mémoire grâce aux oscillations qui ont une fréquence particulière.

Cette activité rythmique et la précision de la synchronie contraignent le nombre d'entités distinctes que le système est capable de maintenir. Ceci constitue l'empan de mémoire à court terme d'INFERNET. Cette étude montre le parallélisme de cet empan avec celui de sujets humains. Le nombre limité de fenêtres de synchronies distinctes contraint également la représentation des prédicats et de leurs arguments. Cette prédiction sera confirmée chez l'humain par des études empiriques. Les différentes fenêtres de synchronies interfèrent entre elles et la force des affectations décroît au fil du temps. Ces deux effets expliquent pourquoi la mémoire à court terme d'INFERNET est soumise à l'effet de primauté ainsi qu'à l'effet de récence comme chez les humains.

L'affectation dans INFERNET est contrainte par le nombre d'étapes intermédiaires nécessaires pour qu'un rôle particulier entre en synchronie avec un objet avec lequel il doit être lié. Cette contrainte permet d'expliquer différents effets observés dans le raisonnement humain.

La dernière contrainte étudiée dans ce travail concerne l'instantiation multiple. Ce problème apparaît dans les réseaux connexionnistes lorsqu'un même symbole doit être représenté plusieurs fois de manière distincte et simultanément. Comme INFERNET utilise une mémoire à court terme qui est l'activation passagère de la mémoire à long terme, il ne peut faire plusieurs copies d'un même concept. La solution d'INFERNET à ce problème implique l'utilisation de superpositions de différentes oscillations affectant les mêmes noeuds. Ce processus est contraint par la période réfractaire des noeuds. Différentes expériences menées sur INFERNET et sur des sujets humains ont permis de montrer la plausibilité psychologique du modèle. Cet effet d'instantiation multiple permet aussi d'expliquer l'effet de similarité sur la mémoire à court terme.

Les systèmes connexionnistes ont montré leurs qualités dans l'explication de nombreux phénomènes psychologiques. Ils tolèrent le bruit, ne montrent pas de fragilité lorsque les situations sont légèrement différentes de celles apprises, ils sont capables de faire émerger de nouvelles représentations, etc. Malheureusement, ils ont des capacités de traitement des symboles assez restreintes. De nombreux travaux ont tenté de résoudre ce problème. Beaucoup y sont arrivés, mais toujours au détriment d'une plausibilité psychologique, de la capacité de généralisation ou de la tolérance au bruit. INFERNET parvient, en maintenant les capacités de généralisation, et les qualités de plausibilité psychologique et neurobiologique, à avoir, non seulement une tolérance au bruit, mais dans certaines circonstances, une capacité accrue face au bruit.

Acknowledgments

I would first like to thank Bob French, whose constant support, help and encouragement, has proved invaluable for both my research and writing. I would also like to thank my dissertation supervisor Véronique De Keyser for her support, advice and confidence. Thanks also to Serge Brédart for his participation in my doctoral committee and for his supportive way of providing answers to my questions. Thanks to the other individuals who agreed to be members of my thesis committee, Axel Cleeremans, Dieudonné Leclercq and Denis Mareschal. Many thanks to all my colleagues with whom I had discussions which contributed to this work. Special thanks to Elizabeth Thomas for help in understanding the neurobiological issues, to Myriam Dupont for her help in the analysis of incomplete designs, to Jim Lindsey for his help in understanding Log-Linear Analysis, and to Jean Pierre Thibaut for his comments on this thesis.

My wife Marie-Françoise provided constant support which helped me to overcome the many difficulties and discouragement on the way to completing this dissertation. My daughter Louise has been a constant source of joy and I thank her for forcing me to return to reality when, all too often, I was lost in the clouds of dissertation issues. Thanks to my son Rémi whose birth gave me the final impetus to finish writing this dissertation. Thanks to my parents, for their education. Thank to my extended family and friends for their encouragement, comprehension, and putting up with my unpleasant moods when things were going wrong.

This work was supported by the Belgian PAI Grant p4/19, and the Belgian FNRS Grant D.4516.93.

1 Introduction

The work presented in this thesis lies at the border of three domains: psychology, artificial intelligence and neuroscience. This cognitive science work uses artificial intelligence as a glue to reconcile work in psychology with neuroscience. The object of neuroscience is brain functioning, while the object of psychology is behavior and mental processes. Even if everybody agrees that the brain controls behavior and mental processes, attempts to link brain functioning to mental processes have not always been accepted by psychologists. For example, the seminal work of Hebb (1949) was completely overshadowed by behaviorism. Many researchers still find that bridging the gap between neuronal phenomenon and mental processes is premature.

This work will make an attempt to bridge the gap. Even if future work shows that this attempt was wrong, it will hopefully have enriched psychological science by generating new and falsifiable hypotheses as about the relationship between low-level neural processes and high level cognition.

1.1 Cognitive modeling

Artificial intelligence and computational simulations provide a means of linking brain studies with psychological studies. Computer programming allows the possibility of constructing a large number of simple neuron-like elements and observing them evolving to produce a complex behavior. Artificial intelligence is here a method.

The work presented here started with programming the computer to simulate a large number of neuron-like elements and their synaptic connections. Of course the neuron model used is a simplification of real neurons. In this context, an attempt was made to model complex mental behavior. This behavior was then compared to empirical psychological data. If there is a match between the simulation and human behavior, it simply means that the model can be considered as a *possible* explanation of the underlying cognitive process. Computer modeling is a method to test *possibility* but not *necessity*. Note that there is no

true necessity in empirical sciences, since an hypothesis can be rejected but never proved. Computer modeling has an advantage over other types of theoretical approaches. A theory can hide several flaws which are not revealed until one tries to program it on a computer. A model that has a computational implementation cannot hide undefined shortcuts. Everything must be defined, otherwise, it simply does not work.

In a second phase, the computational model was used to generate predictions. A computer simulation has the ability to react originally on new tasks. This reaction is taken as prediction that will be tested on humans. If human empirical data are different from those of the simulator, it simply means that the model has to be modified or even abandoned. In the meantime, the method contributed to science in providing new hypotheses. Computer modeling is a powerful means of generating new and original hypotheses. For further discussion about contributions of computational modeling to psychology, see Cleeremans & French (1996) or Defagné, French & Sougné (1999).

In every computational model, some decisions are taken a-priori as axioms. For symbolic models, symbols are taken as axioms, even if nobody has ever seen a symbol in the brain. However, these models can also help science in the understanding of complex behavior. Connectionist modeling uses some neuron-like elements with more or less similarity to real neurons. In this work, some characteristics of neurons have been chosen, others have been rejected. Inputs to the system have been set up to display a particular structure. How this structure is achieved is not the purpose of the study. The specific contribution of the present computer model is: “Given a particular input with a particular structure, using neuron-like elements which have certain properties, can we achieve complex high-level cognitive behavior?” The computer model, called INFERNET, was designed, programmed and tested to generate new hypotheses, which were then tested on human participants.


In order to understand how a neurally inspired model like INFERNET behaves, it is necessary to review a number of facts concerning neural transmission. The next two sections of this chapter describe the neural facts that inspired INFERNET. The fourth section, will provide a short review of computational modeling of neurons. The fifth section will review some of connectionist problems and their potential solutions.

1.2 Neurons and neural networks

The human brain is composed of about 10 to 100 billion neurons like those of Figure 1.1. A neuron is a particular cell composed of a soma with a cytoplasm and a nucleus like other cells. The neuron soma integrates the different signals which come from other neurons. A neuron has dendrites which are extensions of the soma. Both dendrites and the membrane, are the receptors of external signals. There is generally one axon for each neuron which is

prolonged by an axonal tree. Axons are of various length and serve to transmit the signal to the axonal terminals which contains neurotransmitters. Neurotransmitters can be released in the intercellular space. They are the output devices of a neuron. The basic input-output unit for transmission of information between neurons is called the synapse. A neuron is connected to thousands of others, for example a pyramidal cell can have 10,000 inputs, while a Purkinje cell can have 100,000 inputs. A neuron which influences the state of another is called “presynaptic neuron” while the influenced neuron is called “postsynaptic neuron”.


Figure 1.1
How neurons are
connected



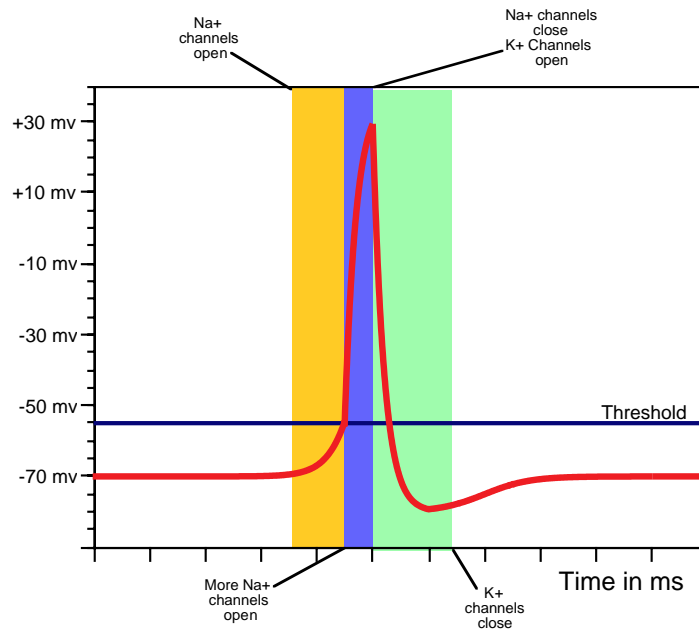
Synapses (see Figure 1.2) are critical devices for communication between neurons. When the presynaptic signal arrives at the axon terminal, vesicles which contain molecules called “neurotransmitters” open and release neurotransmitters into the intercellular space. These neurotransmitters have the ability to fix on some receptor molecules on the postsynaptic neuron dendrite. This process provokes the opening of ion channels. When ion channels are open, ions of a specific kind get into the cell or leave the cell. This flux of ions modifies the electrical potential of the cell. Neurotransmitters have been categorized according to the electrical effect they produce and the speed of their action. Excitatory

neurotransmitters raise the potential: AMPA has a fast but brief action. AMPA channels let Na^+ enter into the cell. NMDA are fast on but switch off slowly. NMDA channels let Na^+ and Ca^{2+} enter the cell. Inhibitory transmitters decrease the potential. GABA_A transmitters are responsible for opening Cl^- channels. They have a fast action in letting Cl^- enter the cell. GABA_B transmitters are responsible for opening K^+ channels. They have a slow action in letting K^+ leave the cell.

Figure 1.2
Synapse



As excitatory channels open, the potential of the membrane slowly increases by the flow of mainly Na^+ which enters the cell. When the potential reaches a certain threshold (around -50 mV), suddenly, more Na^+ channels open and the membrane potential is quickly raised to $\pm 30\text{ mV}$. The Na^+ channels then close, and the K^+ channels open, letting K^+ leave the cell. The potential quickly decreases to a minimum of around -75 mV . When K^+ channels close, membrane pumps expel Na^+ and make K^+ go back into the cell and the membrane potential slowly returns to its resting potential around -70 mV . This process is known as the action potential, as shown in Figure 1.3. This action potential is also called “spike” or a neuron is said to fire. The quick rise and decrease of potential lasts 1 or 2 ms.

Figure 1.3
Action potential

The action potential is a critical event, because it is the only way for a change in the presynaptic cell membrane potential to cross the axonal hillock (see Figure 1.1) and to be transmitted along the axon to the axon terminals. In other words, without an action potential, there is no neurotransmitter release, and no communication between neurons.

When a neuron has fired, it enters into an absolute refractory period in which, no matter what the input, it won't fire again. This absolute refractory period is followed by a relative refractory period. Slowly the membrane resistance decreases, but it will need more neurotransmitter release to obtain the same effect. After ± 10 ms the membrane resistance reaches its resting value. The Figure 1.4 illustrates this process. The X axis represents the time difference in ms between the firing time of the neuron $t^{(f)}$ and the time in consideration t .

What is the time required for a presynaptic action potential to have an effect on the postsynaptic cell potential? The first factor influencing this time course is the length of axon. The speed of the potential transmission along a large myelinated axon is up to 120m/s while for an unmyelinated axon it is about .5m/s. This time is negligible for axons connecting close neurons but for long projections linking distant cortical areas, in the adult human brain, it can be up to 2ms. The second factor is the synaptic delay: the process of neurotransmitter release, crossing the synaptic cliff and binding on a receptor lasts at least 1ms. Finally there is the delay of dendritic integration. The overall process lasts at least 4ms.

Figure 1.4
The membrane
resistance showing
the refractory period

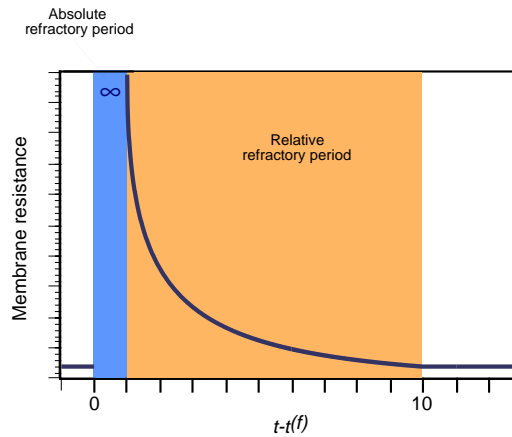
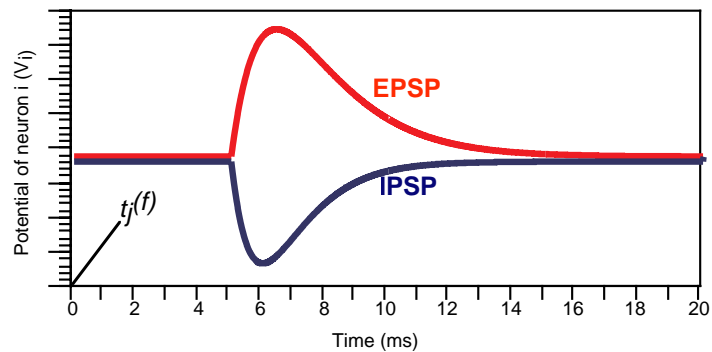


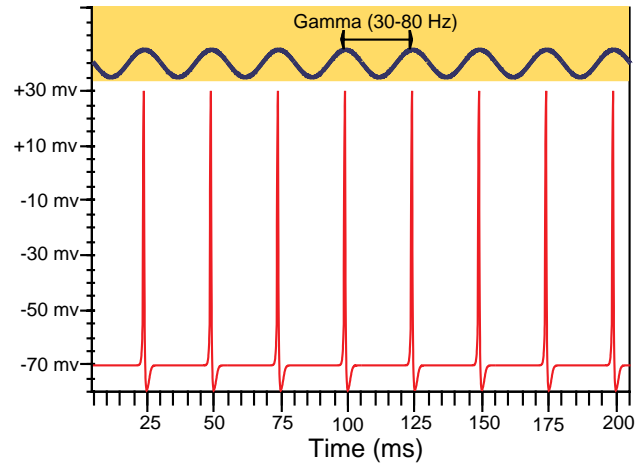
Figure 1.5 shows two curves: Excitatory Post-Synaptic Potential (EPSP) and Inhibitory Post-Synaptic Potential (IPSP). These curves illustrate how the potential of a postsynaptic neuron reacts after a presynaptic action potential. When the presynaptic neuron fires at time $t_j^{(f)}$, a few ms are required for the postsynaptic potential to change (y-axis). This potential rises if the signal is excitatory or decreases if the signal is inhibitory. The signal quickly reaches its maximum and slowly decreases.

Figure 1.5
Function of
Excitatory Post
Synaptic Potential
and Inhibitory Post
Synaptic Potential




There are other temporal properties of neuron firing. Some neurons tend to fire rhythmically, producing oscillations. A well known oscillation shown in Figure 1.6 is called a gamma wave. The lower curve follows the potential of a neuron over time. The oscillation frequency of this neuron is 40Hz, meaning that the neuron fires every 25 ms. The upper curve represents this 40Hz oscillation.

Figure 1.6
Gamma Wave



Oscillations often appear in bursts. A neuron fires rhythmically for a certain amount of time, and then stops firing for a while. This is called a burst. Bursts can reappear rhythmically with a certain period. In that case a second slower wave embeds the faster wave.

Figure 1.7
Gamma sub-cycle in
theta oscillation



The Figure 1.7 shows successive and rhythmic bursts of gamma wave. Each burst of gamma wave is restarted every 250 ms. The upper curve shows a theta wave of 4 Hz in

which appears a gamma sub-cycle of 40 Hz. The lower curve follows the potential of a neuron following these embedded oscillations.

1.3 How does the brain represents the world?

1.3.1 What are the units of representation

How are objects, symbols, concepts represented in the brain? Objects are composed of a large number of features. For example, a soccer ball has a spherical shape composed of alternated black and white hexagons, with a particular texture typical of leather,...

The first hypothesis defended by Barlow (1972), is localist. According to this theory, the soccer ball is represented by a single cell which responds whenever this particular ball is taken into consideration. This unique cell is called a “grandmother cell”. It means that everyone has a particular cell representing her/his grandmother. This kind of representation raises some problems. Imagine a soccer player knocking his head against the goal-post and loosing his football ball cell. What would this player do after getting up? What would he try to kick, the goal-post? A second problem arises from this hypothesis. Think about the enormous number of feature combinations that are necessary for representing the multitude of objects, concepts,... that humans deal with. If every specific combination had a particular corresponding cell, it would require infinitely many neurons. A final problem is generalization. If every new object needs a new representing neuron, similar objects cannot help in representation-building.

The second hypothesis involves fully distributed representations. For Lashley (1950) representation depends on the distribution of neuronal activity and not on particular localization of activity. The brain acts as a whole. A particular object is represented as a pattern of excitation which can affect, potentially, all the neurons in the brain. The problem with this type of fully distributed representation is interference. How can different objects be discriminated if the whole brain reacts to them at the same time? How can the brain learn new material without forgetting old material? Fully distributed networks are prone to catastrophic forgetting (see French, 1999).

The last hypothesis is sparse coding. This hypothesis postulates limited distribution. A particular object is represented by the firing of a number of connected cells. This is Hebb's (1949) description of cell assemblies. A particular cell can participate in different cell assemblies so the distribution quality is maintained. This hypothesis is the most probable one and has received empirical support. A soccer player, waiting near the goal for a corner kick, should be able to recognize the ball, evaluate its movement, its rotation, to decide when to jump and kick or head the ball in an attempt to score a goal. This ability requires the

integration of a series of information. Perceptually, it is well known that cortical area V1 is sensitive to line orientations. The rounded shape of the ball will provoke the firing of a series of orientation cells. But the ball is moving, and area V3 is known to react to moving shapes. Furthermore, motor areas must be synchronized in order for the player to kick the ball at the right time. The whole representation needs successive activations of different cell assemblies. But there is still a problem: how do the different cells participating to a cell assembly, sometimes far apart in the brain, group together to produce a meaningful whole? This will be explained in the next section.

1.3.2 How does the brain communicate

How are neuron assemblies constituted? How can we, as external observers, understand the message involved in the neural activity pattern? What is the code used by the brain? There are two main hypotheses: the rate code and the pulse code.

Rate codes

There are three ways of considering rate. They differ according to the averaging procedure that they use. Rate can be computed on a single cell firing over time. In this case, spikes are counted and divided by the time elapsed, generally between 100 and 500 ms. It has been found that the number of spikes increases as the force applied to some muscle increases (Adrian, 1926). Thomas, Van Hulle & Vogels (1999) recorded several cells in the inferior temporal cortex of a behaving monkey which has to categorize images into tree or non-tree categories. They showed that the majority of neurons responded to both categories. They found that what was crucial is not category specific neurons, but rather those neurons that are responding more to one category than the other (higher rate).


This is an external observer point of view, but do neurons compute similar averages? The first difficulty concerns the fact that behavioral reaction times are sometimes too fast to let the system compute an average. If neurons were firing at regular intervals when oscillating, averages could theoretically be computed after 2 spikes. But there is noise in oscillations. To obtain a good estimation of rate, it is necessary to compute the average over a longer period. In conclusion, this spike rate cannot account for fast processing.

The second procedure for evaluating rate is repeatedly averaging single cell spikes. The experimenter records the spikes of a cell before and after a stimulation, and repeats the experiment. The average is obtained by dividing the total number of spikes by the number of repetitions and the length of the recording intervals. Results are often presented as a peri-stimulus time histogram. Figure 1.8 shows such a peri-stimulus time histogram. The stimulus is presented at time 0. The maximum rate is recorded between 20 and 30 ms after

the stimulus presentation. This measure cannot be the code used by neurons to process information. Much of an organism's reactions ought to be taken consecutively to a single stimulus presentation. For example, a soccer player cannot wait several corner kicks before jumping to try to catch the ball.

The third procedure for computing rate involves recording several neurons before averaging. This rate represents a population activity. Some populations of neurons seem to react to a particular class of stimuli. If after stimulation, an experimenter records the firing of each of these neurons and divides the sum by the number of involved neurons and by the length of the recording time window, a rate measure is obtained. The advantage of this procedure is that it can be calculated from a short time window. The problem raised by this measure concerns the requirement of some kind of rate homogeneity which is not the case in the brain.

Figure 1.8
Peri-stimulus time
histogram



Pulse codes

These codes are based on precise timing of spikes. The first potential code is latency. The idea is that the time separating a stimulus from the first spike of a neuron can carry information. Gawne, Kjaer, & Richmond (1996) recorded activity of striate cortex cells. They showed that the latency was a function of the visual stimulus contrast.


The second code that neurons could use is phases. Oscillation of a global variable like population activity has been found in the hippocampus and cortical areas. This background oscillation can serve as a reference signal (see Figure 1.6). A particular firing of a neuron can be compared to this background oscillation. Its particular position on the oscillation curve can carry a particular code.

This coding scheme has received empirical support. O'Keefe & Recce (1993) for example, showed that phase codes contained spatial information independently of spike rate

in the rat hippocampus. In the olfactory system, some neurons fire preferentially in certain phases of a large scale oscillation (Wehr & Laurent, 1996).

The third potential code is synchrony. A set of neurons that represents micro-features of a stimulus fire at the same time, and thereby allows the representation of the whole stimulus. This hypothesis has received empirical support. Engel, Kreiter, König & Singer (1991) showed that if several objects make up a scene, distinct clusters of synchrony are formed, each associated with a particular object. Engel, König & Singer (1991) observed that individual cells can rapidly change partners of synchrony when visual stimulus changes. Synchrony is regularly measured by cross-correlation. Figure 1.9 shows a cross-correlogram. This measure evaluates the correlation of 2 time series at successive lags. Firing time of two neurons are recorded and constitute 2 time series. If data are correlated at lag 0, as in Figure 1.9, it means that the two neurons are prone to fire synchronously. Synchrony is often associated with oscillation since it has been showed that gamma oscillations enable synchronization (see Singer, 1993).

Figure 1.9
Cross correlogram
sketching relations
between firing times
of two neurons



Rate code and pulse code are not always opposed. Observing synchrony of a population of neurons within a short period of time means also that population firing rate will be high. It is also possible that different information could be coded by different means.

A final problem is how the brain reads the code expressed by neurons. Physiologists as external observers can detect regularities that can account for a potential code in the brain. But people are capable of describing and observing their own thoughts. This would imply that humans have the possibility of actually reading the code. Some hypotheses have been proposed but will be discussed in the chapter 8.

1.4 The 3 generations of artificial neural networks

Maass (1996a) has categorized artificial neural networks into 3 types which correspond to the historical evolution in the domain of connectionist modeling.

1.4.1 McCulloch & Pitts

The first generation is represented by perceptrons (Rosenblatt, 1958), Hopfield networks (Hopfield, 1982), and Boltzmann machines (Ackley, Hinton, & Sejnowski, 1985). These networks are based on McCulloch & Pitts neurons (McCulloch & Pitts, 1943; Pitts & McCulloch, 1947). Each artificial neurons outputs a binary value depending on a threshold value. According to Maass (1996a), these networks are universal for computation with digital input and output, they can compute every boolean function, provided that they include a single hidden layer.

1.4.2 Networks with analogical output

The second generation of artificial neural networks includes those which replace the threshold output function by a sigmoid function or linear function which gives an analogical output. Typical networks of this generation are feedforward or recurrent sigmoidal neural networks. They also can compute boolean functions but mainly, they are universal for computation of functions with analog input and output. They also support powerful learning algorithms like “feedforward backpropagation” (Rumelhart, Hinton & Williams, 1986). One can consider the input and output of such a neural network like the firing rate of biological neuron or population of neurons.

1.4.3 Pulsed Networks

The third generation of artificial neural networks is called pulsed networks (see Maass & Bishop, 1999). They incorporate the properties of neuron firing described in section 1.1.

One of the basic pulsed models is the integrate-and-fire model, it is defined by equation (1.1) which describes how to calculate the state or potential $V_i(t)$ of the node i at time t .


$$V_i(t) = \sum_{j \in \Gamma_i} \sum_{t_j^{(f)} \in F_j} w_{ij} \varepsilon_{ij}(t - t_j^{(f)}) - \sum_{t_i^{(f)} \in F_i} \eta_i(t - t_i^{(f)}) \quad (1.1)$$

A node i fires when its potential $V_i(t)$ reaches the threshold Θ . F_i is the set of all firing times of node i : $t_i^{(f)}$. The set of presynaptic to i node is $\Gamma_i = \{j | j \text{ is presynaptic to } i\}$. F_j is the set of all firing times of node j : $t_j^{(f)}$. Connections weights linking j node to i node is w_{ij} . The kernel $\varepsilon_{ij}(t - t_j^{(f)})$ expresses the postsynaptic potential function as shown in Figure 1.5. While

the kernel $\eta_i(t - t_i^{(f)})$ expresses the refractory period of node i caused by all previous spikes of node i . An example of the function has been given in Figure 1.4. After a node has fired its potential V_i is reset to zero.

This computational neuron acts like a simple electronic device (Figure 1.10) with a capacitor C and two resistances $R_{synaptic}$ and R_{leak} . The resistance $R_{synaptic}$ stands for the connection weight: W_{ij} and the refractory period i.e. the amount of neurotransmitter release and the membrane resistance. The resistance R_{leak} corresponds to a leakage factor, i.e. the decreasing signal in the Post synaptic potential function (Figure 1.5). After a presynaptic action potential, the synaptic strength determines the size of the input current I , it acts as a variable resistance $R_{synaptic}$. When the voltage across the capacitor C reaches the threshold Θ , the circuit is shunted and a pulse is transmitted through the output O to other nodes. This can be compared to the trigger of a camera flash. The capacitor has the power of accumulating potential, but a resistance R_{leak} is responsible for slow loss of potential. The combination of leaky integrator and post-spike resetting are the basic properties of integrate-and-fire neurons.

Figure 1.10
Electronic metaphor
of an Integrate and
fire neuron.



There are many variants of integrate-and-fire neurons: they can take into account the absolute refractory period or not, they can use different approximations for the Post-Synaptic Potential (PSP) function or refractoriness function, they can add noise on threshold, on the reset, or on the integration of inputs.

Networks of integrate-and-fire neurons are a simplification of real neural networks. There are more sophisticated models. Hodgkin-Huxley based models take into account the

difference in behavior of Na and K channels, as well as external and ionic currents. Compartmental models consider a number of dendrite segments which act like a succession of capacitor-resistance circuits.

The more details are taken into account in a computational model, the greater will be the model's computational demands. A compartmental model for example is so demanding, that a computer will not be able to simulate the behavior of a large network. The researcher has to solve the problem of the trade-off between precision in a single neuron model and the network complexity. This can be solved by choosing the appropriate level of detail required for the kind of task to be simulated. If a researcher wants to model a single neuron functioning, she should prefer compartmental models, but if she wants to simulate more coarse grained or more elaborate mechanisms, such a detailed description would probably not be useful. In the present study, the aim is to simulate high-level cognitive behavior which would probably not require many details, but would need large networks (many neuron-like elements). That is why INFERNET is based on simple integrate and fire neurons.

Maass (1996b) showed that spiking neural networks can compute an arbitrary boolean circuit. In addition, Maass (1996c) showed that spiking neural network can perform analog computations.

1.5 Some problems for connectionist networks

What are the difficulties for a connectionist model to display complex cognitive behavior like reasoning? Traditional symbolic systems operate successfully with symbols in a logical fashion, but people's reasoning is not like theorem proving. It more closely resembles a kind of constraint satisfaction (Holyoak & Spellman 1993). Symbolic models lack also of neural plausibility. The brain is not composed of symbols and rules. Symbolic systems also suffer from brittleness i.e. if the input does not match perfectly what is expected, the system cannot produce a response. On the other hand connectionist systems use interconnected neuron-like elements, have a flexible way of functioning, are noise tolerant, have the ability to generalize, and some of them have the ability to construct new representations. However, connectionist systems lack the ability to process symbols. One of the main problems that connectionist system must face is the *binding problem*. There is a need for a solution that solves this problem without losing qualities that give connectionist systems their power.

1.5.1 The binding problem

Humans effortlessly recognize objects, faces, sound, taste,... which are composed of many features. A red ball has a particular shape, color,... Since shape and color are not processed in the same cortical areas, there must be a mechanism able to bind the round shape with the

red color and differentiate this ensemble from a blue cube nearby. A particular face is composed of a set of properties which have to be linked but must also be bound to a particular name and be distinguished from other faces and other names. Representing a predicate and its arguments like “John loves Mary” requires correctly binding the filler “John” to the role of “lover”, and the filler “Mary” to the role of “lovee”, without confusing them. Representing a rule “If a then b” requires correctly binding “a” to the antecedent and “b” to the consequent part of the rule. As these examples show, binding is an essential mechanism used in a varieties of tasks.

Symbolic systems have no problem with binding. They simply define a variable and assign a particular value to it. There is no constraint on the binding process.

However, humans do have problems doing binding in some circumstances. When the time of visual presentation is short, illusory conjunctions are frequent. People bind a feature of one object with another object (Treisman & Schmidt, 1982). Discriminating objects is more difficult if objects share common features (Treisman & Gormican, 1988). There is also more illusory correlations between similar than dissimilar objects (Ivry & Prinzmetal, 1991).

How can an artificial neural network achieve binding? Early critiques of connectionism raised the problem. How could a connectionist network represent the simple fact: “The red rose is on a green table”? Knowing that “red”, “green”, “rose” and “table” have a representation in the system, the problem is to correctly associate “rose” with “red” and “table” with “green” while avoiding “crosstalk”, i.e. avoiding the association between “table” and “red” and between “rose” and “green”. This problem was first introduced by Feldman (1982). The first idea that comes to mind is to increase connection weights between “rose” nodes and “red” nodes and between “table” nodes and “green” nodes while decreasing other connection weights. But once these connection weights have been settled, linked nodes cannot individually participate in other representation. Nodes must be *reusable* for representing another object like “yellow rose”. The binding must occur *dynamically*.

The binding problem is then the inability to correctly associate fillers with roles, values with variable, attributes with concepts, etc. This problem arises as soon as one must discriminate two pairs of associations (in the above example, “red-rose” and “green-table”).

Another problem related to the binding problem is the ability to represent n-ary predicates (predicates with more than one argument). McCarthy (1988) emphasized that connectionist networks can only represent and discriminate unary predicates, (predicates with one argument). How can we represent a predicate and its arguments like “John loves Mary”? If one increases the connection weights between the predicate “love” and its arguments “John” and “Mary”, how could the “lover” be distinguished from the “lovee”. How “John loves Mary” could be distinguished from “Mary loves John”? A solution must involve the discrimination of two pairs of associations: “John-lover” and “Mary-lovee”.

Fodor & Pylyshyn (1988) point out a difficulty that arises from the binding problem. They question the value of connectionism as model of cognition based on the impossibility for connectionism to display what they called “systematicity”. Systematicity is not well defined by these authors. They claim that human cognitive abilities are systematically linked and appear by sets. One can define it as the ability of applying particular knowledge to any content whatsoever. A series of examples are given. These examples make reference to the binding problem. Fodor & Pylyshyn (1988) write that nobody has ever seen people able to think that “John loves Mary” but unable to think that “Mary loves John”. Nobody is able to infer that “John went to the store” from “John, Mary, Susan, and Sally went to the store” but unable to infer that “John went to the store” from “John, Mary and Susan went to the store”. More formally, systematicity can be illustrated by these general examples: the ability to think $P(a,b)$ is linked to the ability to think $P(b,a)$, the ability to think A from $(A \& B \& C \& D)$ is linked to the ability to think A from $(A \& B \& C)$. But as Hadley (1996) states, does systematicity prevail for more complex logical forms? Does the ability to think $\sim A$ from $\sim B$, $(A \supset B)$ is linked to the ability to think A from B, $(\sim B \supset \sim A)$? In reading Fodor & Pylyshyn (1988) we have the impression that they must believe this. But as will be shown in chapter 4, these latter two inferences, even formally equivalent, are not at all as well accepted by people. Fodor & Pylyshyn (1988) do not question the empirical relevance of what they state. Research in cognitive psychology has demonstrated a great many content effects. For example, Ingram (1985) showed that children learning to talk do not display systematicity in their predicate use. The chapter 4 provides examples showing that people’s inferences are sensitive to the content.

While symbolic systems are clearly too efficient in solving the binding problem compared to humans, connectionist systems of the eighties, by contrast, were unable to do binding at all. Some solutions emerged and will be presented next.

1.5.2 Connectionist solutions to the binding problem

Grandmother cells

The first solution to this problem is to use one node for each possible association. Unfortunately, this solution quickly becomes intractable due to combinatorial explosion. This *grandmother cell* solution has already been discussed in section 1.3.1.

Coarse coding

One possible way to avoid combinatorial explosion is to use coarse coding (see Hinton, McClelland & Rumelhart, 1986; McClelland, 1986; Touretsky & Hinton, 1988). Coarse

coding uses a subset of the population of nodes to code a particular value. A node becomes active if the input falls in its receptive field. Coarse coding needs fewer nodes to represent the same amount of information compared to purely localist codes. However, this solution only allows one variable to be bound at a time (see Mc Carthy, 1988).

Tensor product

Smolensky (1987, 1990) proposed the use of a tensor product representation of binding. Tensor products are similar to outer products of vectors. If one has to encode “John loves Mary”, the “John” filler must be bound to the “lover” role. Suppose that “lover” is represented by the vector [1011] and “John” by the vector [0111], their binding will be represented by:

$$\text{lover} \otimes \text{John} = \begin{bmatrix} 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 \end{bmatrix} \quad (1.2)$$

This representation of “John” in the role of “lover” would be clearly different than “John” in the role of “lovee”. Because “lovee” would be represented by a different vector [1010] than “lover”. “John” in the role of “lovee” would lead to this tensor product:


$$[1 \ 0 \ 1 \ 0] \otimes [0 \ 1 \ 1 \ 1] = \begin{bmatrix} 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad (1.3)$$

According to Touretsky (1995), this architecture is a retrieval model and lacks inference ability. However, Halford, Wilson, Guo, Gayler, Wiles & Stewart (1994) modified Smolensky initial model to perform analogical inferences.

Value associated to role and filler

ROBIN (Lange & Dyer, 1989a; Lange & Dyer, 1989b; Dyer, 1991; Lange, 1992) separates roles from fillers. Each fillers has an associated node that outputs a particular constant value (called its signature). Each role has an associated object node or binding node. When a role object node has the same activation as that of a concept signature, this concept is bound to the role. The Figure 1.11 illustrates how ROBIN can successfully encode “John loves Louise” by correctly binding “John” to “Lover” and “Louise” to “lovee”. The signature for “John” has the same value as “lover” object node activation, while the “Louise” signature has the same value as “lovee” object node activation.

Figure 1.11
ROBIN binding is
achieved by a match
between activation
of role object node
and filler signature.



A similar solution has been proposed by Sun (1992, 1995). His CONSYDERR model is a dual architecture, consisting of a localist network and a distributed network. The former is composed of nodes representing concepts and links between these concepts representing rules. Each concept node is linked to several nodes in another network which is distributed. The nodes of the associated distributed network represent the features of the concept. Variable binding is achieved by the use of a particular value which is passed along a link from a role node to a filler node.

Activation similarity & spatial contiguity

COMPOSIT (see Barnden, 1991, 1992, 1994; Barnden & Srinivas, 1991) uses two systems, a Long Term Memory and a Working Memory, both of which are connectionist networks. In this model, Working Memory is composed of several registers filled with activation patterns from Long Term Memory. Fillers and their roles are stored in registers with 2 vectors. One vector represents the filler (the symbol vector) the other vector represents the role (highlighting vector). Predicates, related roles and fillers are stored contiguously and thus constitute a distinguishable ensemble that can be linked to a particular role pertaining to another predicate. A role can be linked to another role by the similarity of their highlighting vectors.

The above solutions have different advantages and limits, but they all lack neural plausibility. We will now explore more neurally plausible solutions.

Binding by temporal frequency

Lange & Dyer (1989) proposed the use of temporal frequency (replacing signatures) for binding. To the best of our knowledge, these authors never actually implemented this solution. Surprisingly, this option has not yet been explored by cognitive scientists, even

though some neurobiological empirical data seem to make it a potential neural mechanism for binding (see Gerstner, Kreiter, Markram & Herz, 1997). If the activation of a node is considered to be its firing rate, binding by activation as done in CONSYDERR and ROBIN can fall into this category. However, firing rate and activation are two different parameters and pulse networks could be used to test binding by temporal frequency.

Binding by synchrony

For temporal synchrony variable binding systems, nodes can be in two different states: they can fire (“on”), or they can be at rest (“off”). A node fires at a precise moment and transmits activation to other connected nodes. When a node activation reaches threshold, it fires. A concept, a predicate, an object or a role is activated when the nodes representing it fire.

Whenever two nodes (or two sets of nodes) representing two objects fire simultaneously, these objects are temporarily associated. On the other hand, if two nodes (or two sets of nodes) fire in succession, they are discriminated. This is how these systems solve the binding problem.

A number of connectionist models have used synchrony for binding, among them: BROADCAST-NET (Clossman, 1988), SHRUTI (Shastri & Ajjanagadde, 1993), LISA (Hummel & Holyoak, 1997), and INFERNET (Sougné, 1996, 1997, 1998a 1998b, Sougné & French, 1997). ART has also been adapted to bind by synchronization (Grossberg & Somers, 1992). Here follows the descriptions of the two main competitors of INFERNET: SHRUTI and LISA.

SHRUTI


SHRUTI (Shastri & Ajjanagadde, 1993; Shastri, 1999) is a model designed to simulate the kind of inferences that people make effortlessly and rapidly. For example, while reading that “John gave a book to Mary”, one can infer that “Mary owns a book”. Representation in SHRUTI is localist, a node represents an object, a role, a fact, a predicate, etc. SHRUTI uses different kinds of nodes which have different behaviors of which the main ones are described below. Predicates are represented by two nodes which are called enabler nodes (e:). One of these nodes serves to affirm the predicate, the other to negate the predicate. When these node are on, they are on for a period of time covering the firing of successive role node firings. Arguments of a predicate are represented by nodes that fire for a shorter time. They fire after the starting of firing of the predicate enabler. Different predicate role nodes fire in succession. The whole process is repeated over and over for a while. This serves to maintain the representation. Figure 1.12 shows the representation of “John loves Louise” in SHRUTI. The affirmative enabler predicate node e:Loves is on for a period of time. Within this period, role nodes “Lover” and “Lovee” will fire in succession. Binding

occurs by synchrony of “John” node firing with “Lover” node firing and synchrony between “Louise” node firing with “Lovee” node firing.

Inferences are performed by connecting different predicate enabler nodes and role nodes. Activated predicate and role nodes transmit their activation to other predicate and role nodes. A particular fact is encoded by an enabler node which will be activated only when role nodes of a predicate are bound. This is done by a set of excitatory and inhibitory connections.

SHRUTI has been used to simulate different tasks like implicit reasoning (Shastri & Ajjanagadde, 1993), planning (Shastri, Grannes, Narayanan, & Feldman, 1998), parsing (Henderson, 1994), etc. However, the use of “grandmother” cells is questionable. Even if Shastri and Ajjanagadde (1993) describe a (non implemented) version of SHRUTI in which each concept is represented by a set of nodes, the sets cannot be overlapping (i.e. one object cannot share a common node with another object).

Figure 1.12
Node firing behavior
in SHRUTI
representing “John
loves Louise”



Many features of SHRUTI seem redundant and this has some disadvantages. SHRUTI makes use of a large number of different purpose nodes. According to Ajjanagadde (1994) himself it is not necessary. But there are other problems related to knowledge abstraction and interaction. For example:

- Each predicate has its own role nodes. If predicates cannot share the same role node, how can phenomena like human analogy-making be explained? Analogy has been observed to be facilitated by objects that share the same role in different predicates.
- The distinction between role and object raises the question of how an object can become, in certain contexts, a role. For example: “gift” as role in Offer(Gift, giver, recip) and as object in “Mary gave John a gift”.
- Each fact is encoded by a particular node, how can SHRUTI explain interference among a number of facts (leading to generalizations or even confabulations).

- The solution to the problem of negation raises some questions. There is no abstraction of the negation because it is specific to a predicate. How can SHRUTI could treat double negations? How could it apply negation to an object instead of a predicate?

These remarks raise the question about how to maintain advantages of connectionist models over symbolic models in a model like SHRUTI.

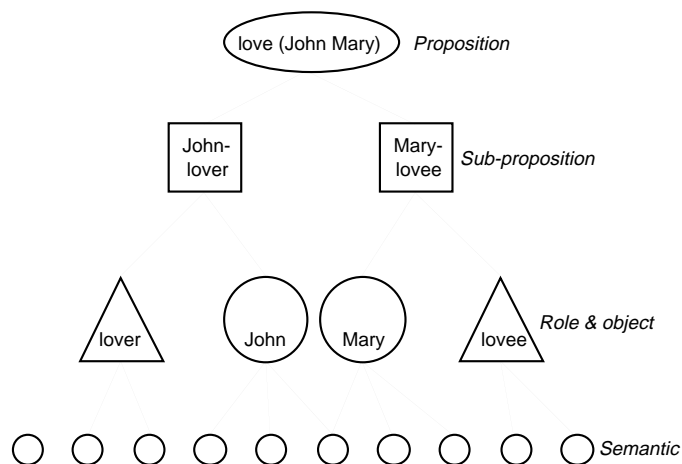
SHRUTI cannot be considered as a pulse network mainly because it does not use a Post-Synaptic Potential function. Transmission along a connection can be instantaneous and there is no activation leakage.

LISA

LISA (Hummel & Holyoak, 1997) is a computational model of analogical inference and schema induction. It represents a proposition by a Long Term Memory (LTM) structure as showed in Figure 1.13.

Each object and role is represented in Working Memory (WM) by an activation pattern of semantic units. These units are connected to role (lover, lovee) and object units (John, Mary). Semantically related roles or objects share semantic units: “John” and “Mary” are human. Sub-proposition units link roles to objects in LTM (John as lover and Mary as lovee). Proposition units link sub-propositions to encode a particular fact, such as, “John loves Mary” in LTM. Connections are bi-directional, enabling encoding a fact as remembering a fact. There is a learning algorithm that modifies connection weights in order to encode a particular fact.

Figure 1.13
LISA encoding of
the fact “John loves
Mary”



The network can also map connections between units of the same type (Proposition, sub-proposition, Role, Objects) in different structures. These connections permit analogy-making to be performed in a structured fashion.

Variable binding in this model is achieved by synchrony. When a particular object node fires in synchrony with a role node they are bound.

LISA has the advantage of using simpler elements than SHRUTI and also uses distributed representations. Finally, it permits fact encoding via a learning algorithm. But LISA is not a pulsed network. There is no Post-Synaptic Potential function nor refractory period.

Neurobiological justifications for synchrony as a binding mechanism

The advantage of synchrony variable binding lies in its neurobiological plausibility. There is considerable evidence for considering synchrony as the most plausible binding mechanism in the brain (see Singer, 1993; Singer & Gray, 1995; Roelfsema, Engel, König & Singer, 1996; Singer, Engel, Kreiter, Munk, Neuenschwander & Roelfsema, 1997). Spatially separated neurons driven by the same stimulus have been found to synchronize. This kind of synchrony has been observed between distant cells in the same cortical area (Gray, König, Engel, & Singer, 1989; Gray, Engel, König, & Singer, 1992; Engel, König, Gray & Singer, 1990), between cells in different cortical areas (Eckhorn, Bauer, Jorden, Brosch, Kruse, Munk & Reitboeck, 1988; Engel, Kreiter, König & Singer, 1991; Murthy & Fetz, 1992; Nelson, Salin, Munk, Arzi, & Bullier, 1992; Roelfsema, Engel, König & Singer, 1997), and between cells in different hemispheres (Engel, König, Kreiter & Singer, 1991).

Synchrony seems to appear if neurons are driven by the same stimulus, and disappears in the presence of two different stimuli. Engel et al. (1990) have shown that individual cells can rapidly change partners of synchrony if the stimulus changes. If several objects make-up a scene, distinct clusters of synchrony are formed, each associated with a particular object (Engel, Kreiter, König & Singer, 1991).

While most of the empirical data involving synchronization have focused on the visual cortex, synchronization has also been found outside the visual cortex. Murthy & Fetz (1992) observed synchronization between cells in somatosensory and motor cortex (see also Riehle, Grün, Diesmann & Aertsen, 1997). It has also been found in acoustic and frontal cortex (Aertsen, Vaadia, Abeles, Ahissar & Bergman, 1992; Vaadia, Ahissar, Bergmann, & Lavner, 1991), as well as in the hippocampus (Bragin, Jando, Nadasdy, Hetke, Wise, & Busaki, 1994). Roelfsema, Engel, König & Singer (1997) observed stimulus-dependent synchronizations between visual and parietal cortex and between parietal and motor cortex of an awake cat.

Some specific connections seem to be responsible for synchronization. Engel, König, Kreiter & Singer (1991) showed that sectioning the corpus callosum abolished response synchronization across area 17 of the two hemispheres. Cortico-cortical connections are probably those enabling synchronization. These connections are established mostly post-natally. When strabismus is provoked in young kittens, cortico-cortical connections, which normally connect areas irrespective of the dominant eye, develop to link only areas served by the same eye (Löwel & Singer, 1992). In that case, these authors have shown that synchrony disappeared between areas served by the two different eyes, and that coordination of binocular signals vanished.

According to Singer (1993) and Singer et al. (1997) specific oscillations in the gamma range seem to be the best candidate for enabling synchronization.

It would be erroneous to believe that all neural activity displays some sort of synchronization. Synchrony occurs between cells but this synchrony is also accompanied by an unstructured activity which resembles noise. We will see in chapter 7, that contrary to common belief, noise can enhance convergence in complex dynamic systems.

1.5.3 Other problems for connectionism

Multiple instantiation

Multiple instantiation (see Sougné, 1998 for a review) involves the simultaneous use of the same parts of the knowledge base in different ways. Knowing that “John is in love with Louise” and that “Louise is in love with John”, one can readily infer that they should be happy. To arrive at this conclusion, one must instantiate the predicate “is in love with” and the objects “John” and “Louise” twice. Precisely how this is done is the problem of multiple instantiation. This problem will be discussed in detail in chapter 6.

Recursive structure

Understanding the sentence: “The boy who kissed the girl who kissed the cat was my friend” requires recursive predication. A particular predicate must be bound to a particular role of another predicate. This point will be discussed in the Chapter 8.

1.6 Overview of the thesis and contribution

INFERNET is a pulsed network using neuronal elements of the integrate-and-fire type. This model is clearly inspired from real neuronal functioning. In INFERNET, nodes can be in two different states: they can fire (“on”), or they can be at rest (“off”). A node fires at a

precise moment and transmits activation to other connected nodes with some time course. When a node activation or potential reaches a threshold, it emits a spike. After firing, the potential is reset to some resting value. Chapter 2 describes in detail INFERNET structure and functioning.

In INFERNET variable binding is achieved by synchrony. How do synchrony and related parameters constrain INFERNET's cognitive abilities? A series of tasks were given to INFERNET. Its reactions have been collected. These provide predictions that were tested with humans in an experimental setting. INFERNET predictions were then compared to human empirical data.

The following chapters compare INFERNET simulation results with human data.

INFERNET implements a short term memory as the activated part of long term memory. By taking realistic neurobiological parameters, we show that short term memory capacity must be limited. Chapter 3 shows how INFERNET can display several features that human short term memory studies reveal.

INFERNET is able to perform certain reasoning tasks. But it is constrained by the length of the reasoning chain, and, consequently, by double negation also. Chapter 4 examines how INFERNET is able to deal with conditional statements in a plausible manner.

Predicates and their arguments are coded by phases in INFERNET. The more arguments a predicate has, the more it will require discriminations among the phases. Chapter 5 will compare INFERNET and human performance on the processing of n-ary predicates.

A particular solution to the problem of multiple instantiation is proposed in INFERNET. This solution puts constraints on the treatment of multiple instantiations. Chapter 6 will compare INFERNET and human performance on multiple instantiation.

In the brain many spikes apparently do not participate in any particular code, and could therefore be considered as noise. Chapter 7 presents several experiments with noise in INFERNET. It will be shown that the introduction of noise can improve INFERNET performance.

Chapter 8 is devoted to conclusions and future research.

This work shows that taking neurobiological constraints and implementing them in a computer program can provide potential explanations of human behavior and mental processes, in addition to generating fruitful hypotheses for future research.


2 INFERNET

INFERNET is a pulsed network which implements a number of neurobiological constraints. In order to understand how a neurally inspired model like INFERNET behaves, it is necessary to keep in mind the introductory discussion in chapter 1, especially sections 1.2 and 1.4.3.

2.1 Connectivity and spikes


INFERNET is not a fully connected network, its structure is organized by clusters of nodes which constitute subnets. Each subnet is fully connected. From each node of a subnet there is a connection to every other node in the subnet. Some of subnet nodes possess connections to external subnet nodes. This organization is shown in Figure 2.1. It not only reduces the computational demands of the program, but also correspond better to the organization of the brain. The brain is not fully connected, in particular, there are more intra-cortical area connections than inter-cortical connections.

Figure 2.1
Subnets are intra
fully connected and
inter partly
connected



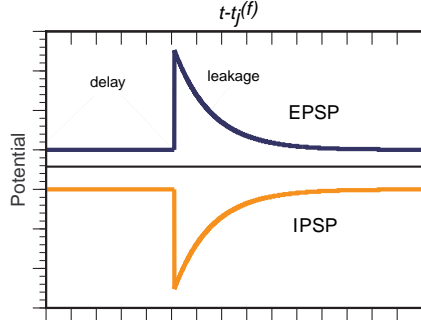
Each connection is either excitatory or inhibitory. Excitatory connections raise the potential of postsynaptic nodes, while inhibitory connections decrease the potential of postsynaptic nodes. Excitatory connections correspond to transmitters which open AMPA and NMDA channels, while inhibitory connections correspond to transmitters which open GABA channels. Two variables affect each connection: weight and delay. They are shown in Figure 2.2. The weight corresponds to synaptic strength between a presynaptic and postsynaptic cell. The weight between a presynaptic node j and a postsynaptic node i is designated by w_{ij} . Noise is added to this value and resulting noisy connection is denoted by \hat{w}_{ij} . The delay d of a connection determines when the effect of the presynaptic node firing will be maximum on the postsynaptic node. There is also a random factor on the delay. The noisy delay is denoted by \hat{d} . This delay corresponds to the axonal, synaptic and dendritic delays of real neurons.

Figure 2.2
Weight and Delay



A signal, whether excitatory or inhibitory will be affected by a leakage factor. When the signal has reached its maximum, at each following step of 1 ms, the signal will be divided by 2. Delays and leakage factor define EPSP or IPSP curve described in Figure 2.3. The y-axis refers to the postsynaptic node potential V_i . The x-axis is the time difference in ms between the time under consideration t and the time of the presynaptic node firing $t_j^{(f)}$.

Figure 2.3
EPSP and IPSP
function in
INFERNET



The resulting postsynaptic (PSP) equation $\varepsilon_{ij}(x)$ where x is the difference between the time in consideration, the time of the presynaptic node firing and the noisy delay on the connection $x = t - t_j^{(f)} - \hat{d}$ is :

$$\varepsilon_{ij}(x) = \frac{1}{2^x} \mathcal{H}(x) \quad (2.1)$$

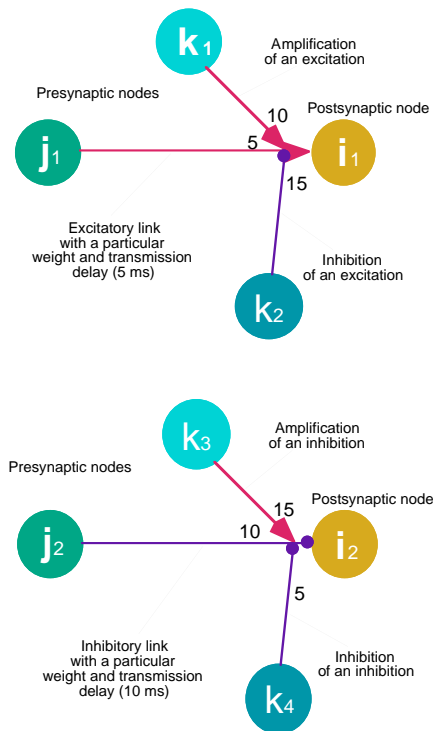
with:

$$\mathcal{H}(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \quad (2.2)$$

In INFERNET, other kinds of connections are used. Unlike most links, these latter links act on *connections* rather than nodes (Hofstadter, 1984; French, 1995). Moreover, each of these connections can be excitatory or inhibitory. There are six types of connections: postsynaptic excitation, postsynaptic inhibition, presynaptic amplification of an excitation, presynaptic inhibition of an excitation, presynaptic inhibition of an inhibition and presynaptic amplification of an inhibition. Figure 2.4 shows the entire set of connections. In INFERNET, for the simplicity of computation excitatory weights have a integer value between 0 and 127 while inhibitory weights have a value between -127 and 0. Moreover, an excitatory connection cannot be modified to become inhibitory and vice-versa.


There are neurobiological justifications for these connections. Some axon terminal neurotransmitter releases act on synapses rather than dendrites. Presynaptic inhibition acts on axonal terminals by opening K^+ or Cl^- channels thus reducing the size of the voltage delivered at the end of the axonal terminal. Ca^{2+} entrance is reduced and neurotransmitter release is decreased. Presynaptic facilitation can occur when serotonin is released near an axon terminal. A second messenger cAMP then enters the terminal and increases the amount of Ca^{2+} entering which increases neurotransmitter release.

Figure 2.4
The six types of connection



When a node potential: V_i reaches a threshold Θ , it emits a spike. Thereafter, the potential is reset to its resting value. Figure 2.5 illustrates a spike of INFERNET node. This is a simplification of the neuron action potential function presented in Figure 1.3.

Figure 2.5
A spike in
INFERNET. (the
potential is reset
to the resting value
after the spike).



After emitting a spike, a node enters a refractory period. This corresponds to the membrane resistance of real neurons which increases after a spike. In INFERNET, the


function is simplified. The resistance of node i depends only on the last spike of the node i : $\hat{t}_i^{(f)}$. In more detail, the equation expressing the refractory period $\eta_i(u)$, (where u is the difference between the time in consideration and the time of the last spike of the node: $u = t - \hat{t}_i^{(f)}$), is:

$$\eta_i(u) = -\left[-e^{-[u^a-b]}\right] \mathcal{H}'(u) \Theta \quad (2.3)$$

$$\text{with} \quad \mathcal{H}'(u) = \begin{cases} \infty, & \text{if } 0 \leq u < 1 \\ 1, & \text{otherwise} \end{cases} \quad (2.4)$$

The parameter a is approximately equal to .8 and b to 4.5. However, in the computer program a list of integers was used to simplify computations. This resistance is plotted in Figure 2.6. The absolute refractory period lasts 1 ms and the relative refractory period ends 10 ms after the spike.

Figure 2.6
Threshold function
in INFERNET



All variables affecting the potential of a node have now been defined. Equation (2.5) express how the potential $V_i^{(t)}$ of node i is calculate at each time step.

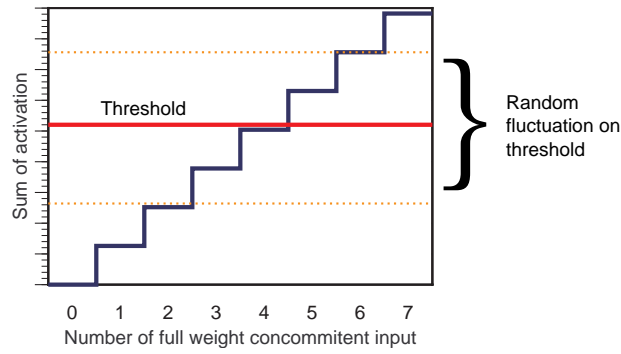
$$V_i^{(t)} = \sum_{j \in \Gamma_i} \sum_{t_j^{(f)} \in F_j} \left[\hat{w}_{ij} + \sum_{k \in K_{ij}} \sum_{t_k^{(f)} \in F_k} \hat{w}_{k \rightarrow ij} \varepsilon_{k \rightarrow ij}(x) \right] \varepsilon_{ij}(x) - \eta_i(u) \quad (2.5)$$

A node i fires when its potential $V_i(t)$ reaches threshold Θ . This potential is affected by connection weigths coming from presynaptic node \hat{w}_{ij} but also by the connections weigth that modify this connection $\hat{w}_{k \rightarrow ij}$. The set of presynaptic to i node is $\Gamma_i = \{j | j \text{ is presynaptic to } i\}$. F_j is the set of all firing times of node j : $t_j^{(f)}$. Noisy connection weights linking j node to i node are \hat{w}_{ij} . The set of presynaptic to ij synapse is $K_{ij} = \{k | k \text{ is presynaptic to } ij \text{ synapse}\}$. These are connections that act on connection ij . F_k is the set of all firing times of node k : $t_k^{(f)}$. These are the nodes from which start a connection acting on the connection ij . Noisy connection weight linking k node to ij synapse is $\hat{w}_{k \rightarrow ij}$. The equations $\varepsilon_{ij}(x)$ and $\varepsilon_{k \rightarrow ij}(x)$ expresses the postsynaptic potential function as shown in Figure 2.3.

While the equation $\eta_i(u)$ express the refractory period of node i caused by the previous spikes of node i .

A comment is necessary about threshold Θ . Random noise is added to its value. The firing of a node requires at least 3 concomitant full excitatory inputs and at most 7 concomitant full excitatory inputs, if the node is out of refractory period. Figure 2.7 illustrates this process. In agreement with real neural networks, a single synapse cannot provoke a postsynaptic action potential. In reality, a single cortical synapse raises the potential of a postsynaptic neuron by 0.1 to 1 mV (Abeles, Prut, Bergman, Vaadia, & Aertsen, 1993).

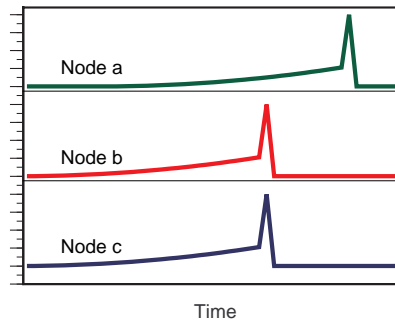
Figure 2.7
Number of full weight
concomitant input to
cross the threshold
of a node out of
refractory period



2.2 Synchrony

As described in the previous section, node firing occurs at a precise moment in the ms range. One can record firing time of several nodes. In Figure 2.8, three node spikes are displayed. The spikes of nodes b and c are synchronous while node a spike is asynchronous with node b and c spikes. Synchronous activity in INFERNET is the signal of a particular code. It simply means that these nodes belongs together.

Figure 2.8
Synchrony: Nodes b
and c spikes are
synchronous




2.3 Phases

INFERNET has the ability to display rhythmic patterns of node-firing. A node can fire regularly for a period of time. It displays a particular oscillation pattern. In Figure 2.9, three nodes fire regularly at an identical frequency. The oscillation displayed on the upper part of the figure is the oscillation corresponding to node *c* firing. This oscillation can be taken as a reference. Nodes *a* and *b* fire at different phases comparing to node *c*. In INFERNET, this is a sign of discrimination.

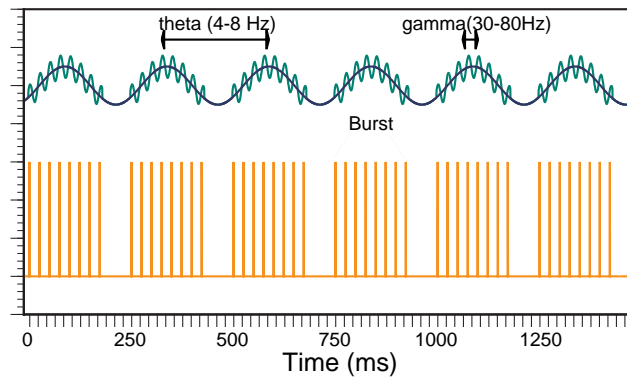
This rhythmic pattern corresponds to well known 40 Hz gamma waves which have been found to enable synchronization by experiments in slices. (see Singer, Engel, Kreiter, Munk, Neuenschwander, & Roelfsema, 1997). Eckhorn, Bauer, Jorden, Brosch, Kruse, Munk & Reitboeck (1988) recorded such oscillations with micro-electrodes in the visual cortex of the anesthetized cat when stimuli were presented. These gamma waves have also been observed to be associated with attention (Wang & Rinzel, 1995) and with associative memory (Wilson & Shepherd, 1995). The lower limit of gamma wave frequency is 30Hz and the upper limit varies according to various authors from 70Hz (Abeles, Prut, Bergman, Vaadia, & Aertsen, 1993) to 100 Hz (Wilson & Shepherd, 1995). The temporal gap between 2 spikes of a node is therefore from 10-14 to 33 ms.

Figure 2.9
The 3 nodes belongs
to different phases
with respect to the
oscillation



In INFERNET, this rhythmic pattern lasts a limited period of time. After while the node enters a resting period. After the resting period, the rhythmic pattern can restart. Such a rhythmic activity corresponds to bursts of a gamma wave. Such bursts are shown in Figure 2.10 (lower curve). The resulting oscillation pattern shows gamma subcycles in a theta cycle (upper curve). Theta waves have been observed to be associated with visual short term memory tasks in monkeys (Nakamura, Mikami & Kubota, 1992). This wave was maintained as long as attention was required. Theta wave duration can exceed 10 seconds. In INFERNET, this wave serves to maintain information.


Figure 2.10
INFERNET gamma
subcycles in a theta
oscillation



2.4 Symbols

Symbols are represented by a set of nodes that can be compared to Hebb (1949) cell assembly. As shown in Figure 2.11, a particular node can participate in different symbols or cell assemblies (different node colors in the figure). But a node cannot participate in every cell assembly. The representation is distributed, but not fully distributed (In fact, it is quite sparsely distributed). Nodes participating in a particular symbol do not need to be in the same topographical area, but can belong to different subnets.


Figure 2.11
Each color stands
for a different
assembly. A node
can participate in
different assemblies



The activation of a symbol is achieved by the firing of nodes which participate in this symbol. But even if every node in an assembly fires, it does not mean that this symbol is active. These nodes have to fire *in close synchrony* for the symbol to be active. Figure 2.12

illustrates three cases by displaying histograms representing the percentage of nodes, in an assembly, firing within a certain amount of time. In the first case, nodes belonging to a particular symbol fire in close synchrony. This symbol is highly activated. In the second case, nodes belonging to a symbol fire less synchronously. This symbol is less active than the first one. In final case, nodes belonging to the symbol do not fire in synchrony, the distribution of node firing times is too spread out in time. This symbol is not activated.

Figure 2.12
A symbol is
activated when
nodes participating
in the assembly fire
in close synchrony



A representation of symbols by synchronous firing of nodes belonging to an assembly may be considered very speculative. Suffice it to say that this hypothesis has received neurobiological empirical support. Engel, Kreiter, König & Singer (1991) showed that if several objects make-up a scene, distinct clusters of synchrony are formed, each associated with a particular object. Engel, König & Singer (1991) observed that individual cells can rapidly change partners of synchrony when the visual stimulus changes.


2.5 Binding, discrimination and predicates

We have seen that each symbol is represented by a cluster of nodes firing in synchrony. Symbols are then bound together by synchronous firing too. For example, to represent associated symbols like “red rose”, nodes belonging to “red” must fire synchronously with nodes belonging to “rose” (Figure 2.13). This is how INFERNET represents attribute binding.

This hypothesis was first proposed by von der Malsburg (1981). There is now neurobiological evidence for considering synchrony as a possible binding mechanism in the brain. In particular, synchrony has been observed between distant cells in the same cortical

area (Gray, König, Engel, & Singer, 1989; Engel, König, Gray, & Singer, 1990), between cells in different cortical areas (Eckhorn, et al., 1988; Engel, Kreiter, König & Singer, 1991; Murthy & Fetz, 1992; Nelson, Salin, Munk, Arzi, & Bullier, 1992), and even between cells in different hemispheres (Engel, König, Kreiter, & Singer, 1991). (See also section 1.5.2.)

Figure 2.13
Binding and
discrimination

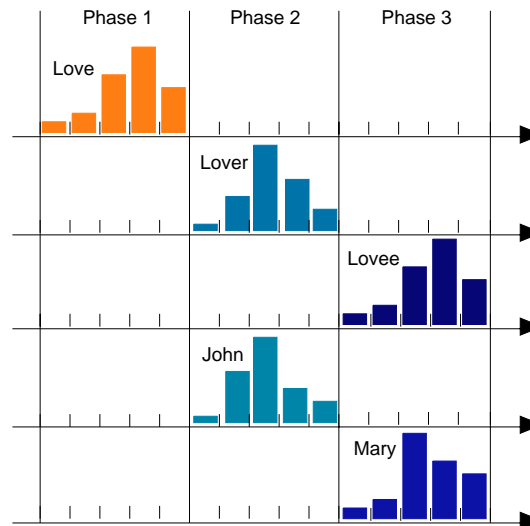


In INFERNET discrimination is achieved by successive synchronies, for example, to discriminate a red rose from a green lawn. The nodes belonging to “red” and “rose” must fire in synchrony and those corresponding to “green” and “lawn” must also fire in synchrony. Further, these two sets of nodes must fire asynchronously for “the red rose” to be discriminated from “the green lawn” (see Figure 2.13).

Neurobiologists have found that if a number of different objects make up a scene, distinct windows of synchrony are formed, each associated with a particular object (Engel, Kreiter, König & Singer, 1991).

In INFERNET, predicates and roles are linked by a specific temporal order. The activation of a predicate is always followed by the successive activation of its different roles, each of which is assigned to a particular window of synchrony. Figure 2.14, shows how INFERNET codes “John loves Mary”. The firing of the predicate “love” is followed by the firing of the arguments roles “lover” and “lovee”. Nodes representing “John” fire in synchrony with those of “lover” while nodes representing “Mary” fire in synchrony with “lovee” nodes.

Figure 2.14
Predicate and role
binding



A number of neurobiological parameters are involved in a representation that relies on clusters of nodes firing simultaneously. The first is the frequency of oscillation. As already stated, some specific oscillatory activities seem to facilitate synchronization (Singer, 1995, Roelfsema, Engel, König, & Singer, 1996, Singer, 1993). In INFERNET once a node is activated, it tends (but not necessarily) to begin oscillating at a gamma frequency range, whose lower limit is 30Hz and the upper limit varies according to various authors from 70Hz (Abeles et al., 1993), 80 Hz (MacKay, 1997) to 100 Hz (Wilson & Shepherd, 1995). The temporal gap between 2 spikes of a node is therefore from 10 to 33 ms. These gamma waves seem to be the best candidate for enabling synchronization and binding (Singer, 1993). The second key parameter is the precision of the synchrony at this frequency range. According to Singer and Gray (1995) this precision is between 4 to 6 ms., while for Abeles et al. (1993), it is about 5 ms, sometimes less, and depends on the frequency of oscillation. This allows us to approximate the number of windows of synchrony or phases that can be differentiated, i.e., $25/5 = 5$, based on a typical frequency of 40Hz. This constraint is very important because the more the system needs to discriminate objects at a particular moment, the more precise the synchrony should be. Since this parameter is bounded, it can lead to overload in which windows of synchrony are no longer distinguishable. Therefore, the number of distinct items and the number of predicate arguments are limited. The precision of the synchrony in INFERNET is regulated by the noise on delays (see section 2.1). This common effect of oscillation frequency and precision of the synchrony will lead INFERNET to make predictions about Short term memory that will be treated in chapter 3 and about the number of predicates arguments that will be treated in chapter 5. Similar explanations for the

brain's ability to store a limited number of short-term memory items can be found in Lisman & Idiart (1995), Jensen & Lisman (1996), Shastri & Ajjanagadde (1993), Sougné (1996), Sougné & French (1997), Jensen & Lisman (1998), Luck & Vogel (1998)

After receiving a premise like “John loves Mary” INFERNET should be able to answer a question like “Who is the lover?”. Therefore a mechanism should be able to link roles like “lover” and fillers like “John” for a limited amount of time. This is done in INFERNET by the modification of connection weights and delays. During the premise-encoding phase, connection weights and delays will be modified by a Hebbian learning rule to reproduce synchronies. In Figure 2.14, “John” is synchronized with “Lover”. The connections strength between the filler nodes representing “John” and role nodes representing “lover” will be increased. The delays will be adjusted to enable synchronous firing. After a number of connection modifications, the connections from “John” nodes will be sufficient for the “lover” nodes to fire in synchrony whenever “John” nodes fire. Similarly, connections from “lover” nodes will be sufficient for the “John” nodes to fire in synchrony whenever “lover” nodes fire.

The increase of connection strength between node i and node j is defined by equations (2.6). This equation expresses the fact that at every synchronous firing of nodes i and j the weights w_{ij} as well as w_{ji} will increase by a value Δw_{ij} which takes the value of a constant c in that case.

$$\Delta w_{ij}^+ = \begin{cases} c, & \text{if } (t_j^{(f)} - t_i^{(f)}) = 0 \\ 0, & \text{otherwise} \end{cases} \quad (2.6)$$

The connection delays are modified according to equation (2.7). This equation expresses the following mechanism. At every synchronous firing of nodes i and j the delays d_{ij} as well as d_{ji} will take the value $d_{ij}^{new} = \Delta t_\gamma$. Δt_γ is defined by the gamma wave frequency. If the oscillation has a frequency of 40Hz, Δt_γ is 25.

$$d_{ij}^{new} = \begin{cases} \Delta t_\gamma, & \text{if } (t_j^{(f)} - t_i^{(f)}) = 0 \\ 0, & \text{otherwise} \end{cases} \quad (2.7)$$

INFERNET forgets also bindings. It is necessary to make the system ready for encoding a new fact. For example, if in a first phase the premise “John loves Mary” is presented to INFERNET, and in a second phase a new premise is presented: “Mark loves Gill”. The system ought to have forgotten bindings between “lover” and “John” and between “lovee” and “Mary” to avoid spurious representations like “John loves Gill”. This process is done by decreasing the value of connection strength. The decrease of connection weights is defined by equation (2.8). There will be no decrease if both node i and j fire inside a theta wave in

question. If one or both do not fire inside the same period, a constant e will be subtracted from their connections weight w_{ij} and w_{ji} .

$$\Delta w_{ij}^- = \begin{cases} 0, & \text{if } \mathcal{T}_j^\theta \text{ and } \mathcal{T}_i^\theta \neq \emptyset \\ -e, & \text{otherwise} \end{cases} \quad (2.8)$$

\mathcal{T}_j^θ is the set of firing times of node j in a period defined by a theta wave, while \mathcal{T}_i^θ is the set of firing times of node i in the same period.

There is, however, a further detail to be noted. Modifications in connections strength and associated delays will never happen if the connection weights are 127, 0 or -127, i.e., values reserved for the encoding of Long Term Knowledge. All other connections can be modified but are bounded. These latter connection are modifiable. Modifiable excitatory connection cannot be greater than 126 or less than 15. Modifiable inhibitory connections cannot be greater than -15 and less than -126. Finally, since INFERNET network is not fully connected, some connections do not exist and cannot be affected by learning.

This fast and transient modifications of weights is comparable to neurobiological short-term potentiation by Hebbian learning.

2.6 Temporal Gates and coincidence detection

In the previous section, it was noted that some connection weight values are reserved for the encoding of Long Term Knowledge. This knowledge is encoded by hand. This enables the experimenter to encode rules and facts. Examples of Long Term Knowledge Base will be given in chapters 3, 4, 5, and 6. In this section, the basic types of rules will be presented. They are called temporal AND-gate, OR-gate, and XOR-gate. These gates permit the encoding of complex types of knowledge, like those presented in chapter 3 to 6.

The first type of gate, presented in Figure 2.15, is the AND-gate. This gate only opens when two events happen. In the figure, node C only fires if A nodes fire 10 ms before B nodes. This gate serves to detect the conjunction of events in a particular order. This is a temporal gate. If the B nodes fire alone, the C node will get two excitatory signals and one inhibitory signal. Even if all B nodes fire together, this will not be sufficient for the node C to fire. The same lack of firing of the C node will occur if A nodes fire solely. However, if B nodes fire 10 ms later than A nodes, the inhibitory connection from one of the B nodes will be inhibited (by the firing of A nodes). Excitatory connections from B nodes to C node will be amplified, node C will get two additional excitatory inputs from one of the A nodes which will be amplified by one of the B nodes, and the inhibitory connection from one of the A nodes will be inhibited. It is equivalent to the C node getting 8 simultaneous excitatory inputs. Under these circumstances, the C node will necessarily fire.

Figure 2.15
An example of AND-gate, node C will fire whenever nodes A fire 10 ms before nodes B




Figure 2.16 presents a second type of gate, the OR-gate. It opens whenever at least one of the two input events occurs. In the figure, node C will fire, if the A nodes fire, or if the B nodes fire or if both fire.

Figure 2.16
An example of OR-Gate, node C will fire, if A nodes fire, if B nodes fire or if A and B nodes fire





Figure 2.17 presents XOR-gate. This gate only opens if one event happens independently of another. In the figure, if the A nodes fire, node C will fire 10 ms later, unless the B nodes fire in synchrony with the A nodes. If the B nodes fire, node C will fire 10 ms later unless the A nodes fire in synchrony with B nodes.

Figure 2.17
An example of XOR-Gate, node C will not fire if the A nodes fire in synchrony with the B nodes.



Gates are able to detect the coincident occurrence of events. There is evidence that neurons can be sensitive to temporal relations among inputs. According to Abeles et al. (1993), people hearing a sound from a source situated 1 degree to the right of their face, at a distance of 1 meter, will receive the acoustic signal at their right ear 12 μ s before the left ear. This means that detecting the source of a sound requires very accurate temporal detection. According to Calvin (1983), throwing a projectile at a 20cm wide target, from a distance of 7 meters, requires an accuracy in the release time of less than 1 ms. We saw in section 2.5 empirical evidence to support precise synchronization of neuronal response. Moreover, Abeles et al. (1993), Abeles (1991) findings on syn-fire chains indicate very precise temporal behavior of neuron firings. These researchers have recorded spike timing of different cortical cells. In particular, they observed the following kind of pattern: when a neuron A fired, neuron B would fire 151ms later while neuron C would fire 289ms later with a precision across trials of 1 ms! Such long delays, would require dozens of transmission delays from presynaptic (A) to postsynaptic (C) neuron. The mechanism proposed by Abeles et al. (1993) for generating such precise delayed synchronization has been called *syn-fire chains*. Since cortical synapses are relatively weak, inputs to cells must arrive at the same time, and so make a very precise temporal filter. This mechanism is very close to a succession of temporal gates. More details about coincidence detection in the brain can be found in Konnerth, Tsien, Mikoshiba, & Altman (1996).

Each gate in INFERNET has a probability of malfunction. As the number of successive required gates to perform a task increases, the probability of successfully performing the task decreases. This hypotheses will be tested and discussed in chapter 4.

2.7 Multiple instantiation

Classically, in a connectionist network there is no separation, as there is in symbolic artificial intelligence models, between the Long Term Knowledge Base (LTKB) and a temporary store (or a working area) in which copies of pieces of LTKB are loaded before transformation. In these models, activation of the LTKB *creates* Short Term Memory (STM). For systems that do separate LTKB and STM (most traditional AI models), multiple instantiation is not a problem since the system can make as many copies of LTKB facts as needed and load them into STM. Without this copy-and-load process, neural nets suffer from “crosstalk.” (Feldman, 1982). Simultaneous processing of “John loves Mary” and “Gary loves Rita” can lead to pseudo-memories (Dyer, 1991) like “John loves Rita”. Even if we assume that John and Gary are correctly bound to the role of lover, and Mary and Rita to the role of lovee, both men and both women remain bound to the same respective roles. The system need to distinguish the two facts by separating the two identical predicates and their respective roles bindings.


The problem of multiple instantiation arises in localist networks if two instantiations differ by more than the value of one argument. For example, “Jack eats eggs and Jack eats fish” does not require separate instances of the predicate “eats” since this statement can be reduced to “Jack eats eggs and fish”. However, when two sets of two contents must be bound to identical pairs of roles, the system must be able to handle two copies of the predicate and argument slots. For example, “Jack eats eggs and Mary eats fish” cannot be reduced to “Jack and Mary eat eggs and fish,” otherwise one cannot distinguish who eats what.

The problem for distributed representations is even more difficult. Multiple instantiation problems appear as soon as one node must be shared by entities that have to be differentiated. If an n -ary predicate must be represented where either predicate roles or their fillers need to use a common node, this node will have to be linked to different entities. In systems of distributed representations, the problem will be a function of the proportion of shared nodes.

Multiple instantiation will be discussed extensively in chapter 6. Here we will present the basic mechanism of multiple instantiation in INFERNET.

In INFERNET, nodes pertaining to a doubly instantiated object will be associated with two oscillations while those objects that are singly instantiated will be associated with one oscillation. This makes doubly instantiated object nodes fire twice while singly instantiated object nodes fire once. This means that each new instance will occupy a new phase to avoid crosstalk. Figure 2.18 illustrates how INFERNET deals with multiple instantiations. In this figure, a particular node activity is plotted (lower curve) with reference to its oscillation pattern (upper curve). Initially, we can imagine that this node is bound to a particular node firing in synchrony and having the same oscillation pattern. After three spikes, another pattern of firing is added, following an additional oscillation pattern which is represented in the upper curve. The additional firing of this node enables it to also bind with a second, different node.

Figure 2.18
Multiple instantiation
by period doubling



This solution can be compared to the neural phenomenon of bifurcation by period doubling (Canavier, Clark, & Byrne, 1990; Glass, & Mackay, 1988), whereby a stable

oscillatory state can lose its stability, giving rise to a new stable state with doubled period. Figure 2.18 shows the point at which the bifurcation occurs. The number of instantiations that INFERNET can support is, limited to 2 or 3 by the precision of the synchrony and refractory period. This puts processing constraints on multiple instantiation which will be tested in chapter 6. INFERNET's ability to process a greater number of instantiations will be explained in section 6.3

2.8 Treatment of Simulator Data

When the simulator runs, it records in a file for each ms in a gamma cycle, for each gamma cycle in a theta cycle and for each theta cycle, which nodes are firing. This file contains all the data concerning node firing times.

These data can be considered as a time series in which successive values represent successive measurements taken at equal lag intervals. Since INFERNET represents binding by synchrony, and relations between predicate and arguments by delayed synchrony, the measure that we need to use must be sensitive to synchrony and to lags among nodes firing. In neuroscience, such measurements are quite common. One of the most popular techniques is cross-correlation. This technique has been in use since the early 1970's (Moore, Segundo, Perkel & Levitan, 1970) and is still used (Singer, Engel, Kreiter, Munk, Neuenschwander & Roelfsema, 1997). This measure evaluates the correlation of 2 time series at successive lags. Firing times of two neurons are recorded and constitute 2 time series. If data are correlated at lag 0, it means that the two neurons are prone to fire synchronously. If data are correlated around a lag of 10 ms, it means that the firing of the first neuron tends to precede the firing of the second neuron by 10 ms.

There is however a problem with this technique. Correlation is known to be sometimes a spurious measure. Imagine that we have collected two node firing times and obtained these 2 time series.

[0 0 0 0 0 0 0 0 20 80 20 0 0 0 0 0] for the i node, and

[0 0 0 0 0 0 0 0 2 8 2 0 0 0 1 0] for the j node.

Their correlation $r = .992536$

This almost perfect correlation does not take into account the fact that the j node almost never fires, it only accounts for the moment of firing.

For this reason, some researchers doubt the relation between correlation and synchrony, (see Brody, 1998, 1999a, 1999b) for a recent review of this topic). To eliminate the problem, some researchers have proposed considering the relative number of spikes (see Golomb, Wang & Rinzel, 1994).

Data presented in previous studies (Sougné, 1998b; Sougné, 1999) were based on some arbitrary threshold. Two nodes were considered synchronous if they fired more than 4 times in a theta cycle, within an interval of less than 3 ms. The measure used in the present study is a correlation corrected by the ratio of the two nodes' spike density: the observed density divided by the ideal density or the density of input nodes (this ensures that the denominator is \geq than the numerator). This correction provides a normalized correlation. For the example provided above, the normalized correlation is: $\sqrt{.992536^2 \frac{13}{120}} = .326684$, which is a far more realistic.

In the following sections are described how proportions of response and reaction times are computed. Two files are saved by INFERNET: the first contains data when learning occurs, the second contains the data of INFERNET's response to a question. The first file will be used to calculate reaction times, the second to calculate proportions of responses.

2.8.1 Proportion of response

To obtain the proportion of a particular response, the ideal pattern of response is defined or the input pattern is observed this is the reference pattern. This pattern is compared to the observed pattern of response. This observed pattern is obtained by counting the precise moment of a particular node firing within each gamma cycle. A sequence of length corresponding to the number of ms in a gamma cycle is obtained. The correlation between the reference pattern and the observed one is computed. This correlation is normalized by multiplying it by the ratio of the number of observed spikes and the number of spikes in the reference pattern. This process is repeated for each node that participates to a particular response. The combination of obtained correlations is performed by computing the square root of the mean of squared correlations (with taking care of signs). This results is taken as the proportion of response. The details are following.

Define Δt_γ as the delay in ms corresponding to a gamma wave. If the oscillation has a frequency of 40Hz, Δt_γ is 25. For each node, and for each gamma cycle, we can construct a vector of binary values of length Δt_γ , the firing vector of node i is $\vec{F}i_\gamma$:

$$\vec{F}i_\gamma = \left[f_i^{(1)}, f_i^{(2)}, \dots, f_i^{(\Delta t_\gamma)} \right] \quad (2.9)$$

$$\text{where } f_i^{(t)} = \begin{cases} 1, & \text{if node } i \text{ fires at time } t \text{ (in ms)} \\ 0, & \text{otherwise} \end{cases}$$

This vector contains the information about a node firing times in a gamma cycle from the first ms to the Δt_γ ms.

The sum \vec{X}_i of these vectors across every gamma cycle and theta cycle for the node i is:

$$\vec{X}_i = \sum_{p=1}^{n_\theta} \sum_{q=1}^{n_\gamma} \vec{F}i_{\gamma} \quad (2.10)$$

Where n_θ is the number of theta cycles, and n_γ is the number of gamma cycle in a theta cycle.

The sum vector of i node contains Δt_γ integers z_t :

$$\vec{X}_i = [z_1, z_2, \dots, z_{\Delta t_\gamma}] \quad (2.11)$$

The sum vector of j node contains Δt_γ integers y_t :

$$\vec{X}_j = [y_1, y_2, \dots, y_{\Delta t_\gamma}] \quad (2.12)$$

The correlation between the two sequences \vec{X}_i and \vec{X}_j is

$$r = \frac{\Delta t_\gamma \sum_{t=1}^{\Delta t_\gamma} z_t y_t - \sum_{t=1}^{\Delta t_\gamma} z_t \sum_{t=1}^{\Delta t_\gamma} y_t}{\sqrt{\left[\Delta t_\gamma \sum_{t=1}^{\Delta t_\gamma} z_t^2 - \left(\sum_{t=1}^{\Delta t_\gamma} z_t \right)^2 \right] \left[\Delta t_\gamma \sum_{t=1}^{\Delta t_\gamma} y_t^2 - \left(\sum_{t=1}^{\Delta t_\gamma} y_t \right)^2 \right]}} \quad (2.13)$$

z_i = i th element of the sequence \vec{X}_i

y_i = i th element of the sequence \vec{X}_j

Δt_γ = The delay in ms defined by a gamma wave

S is the set of binding pairs that are required to consider that a response belongs to a category.

$$S = \langle (A_1, B_1), (A_1, B_1), \dots, (A_c, B_c) \rangle \quad (2.14)$$

The normalized signed square correlation (or signed coefficient of determination) between a set of paired cell assemblies S is:

$$\rho_S = \prod_{AB \in S} \left[\frac{\sum_{\vec{X}_i \in A} \sum_{\vec{X}_j \in B} r_{\vec{X}_i, \vec{X}_j}^2 \mathcal{H}(r_{\vec{X}_i, \vec{X}_j}) \frac{\sum y_t}{\sum z_t}}{nA nB} \right] \quad (2.15)$$

$$\text{where } \mathcal{H}(r_{\vec{X}_i, \vec{X}_j}) = \begin{cases} 1, & \text{if } r_{\vec{X}_i, \vec{X}_j} \geq 0 \\ -1, & \text{if } r_{\vec{X}_i, \vec{X}_j} < 0 \end{cases}$$

nA is the number of node in A cell assembly

nB is the number of node in B cell assembly

The normalization term: $\frac{\sum y_t}{\sum z_t}$ is the ratio of the two sets of nodes spike density. The numerator correspond to the observed density and the denominator to the reference density.

Note the use of the mean of the signed square normalized correlation. The reason is that correlations are not additive. The correlation formula contains a square root. Squaring correlation removes it and permits addition and mean calculation.

The proportion of response containing S pairs of associations is:

$$\frac{\mathcal{H}(\rho_s) \sqrt{\mathcal{H}(\rho_s) \sum_n \rho_s}}{n} \quad (2.16)$$

where n is the number of simulator runs and $\mathcal{H}(\rho_s)$ is a stepwise function which take care of the sign:

$$\mathcal{H}(\rho_s) = \begin{cases} 1, & \text{if } \rho_s \geq 0 \\ -1, & \text{if } \rho_s < 0 \end{cases}$$

2.8.2 Reaction times

The principle for calculating reaction times is the same as for proportions of responses. The difference is that correlations are continuously monitored after each theta cycle. When the correlation reaches a threshold, a reaction time is returned. The second difference is that the data that are taken into account are those when INFERNET learns bindings. Recall that when INFERNET learns bindings, after each theta cycle containing the premises, a question is presented with all learning turned off. This process enables monitoring of learning.

The first step is to compute ρ_s for each theta cycle, and then add them successively. When the sum reaches a threshold (1.5), the number of theta cycles n_θ explored so far is returned. The ρ_s is then computed for each gamma cycle inside this theta cycle. They are added successively and when this addition reaches a threshold (.25) the number of gamma cycles explored so far n_γ is returned. The reaction time is then:

$$RT = n_\theta \Delta t_\theta + n_\gamma \Delta t_\gamma \quad (2.17)$$

where Δt_θ is the delay in ms defined by a theta wave, and Δt_γ is the delay in ms defined by a gamma wave.

If the above thresholds are not reached, it means that the simulator could not find any response to the question, a constant is then returned.

3 **Short term memory as the activated part of long term memory**

Cognitive psychology distinguishes sensory memory, short term memory and long term memory. Sensory memory is the very short term maintenance of the perceptual signal. Short term memory (STM) refers to the memory that is maintained during the human psychological present, and long term memory (LTM) is the storage of the past experiences. The concept of short term memory has generally been replaced by the concept of working memory (WM), (Baddeley, 1986). The idea is that the transient and present memory is not a passive store; rather, it permits and enables information processing. When referring to working memory we will adopt the definition of Cowan (1998): *“the collection of mental processes that permit information to be held temporarily in an accessible state, in the service of some mental task.”* Working memory does involve the storage of information but also all mechanisms that help this information to be used or to be maintained, in particular, the process of attention. When the storage itself will be referred to, the term “short term memory” will be used.

The traditional computer metaphor of working memory is still alive (Luck & Vogel, 1998). Long term memory is compared to data stored in Random Access Memory (RAM). Portions of RAM can be loaded (or copied) into the memory registers of the processor to constitute a Short-Term Memory. The capacity of these registers is relatively small like the capacity of Short Term Memory. The capacity of the registers limits the power of the computer just as the capacity of Short term memory limits human processing capabilities. The registers also contain the code that will process the information or the data stored in them. In other words, these registers have the ability of processing information like Working memory. But this metaphor is only based on surface similarity. Human memory is more limited, is subject to forgetting, to interference, is addressable by its content and not by an “address” (Kanerva, 1990), etc. Even if this metaphor is not taken completely seriously,

some of the basic ideas still form the dominant paradigm. The separation of LTM and STM into two systems, the idea that information can be loaded from LTM to fill WM, and that LTM can be constituted by loading the WM content into LTM are the basic ideas of the multiple store model. This section will show how INFERNET can exhibit characteristics of WM and LTM without separating them. The first part of this section will first be devoted to presenting some of the main findings about memory.

3.1 The capacity of short term memory

In a series of experiments, Miller (1956) found that people can recall at most seven independent sets at a time. These sets are called chunks. They can involve an unlimited amount of information. The capacity of STM is generally measured by a span task. The experimenter presents a collection of items to the participant then removes the stimulus and asks the participant to recall all items present in the original stimulus in order. When the participant succeeds, a longer collection of items is presented and so on until the participant fails. The number of items contained in the last successful trial of the participant is the span of that participant. Note that it is assumed that if a participant fails at 7 items she/he will fail at 8, 9, 10.

The span task has been used for different stimuli: digits, words, non-words, visual features etc. One of the important findings is that this limit depends on the stimuli. For example: the span is higher for digits than words, it is higher for short words than long words, it is higher for quickly pronounced words than for slowly pronounced words...

It has been argued that the true capacity of STM was actually lower than seven. Broadbent (1975) argued that the capacity is about 3 or 4 chunks. Fisher (1984) found a maximum of 4 channels in visual search. Yantis (1992) found that people can track 4 moving visual items at a time. Luck & Vogel (1998) reported that people can hold an average of four visual features, chunks or combination of features.

The problem with span is that it is often considered as a clear cut threshold, and it seems that researchers hope to find a “magical number” that limits the capacity of STM. It would be better to describe this phenomenon as a decreasing function describing the lower frequency of success as the number of items increases. But this is rarely how researchers present results. STM span studies seem to favor the belief in a threshold function, whereas data seems to favor a sigmoid function.

3.2 Forgetting

Short term memory content is evanescent, after something like 20 seconds the initial content of STM has disappeared. Peterson & Peterson (1959) used the Brown-Peterson technique to

study forgetting in STM. Participants received a trigram to be memorized, then they had to count backward from a particular number by steps of 3 (e.g. 101, 98, 95, 92, ...). After a certain time interval, not known by participants, they had to report the trigram. Results show an exponential decrease in the recall frequency. After 18 seconds the proportion of correct reports is near 0. These authors interpreted this forgetting in terms of trace decay. Other researchers found that by manipulating variables affecting proactive inhibition, they could explain forgetting in terms of interference. Proactive inhibition is the reduced recall of newly presented material due to what has been learned before. Keppel & Underwood (1962) showed that proactive inhibition was affected by the number of previously learned items. Wickens, Born, & Allen (1963) showed that the similarity between previously learned items and newly learned items affected proactive inhibition. They argue for interaction between a decay and interference. Turvey, Brick and Osborn (1970) showed that the influence of previously learned items on newly learned items was affected by the time delay between their presentation.

3.3 Arguments for STM LTM distinction

3.3.1 Serial effect

There seems to be two storage systems: one for LTM and one for STM. Some early studies (Deese & Kaufman, 1957; Murdock, 1962) found that in free recall the first items and the last items were better remembered. In free recall, people memorize a series of 10 or 15 or even 40 items (more than the span). The participants' task is to recall as many items as possible, regardless of order. Murdock (1962) showed that for lists from 10 to 40 items people were better at remembering the first items in the list and even better at recalling the last items of the list. The fact that the last items are recalled better is called the recency effect, while the fact that the first items are recalled better is called the primacy effect. Postman & Phillips (1965) showed that if one increases the time between the presentation of items and the beginning of the recall to approximately 15 seconds, the recency effect disappears while the primacy effect remains. The multiple store model explanation is that the first items are stored in LTM, while the last items of the list are stored in STM. Deferring the recall has an effect on STM, but no effect on LTM. The conclusion is that STM must be separated from LTM, since the recency effect is a STM effect. However, Baddeley & Hitch (1977) asked rugby players to recall all the teams with whom they played since the beginning of the season. They found a recency effect in this purely LTM task.

3.3.2 Coding systems

The phonological similarity between letters or between words produces replacement errors or impaired immediate serial recall in STM tasks (Baddeley, 1966a). Semantic similarity, on the other hand, does not produce as many errors. Visual similarity among items has also been found to impair STM (Longoni & Scalisi, 1994; Walker, Hitch, & Duroe, 1993). In LTM, Baddeley (1966b) showed that semantic similarity does induce more errors, while phonological similarity does not. The multiple store model proponents conclude that STM and LTM do not use similar coding systems and for this reason must be separated.

3.3.3 Double dissociation

Some neuropsychological data indicate a double dissociation between STM and LTM. Baddeley & Warrington (1970) describe a patient with normal STM and defective LTM while Shallice & Warrington (1970) describe a patient with normal LTM and impaired STM. Again, the multiple store model proponents conclude that this supports the conclusion that STM and LTM should be separated.

3.4 INFERNET as a model of STM

In this section, INFERNET abilities to model STM as the activated part of LTM will be described and illustrated by empirical support from simulations.

3.4.1 INFERNET STM capacity

As indicated in chapter 2, a number of neurobiological parameters are involved in representations that rely on clusters of nodes firing simultaneously. The first is the frequency of oscillation. In INFERNET, once a node is activated, it tends (but not necessarily) to fire rhythmically between 30 and 100 Hz. The temporal gap between 2 spikes of a node is therefore from 10 to 33 ms. This corresponds to the observed 30-100 Hz (γ wave) oscillations of certain types of neurons. These γ waves have been observed to be associated with attention (Wang & Rinzel, 1995) and with associative memory (Wilson & Shepherd, 1995). The second key parameter is the precision of the synchrony. According to Singer & Gray (1995) this precision is between 4 to 6 ms. For Abeles, Prut, Bergman, Vaadia, & Aertsen (1993), the precision is about 5 ms and depends on the frequency of oscillation.

This allows us to approximate the number of windows of synchrony that could be differentiated, i.e. $25/5 = 5$. If we assume that a window of synchrony corresponds to an item or a chunk in STM, then this puts STM span at approximately 5, with a small amount of variance since precision is proportional to oscillation frequency. This corresponds to current estimates of human STM span. It has been suggested (Cowan, 1998), that the traditionally

accepted size of STM (i.e., 7 ± 2 items) may be too high. An item can be a word, an idea, an object in a scene or a chunk, i.e., a grouping of items. Similar explanations for the brain's ability to store approximately 7 short-term memory items can be found in Lisman & Idiart (1995), Jensen & Lisman (1996), Shastri & Ajjanagadde (1993), Sougné (1996), Sougné & French (1997), Jensen & Lisman (1998), Luck & Vogel (1998)

How can representations be maintained in STM? The problem with γ waves is that they persist only a few hundred milliseconds. This is not long enough to reflect the time taken by people to draw inferences, nor does it correspond to standard estimates of STM retention time (10 to 20 seconds). For this reason, following Lisman & Idiart (1995), γ waves in INFERNET occur in bursts which restart every 146 to 333 ms. This corresponds to θ waves [3 - 7 Hz] whose duration can exceed 10 seconds. The resulting temporal firing pattern for a single node is a set of firing at 40 Hz for the seven spikes that constitute a burst. This is followed by a resting period of, say, 75 ms. Thereafter, the burst begins again. The burst interval is, therefore, about 250 ms (4 Hz). There is neurobiological evidence for this rhythm in STM. θ waves have been observed to be associated with visual short term memory task on a monkey (Nakamura, Mikami, & Kubota, 1992). This wave was maintained as long as attention was required.


STM capacity is limited, and chunking increases the amount of information it can contain. In INFERNET, chunking is achieved by two processes:

1. Increasing the number of nodes — and, as a result, increasing the number of objects or symbols firing in synchrony. This is achieved by means of spreading activation.
2. Replacing the content of two or more windows of synchrony by a single one that sums them up. This is achieved by the use of excitatory and inhibitory connections.

Examples of chunking will be provided in sections 5.3 and 6.3.

We will run simulations on INFERNET to explore STM capacity. In this experiment, 84 nodes fire rhythmically with a delay $\Delta t_\gamma = 30\text{ms}$. They are distributed over the 30ms to cover the entire interval. This represents the context or task nodes. In the binding-learning phase, as input, each object node will only fire twice per theta cycle and will respect a particular sequence (see Figure 3.1). “Object 1” nodes will fire twice, then “Object 2” will fire twice etc. The phase assigned to each object is randomly chosen. In the binding-learning phase, object nodes will fire at the same time as particular context nodes, thereby binding these object nodes to these context nodes. The connections linking context and object nodes will be modified. In the question phase, context nodes will fire and the simulator memory will be evaluated by looking at the firing of object nodes.

Figure 3.1
An example of input
for the binding-
learning phase for
12 objects. The
particular phase is
randomly assigned
for each object.




In this experiment, there were 7 nodes per object, and objects inhibits each other (Inhibiting connection with a delay of Δt_γ). There is a probability for noise on the delay, and therefore the nodes will not fire precisely at 30 ms intervals. The more windows of synchrony that are required, the more competition there will be between nodes, since there is an increased probability that 2 nodes pertaining to different objects will fire in synchrony. In that case, we would expect that the proportion of recall will decrease as a function of the number of items to memorize. This is, in fact what happens.

Figure 3.2 displays the proportion of correct recall for different list lengths. Twenty trials for 4, 5, 6, 7, 8, 9, 10, and 12 items were tested with the INFERNET simulator. Normalized correlation (as described in section 2.8) was computed for each object. This is a correlation between data observed and data that we would obtain for perfect recall. To obtain the proportion of recall for a particular set length, normalized correlations for each of the items were multiplied together. The reason is that we are looking for the cases where all items are recalled in the correct order. In span task, if one item is missing, the response is considered incorrect.

In order to compare these results with human data, we plotted the data of different studies for which the authors have collected the frequencies of correct recall for different list lengths. Luck & Vogel's (1998) study on visual short term memory analyzes visual span by successively showing two stimulus arrays. Subjects are asked to judge if the two stimuli are equivalent or differ by a a single feature or a combination of features. The set size is

determined by the number of objects present in the array. The data published by Baddeley, Thomson & Buchanan (1975) on short words are somewhat biased. When a participant failed on the 8 trials of a list of items, the test was stopped, which means that proportion of correct recall was underestimated for higher numbers of item. Moreover, this study involved only 8 subjects. Baddeley & Warrington's (1970) results should also be taken with caution, since they were based on only 6 participants. Unfortunately, we did not find an extensive study which presented frequency of recall for different list lengths and verbal material.



Data reported in Figure 3.2 show a decrease of correct recall as the number of item increases. Moreover, the decreasing function is sigmoidal in shape. INFERNET seems to be somewhat better than humans, but the important point is that STM capacity is limited and that the decreasing functions are parallel.

3.4.2 Forgetting

Forgetting in INFERNET is caused by both trace decay and interference. As described in chapter 2, connections that were strengthened during a binding-learning phase, will decay slowly over time. This is compatible with a trace decay. When there is a high number of objects to be distinguished, competition for windows of synchrony increases and, as we showed in the previous section, the proportion of correct recall decreases. This is compatible with what would be produced by interference. The data presented in the next section are a result of these two processes, i.e. decay and interference.


3.4.3 Serial effects

What happens when INFERNET has to memorize a number of items beyond its span? The following experiment examined which items were recalled. Forty trials were performed by

INFERNET faced with the task of memorizing 12 items. The procedure was the same as the preceding experiment. For each item, the proportion of recall was computed as described in section 2.8. The results are shown in Figure 3.3. We also report data that was inferred from Murdock (1962). In this study, there was no data for list length of 12 items. Data were obtained by comparing 10-item results with 15-item results.

As we see, the first item in INFERNET is better recalled than successive ones. The five last items are also better recalled. INFERNET displays a primacy and a recency effect without separating LTM from STM. Why this happens? In short, the primacy effect is caused by interference, and the recency effect by trace decay. In Figure 3.1, we plotted the input that is provided to INFERNET in the binding-learning phase with 12 items. Each object node is externally excited twice, but these nodes will continue to oscillate independently at the gamma frequency. The only object nodes that are not in competition with others are the nodes associated with the object presented first. The others will always appear when other object nodes are already oscillating. This is why the object presented first is recalled better than subsequent ones. Binding-learning happens when independent excitation is synchronous with external excitation. For the objects presented first, at the end of a gamma cycle, the time elapsed from this moment will be longer than for the items presented last. If the nodes associated with an object continue to oscillate, nothing will happen, but if a node stops oscillating, the connection strength from the bound context nodes will decrease. The more time elapsed from the external excitation, the more the chance of having a node that stops firing. That is why the items presented first have a greater chance that their connections from context nodes will decay.

Figure 3.3
INFERNET serial
position effect




The data reported in Figure 3.3 shows that INFERNET performs better than humans, and that the INFERNET trace decay takes more time. Nevertheless, both primacy and recency effects occur with a unique system for STM and LTM. It should be possible to obtain a better match with different settings of INFERNET parameters.

3.4.4 Similarity

The following simulation will show how INFERNET reacts when the objects to memorize are similar. The similarity between objects was simulated by assigning to the different objects a common node. Five objects were created. Each contained 7 nodes. In the dissimilar group, the 5 objects were represented by 7 different nodes; for the similar group, the five objects had 6 of their own nodes and one common node (i.e. a degree of similarity of 1/7). The frequency of correct recall was computed as before. Data were collected on 40 trials for each condition. The results are displayed in Figure 3.4, and are compared with Baddeley (1966a), Experiment 1.

Figure 3.4
Effect of item
similarity on STM
performance



The INFERNET results show an effect of similarity between items. The slope of the decreasing function falls between the effect of acoustic similarity and semantic similarity, according to Baddeley (1966a) study. The important point is the existence of a similarity effect. The number of nodes shared by objects could have been increased in order to obtain an effect similar to that observed for “acoustic similarity”.


Since the current version of INFERNET does not separate phonemic from semantic representations, no data are available to contrast with this difference obtained in Baddeley (1966a).

3.4.5 Double dissociation between LTM and STM

In this simulation, the ability of INFERNET to display double dissociations between STM and LTM was explored. Three settings were tested with 20 trials per settings. The STM was tested as before, noting the proportion of correct recall for a list of 5 objects. A LTM was included in the system: each of the 5 objects were linked with a particular object. So, the experiment involved 10 objects. This simulation can be viewed as if the network were

maintaining a list of exemplar words like *robin*, *siamese*, *griffin*, *trout*, *oak* in STM and the LTM task involved correctly associating each word with a particular category: *robin-bird*, *siamese-cat*, *griffin-dog*, *trout-fish*, *oak-tree*. Specifically, the LTM score is the total category node firings divided by the total exemplar node firings. If the LTM is good, whenever an exemplar is activated, the corresponding category will be activated.

Figure 3.5
INFERNET
simulator results on
double dissociation
between LTM and
STM (20 trials)



For the first group (“full”) the noise on delay probability was set to 20% and the links between our five exemplar-category pairs were maintained. For the impaired LTM group, the noise on delay probability was maintained at 20% but 5/6 of the links between exemplars and category were destroyed. For the impaired STM group the noise on delay probability was raised to 50%, and the content of LTM was the same as in the “full” group.

INFERNET results are displayed in Figure 3.5. These data clearly show a double dissociation between STM and LTM.

Other variables could have been used to obtain this double dissociation. Noise decreasing connection strength should affect more LTM than STM, binding-learning rate should afflict more STM than LTM, etc. This experiment illustrates the fact that for obtaining double dissociations only two variables suffice.

3.5 Discussion

A large number of data have been accumulated in this half of century on memory. A number of phenomena have been discovered and assessed. INFERNET has not the pretension to explain all of these phenomena. Our goal was simply to review a number of classic phenomena and to test INFERNET to see if it can account for them. The overall picture that INFERNET draws is quite close from reality. This success is more especially valuable as INFERNET (an unique store model) simulates data that are considered by many authors to be the main arguments for the multi-store model. Supporters of the multi-store model could

always argue that even if STM is the activated part of LTM, STM is functionally distinct from LTM in INFERNET. But since they both share the same substrate, they are not functionally independent in INFERNET. This is important in explaining data involving the relationship between LTM and STM. For example: Hulme, Maughan & Brown (1991) showed that the STM span is higher for words than for non-words (which have no semantic representation in LTM). Memory span is worse for foreign language words than for mother-tongue words (Hulme, Maughan & Brown, 1991). Experts have a better STM than novices (Chase & Simon, 1973). High frequency words are better remembered than low frequency words (Watkins & Watkins, 1977). To explain these data, Logie (1996) has proposed that information in working memory also activates LTM traces. These kind of dual STM-LTM systems seem ad-hoc, uneconomical, and probably unnecessary. Another advantage of a one store model over multiple store models is related to the “Occam’s razor” principle. This principle was enunciated by the philosopher William of Occam in the fourteenth century. The less a theory postulates elements in order to explain a phenomenon, the better it is. In other words, the more parsimonious an hypothesis is, the more likely to be true.

The reader may have remarked that the term “working memory” has almost never been used. The reason is simple: working memory involves a set of process control mechanisms like attention and even consciousness. INFERNET was not designed to simulate these kinds of phenomena. Even if certain colleagues have built simulators which also use oscillation and synchrony to model attention (Baird, 1996; 1997), and even if synchrony has been hypothesized to be the neural correlate of consciousness (Crick & Koch, 1990), we are still very far from making a simulator which displays the process control of working memory.

The contribution of INFERNET over “box” models like those of Baddeley (1986) or Atkinson & Shiffrin (1968) is the level of processing details. These “box” models provide a gross description of phenomena. Consequently, one would have great difficulty building a computational model based on these descriptions without building in a large number of special purpose mechanisms. For example, in Baddeley’s model, the phonological loop is responsible for maintaining verbal information. What are the mechanisms of looping, why does it start, why does it stop? Why is the phonological store impaired by phonological similarity? Many magical processes have been attributed to the central executive ensuring the model fits empirical data, but nothing is said about how this central executive actually works. Some researchers even doubt its very existence (see Parkin, 1998).

INFERNET is a theoretical and computational model that gives detailed and falsifiable predictions, and provides detailed mechanisms of cognitive phenomena grounded in neurobiology. Its limits in explaining some data also provide perspectives for future work.

What are the advantages and disadvantages of INFERNET over other connectionist models of memory which attempt to model the relationship between STM and LTM? Levy & Bairaktaris (1995) distinguish three types of connectionist solutions: the dual component

architecture, the dual weight architecture and the dual trace architecture. In the dual component architecture, there are two networks one for the STM one for LTM. An example is Barnden's COMPOSIT model (Barnden, 1991, 1992, 1994; Barnden & Srinivas, 1991). In this model, STM is composed of several registers filled with activation patterns from LTM. The problem is finding a neurally plausible justification for the rapid transfer of information from LTM to STM. In the dual weight architecture, there are two types of weights, fast weights are modified for encoding STM and slow weights for LTM. Examples of dual weight connectionist models are Hinton & Plaut (1987), Cleeremans & McClelland (1991), Levy & Bairaktaris, (1995). There is a neural justification for this dichotomy — namely, some synapses change slowly, others quickly, some are long lasting changes, others transient (see Markram, Wang & Tsodyks, 1998; Markram, Pikus, Gupta & Tsodyks, 1998). For the dual trace architecture, LTM is stored in connection weights, while STM is the activation of the network. These models derive from the ideas of Hebb (1949). This solution has been applied to model different aspects of STM by Shastri & Ajjanagadde (1993) and Jensen & Lisman (1998).

None of these solutions alone are sufficient for modeling STM and LTM: the problems of binding and multiple instantiation must be solved, the capacity of STM must be constrained, etc.

INFERNET falls into the category of dual trace architectures, since STM is the activated part of LTM, but it also has some characteristics of dual weight architecture. Binding in STM is performed by fast weight changes while LTM weight cannot be modified, they have fixed values.

INFERNET does not separate objects, symbols, and phonemes and, as a result cannot model effects related to different codings involving word-length effects, semantic vs. phonemic similarity, and speech rate effects. However, such distinctions could likely be done as extensions of the present INFERNET architecture. Some connectionist work has already been done on this subject. Hartley & Houghton (1996) built a connectionist model able to simulate errors human make with unfamiliar phoneme sequences. Grossberg, Boardman & Cohen (1997) simulate speech rate effects. However, these models in general, simulate fewer STM characteristics than INFERNET.

Sougné & French (1997) provide a theoretical explanation of Sternberg (1966) memory scan effect within the framework of INFERNET model. The memory matching capabilities necessary to do this task have not yet been implemented in INFERNET. Jensen & Lisman (1998) provide an elegant explanation and simulation data on this problem with a model very similar to INFERNET. The difference lies in the fact that the different items are encoded in different phases inside a theta cycle instead of gamma cycle in INFERNET.

The final ability involving LTM that INFERNET does not have is the ability to construct new memories and to integrate them in LTM. There are two difficulties with this process. The first is developing a learning algorithm capable of learning long chains of activation. This point will be discussed in chapter 8. The second problem is avoiding catastrophic interference. Recent studies French & Mareschal (1998), French, (1997a, 1997b, 1999) and MacClelland, McNaughton, & O'Reilly (1995), Ans & Rousset (1997), provide insights on that problem with a good neurobiological justification.

Overview of empirical studies

The preceding chapter showed that INFERNET's short term memory (STM) is constrained in a similar manner as human STM. The following three chapters will describe additional INFERNET constraints, will explain how the INFERNET simulator behaves and will describe human empirical studies that test INFERNET's predictions.

Chapter 4 presents constraints on the length of propagating chain. Bindings in INFERNET are constrained by the number of intermediate steps required for particular role nodes to enter into synchrony with the filler nodes. This constraint is shown to provide a plausible explanation of negation effects in human reasoning. *Experiment 1* in chapter 4 presents human reactions to negated conditionals. INFERNET predicts a lower rate of inferences for cases involving double negations which require a longer chain of propagation.

Experiment 4 in chapter 6 provides further empirical evidence for the effect of the number of intermediary steps. Two situations were compared, one involving a symmetry and the other involving an asymmetry. Since dealing with asymmetry requires a longer propagation chain in INFERNET, it was predicted that dealing with asymmetry would be more difficult than dealing with symmetry. *Experiment 4* tests this hypothesis on humans.

Chapter 5 explores constraints on predicate representation. In INFERNET the more arguments a predicate has, the more windows of synchrony will be needed to represent it. *Experiment 2* tests this hypothesis on humans. *Experiment 3* explores how humans can reduce the number of windows of synchrony when arguments are too numerous.

Chapter 6 presents the multiple instantiation constraint. This problem arises in connectionist networks as soon as a symbol has to be simultaneously used twice in different ways. Since INFERNET's short term memory is the transient activation of parts of long term memory, it cannot make multiple copies of a symbol, in the same way, for example, that a symbolic system can. The INFERNET solution to the multiple instantiation problem involves superposition of different node oscillations. This process is constrained by the refractory period of the nodes. *Experiment 4* explores the feasibility of the INFERNET

solution. *Experiment 5* provides evidence for a relation between human reaction time and the number of instances of a same predicate. Finally, in INFERNET, a symbol is represented by a set of nodes firing in synchrony. The distributed nature of each symbol implies that closely related symbols have nodes in common. If two related symbols are needed simultaneously, and if they cannot belong to the same window of synchrony, the nodes that they share must be instantiated twice. *Experiment 6* provides evidence for a relation between human reaction time and the level of similarity between elements involved in premises.

In these experiments when the premises involved in the compared situations contained different words, we wanted to be sure that effects were not due to the use of particular words. One way to compare two situations involving different words is to consult a word frequency table. We preferred to perform a lexical decision task with all words present in those experiments.

The experiment had a within-subject design. The 20 participants were undergraduate psychology majors (14 females 6 males) mean age: 21.8 SD: 2.09. All the words involved in the different experiments were collected: (*bûcheron, coupe, chêne, apprenti, menuisier, clouera, planches, tracteur, fermier, pourra, passer, chemin, si, le, du, sur, des, il, faut, portier, payer, francs, 100, au, présenter, ou, 2, filles, patron, pour, rentrer, dans, bar, avoir, plus, 14, ans, être, accompagné, adulte, entrer, volé, argent, donné, bonbon*). They were added to a collection of non words obtained by replacing a letter in a real French word: (*contactir, débli, entilé, fertale, gralier, inutale, livade, manelle, mochir, orgone, pinle, pochir, quantare, relatir, sapon, serter, stralique, tambard, telon, tesu, tirte, usoge, vengir*). Participants were asked to decide as fast as possible if the presented string was a word or not. If they thought it was a word they had to press a key on the right side of the keyboard if not they had to press a key on the left side of the keyboard. The computer recorded the time required for them to respond.

Whenever comparison involved dissimilar words, respecting word reaction times were combined. The resulting mean word reaction times (for every situations of an experiment) were then compared with appropriate statistical measure. This indicated if mean word reaction times were significantly different (or not) in the situations compared.

4 Constraints on the length of the reasoning chain

4.1 Introduction

In INFERNET the number of steps required for transmitting activation from one point to another is constrained by two factors. The first is the probability of errors in transmission which increases as the number of steps increases. In Figure 4.1, for the role 1 nodes to fire, the transmission must go through intermediate nodes 1. For the role 2 nodes to fire, the transmission must go through intermediate nodes 1, 2 and 3. At each intermediary step there is a chance of error. The second constraint concerns the lag between the starting point of the filler's oscillation and the starting point of the role's oscillation. The more steps that are required, the greater the lag between roles and fillers. As the lag increases, the number of synchronizations between the firing of role and filler nodes decreases.

Figure 4.1
Illustration of lag
between first filler
nodes spikes and
first role node
spikes.

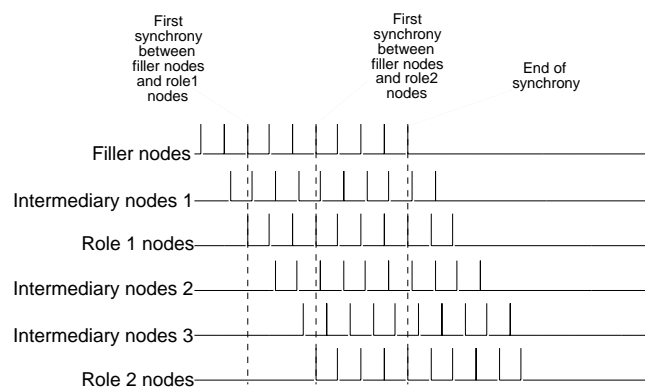


Figure 4.1 illustrates this second constraint. The synchrony between “Filler nodes” and “Role 1 nodes” starts before the synchrony between “Filler nodes” and “Role 2 nodes” but

both end at the same time. Since in the process of learning bindings, each synchrony provokes a modification of connection strength, the more synchronies that occur, the more learning will be fast and reliable. In this case filler nodes will be bound to role 1 nodes faster and more reliably than to role 2.

This chapter will illustrate this constraint by having INFERNET solve conditional reasoning tasks. The behavior of INFERNET will produce hypotheses that will be verified on humans.

4.2 Conditional reasoning

Many psychological studies in the area of deductive reasoning have focused on conditional reasoning of the type “if p then q”. This kind of statement is common in ordinary language. For example, the reader has surely heard a mother promising her daughter: “If you drink your soup, you will get a desert”. For the logic point of view “If-then” forms refers to material implication. Of course, some logicians would deny that material implication “ \supset ” is really what humans mean by “if ...then”. Nonetheless, here are the standard logical rules related to material implication: modus ponens (MP): *if p then q, p, therefore q*, $\frac{p \supset q, p}{q}$, and modus tollens (MT): *if p then q, not q, therefore not p*, $\frac{p \supset q, \sim q}{\sim p}$. While most humans follow modus ponens, it is different for modus tollens. People also use two inappropriate rules: Denial of the antecedent (DA): *if p then q, not p, therefore not q* is an error, we write this: $\frac{p \supset q, \sim p}{\sim q} \times$, and Affirmation of the consequent (AC): *if p then q, q, therefore p* is also an error, we write this: $\frac{p \supset q, q}{p} \times$. Table 4.1 summarizes the application of MP, DA, AC, and MT. Throughout this chapter the “if p then q” form will be called the “major premise”, p the antecedent, q the consequent.

Table 4.1
The result of
applying the four
inference rules

	MP		DA		AC		MT	
	given	infer	given	infer	given	infer	given	infer
$p \supset q$	p	q	$\sim p$	$\sim q$	q	p	$\sim q$	$\sim p$

Psychological studies on conditional reasoning generally have used 3 different paradigms: Wason selection task, Inference paradigm, and truth table task. The “Wason selection task” was invented by Wason (1966). Wason's problem consists of giving subjects a series of cards that have, for example, on one side a number and on the other a letter. Participants receive a major premise, for example: *if there is a vowel on the card, then there is an even number on the other side*. A series of cards is presented, for example: A, B, 4 and 7 and the participants are asked to only turn over the cards which will allow them to determine the validity of the given major premise. Turning over card A corresponds to applying modus ponens, turning over card B corresponds to applying the deny of the antecedent, turning over

card 4 corresponds to applying the affirmation of the consequent and turning over card 7 corresponds to applying modus tollens. In the inference paradigm, people receive a conditional statement and a minor premise and they have to state the conclusion or to choose a conclusion from a set of different conclusions. In a truth table task, participants receive a major premise e.g.: “If there is an A then there is a 3” along with instances which combine truth or falsity of antecedent and consequent, for example: A-3 (true antecedent - true consequent), A-4 (true antecedent - false consequent), B-3 (false antecedent - true consequent) and C-5 (false antecedent - false consequent). The participant’s task is to judge whether each instance contradicts or conforms to the major premise.

Early studies with the Wason selection task (Wason, 1966; 1968; 1969) showed that very few (around 10%) of undergraduate students were applying Modus Tollens. Wason explanation was in term of “verification bias”. When given a major premise like “if p then q”, people try to prove that the statement is true instead of false. They choose a combination that confirms the major premise (true antecedent, true consequent) and avoid the falsifying card (false consequent) with a “defective truth table”. In this “defective truth table” the false antecedent cases (i.e., cases with a true consequent and false antecedent or a false consequent and false antecedent) are considered as undetermined.

Studies with the inference paradigm (Taplin, 1971; Evans 1977, Wildman & Fletcher, 1977, Marcus & Rips, 1979) reveal a greatly improved performance on application of modus tollens (more than 60%).

People’s performance on the truth table task is closer to that observed with Wason selection task (Johnson-Laird & Tagart, 1969).

4.2.1 Models of conditional reasoning

Current models of conditional reasoning tend to be symbolic. They rely on three theoretical currents: general purpose rule theory, Mental model theory and pragmatic schema theory.

According to authors like Inhelder & Piaget (1955), Beth & Piaget (1961), Beth, Grize, Martin, Matalon, Naess & Piaget (1962), Braine (1978), Braine & O'Brien (1991) or Rips (1983, 1990, 1994), humans possess a “natural” mental logic. This natural formal logic describes and determines the nature of human reasoning. When faced with a problem, the subject translates its content into an abstract representation to which inference schemas are applied. Once the inference is generated, it is channeled towards the real domain. A computational model PSYCOP has been presented by Rips (1994). It provides simulator results but not for simple conditional reasoning.

The theory of mental models of Johnson-Laird specifies that a logical reasoning can be done without recourse to either general or specific rules (see Johnson-Laird 1983;

Johnson-Laird & Byrne 1991). According to this theory, logic is a set of procedures allowing us to establish the validity of a given inference. However, an inference system can behave in an entirely logical manner without using inference rules or any other formal machinery. It would thus be useless to postulate the necessity of some mental logic. When a subject reasons, he/she only tests if the conclusion is true, assuming that the data are true. A mental model is a structure analogous to the world which allows the possibility of testing its veracity. When a subject reads the premises of a problem, he/she constructs a mental model to represent the possible states of the world which are consistent with the available information. The subject constructs a provisional inference based on true propositions of the model. In order to ensure the validity of the inference, the subject looks for counter-examples by the construction of alternative models for which the data remain true, but not the conclusion. If no counter-example is discovered, the inference is held to be true. A computational symbolic model, PropAI, is briefly described in Johnson-Laird & Byrne (1993). No comparison between PropAI and human performance is provided. PropAI is available on the Web, which allowed us to test it on a simple conditional reasoning “If p then q” and each puzzle (MP, DA, AC and MT) were tested 20 times. This program always returns the same response for a given input. It leads to 100% of MP and MT acceptance and 0% of DA and AC acceptance. Clearly, this model is more a computational model of a particular logical derivation than a model of human reasoning.

Cheng and Holyoak's theory of pragmatic reasoning schemas (Cheng & Holyoak, 1985, 1989, Holyoak & Cheng, 1995) describes schemas which are structures that are more abstract than knowledge-specific, but are, nonetheless, more knowledge-specific than the general-purpose inference rules. Reasoning is most likely neither based on rules independent of context as in Piaget, Braine or Rips, nor on the memory of specific experiences as Griggs & Cox (1982) suggested in their “memory cueing hypothesis”. Subjects instead use abstract structures of knowledge induced by everyday life, which are called pragmatic schemas of reasoning. Cheng and Holyoak (1985, p. 395) define them as follows: “A *pragmatic reasoning schema* consists of a set of generalized, context-sensitive rules which, unlike purely syntactic rules, are defined in terms of classes of goals (such as taking desirable actions or making predictions about possible future events) and relationships to these goals (such as cause and effect or precondition and allowable action)”. As we see, the intention, the goal and their relationships make up the organizing structure of the application of a schema. Holyoak did not published a specific computational model but much symbolic program in AI would fit this theory.

4.2.2 Content effect on Conditional reasoning

Content effect

Wason and Shapiro (1971), using a “Wason selection task”, discovered that making the content of the problem more concrete raised the proportion of sound answers. The name of a city was written on one side of the cards, and, on the other side, a means of transportation. The participants were informed that each of these cards represented a day of travel for the experimenter. They were asked to verify the assertion that “Every time I go to Manchester, I travel by train”. The proposed cards were: Manchester, Leeds, Train and Car. A significantly higher proportion of participants gave the right answer with this concrete content than with the abstract data of letters and numbers. However, this “concrete content” effect has not been replicated for this transportation rule. Johnson-Laird, Legrenzi, & Legrenzi (1972) found that subjects’ performance significantly improved with respect to the standard Wason task if the major premise was: “If a letter is sealed, it has a 50 lire stamp on it”. The facilitating effect of concrete content has nonetheless been challenged by Manktelow & Evans (1979). These authors discovered concrete contents for which no facilitation was observed. More recently, Pollard & Evans (1987) stated that a concrete and adequate scenario should be used to produce facilitation. It seems that an adequate context and content are required in order for a facilitation to appear. According to Evans (1989), content and context must be rather coherent in order for the participant to apply actions that would be appropriate in real life.

Deontic effect

The most studied content effect on reasoning concerns facilitation subsequent to the use of statements which make reference to deontic factors. Deontic reasoning is reasoning about what people may, ought or must do in a particular context. Griggs & Cox (1982) found a facilitation in Wason selection task. They gave participants a major premise: “*If a person is drinking beer, then the person must be over 19 years of age*” and asked them to imagine that they were a police officer and that their job was to ensure that the law was respected. People then performed significantly better than in the abstract Wason task using letters and numbers. Griggs & Cox (1982) explained this facilitation by the fact that the situation involved a particular piece of memory that was relevant. The Wason selection task performance would be facilitated when the presentation of the task allows people to recall a real experience similar to the one they are being tested on.

Cheng & Holyoak (1985) postulated another explanation for the phenomenon. They described pragmatic reasoning schemas which are a set of rules that are invoked in particular situations. These schemas include a permission (an action requires a precondition to be

satisfied), or an obligation (an action must be put forward if a precondition is present). In their 1985 study they described the permission schema as:

- Rule P1: If the action is to be taken, then the precondition must be satisfied.
- Rule P2: If the action is not to be taken, then the precondition need not be satisfied.
- Rule P3: If the precondition is satisfied, then the action may be taken.
- Rule P4: If the precondition is not satisfied, then the action must not be taken.

When a conditional statement matches rule 1 of the permission schema, the solution that follows from the application of the schema follows the rules governing logical material implication. They replicated the Griggs & Cox (1982) experiment. They also showed that improved performance occurs when the permission rule involved abstract content like “If one is to take action A, then one must first satisfy precondition P”.

Cosmides (1989) provided another explanation of the phenomenon which involves a “social contract”. This evolutionary theory postulated that people possesses a rule: “*If you take the benefit, then you pay the cost*”. This is a kind a social contract which would have been shaped and included in human genes by evolution. Evolution would have also shaped humans to look for cheaters. With the rule “*If you take the benefit, then you pay the cost*”, looking for cheaters is looking for people who benefited without then paying a cost (MP cases) and people who did not incur a cost but then took a benefit (MT cases). Cosmides (1989) invoked the social contract rule in a Wason selection task by an introductory story which leads people to look for violations of the rule. She obtained a facilitation with unfamiliar content like “*If a man eats cassava root, then he must have a tattoo on his face*”. She also did an experiment with a “switched” rule of the form “*If I pay a cost to you, then you must benefit me*” which according to the social contract theory should trigger looking for cheaters who did not pay a cost (DA) and cheaters who benefitted without paying a cost (AC). With an unfamiliar content: “*If a man has a tattoo on his face, then he eats cassava root*”, Cosmides found more DA and AC than MP and MT. Cosmides (1989) also compared permission rules with or without the social contract context, and found better correct performance when the social contract was included in the context. When these latter rules were switched, she found more DA and AC for rules accompanied by a social contract context.

Cheng & Holyoak (1989) criticized Cosmides (1989) on the grounds of the differences of the phrasing of the introductory context which were confounded for cases without social contract context. They also provided results of an experiment which compared an identical rule with an introductory context that did or did not involve a cost. They found no difference and concluded that the facilitation observed by Cosmides (1989) was a consequence of the invocation of the permission schema. Concerning the reverse effect of a switched rule: e.g. “*If a man has a tattoo on his face, then he eats cassava root*” it can be stated that it could be interpreted to be like rule 3 of the permission schema: “*If the precondition is satisfied, then the action may be taken.*” and lead to an increase of DA and AC.

Platt & Griggs (1993) showed that if cost-benefit context does increase facilitation, the use of the word “must” in the major premise also increases facilitation.

Gigerenzer & Hug (1992) showed that when there is a bilateral cheating option, the switching effect was more striking than when the cheating option is unilateral. Unilateral cheating can be illustrated by this major premise: *“If the envelope is sealed, then it must have a 1 mark stamp.”* The only type of cheaters are those who used a sealed envelope with a stamp whose value was less than 1 mark. Bilateral cheating is present in this major premise: *“If an employee works on the week-end, then that person gets a day-off during the week”*. The difference with the first type of cheating is the number of possible types of cheater. In bilateral cheating there are two. The first type of cheaters are those employers who do not give a day off when an employee worked on week-end, and the second type of cheaters are those employees who get a day-off without working on week-end. They concluded that only social contract theory explain “cheater” effects. But Liberman & Klar (1996) showed that cheating and not cheating contexts were equivalent when the understanding of the task was controlled for.

Jackson & Griggs (1990) began a new debate in the field. They found that using implicit negations in the DA and MT minor premises reduced the facilitation effect in the permission context. Consider a major premise like: *“If one is to take action A, then one must first satisfy precondition P”* (this major premise has the form of the rule-1 of the Cheng & Holyoak’s permission schema). Explicit negation for the DA question is presenting the minor premise: *“Has not taken action A”* while using implicit negation is presenting the minor premise: *“Has taken action B”* (implicitly the action A is not taken). Explicit negation for the MT question is presenting the minor premise: *“Has not fulfilled precondition P”* while using implicit negation is presenting the minor premise: *“Has fulfilled precondition Q”* (implicitly the precondition P is not fulfilled). Jackson & Griggs (1990) observed a better performance with explicit negations than with implicit negations. However, when permission was removed (standard Wason selection task), the use of explicit negation did not produced a better performance. Finally, when participants were asked to check for violation, facilitation occurred with explicit negation in the standard Wason selection Task. This last effect has not been replicated in the study of Kroger, Cheng & Holyoak (1993). Girotto, Mazzocco & Cherubini (1992) replicated the implicit negation effect but also found that when the implicit negation is deontically relevant, facilitation occurs. Griggs & Cox (1993) found similar results and concluded that failing to apply MT with implicit negation was probably due to ambiguity in the MT minor premise which make the MT application irrelevant (see also Noveck & O’Brien, 1996).

Manktelow & Over (1991) introduced the study of point of view in deontic reasoning. They presented a context in which a mother tells her son: *“If you tidy your room, then you may go out to play”*. In terms of Cheng & Holyoak’s permission schema, this major

premise has the form of rule 3 and should increase the proportion of DA and AC and decrease MP and MT. Manktelow & Over (1991) observed this result when participants were asked to verify if the son had broken the rule (mother's point of view), but observed the contrary (increase of MP, MT and decrease of DA, AC) when participants were asked to check if the mother had broken the rule (son's point of view). To explain these data, a theory should explain how utilities of outcomes of different actions taken by oneself or others can be evaluated and used to perform inferences.

To assess the effect of expected utility on deontic reasoning, Kirby (1994) used a Wason selection task with a drinking age major premise: *"If a person is drinking beer, then the person must be over 19 years of age."* Kirby (1994) used different MT minor premises with different ages: *"the person is 19 years old"*, *"the person is 12 years old"*, and *"the person is 4 years old."* These different versions lead to different frequencies of MT application reflecting the different utilities of outcomes. It is more probable that a 19 year old person would break the rule than a 4 year old child. This experiment provides empirical evidence that utilities need to be incorporated in a theory of reasoning. The effect of utilities has been replicated by Manktelow, Sutherland & Over (1995), Green, Over & Pyne (1997). Finally, Oaksford & Chater (1994), Chater & Oaksford (1999) have developed a probabilistic approach based on Bayesian statistics. They suggest that people select cards in order to maximize information gain. People first select the data that has the greatest informativeness. They do not tend to care about rare events. Since there is a cost of examining data, the less possible examination is performed. Combining cost and rarity, people's inferences will lead to an ordering of information gain associated with turning over each card: MP>AC>MT>DA. So we should observe more MP application than AC, more AC than MT and more MT than DA. In the Kirby (1994) experiment the MT minor premise *"the person is 19 years old"*, is more probable (less rare) than the MT minor premise *"the person is 4 years old"*.

Holyoak & Cheng (1995) extended the pragmatic schema theory to explain these point of view effects. They first presented the schema of obligation:

- Rule O1: If the precondition is satisfied, then the action must be taken.
- Rule O2: If the precondition is not satisfied, then the action need not to be taken.
- Rule O3: If the action is to be taken, then the precondition may have been satisfied.
- Rule O4: If the action is not to be taken, then the precondition must not have been satisfied.

The theory predicts that changing point of view encourages the mapping of a particular rule in the permission schema to another rule in the obligation schema. This leads to the reverse of facilitation effect. The permission schema involves rights and the obligation schema involves duties. They are complementary. Changing point of view leads to a switch from rights to duties and vice-versa.

For Oaksford & Chater (1995b) the pragmatic reasoning schema explanation is less parsimonious than their probabilistic model.

The ability to correctly solve deontic problems seems to be a universal human ability. Dellarosa-Cummins (1996), showed evidence of deontic reasoning in 3 years old children. Stanovich & West (1998) showed that, unlike nondeontic tasks, deontic tasks were not related to cognitive ability.

The above empirical data were collected using the Wason selection task paradigm. Studies of deontic reasoning with other paradigm are rare. Newstead, Ellis, Evans & Dennis (1997) however, used truth table task. This paradigm, like the Wason selection task, does not involve open answers.

4.2.3 Effect of negated constituents

What happens when negations are introduced into the major premise? Negation can affect the antecedent or the consequent. It produces four forms of major premises. Table 4.2 shows these four forms and the inferences resulting from the application of the four rules (MT, DA, AC, MT).


Using the inference paradigm some studies showed an effect of the negation (Evans, 1977, Wildman & Fletcher, 1977, Pollard & Evans, 1980). The observed frequencies are showed in Figure 4.2. Pollard & Evans (1980) explain these data with what they call a “negative conclusion bias” which is *a tendency to prefer accepting a conclusion in the negative form*. This is effectively the case for DA and MT. This is not the case for MP, but one could invoke a ceiling effect. Finally the effect is not clear for AC. As stated by Evans, Newstead & Byrne (1993), this bias could be explained by people’s caution. Concluding, for example, in the absence of other information that “the letter is not an X” would have a higher probability (25/26) than concluding that “the letter is X” (1/26). Oaksford & Chater (1994) provide a similar explanation.

Table 4.2
Combination
between form of
major premises and
the result of
applying the four
inference rules

	MP		DA		AC		MT	
	given	infer	given	infer	given	infer	given	infer
$p \supset q$	p	q	$\sim p$	$\sim q$	q	p	$\sim q$	$\sim p$
$p \supset \sim q$	p	$\sim q$	$\sim p$	q	$\sim q$	p	q	$\sim p$
$\sim p \supset q$	$\sim p$	q	p	$\sim q$	q	$\sim p$	$\sim q$	p
$\sim p \supset \sim q$	$\sim p$	$\sim q$	p	q	$\sim q$	$\sim p$	q	p

There is also an interpretation of the negation effect in terms of a “Matching bias”: a tendency to verify cases that are stated in the major premise. If the major premise is “*If p then not q*” people would prefer turning *p* and *not q* cards because they match the major premise content. However, this bias concerns only procedures like the “Wason Selection Task” or the “Truth Table Task” in which participants have to test or verify a major premise instead of applying it. Moreover, the matching bias is closely related to implicit negation

(Evans, 1998) and seems to disappear with the use of explicit negations (Evans, Clibbens & Rood, 1996). While negation in conditionals is known to create difficulties (Oaksford & Stenning, 1992), little is said about double negation (for an exception, see Sperber Cara, & Girotto, 1995, Evans, Clibbens & Rood, 1995). The present study focuses on explicit negation with a production task.



In this study, the INFERNET simulator's performance will be compared with human data. INFERNET provides an interesting hypothesis, related to the difficulty of removing double negations. A replication of previous experimental work was carried out in order to collect reaction times which are not available in previous related studies, like Evans, Clibbens & Rood (1995).

4.2.4 Connectionist models of conditional reasoning

Bechtel & Abrahamsen (1991) presented a connectionist model capable of treating material implication, (see also Berkeley, Dawson, Medler, Schopflocher, & Hornsby, 1995). This is a back propagation network able to test the logical validity of a given conclusion. It is a model of logical proof but not a model of human reasoning. Moreover, this network is not capable of solving the binding problem and consequently has limited capabilities.

Connectionist systems that solve the binding problem (see chapter 1) are theoretically capable of conditional inference. However, none of them reports comparison with human data.

4.3 Effect of the length of reasoning chain: the case of abstract content with negations

In this section conditional reasoning in INFERNET will be compared to human performance on abstract major premises with and without negations.

The task that will be simulated on INFERNET corresponds to the one already described in Table 4.2.

INFERNET has a Long Term Knowledge Base that is used for encoding premises and answering queries. Figure 4.3 shows the knowledge necessary to make conditional inferences with negations. Arrows represent connections; they are tagged with numbers that indicate the time required to propagate activation. Specifically, in this example, a delay of 30ms corresponds to the lag between two spikes of a node oscillating at 33Hz. This delay ensures that these symbol-node spikes will synchronize after 30ms. The knowledge encoded, as shown in Figure 4.3, can correctly answer the query related to material implication.

The diagram illustrates a fuzzy inference process. It begins with four input nodes: two 'Content' nodes (purple) and two 'If'/'Then' nodes (yellow). These inputs feed into 'Antecedent' and 'Consequent' nodes (yellow). The 'Antecedent' node has a 'binding' value of 10. The 'Consequent' node has a 'binding' value of 10. The 'Antecedent' and 'Consequent' nodes feed into a 'Negation' node (yellow). The 'Negation' node feeds into four 'AND Gate' blocks (blue) and two 'XOR gate' blocks (green triangles). The 'AND Gate' blocks feed into 'AND Gate 5' (blue boxes). The final output is a numerical value, 15, shown at the bottom.

The first ability that INFERNET must have is that it must be able to detect negations in the major premise. AND-gate 2 detects when the antecedent is negated in the major premise and AND-gate 3 detects a negated consequent. During the premise-encoding phase, if an antecedent is negated, for example: $\sim A \supset B$, the connection between the AND-gate 2 and “A” will be strengthened as well as the connections between “A” and “Antecedent”. After this phase, the firing of “A” nodes will be sufficient to induce the synchronous firing of the nodes of AND-gate 2. The second ability of INFERNET is to detect whether in the question (minor premise) the antecedent or the consequent (as stated in the major premise) is negated and that is done by AND-gate 1 and AND-gate 4.

By following the diagram carefully, one can see that AND-gate 1 detects the denial of the antecedent, and AND-gate 4 detects the denial of the consequent. If the antecedent or the consequent has a negative form in the major premise (e.g. $\sim A \supset \sim B$), and if the minor premise is in the affirmative form (e.g. “A”), AND-gate 1 will be activated by AND-gate 2 by the means of an XOR gate. The same principle activates AND-gate 4. The role of AND-gate 5 is to detect double negations. This gate will be active whenever AND-gate 1 and 2 or AND-gate 3 and 4 are active. This gate prevents nodes representing “negation” from firing, therefore solving double negation. In order to do correct inferences, “Antecedent” and “Consequent” must be linked. The detection of the “Antecedent” in the question must enable firing of “consequent” nodes, unless AND-gate 1 is active (thereby avoiding Denying the Antecedent). The detection of the “Consequent” in the question must enable firing of “Antecedent” nodes if AND-gate 4 is not active (it avoids Affirming the Consequent). Finally, if AND-gate 1 is active, AND-gate 4 will be activated, and vice-versa.

In section 2.6, it was argued that circuits of successive temporal gates were a possible explanation of syn-fire chains (Abeles et al., 1993). Consequently, the diagram of Figure 4.3 has some neurobiological plausibility.

Hypotheses

Classical explanations of negation effects in conditional reasoning rely on the notion of “negative conclusion bias”: a tendency to prefer inferences in the negative form with the exception of MP (Pollard & Evans, 1980).

The first hypothesis that follows from INFERNET is that it should be easier to apply Modus Ponens than any other rule. This effect is attributed to the stronger links from antecedent nodes to consequent nodes. The second hypothesis states that any time AND-gate 5 (see Figure 4.3) is needed, a decrease in performance should occur. This effect is due to an increase in the number of steps required to propagate the activation, and to this gate’s role of blocking the oscillation of “negation” nodes. This AND-gate 5 is required to treat double

negations. Therefore this hypothesis predicts a decrease in the number of incorrect DA answers for major premises “ $p \supset \sim q$ ” and “ $\sim p \supset \sim q$ ” and a decrease in the number of correct MT answers for major premises “ $\sim p \supset q$ ” and “ $\sim p \supset \sim q$ ”. In order to contrast classical and INFERNET hypotheses frequencies of inference and reaction times will be used.

Material

Four major premises were constructed. They differ by the use of negation in their antecedent or consequent parts. The first major premise involves affirmative antecedent and consequent ($p \supset q$). The second major premise involves affirmative antecedent and negative consequent ($p \supset \sim q$). The third major premise involves negative antecedent and affirmative consequent ($\sim p \supset q$). The last major premise involves negative antecedent and consequent ($\sim p \supset \sim q$). Four minor premises or questions were constructed: p , $\sim p$, q and $\sim q$.

Procedure and parameters

Each object in the experiment is composed of 12 nodes connected to each other with a delay of Δt_γ . In the learning phase, the major premise is presented to INFERNET by making corresponding nodes fire in a particular order (Figure 4.4). This input is repeated 10 times which correspond to the number of gamma cycles in a burst. An input like “If $\sim p$ then $\sim q$ ” is presented to the system with a particular a priori fixed structure (Figure 4.4) (how this structure is achieved is beyond the scope of this work.) The different parts of this input are split in different windows of synchrony. The first is attributed to “If” nodes, the second, to “Negation” and “p” nodes, the third, to “then” nodes, and the last to “q” and “negation” nodes. When “negation” nodes fire in two windows of synchrony, it is said to be “doubly-instantiated” (see chapter 6 for more details).

The question is then presented to the system (Figure 4.5). The response of the system will be used to evaluate reaction time. This process is repeated 10 times (theta wave). For every ms interval, firing nodes are recorded. The question phase follows. The four questions are presented successively to the network (a “within-subjects” design). The “within-subjects” design is simulated by saving random states, every time the system receives a new premise it reloads this saved random state. It changes random state when it simulates another subject (or trial). This ensures that the difference between simulated subject questions is not an effect of a modification in the initial state of the network. This input is presented once every theta cycle. The reaction of the network is recorded and the resulting data files will be used to calculate the proportion of each INFERNET responses.

Figure 4.4
Major premise input


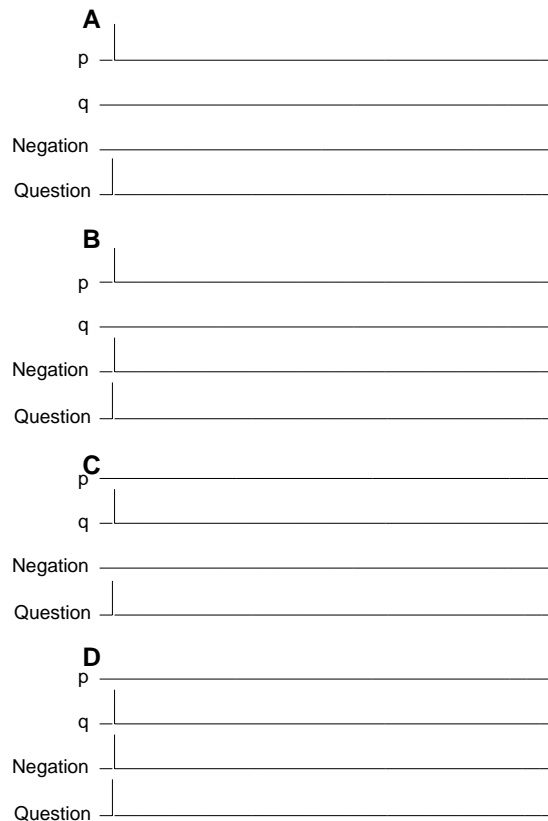


Figure 4.5
The four minor
premises input. A: p,
B: $\sim p$, C: q, D: $\sim q$

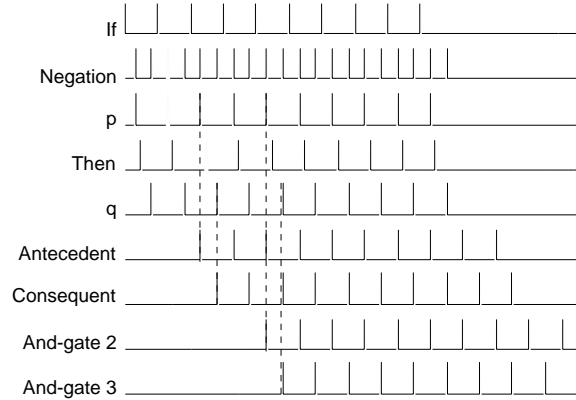


INFERNET simulation results

INFERNET uses a two-step process for drawing inferences. The first is to encode premises by temporarily increasing the weight of connections between objects and their respective roles. For example, giving the major premise “If $\sim p$ then $\sim q$ ”, this first step will increase connection weights between the object “p” and the role “Antecedent”, between the object “p” and the nodes associated with gate 2, between the object “q” and the role “Consequent” and between the object “q” and the nodes associated with gate 3. Imagine that when a gate is opened some nodes fire. Figure 4.6 shows an example of these synchronies as they happen as a result of the connections shown in Figure 4.3 and of an input: “If $\sim p$ then $\sim q$ ”. Figure 4.6 shows the ideal distribution of node firing times for each object that follows from the major premise input “If $\sim p$ then $\sim q$ ”. Each object is represented by a set of nodes and if a majority of these nodes fire within a short period, the object is said to be activated. The firing of “If” nodes is responsible for the firing of “Antecedent” nodes 10 ms later (see the connections in Figure 4.3). The firing of “Then” nodes is responsible for the firing of

“Consequent” nodes 10 ms later. Shortly after, “Antecedent” nodes synchronise with “p” and with “Negation” nodes, while “Consequent” nodes fire in synchrony with “q” and “Negation” nodes. From the firing synchrony of “Antecedent” and “Negation” nodes, AND-gate-2 will open and its associated nodes will fire in synchrony with “p” nodes. From the synchrony of “Consequent” and “Negation” nodes, AND-gate-3 will open and its associated nodes will fire in synchrony with “q” nodes. These synchronies enable modifications of connections between “p” and “Antecedent” nodes and between “p” and AND-gate-2 nodes. After a certain amount of repetition, the firing of “p” nodes will enable the firing of “Antecedent” and AND-gate-2 nodes, in the same manner, the firing of “Antecedent” or AND-gate-2 nodes will enable the firing of “p” nodes. The same thing will happen to connections between “q” and “Consequent” nodes and between “q” and AND-gate-3 nodes.

Figure 4.6
Distribution of node
firing time following
the input “If not q
then not p”.




The second step is the network’s response to a query. In the above example (“If $\sim p$ then $\sim q$ ”), giving as minor premise the object “q”, we expect the system to apply the “modus tollens” rule and answer “p”.

Figure 4.7 shows the network’s response to the query “There is a q, what do you conclude?” which is represented by a binding of “q” to “Question”. When “q” nodes fire, they enable the firing of “Consequent” and AND-gate 3 nodes. This follows from the strengths modifications from the encoding phase. Nodes belonging to AND-gate 3 and “Consequent” activate AND-gate 4 nodes in the absence of inhibition from “negation”. The conjunction of AND-gate 3 and AND-gate 4 nodes firing opens AND-gates 5, whose firing prevents “negation” nodes to fire near “Consequent” nodes firing time. Meanwhile, “Consequent” nodes have activated “Antecedent” nodes. Because of the rule encoding phase, “Antecedent” nodes activates “p” nodes which, in turn, activates AND-gate 2 nodes. This last gate nodes in conjunction with AND-gate 1 nodes that have been activated by AND-gate

4 nodes, enable AND-gate 5 nodes firing in a second window of synchrony. The nodes following AND-gate 5 inhibit the firing of “Negation” nodes around “Antecedent” nodes firing time. This mechanism enables INFERNET to respond “p” to the query.

This graph is a theoretical and ideal representation. INFERNET’s performance is more noisy. Negation may, for example, fire before AND-gate 5 inhibits it.

Figure 4.7
Distribution of node
firing time following
the Query “There is
a q, what do you
conclude?”



INFERNET’s efficiency was measured as described in section 2.8. The frequency of inferences (MP, DA, AC, and MT) are reported in Table 4.3. INFERNET was run 40 times on each of the 16 problems. As expected, MP is more often applied than any other rule. There is also an effect of double negation which is responsible for the low frequencies of DA when the consequent is negated and of MT when the antecedent is negated.

Table 4.3
INFERNET
frequencies of
inference

	MP	DA	AC	MT
$p \supset q$	39	28	25	34
$p \supset \sim q$	37	11	27	31
$\sim p \supset q$	39	30	27	13
$\sim p \supset \sim q$	34	15	26	20

The data were arranged to show the contrast between “negative conclusion bias” hypothesis and INFERNET double negation hypothesis. Figure 4.8 shows the hierarchical organization of the data. The first level splits data in two groups: Is the conclusion negative or positive? The second splits these two groups into two subgroups: is the inference expected to be negative or positive? MP-AC is grouped on one side and DA-MT on the other side. The combination of “conclusion sign positive” and “expected sign negative” involves

all the cases that require double negation. The third factor splits the data in two levels: infer a consequent “forward” (MP-DA) and infer an antecedent “backward” (AC-MT). The combination of “expected sign positive form” and “infer forward” involves all the MP cases. A final factor splits the data into those cases where the minor premise (question) is in negative or positive form.

Figure 4.8
Arrangement of the
data for the analysis



The log-linear analysis results are displayed in Table 4.4. The G^2 are underestimated because data were analysed with a between design. The Rasch model could not be applied because it would require a 2^{16} cases contingency table and no statistical software has been found that can compute such a huge table. (Even if widely used, an ANOVA on categorical data is not appropriate.) Nonetheless, an ANOVA has been performed and Table 4.4 displays results on 4 within factors with 2 levels each. The sphericity assumption was violated, so a box correction ($\hat{\epsilon}=0.67235$) on the number of degrees of freedom has been performed.

The “negative conclusion bias” effect (*a tendency to prefer accepting a conclusion in the negative form*) is measured by the effect conclusion sign. This effect is significant. There are 199 positive conclusions and 257 negative conclusions. The INFERNET hypothesis is tested by the interaction between “conclusion sign” and “expected sign” which is significant. Among positive conclusions, there are 130 conclusions where the expected conclusion is positive and only 59 conclusions where they are expected negative, while among negative conclusions the expected positive conclusions (124) are approximately as numerous as the expected negative conclusions (133). The other predicted effect of INFERNET is the interaction “Expected sign” * “Infer forward vs. backward” which is also

significant. Inside expected positive conclusions cases, there is more MP (149) than AC (105) while inside expected negative conclusion cases, DA (94) and MT (98) are approximately equal. One way to decide between the two hypotheses is to test their model respective fitting. The best fitting model is: Conclusion sign * Expected sign * Response + Expected sign * Infer Forward Backward * Response $G^2_{(20)}=10.90787$ $p=.9485464$ which is what was expected. In summary, there is more MP inferences than any others, and when an inference involves a double negation, it is hardly stated.

		Effect	DF	G ²	p	F	DF	p
Table 4.4 Log-linear analysis and ANOVA on INFERNET frequencies of inference		Independence	15	145.5333	1.85 E-23			
		Premise sign	8	1.15999	.2814676	1.45643	1,26	.2383648
		Infer Forward Backward	1	7.20272	.0072793	10.76687	1,26	.0029430
		Expected sign	1	39.98231	2.56 E-10	41.29412	1,26	.0000008
		Conclusion sign	1	26.35683	.0000003	31.56785	1,26	.0000066
		Cs*Es	1	29.64906	.0000001	36.89189	1,26	.0000020
		Cs*I*	1	1.842829	.1746197	.66102	1,26	.4235830
		Cs*Ps	1	.229631	.6317976	.26773	1,26	.6092331
		Es*I	1	36.69561	1.38 E-9	43.21344	1,26	.0000006
		Es*Ps	1	.089985	.7641962	.54545	1,26	.4667978
		I*Ps	1	.436194	.5089652	.31593	1,26	.5788779
		Cs*Es*I	1	2.245933	.1339661	.60938	1,26	.4420711
		Cs*Es*Ps	1	2.692276	.1008357	2.92771	1,26	.0989780
		Cs*I*Ps	1	.9941513	.3187299	.37500	1,26	.5456104
		Es*I*Ps	1	.0130477	.9090582	.05424	1,26	.8176668
		Cs*Es*I*Ps	1	.0808797	.7761106	.39394	1,26	.5357097

Table 4.5 reports mean and SD reaction times for the four premises and the four inferences.

			MP	DA	AC	MT
Table 4.5 means and Standard deviations INFERNET data (in ms)	p>q	mean	2402	5808	3794	3472
		SD	383	2773	1130	816
	p>~q	mean	2677	4970	3764	3809
		SD	217	2867	924	1391
	~p>q	mean	2351	4834	5069	5384
		SD	416	2773	1411	858
	~p>~q	mean	2921	4309	5307	5484
		SD	443	2316	1458	1586

Data were analyzed by an ANOVA, 4-way within-subjects with 2 levels. Normally with 2 levels the sphericity assumption is not required, but the organization of the data is somewhat artificially nested. That is the reason why we took sphericity into account. Since the sphericity assumption is violated and no transformation met it, the Box correction was used (Box correction $\hat{\epsilon}=.4062180$). ANOVA results are displayed in Table 4.6.

The “negative conclusion bias” effect (i.e., *a tendency to prefer accepting a conclusion in the negative form*), (conclusion sign) is not significant. The “expected sign” (expected positive: MP+AC, expected negative: DA+MT) and “infer forward vs. backward” (forward: MP+DA, backward: AC+MT) are both significant. The mean reaction time for expected

positive cases is 3548ms, while for expected negative cases it is 4759 ms. Forward inferences are faster (mean: 3796 ms) than backward inferences (mean: 4510 ms).

Table 4.6
ANOVA on
INFERNET Reaction
times

Effect	DF	F	p
Conclusion sign	1,16	2.8021	.1135742
Expected sign	1,16	172.8939	5.41 E-10
Infer Forward Backward	1,16	47.3611	.0000037
Premise sign	1,16	2.0023	.1762287
Cs*Es	1,16	32.6280	.0000321
Cs*I	1,16	7.6755	.0136471
Cs*Ps	1,16	.3011	.5907693
Es*I	1,16	81.7152	.0000001
Es*Ps	1,16	.5128	.4842507
I*Ps	1,16	2.8859	.1087156
Cs*Es*I	1,16	39.7897	.0000104
Cs*Es*Ps	1,16	.1851	.6727654
Cs*I*Ps	1,16	.2496	.6241571
Es*I*Ps	1,16	5.3560	.0342673
Es*C*I*Ps	1,16	.3924	.5398757

The INFERNET hypothesis is tested by the interaction between “conclusion sign” and “expected sign” which is significant. Among positive conclusions cases, the mean reaction time for expected positive cases is 3078 ms and for expected negative cases it is 5037 ms. While among negative conclusions cases, the mean reaction time for expected positive cases is 4018 ms and for expected negative cases it is 4481 ms, this difference is less important. Post hoc Tukey contrasting double negation cases with others are all significant.

The other predicted effect of INFERNET is the interaction “Expected sign” * “Infer forward vs. backward” which is also significant. The MP cases mean reaction time is lower (2213 ms) than any other (DA: 4980, AC: 4483, MT: 4537). Post hoc Tukey contrasting MP with others are all significant.

4.3.2 Experiment 1: Conditional reasoning with negated constituents in humans

The following experiment was conducted to provide a comparison between INFERNET and humans.

Participants and Design

The experiment has a within-subjects design. Forty participants received four major premises in a random order and had to answer four questions for each major premise in a random order. The 40 participants were undergraduate psychology majors, 31 females 9 males, mean age was 21.32 and SD was 2.1.

Material

Four major premises were constructed, alternating positive and negative antecedent and consequent. Positive antecedent, positive consequent: “If the number is 3 then the letter is X”, Positive antecedent, negative consequent: “If the number is 3 then the letter is not X”, Negative antecedent, positive consequent: “If the number is not 3 then the letter is X”, and Negative antecedent, negative consequent: “If the number is not 3 then the letter is not X”. Each major premise presentation was followed by four questions: “The number is 3, what do you conclude?”, “The number is not 3, what do you conclude?”, “The letter is X, what do you conclude?”, “The letter is not X, what do you conclude?”.

Procedure

Each participant was seated approximately 50 cm in front of the monitor. One of the randomly chosen major premises appeared on the screen. Participants were asked to read it and to indicate when they had understood it. The major premise stayed on the screen during the entire experiment. Questions appeared on the screen, one at the time and in random order. Participants had to answer each question. The computer recorded the time required for them to respond. The experimenter recorded the response. When the participant answered the four questions, the next major premise appeared on the screen with the same procedure until the four major premises had been presented. Before presenting the experimental material, participants received training exercises with the same procedure, but with an arithmetic content. The refreshing frequency of the computer monitor was 75 Hz so the timing precision is about 13 ms. In the program made for the experiment, we took care to block all processor interrupts to be sure that the computer waits only for signal of the time measure (a click on the mouse).

Results

Frequencies of stating each inference are shown in Table 4.7. According to the “Negative Conclusion Bias” hypothesis, there should be more DA type inferences for major premises $3 \supset X$ and $\sim 3 \supset X$ (applying DA for these major premises lead to $\sim X$ conclusion), more AC type inferences for major premises $\sim 3 \supset X$ and $\sim 3 \supset \sim X$ (applying AC for these major premises lead to ~ 3 conclusion), more MT type inference for major premises $3 \supset X$ and $3 \supset \sim X$ (applying MT for these major premises lead to ~ 3 conclusion), and finally more MP. According to the INFERNET prediction, there should be more MP type inferences than any other, fewer DA type inferences for major premises $3 \supset \sim X$ and $\sim 3 \supset \sim X$ (double negation cases), and fewer MT type inferences for major premises $\sim 3 \supset X$ and $\sim 3 \supset \sim X$ (double negation cases).

Table 4.7
Humans frequencies
of inference

	MP	DA	AC	MT
$3 \supset X$	40	31	28	31
$3 \supset \sim X$	37	14	23	32
$\sim 3 \supset X$	38	29	31	17
$\sim 3 \supset \sim X$	40	18	34	16

The INFERNET model seems to be a better explanation than “negative conclusion bias”. Data were analyzed by a log-linear analysis. The G^2 are underestimated because data were analysed with a between-subjects design. The Rasch model could not be applied because in that case it would require a 2^{16} cases contingency table and no statistical software has been found that can compute such a huge table. The Table 4.8 also displays ANOVA results for 4 within-subjects factors with 2 levels each. The sphericity assumption was violated so a Box correction ($\hat{\epsilon}=0.55088$) on the number of degrees of freedom was performed.

In addition to the effect of the conclusion sign being significant (194 positive and 265 negative conclusions), other effects are also. The effect of expected sign is significant, which means that DA + MT (188 inferences) are less often applied than MP + AC (271 inferences). Forward inferences (MP+DA) are more often done (247 inferences) than backward inferences (AC+MT) (212 inferences). The interaction between the expected sign and the conclusion sign is also significant: among the positive conclusions those which involve a double negation are less often inferred (65 inferences) than others (129 inferences) while for negative conclusions cases, expected positive cases (142 inferences) are more comparable with expected negative cases (123 inferences). There is also an interaction between the expected sign and Forward and backward inferences. MP are more often applied (155 inferences) than AC (116 inferences) while DA (92 inferences) and MT (96 inferences) are sensibly equal. The best fitting model is: Conclusion sign * Expected sign * Response + Expected sign * Infer Forward Backward * Response $G^2_{(20)}=10.87807$ $p=.9492933$. This is what INFERNET predicted. ANOVA results confirm these effects.

Table 4.8
Log-linear analysis
and ANOVA on
human frequencies
of inference

Effect	DF	G^2	p	F	DF	p
Independence	15	157.6479	7.25 E-26			
Premise sign	1	.0824	.7740710	.17991	1,21	.6757641
Infer Forward Backward	1	11.09174	.0008671	20.46038	1,21	.0001860
Expected sign	1	59.35807	1.31 E-14	42.84341	1,21	.0000017
Conclusion sign	1	44.13527	3.06 E-11	28.74675	1,21	.0000257
Cs*Es	1	4.893249	.0296190	22.85818	1,21	.0001007
Cs*I*	1	.785947	.3753285	5.66097	1,21	.0269077
Cs*Ps	1	2.658119	.1030225	2.88187	1,21	.1043537
Es*I	1	30.22554	3.85 E-8	12.99063	1,21	.0016667
Es*Ps	1	.003811	.9507752	.01991	1,21	.8891335
I*Ps	1	.002309	.9616748	.01711	1,21	.8971744
Cs*Es*I	1	.655936	.4179984	2.39586	1,21	.1365958
Cs*Es*Ps	1	3.839037	.0507227	1.43656	1,21	.2440483
Cs*I*Ps	1	2.489859	.1145820	.30327	1,21	.5876509
Es*I*Ps	1	.045465	.8311513	.43237	1,21	.5179715
Cs*Es*I*Ps	1	2.686908	.1011859	2.56607	1,21	.1241151

The Table 4.9 shows mean and SD reaction times for the four major premises and the four questions. The double negation cases seems to be slower.

Table 4.9
Human means and
standard deviations
reaction times

		MP	DA	AC	MT
$3 \supset X$	mean	3083	5345	4074	4739
	SD	1749	5033	3414	2998
$3 \supset \sim X$	mean	3486	6588	7564	5739
	SD	1483	5047	11470	3704
$\sim 3 \supset X$	mean	3948	6858	6546	12364
	SD	2037	5758	15630	11058
$\sim 3 \supset \sim X$	mean	4351	6352	5293	8211
	SD	2838	4147	5808	10043

Data were analyzed by an ANOVA as a 4-way within-subjects design with 2 levels. Normally, with 2 levels the sphericity assumption is not required, but the organization of the data is somewhat artificially nested. That is the reason why we took sphericity into account. Since the sphericity assumption is violated and no transformation met it, the Box correction was used (Box correction $\hat{\epsilon}=.236524$). ANOVA results are displayed in Table 4.10.

Table 4.10
ANOVA on human
Reaction times

Effect	DF	F	p
Conclusion sign	1,9	11.01899	.0089476
Expected sign	1,9	18.23539	.0020794
Infer Forward Backward	1,9	6.27900	.0335415
Premise sign	1,9	3.68833	.0869831
Cs*Es	1,9	12.79157	.0059630
Cs*I	1,9	7.73862	.0213366
Cs*Ps	1,9	4.48006	.0633851
Es*I	1,9	.38741	.5491067
Es*Ps	1,9	.26015	.6222895
I*Ps	1,9	3.48920	.0946088
Cs*Es*I	1,9	7.80778	.0209085
Cs*Es*Ps	1,9	.29797	.5984255
Cs*I*Ps	1,9	2.26423	.1666501
Es*I*Ps	1,9	.97017	.3503749
Es*Cs*I*Ps	1,9	.13099	.7257598

The effect conclusion sign is significant: when the conclusion is negative people respond faster (mean: 5295 ms) than when the conclusion is positive (mean: 6523 ms). The effect of expected sign is significant, this means that DA + MT (mean: 7025) are slower than MP + AC (mean 4793 ms). Forward inferences (MP+DA, mean: 5001 ms) are faster than backward inferences (AC+MT, mean: 6816).

The interaction between the expected sign and the conclusion sign is significant: among the positive conclusions those which involve a double negation are slower (mean: 8379) than expected positive cases (mean: 4667 ms). While for negative conclusion cases, expected positive cases (mean: 4919) are approximately equal to expected negative cases (mean: 5770 ms). Post hoc Tukey test indicates that double negation cases are significantly slower than other cases.

There is also an interaction between the conclusion sign and Forward and backward inferences. Among positive conclusion cases, forward inference cases (mean: 4993 ms) are faster than backward inference cases (mean: 8054 ms). While among negative conclusion cases, forward inference cases (mean: 5010 ms) are approximately equal to backward inference cases (mean: 5579 ms). Post hoc Tukey test indicates that positive conclusion + backward inference (AC for rules $3 \supset X$, $3 \supset \sim X$ and double negation MT) cases are significantly slower than other cases.


The interaction conclusion sign * Expected sign * Forward and backward inferences is also significant. Post hoc Tukey reveals that the double negation MT cases are significantly slower than any other cases.

While MP cases (mean: 4667 ms) are faster than other cases, Post hoc Tukey indicates that these case are only significantly faster than DA cases.

These data show departures from both “Negative conclusion bias” and INFERNET hypotheses. MP cases are faster than other cases but not always significantly contrary to both hypotheses. Negative conclusion cases for DA and AC are not always faster than positive conclusion cases, contrary to the “Negative conclusion bias” hypothesis. Double negation cases for DA are not always slower than other DA cases, contrary to INFERNET hypothesis.

4.3.3 Conditional reasoning with negated constituents comparison between INFERNET and humans

The Figure 4.9 compares the frequency of inference between INFERNET simulator and human data. These proportions are similar but a statistical measure is necessary.




In Table 4.11 are displayed the statistical analysis concerning frequencies of inference. A Log-linear analysis was performed to test the effect of the contrast INFERNET-human (called

“Group”) on response frequency. The G^2 are underestimated because data were analysed with a between-subjects design. The Rasch model could not be applied because of the between-within design. Statisticians have not yet studied the application of log-linear analysis to this type of design (Lindsey, personal communication). The Table 4.11 also displays ANOVA results on 1 between-subjects and 4 within-subjects factors with 2 levels each. The sphericity assumption was violated so a Box correction ($\hat{\epsilon}=0.73326$) on the number of degrees of freedom was performed. Variances for the 16 measures are all homogeneous across groups. The Table shows the very close fit of INFERNET to human data.

The best fitting model is Conclusion sign * Expected sign * Response + Expected sign * Infer Forward Backward * Response $G^2_{(52)}=26.40158$ $p=.9988014$. This is the prediction of INFERNET resulting of the difficulty of double negation cases and the easiness of MP cases.

Effect	DF	G^2	p	F	DF	p
Group	1	2.0617	.1510416	1.03658	1,57	.3129203
G*Conclusion sign	1	.76531	.3815786	.59945	1,57	.4419876
G*Infer Forward Backward	1	1.01377	.3140014	.17426	1,57	.6779215
G*Expected sign	1	1.73554	.1877049	.42258	1,57	.5182640
G*Premise sign	1	.19013	.6628084	.54400	1,57	.4638053
G*I*Ps	1	.20816	.6482134	.28425	1,57	.5960018
G*Es*Ps	1	.04432	.8332594	.29243	1,57	.5907746
G*Es*I	1	.00699	.9333695	1.02189	1,57	.3163432
G*Cs*Ps	1	.66852	.4135681	.51546	1,57	.4757172
G*Cs*I	1	.00323	.9546782	.54777	1,57	.4622680
G*Cs*Es	1	3.00937	.0827845	2.82281	1,57	.0980689
G*Es*I*Ps	1	.138421	.7098558	.37244	1,57	.5441029
G*Cs*I*Ps	1	3.778989	.0519000	.66771	1,57	.4172535
G*Cs*Es*Ps	1	6.910677	.0085683	4.36393	1,57	.0411835
G*Cs*Es*I	1	.071965	.7884969	.14240	1,57	.7073081
G*Cs*Es*I*Ps	1	1.669577	.1963245	1.72220	1,57	.1946702

Figure 4.10 displays the comparison between INFERNET and human reaction times. Reaction times seem to be longer for humans but the overall shape is similar.



The reaction times were analyzed by an ANOVA with 1 between-subjects factors (INFERNET vs. humans) and 4 within-subjects variables with 2 levels each. The sphericity assumption was violated, so a Box correction $\hat{\epsilon}=.2682286$ was applied. ANOVA results are displayed in Table 4.12.

There is a significant main effect of group: reaction times are faster in INFERNET (means: 4153 ms vs. 5909 ms for humans). The significant effect of the interaction group*conclusion sign means that INFERNET is faster for positive conclusion (mean: 4057 ms) than for negative conclusions (mean: 4249 ms) while for human it is the contrary (positive conclusions mean: 6523 ms, and negative conclusion mean: 5295 ms). This last effect is probably due to the greater increase of reaction time for double negation cases MT in human compared to INFERNET. Reducing the learning constant in INFERNET would probably increase all reaction times and the difference between double negation cases and other cases.

Table 4.12
ANOVA on
comparison between
human and
INFERNET Reaction
times

Effect	DF	F	p
Group	1,20	12.07717	.0023878
G*Conclusion sign	1,20	13.44253	.0015325
G*Expected sign	1,20	3.70113	.0687256
G*Infer Forward Backward	1,20	2.26443	.1480045
G*Premise sign	1,20	1.8171	.1927298
G*Cs*Es	1,20	2.84294	.1073147
G*Cs*I	1,20	3.37645	.0810386
G*Cs*Ps	1,20	4.75268	.0413727
G*Es*I	1,20	2.16598	.1566559
G*Es*Ps	1,20	.48607	.4937110
G*I*Ps	1,20	5.71672	.0267525
G*Cs*Es*I	1,20	.33985	.5664317
G*Cs*Es*Ps	1,20	.19024	.6673888
G*Cs*I*Ps	1,20	1.94634	.1782827
G*Es*I*Ps	1,20	2.22890	.1510585
G*Es*Cs*I*Ps	1,20	.26254	.6139927

4.4 Discussion

Connectionist modelling of human reasoning is a difficult challenge. Even though Holyoak & Spellman (1993) have described human reasoning in terms of constraint satisfaction, no connectionist system has previously been done for modeling human reasoning. INFERNET shows how reasoning might be possible based on certain low-level neurobiological mechanisms. These properties constrain the reasoning process and explain human limitations. People are sensitive to negated conditionals. INFERNET's account of the phenomenon involves the type of inference and double negation effects and challenges classical explanations that rely on the notion of "Negative Conclusion Bias".

It was predicted that the number of steps required to perform an inference constrained the reasoning process. Removing double negations requires a long chain of gates opening. The longer the chain of successive gates, the higher the number of errors, and the less opportunity

for binding fixation. This section presented INFERNET's predictions and results. These results confirmed that INFERNET is sensitive to double negations. A similar experiment has been conducted on human participants. Results confirmed INFERNET's prediction and showed that the INFERNET explanation is better than classical explanation in terms of "negative conclusion bias". Finally, INFERNET and humans data were compared and there is a high degree of qualitative similarity between the two.

5 The number of predicate's arguments constraints processing

As we saw in chapter 2, the nodes representing the arguments of a predicate fire after the nodes representing the predicate of these arguments. Each predicate is assigned to a particular window of synchrony inside a gamma cycle. Thus, the more arguments a predicate has, the more windows of synchrony will be needed to represent it. For example “John buys a car for Mary” needs 4 windows of synchrony, while “John loves Mary” needs 3 windows of synchrony. This constraint should have an effect on INFERNET's performance. This section will present the psychological literature on the effect of the number of argument, then will present INFERNET's performance, and will finally compare human empirical data with INFERNET.


5.1 The effect of the number of predicate arguments on cognition

It is well known that people have a limited information processing capacity in terms of the number of chunks concurrently available in STM (Miller, 1956). Is there a similarly limited capacity regarding the number of a predicate's arguments?

Halford & Wilson (1980) have defined cognitive complexity on the basis of mathematical structure. The complexity of a statement depends on its dimension or on the number of independent units that are involved in it. This is related to the number of arguments that a predicate has. A predicate with one argument has one dimension of variation, a predicate with two arguments has two dimensions of variation and each of them may vary independently. Since each predicate's arguments ought to be mutually distinguished, they should occupy a particular chunk in memory. For a predicate with 2 arguments like Love(John,Mary), a cognitive system must be able to distinguish the “lover” and the “lovee”.

Halford & Wilson (1980) present two experiments that test how the number of predicate's arguments is related to developmental aspects of analogy making. Children were asked to move a truck around a square similar to Figure 5.1. In group A children had these 3 operators: (remain in the same place, move clockwise and move counter-clockwise) while in group B (remain in the same place, move clockwise and move diagonal). Each of the moves was represented by a geometrical figure. In the first phase, children had to learn by trial and error which move corresponded to the 3 geometrical forms presented by the experimenter. When they succeeded, a new set of 3 geometrical forms was randomly assigned to each of the permissible moves. Children had to learn the new set of correspondences.

Figure 5.1
Environment used in
Halford & Wilson
(1980) experiment 1
first phase.



In the second phase, children had to relate the system they learned, to a problem with the same structure but with new elements. Three new geometrical forms were presented, and the four states (each square corners) were represented by a colored house. The experimenter placed one of the four houses on a corner and the children had to place the houses around the square in an appropriate way (consistent with permissible moves). Children had no feedback for this. They only had feedback for the move corresponding to the cards like “move to the green house”. The group B permissible moves require more arguments because “move clockwise” and “move diagonal” are not symmetrical while in group A “move clockwise” and “move counter-clockwise” are. In group B, children had to keep in mind 2 preceding moves to decide which place to assign to a particular colored house. Figure 5.2 shows that with the group A operators, when the first card comes (move to green) the two alternatives are equivalent and will not lead to inconsistency. The maximum load is \Rightarrow (yellow, green) and \triangleright (green, yellow). Figure 5.3 shows that with the group B operators, when the first card comes (move to green) the two alternatives are possible, one will lead to an inconsistency. The maximum load is now \Rightarrow (yellow, green) \triangleright (green, red) and \Rightarrow (red, yellow).

Figure 5.2
Diagram of Group A
operators analogy
task





Figure 5.3
Diagram of Group B
operators analogy
task



Since the memory load for group B is higher than for group A, there should be more errors in the group B. Halford & Wilson (1980) found that 4 year old children succeeded with system group A, but few 4 year old children succeeded with system group B while 6 year old children had a better performance.

Halford & Wilson's (1980) second experiment involved a hexagon instead of a square truck circuit. Thus the information to be held in memory can involve more steps. The authors found that children could perform the task successfully around 11 years of age.

This empirical data seem to show that the ability to deal with a growing number of arguments is a developmental function. This hypothesis is also corroborated by Gentner & Toupin (1986). They defined systematicity as the potential use of embedded predicates like Cause [Push (a,b), Collide(b,d)]. They observed that 10-year-old children were able to base their analogies on the use of systematicity while 6 year old children had a tendency to base them on surface similarity. However, Goswami & Brown (1989) have shown that when a causal relation linking 2 predicates is simple, younger infants can use it to derive an analogy. This kind of two-level, embedded predicates can be used by trained chimpanzee also (see Premack, 1983). But higher order embeddings are more difficult.

In another experiment (see Halford, Wilson, Guo, Gayler, Wiles & Stewart, 1994), experimenters presented to participants graphical representations and verbal descriptions of 3- and 4-way statistical interactions. The participant's task was to map a verbal description to the corresponding graphical representation. Results show that mapping can usually be done with 3-way interactions, but rarely with 4-way interactions. Understanding interactions requires the ability to embed predicates within each other.

Maybery, Bain & Halford (1986) compared transitive inferences requiring integration with premise verification. For example, deducing that "John is taller than Tom" from the premises: "John is taller than Mike" and "Mike is taller than Tom" requires integration and processing of two relations. Verifying that "Mike is taller than Tom" requires the processing of a single relation. Reaction times are longer when integration is needed (see also Halford Mayberry & Bain, 1986).

Halford, Bain & Maybery (1984) compared arithmetic problems in which participants had to find the operator that was removed from an expression. They showed that when two operators were removed participants had to combine two operations. In that case, a secondary task interfered more with performance.

More recently, Kroger & Holyoak (1997) asked participants to detect changes in two successive stimuli. The rules for considering changes were involving three levels of relational complexity. At a perceptual level any modification of color was considered as a change. At the relational level a change was defined as a relation between color and 2 pairs of squares. At a system level, a change was defined as a relation between the sameness relation

for two pairs. The authors found that reaction time increased as the level became more complex in terms of nesting of predicates.

Some connectionist computational models are sensitive to the number of predicate's arguments. STAR (Halford, et al. 1994) uses a tensor product representation. Each of the predicate's arguments is assigned to a particular dimension in the tensor product. Representing a predicate with 2 arguments requires an array of 3 dimensions. If a predicate has 5 arguments it will require an array of 6 dimensions. As the number of arguments increases, the number of cells in the array increases by an exponential factor, so capacity to process arguments is constrained. Synchrony and phase-based models like LISA (Hummel & Holyoak, 1997) , SHRUTI (Shastri & Ajjanagadde, 1993) or INFERNET represent each argument binding with a particular phase or window of synchrony. In these models, the number of predicate arguments that can be represented depends on the oscillation frequency and the precision of synchrony.

All of the preceding empirical studies compare situations that vary according to the number of arguments needed but also with the level of nesting of these predicates (predicates whose arguments are predicates whose arguments are predicates etc.). In the following section, we present empirical data that vary the number of arguments. The task chosen is disjunctive reasoning in which the number of arguments is manipulated. Participants and INFERNET will be compared with regard to the way they can apply disjunctive syllogisms. This rule is generally well understood. Imagine that you receive this major premise: "People must be over 14 years old or be accompanied by an adult to get into a bar". Being confronted with this fact "Mary is 10 and drinks a Coke in the bar", you should easily conclude that there is an adult with Mary. Disjunctive syllogism is represented by this logical formulation:

$$\frac{[A \vee B] \sim A}{B} \quad (5.1)$$

In the context of psychology of reasoning, this task is called "Denial inference". People seem to derive the correct conclusion $\pm 80\%$ of the time (Roberge, 1976a, 1976b, 1977, 1978) . This task seems also to be a little bit easier when the disjunction is considered as exclusive (with the exception of Johnson-Laird, Byrne, & Schaeken (1992) study in which the proportion of correct deduction was .48). The discrepancy between Roberge and Johnson-Laird et al. data can reasonably be attributable to a difference in the participant sample. Roberge participants were undergraduate students, while Johnson-Laird et al. participants were drawn from the general population.


5.2 Effect of the number of Predicate's arguments

5.2.1 Effect of the number of predicate's arguments on INFERNET

INFERNET LTM


Figure 5.4 shows the knowledge encoded in INFERNET which deals with disjunctive syllogism. When one term of the disjunction is negated the other term is activated. Remember that running INFERNET involves two phases. The first phase is necessary to learn the bindings. A particular content will be linked to a role. For example, a major premise is presented Or (A B) with 10 ms separating “Or” and “A” and 25ms separating “Or” and “B”. From the diagram, it can be seen that “Or” will activate “1st Term” 10 ms later and “2nd Term” 25 ms later. Nodes composing the object “A” will fire in synchrony with nodes composing “1st Term” and their connection weights will increase with a delay corresponding to Δt_γ the delay in ms corresponding to a gamma wave. The same process will happen between “2nd Term” and “B”. The second phase in INFERNET is a question-and-answer phase. A minor premise is presented to INFERNET and a response will follow. If negation precedes “A” by 5 ms, the “A” nodes will excite “1st Term” nodes. After a while these “1st Term” nodes will fire in synchrony with “A” nodes. From the conjunction of Negation nodes and “1st Term” nodes, an AND-gate will open and “2nd Term” nodes will receive activation. These nodes will oscillate 20 ms later than A nodes. Since “2nd Term” nodes have been associated with “B” nodes, “B” nodes will fire also and will constitute the network response.

Figure 5.4
The encoded
INFERNET
knowledge able to
deal with disjunctive
syllogism.



In this experiment, the object which will constitute the terms of the disjunction will be predicates with 1 or 2 arguments. These predicates are encoded as Figure 5.5 shows. Each argument of a predicate must be linked with the predicate. For example, the predicate “Pay” excites 2 role arguments: the amount to be paid and the receiver of the money.

Figure 5.5
Encoding of
predicates with their
role arguments



Hypothesis

The hypothesis that follows from INFERNET's structure is as follows: The more arguments predicates involved in a reasoning episode have, the poorer the performance. Since every predicate and its arguments occupy a distinct phase or window of synchrony, the more arguments a predicate has, the more phases will be required. If more phase are required, the chances for overlapping of phases will increase. The result will be that the different objects to be distinguished will be less separated from each other. It will happen that some nodes constituting an argument will fire in the wrong window of synchrony and will be bound to the wrong object.

Design

The design of INFERNET experiment has 1 between-subjects and 1 within-subjects variable. The between-subjects variable is the number of predicates' arguments. The within-subjects variable is the 2 questions asked to INFERNET (negation of the first term and negation of the second term).

Material


Two major premises were constructed. They differ according to the number of predicate arguments involved. The first major premise: *Or [$> (14)$, accompany (adult)]* (i.e., "It is required to be over 14 years old or be accompanied by an adult to get into a bar") is assigned to the "5 phases" group because this major premise requires 5 windows of synchrony. The second major premise is: *Or [Pay (100, gatekeeper), Present (clients, owner)]* (i.e., "It is required to pay 100 francs to the doorman or to present two clients to the owner to get into the bar") is the "7 phases" group rule. It requires 7 windows of synchrony. Two questions by major premise were constructed. When the question involves the negation of the first term, it will be called *forward*, when the second term is negated, it will be called *backward*. The forward questions are for the 5-phases group: *negation $> (14)$* (i.e., "Mary has not

more than 14 years old”), for the 7 phases group: negation *Pay* (100, *doorman*) (i.e., “Mary did not paid 100 francs to the doorman”). The backward question are *negation accompany* (*adult*) (i.e., “Mary do not accompany an adult”) for the 5 phases group and *negation Present* (*clients*, *owner*) (i.e., “Mary did not present 2 clients to the owner”) for the 7 phases group.

Procedure and parameters

Each object in the experiment is composed of 12 nodes connected each other with a delay of Δt_γ . In the learning phase, a rule is presented to INFERNET by making corresponding nodes fire in a particular order (Figure 5.6A). This input is repeated 10 times, which corresponds to the number of gamma cycles in a burst. The question is then presented to the system (Figure 5.6B). This serves to monitor the evolution of learning bindings. This response of the system will be used to evaluate reaction time. This process is repeated 10 times (theta wave). The firing nodes are recorded for every ms interval. The question phase follows. In this phase, one of the two terms is presented in input with negation (Figure 5.6B). The firing nodes are recorded for every ms interval. The network reaction to negation of the other term is recorded in the same manner. These data files will be used to calculate the proportion of INFERNET's correct responses.

Figure 5.6
A. Binding learning
phase input. B.
Question phase
input



Results

When the network has learned the bindings, it will be able to answer the questions correctly. Figure 5.7 shows a histogram of object node firing time following a question for the “5 phases” group. “Negation”, “>” and “14” nodes fire as input and begin to oscillate.

Activation is propagated from ">" to "1st Term" which begin to oscillate one cycle afterward in synchrony. In the meantime, activation is propagated from ">" to "age". From "negation" and "1st Term" the AND-gate opens and the activation is transmitted to "2nd term". "2nd Term" excites "Accompany" which begin to fire in synchrony. "Accompany" nodes excite "A-obj" nodes which activate "Adult" nodes in synchrony.

Figure 5.7
INFERNET reaction
to the question
"negation > 14"

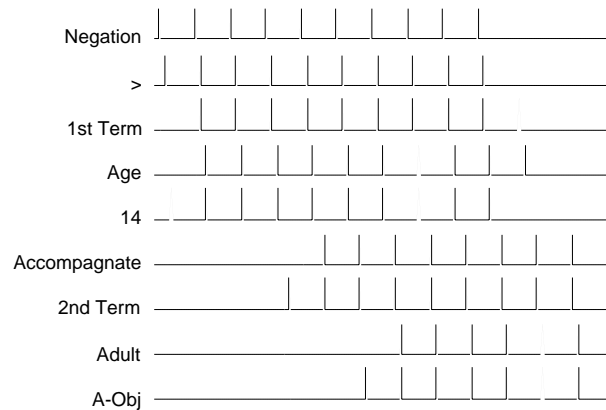
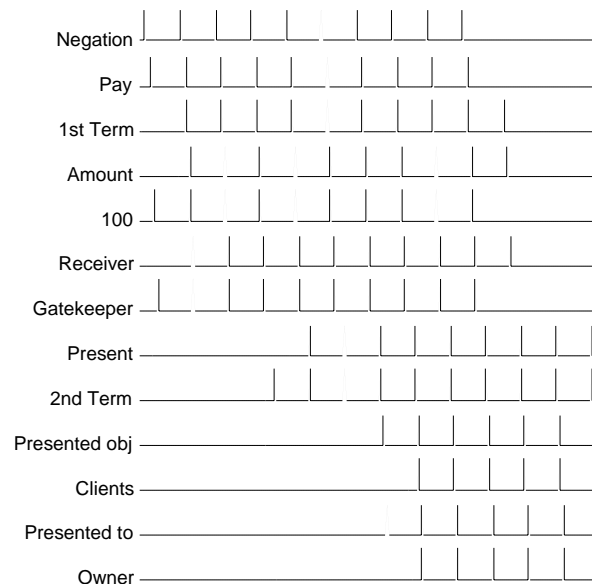


Figure 5.8 shows INFERNET's response following "negation pay 100 doorman" in the "7 phases group". The principle of activation propagation is the same as in the other group. Here, the reader can see that the lag around the different arguments of predicates are shorter. So the likelihood of confusion increases.

Figure 5.8
INFERNET reaction
to the question
"negation pay 100
gatekeeper"



The performance of INFERNET was measured as described in section 2.8. The frequency of correct inferences is reported in Table 5.1. INFERNET was run 20 times on each of the problems (5 phases and 7 phases). The overall performance of INFERNET is quite high, with a slight advantage for the 5-phase group.

Table 5.1
Frequency of
correct response in
INFERNET (n=20)

	5 phases	7 phases
Forward	14	13
Backward	14	13

Is this difference significant? According to a log-linear analysis, reported in Table 5.2, there is no significant difference between the performances of the two groups. There is no difference between forward and backward questions. Finally there is no interaction effect. Since there are only 2 questions for each trial, it is not necessary to take the repeated measure into account (see appendix B).

Table 5.2
Max likelihood chi
square for
INFERNET data

	DF	G ²	p
Independence	3	.228088	.9729319
Direction	1	0	1
Phase	1	.2280880	.6329452
Direction*phases	1	0	1

Table 5.3 reports mean and SD reaction times for both groups on the only forward question. The "7 phases" seems to be slower than the other group.

Table 5.3
Means and Standard
deviations
INFERNET

	5 phases	7 phases
Mean	4991	5993
SD	762	1182

Effects of the number of phases was analyzed by a t-test for independent samples. After an inverse (1/x) transformation of reaction times, variables fit the normal distribution and variances are homogeneous across groups. This transformation can be justified by the fact that the inverse of reaction time is speed. Table 5.4 shows the t-test results.

Table 5.4
effect of the number
of phase on
INFERNET speed of
processing.

Effect	DF	t	2 tailed p
Phases	38	3.472401	.0013028

As we see in Table 5.4, the difference of speed between “5 phases” and “7 phases” group is significant. We can conclude that INFERNET is sensitive to the number of predicate's arguments. In the following section a similar task will be proposed to human.

5.2.2 Experiment 2: Effect of the number of Predicate's arguments

This experiment studies the effect of the number of predicate arguments on disjunctive reasoning.

Participants and design

The experiment has a mixed design with 1 between-subject variable and 1 within-subjects variable. Forty participants received one of the two disjunctive premises in a random order and had to answer two questions in a random order. The 40 participants were undergraduate psychology majors, 22 females 18 males, median year: 3rd (around 21 years old).

Material

Two disjunctive premises were constructed. The first contained predicates with one argument: “*It is required to be over 14 years old or be accompanied by an adult to get into a bar.*”, the second involved predicates with two arguments: “*It is required to pay 100 francs to the doorman or present two clients to the owner to get into the bar.*” (In French, the two major premises involved the same number of words.) These two major premises were attributed respectively to the “5 phases” and “7 phases” groups. Two minor premises or questions were constructed for each group denying the first term and the second term. For the first group, the negated first term question was: “*Helen is 10 and is in the bar, what do you conclude?*”, while the negated second term question was: “*Mark is alone in the bar, what do you conclude?*”. For the “7 phases” group the negated first term question was: “*Helen gave 10 francs to the doorman, what do you conclude?*”, while the negated second term question was: “*Mark presented one person to the owner, what do you conclude?*”.

To verify further that the two groups' stimuli were comparable except for the number of arguments, a lexical decision task was performed on the words used in both groups. The participants were undergraduate psychology majors (14 females 6 males) mean age: 21.8 SD: 2.09. The mean and SD reaction times are displayed in Table 5.5.

Table 5.5
Means and SD
reaction times by
word present in each
stimulus

	5 phases	7 phases
Mean	765.7535	779.4844
SD	269.6622	255.0238

An ANOVA with one within-subjects factor on mean word latency for each sentence has been performed. Variances are homogeneous. There are only two levels in the within-subject factor, so the sphericity and compound symmetry assumptions are not taken into account. Table 5.6 shows the ANOVA results. As one can see, there is no significant difference between the two groups regarding the mean reaction time to decide if strings are words or not. We can be reasonably confident, therefore, that the two groups stimuli are comparable.

Table 5.6
Repeated measure
Anova on word
latency

Effect	DF	F	p
Repeated measure	1,19	.9981411	.3303154

Procedure

Each participant was seated approximately 50 cm in front of the monitor. The major premise appeared on the screen. Participants were asked to read it and to indicate when they understood it. The major premise stayed on the screen during the entire experiment. Questions appeared on the screen, one at the time and in random order. Participants had to answer each question. The computer recorded the time required for them to respond. The experimenter recorded the response. Before presenting the experimental material, participants received training exercises with the same procedure, but with an arithmetic content. The refreshing frequency of the computer monitor was 75 Hz so the timing precision is about 13 ms. In the program which was written for the experiment, all processor interrupts were blocked to be sure that no other time-shared process interfered with time measurements.

Results

Frequencies of sound responses are displayed in Table 5.7. As the reader can see, there is little difference between groups or between forward and backward inferences.

Table 5.7
Humans'
frequencies of
correct responses

	5 phases	7 phases
Forward	17	16
Backward	15	14

Are there any significant differences? Table 5.8 displays the log-linear results. None of the differences is significant: there is no effect of the number of arguments nor of the direction (forward vs. backward), the interaction is not significant. To confirm all, independence is not rejected.

Table 5.8
Max likelihood chi
square for human
data

	DF	G ²	p
Independence	3	1.453773	.6929822
Direction	1	1.1587794	.2817188
Phase	1	.2913863	.5893335
Direction*phases	1	.0077896	.9296718

Its time now to look at reaction times. Table 5.9 shows mean and SD reaction times for the two groups and the two questions. The 7 phases group appears to be slower and the backward question seems to be answered faster.

Table 5.9
means and Standard
deviations Human
data

		5 phases	7 phases
Forward	Mean	5321	7821
	SD	2770	2966
Backward	Mean	4927	6999
	SD	2441	4923

Data were analyzed by a 1 between-subjects (phases) and 1 repeated measure (direction) ANOVA. Variables did not fit the normal distribution and variances and covariances were not homogeneous. After transforming data into speed (inverse transformation), the homogeneity of variances and covariances assumption was met. Normality assumptions were also met for both forward and backward responses. There are only two levels in the within-subjects factor, so the sphericity and compound symmetry assumptions are not taken into account. ANOVA results are displayed in Table 5.10.

Table 5.10
Anova on human
reaction times

Effect	DF	F	p
Phases	1,38	18.74553	.0001050
Direction	1,38	3.30003	.0771730
Phases*Direction	1,38	.24498	.6234805

The results shown in Table 5.10 indicate a strong significant main effect of the number of predicate's arguments on the speed of processing. This is in agreement with INFERNET's results.

5.2.3 Comparison of simulator and humans data

Figure 5.9 compares the proportion of correct responses between the INFERNET simulator and human data. These proportions are similar but a statistical measure is necessary.

Figure 5.9
Comparison of
frequency





Table 5.11 shows the statistical analysis of the frequencies of correct responses. A Log-linear analysis was performed to test the differences between INFERNET and human. The Table shows that no effect involving the variable Group (which can takes two values INFERNET or human) is significant.

Table 5.11
Max likelihood chi
square comparison
between INFERNET
and human data

	DF	G ²	p
Independence	7	3.696693	.8139659
Group	1	2.027644	.1544599
Group*Direction	1	.6462109	.4214706
Group*Phase	1	.008543	.9263577
Group*Direction*phases	1	.0044276	.9469479

Figure 5.10 compares INFERNET and human reaction times. Reaction times appears to be longer for humans, especially for the “7 phases” group.

Figure 5.10
Comparison of
reaction times



The reaction times were analyzed by an ANOVA with 2 between-subjects factors (number of phases, contrast INFERNET-humans). Variables did not fit the normal distribution and variances were not homogeneous across groups. Since no transformation meets assumptions, the Brown and Forsythe (1974) procedure for multi-way ANOVA was used. Results are displayed in Table 5.12.

Table 5.12
Anova comparing
human and
INFERNET reaction
times
 $f=21.0267$

Effect	DF	DFcorr	F	p
Group	1,76	1,21	4.223011	.0525298
Phases	1,76	1,21	6.192374	.0213029
Group*Phases	1,76	1,21	3.391533	.0797017

The Table shows that effects containing the contrast INFERNET-humans are not significant (no significant group effect, no significant interaction effect).

5.2.4 Discussion

INFERNET's prediction appears to be correct: the more arguments that predicates involved in an episode of reasoning have, the more difficult the task. It is not a prediction that INFERNET alone could make; other models would also predict this effect. Symbolic models are too powerful to simulate people's difficulties with the number of predicate arguments unless one limits the number of arguments by an external mechanism. This kind of ad-hoc constraint, which is very common in cognitive symbolic models, is less plausible than a constraint that emerges from the structure of the model itself. In contrast, connectionist models like INFERNET, LISA, STAR, do possess in their structure something that constraints argument processing. STAR is a tensor-product architecture and the more arguments a predicate has the more dimensions the required array will need, and so the total number of required nodes will increase exponentially. LISA and INFERNET are based on phases and synchrony. In both of these models, predicates and their arguments need their own phases, since the overlapping of phases constrains the binding process, the more phases that are required, the more likely there will be an error. LISA and INFERNET and all synchrony based binding systems have an advantage over tensor product binding based system, in that they are motivated by neurobiological considerations. LISA and STAR have, however, an advantage over INFERNET. They have the capacity to embed predicates (i.e. a predicate with an argument that is also a predicate). The version of INFERNET presented here does not have this capability. A possible enhancement of INFERNET to do this will be discussed in the Chapter 8.

However, we only observed significant effects on reaction time and not on accuracy, where differences were not significant. This would imply that the cognitive system has a

means of simplifying the task, and that this process should take time. This process could be compared to “chunking”. This simplifying process should also be required when too many phases or separate bindings are required. However, sometimes it is not possible, and the cognitive apparatus reaches its limits. This is the case when predicates are embedded. For example in Halford et al. (1994) study, they observed that the understanding of four way statistical interactions were almost impossible.

If humans, when confronted with a large number of predicate arguments use chunking to reduce the number of concurrently activated chunks, this process should take time, and it should take less time when the chunking is obvious. The following experiment manipulates the ease of chunking predicate arguments.

5.3 Experiment 3: Reducing the number of a predicate's arguments

If arguments are coded by phases, when facing predicates with many arguments, a cognitive system may collapse or try to reduce the number of phases required to represent predicate and role bindings. In some situations, it should be easier than in others. For example, If there is two equivalent predicates with the same argument roles but bound to different contents, it should be easy to group the two instances of a particular predicate and to bind argument's roles to two contents.

To test the last prediction, we ran a conditional reasoning experiment comparing two situations. Both situations involve 2 predicates with the same number of arguments. In one situation, there is a possibility of grouping the two predicates and their arguments; in the other, this is much more difficult.

This task of conditional reasoning refers to material implication. People's inferences use four types of logical rules: Modus Ponens $\frac{p \supset q, p}{q}$, Modus Tollens $\frac{p \supset q, \sim q}{\sim p}$ (sound inferences) Denying the antecedent $\frac{p \supset q, \sim p}{\sim q}$ and Affirming the consequent $\frac{p \supset q, q}{p}$ (sound only in material equivalence). For a more detailed introduction to conditional reasoning see chapter 4.

Participants and design

The experiment has a mixed design with 1 between-subjects variable and 1 within-subjects variable. 28 participants (14 by group) received one of the two conditional premises, randomly selected, and had to answer four questions, chosen in a random order. The 28 participants were undergraduate psychology majors, 18 females, 10 males, median year: 3rd (around 21 years old).

Material

Two conditional major premises were constructed. The first contained two identical predicates with three arguments: “*If Tom gives a candy to Gus then Tom has given a candy to Lawrence*”, the second involved 2 different predicates with three arguments: “*If Tom gives a candy to Gus then Tom has stolen money from Mary*”. (In French, the two major premises involved the same number of words.) These two major premises were attributed to the “easy chunking” and “hard chunking” groups, respectively. Four minor premises or questions were constructed for each group each corresponding to the four types of inferences, they are shown in Table 5.13

Table 5.13
The four questions
and answers in each
group

	Modus Ponens	Denying the Antecedent	Affirming the Consequent	Modus Tollens
Easy chunking group	from <i>Tom gives a candy to Gus</i> infer <i>Tom has given a candy to Lawrence</i>	from <i>Tom does not give a candy to Gus</i> infer <i>Tom has not given a candy to Lawrence</i>	from <i>Tom has given a candy to Lawrence</i> infer <i>Tom gives a candy to Gus</i>	from <i>Tom has not given a candy to Lawrence</i> infer <i>Tom does not give a candy to Gus</i>
Hard chunking group	from <i>Tom gives a candy to Gus</i> infer <i>Tom has stolen money from Mary</i>	from <i>Tom does not give a candy to Gus</i> infer <i>Tom has not stolen money from Mary</i>	from <i>Tom has stolen money from Mary</i> infer <i>Tom gives a candy to Gus</i>	from <i>Tom has not stolen money from Mary</i> infer <i>Tom does not give a candy to Gus</i>

To verify that the two group stimuli were comparable excepting for predicates, a lexical decision task was performed on the words used in the 2 groups (only differing words are taken into account). The participants who did the lexical decision task were not the same as those who participated in the final experiment, but were also undergraduate psychology majors (14 females 6 males) mean age: 21.8 SD: 2.09. There was one missing value. The mean and SD reaction times are displayed in Table 5.14.

Table 5.14
Means and SD
reaction times by
words present in
each stimulus

	Easy	Hard
Mean	704.1930	696.2807
SD	151.9474	186.9257

An ANOVA within-subjects on mean word latency for each sentence was performed. Variances are homogeneous. There are only two levels in the within-subject factor, so the sphericity and compound symmetry assumptions are not taken into account. Table 5.15 shows the ANOVA results. There is no significant difference between the two groups for mean reaction times to decide if a string is a word or not. We can be more confident now that the two groups' stimuli are comparable.

Table 5.15
Repeated measure
Anova on word
latency

Effect	DF	F	p
Repeated measure	1,18	.0948364	.7616516

Procedure

Each participant was seated approximately 50 cm in front of the monitor. A major premise appeared on the screen. Participants were asked to read it and to indicate when they had understood it. The major premise stayed on the screen during the entire experiment. Questions appeared on the screen, one at the time and in random order. Participants had to answer each question. The computer recorded the time required for them to respond. The experimenter recorded the response. Before presenting the experimental material, participants received training exercises with the same procedure, but with an arithmetic content. The refreshing frequency of the computer monitor was 75 Hz so the timing precision is about 13 ms. In the program written for this experiment, all processor interrupts were blocked to be sure that no other time-shared process interfered with timing measurements.

Hypothesis

If we consider that predicates and arguments are coded by phases, each of the two major premises used in this experiment would need 8 phases or windows of synchrony (4 for each of the antecedent and consequent part of the “if, then” clause). The possible predicate representation could be for the first major premise: *If-then [give (Tom, candy, Gus), give (Tom, candy, Lawrence)]*, and for the second major premise: *If-then [give (Tom, candy, Gus), steal (Tom, money, Mary)]*. It will be possible to reduce the first major premise to: *give [Tom, candy, If-then (Gus, Lawrence)]* which need 5 phases for encoding the antecedent and consequent parts. This shortening would not be possible for the second major premise. Despite the fact that it would require some time for transforming the first major premise, the reducing of the number of phases should help.

Results

Frequencies of responses are displayed in Table 5.16. As the reader can see, there is a slightly higher rate of inference for the easy chunking group.

Table 5.16
Humans frequencies
of inference n=14

	MP	DA	AC	MT
easy	14	12	12	11
hard	12	9	8	10

Are these differences significant? Table 5.17 displays the log-linear results. The G^2 are underestimated because data were analysed with a between design. The Rasch model could not be applied because of the between-within design. Statisticians have not yet studied the application of log-linear analysis on this type of design (Lindsey, personal communication). Anova between-within results are also displayed in Table 5.17. The sphericity assumption is met. Variances across group for the four repeated measures are all homogeneous.

According to the Log-linear analysis, there is a significant difference between the easy (49 inferences) and hard (39 inferences) chunking groups. However this difference is not confirmed by an ANOVA. Since an ANOVA requires the dependant variable to at least be taken from an interval scale, we are more prone to believe a log-linear analysis, especially since G^2 is underestimated.

		DF	G^2	p	F	DF	p
Table 5.17 Max likelihood chi square and anova for human data	Independence	7	13.26572	.0659313			
	Chunking	1	5.672623	.0172317	3.591160	1,26	.0692589
	Inference	3	5.728730	.1255804	1.938983	3,78	.1301726
	Chunking*Inference	3	2.101684	.5515795	.440678	3,78	.7245547

Table 5.18 shows mean (and SD) reaction times for the two groups and the two questions. MP appears to be faster than other inferences. The easy chunking group also seems to be faster than hard chunking group.

			MP	DA	AC	MT
Table 5.18 Means and SD reaction time humans	Easy	Mean	4573.071	6778.000	5123.214	8169.93
		SD	2293.408	4952.844	3144.627	9149.461
	Hard	Mean	6137.786	9351.500	6960.429	10070.36
		SD	2942.197	5796.474	3488.050	7919.526

Effects are computed by an analysis of variance with 1 between-subjects (chunking) and 1 within-subjects (inference) variable. The sphericity assumption is violated. After transformation ($1/x$) this assumption is no longer violated. Variances across groups are all homogeneous. ANOVA results are displayed in Table 5.19.

		Effect	DF	F	p
Table 5.19 ANOVA on Reaction times		Chunking	1,26	5.186696	.0312242
		Inference	3,78	6.072524	.0009060
		Chunking*Inference	3,78	.428162	.7333792

The effect of chunking is significant. When using predicates that can be easily merged, people are faster (mean: 6161 ms) than when they cannot merge them (mean: 8130 ms). There is also a significant effect of the inference. A Post hoc Tukey test reveals that MP

(mean: 5356 ms) is significantly faster than DA (mean 8065 ms), than MT (mean: 9120 ms), but not significantly faster than AC (mean: 6042 ms) Finally, there is no significant interaction effect.

5.4 General discussion

As is described in section 5.1, people seem to be sensitive to predicate embedding. This section tried to explore whether people are sensitive to the number of predicate arguments. The first experiment showed that there is a relation between the number of arguments and the speed of processing disjunctions. The second experiment examined the possibility of reducing the number of arguments by chunking. We found that when chunking is possible, reaction times decrease.

The experiment 2 explored the number of arguments of a predicate and its effect on disjunctive reasoning. The lack of a significant difference between the frequencies of responses in this experiment can be interpreted in two ways. The processing capabilities does not depend on the number of predicate arguments. In this case, we would not expect to find differences in reaction time. The second possibility is that mechanisms exist to deal with numerous arguments and the application of these mechanisms should take time, thereby increasing reaction times but not affecting response accuracy.

What is the nature of these mechanisms? There could be numerous mechanisms of reduction. In experiment 3, we explored the possibility of phase reduction and we obtained faster reaction times for a case that seemed to be easier to reduce a-priori. This case seems simpler to reduce but also involves a double instantiation of the predicate. As we will see in chapter 6, double instantiation does increase reaction time, but not dramatically. This would suggest that the time lost in processing the double instantiation is more than recovered by the decrease in the number of phases.

6 Constraints on Multiple Instantiation

The problem of multiple instantiation concerns the ability to handle different instances of the same object simultaneously. From these two facts: “Pepin the Short was the son of Charles Martel” and “Charlemagne was the son of Pepin the Short”, one can infer that Charles Martel was the grandfather of Charlemagne. This inference requires two instantiations of “Pepin the Short”: the first in the role of son, the second in the role of father. For a connectionist model that does not use a working area receiving copies of items from a long-term knowledge base, the problem of multiple instantiation is a particularly thorny one. People are able to deal with multiple instances, unlike most connectionist models, but their performance when doing so is nonetheless reduced. On the other hand, there is no decrease in performance for symbolic models doing multiple instantiation. A good cognitive model should reflect both human competence, and human limits. Different connectionist solutions to the problem of multiple instantiation are reviewed and their merits examined. The INFERNET solution is then presented. It does not separate the long term knowledge base from a working area and has no recourse to copies. This solution constrains the process of multiple instantiation in a way that better reflects human data. Empirical data are compared with INFERNET’s performance.

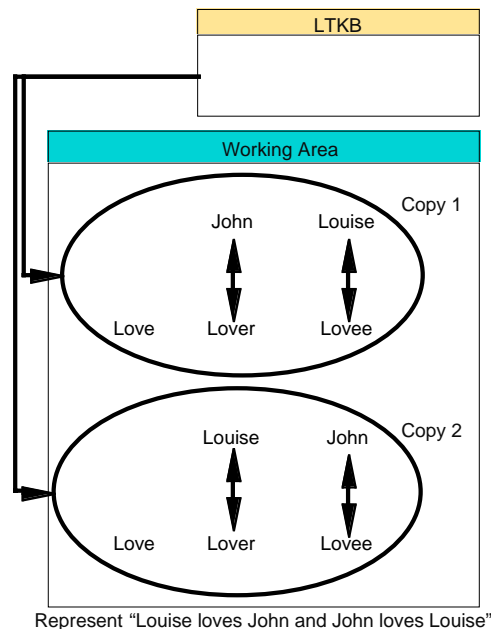
6.1 Introduction

Multiple instantiation involves the simultaneous use of the same parts of the knowledge base in different ways. Knowing that “John is in love with Louise” and that “Louise is in love with John”, one can readily infer that they should be happy. To arrive at this conclusion, one must instantiate the predicate “is in love with” and the objects “John” and “Louise” twice. Precisely how this is done is what certain authors (Barnden, & Pollack, 1991; Mani, & Shastri, 1993) call the problem of multiple instantiation. This problem is also called “the type-token problem” (Norman, 1986; Dyer, 1991). This problem is related to, but not

equivalent to, the binding problem. Even if a connectionist model solves the binding problem, it does not mean that the problem of multiple instantiation is solved. However, solving the problem of multiple instantiation requires the binding problem to be solved.


Symbolic models that load copies of pieces of knowledge into a working area before transforming them have no problem with multiple instantiation. They simply make several copies of the same content from the long-term knowledge base (LTKB) and place them in the working area (see Figure 6.1).

Figure 6.1
Straightforward
symbolic solution to
the problem of
multiple
instantiation.



However, for connectionist models that use the structure of the knowledge base itself as the place where concepts are associated and transformed, and where inferences are drawn, multiple instantiation poses a serious problem (see Figure 6.2). How can the same part of the knowledge base be associated with different roles at the same time without making several copies of the knowledge in question? Multiple instantiation also poses significant problems for distributed representations. Two closely related concepts will, in principle, share nodes. If both concepts are needed simultaneously, their common parts must be associated with two different entities.

Figure 6.2
The problem of
Multiple
Instantiation for
Connectionism



Can be interpreted as:

John loves Louise
 Louise loves John
 Louise loves Louise
 John loves John
 John loves Louise and John
 Louise loves John and Louise
 ...

The ability to handle multiply instantiated concepts assigned to different roles at the same time is required for many cognitive tasks, for example:

- Transitive inferences: Knowing that redwoods are taller than sequoias and that sequoias are taller than baobabs, a cognitive system should be able to infer that redwoods are taller than baobabs. This task requires two instantiations of the same predicate “taller” and two instantiations of the object “sequoias,” each of them assigned to two different roles, “taller object” and “shorter object”.
- Symmetric and non-symmetric inferences: From “John loves Louise” and “Louise loves Gary”, the system should be able to infer that “John is jealous of Gary.” Here again, the task involves two instantiations of the predicate “love” and two instantiations of “Louise” in the role of “lovee” in the first sentence and in the role of “lover” in the second.
- Recursion: Understanding this sentence: “The boy who kisses the girl who kisses the cat was my friend” requires recursive predication. Specifically, it needs two instances of “kiss” and of “girl,” successively assigned to the role of “kissee” and “kisser”.

6.1.1 Connectionism and Multiple Instantiation

As indicated in chapter 2, connectionist networks do not typically separate LTKB from a temporary store (or a working area), inside which copies of pieces of LTKB are loaded before transformation. In these models, transient activation of the LTKB creates a Short Term Memory (STM). For models that do separate LTKB and STM (most traditional AI

models), multiple instantiation is not a problem since the system can make as many LTKB copies as needed in STM. Lacking this copying process, neural nets suffer from “crosstalk.” (Feldman, 1982). Adding “John loves Mary” to “Gary loves Rita” can create pseudo-memories (Dyer, 1991) like “John loves Rita”. Even if we assume that John and Gary are bound to their role of lover, and Mary and Rita to their role of lovee, both men and both women remain associated with the same respective roles. What we need is a way to distinguish the two facts by separating the two identical predicates and their respective roles bindings.

The problem of multiple instantiation arises in localist networks (where each concept is represented by one node) that can represent n-ary predicates (predicates with more than one argument), if two predicate instantiations differ by more than one value of their arguments. For example, “Jack eats eggs and Jack eats fish” does not require separate instances of the predicate “eats” since this statement is equivalent to “Jack eats eggs and fish”. However, when two sets of two fillers must be associated with identical pairs of roles, the system must be able to handle two copies of the predicate and argument slots. For example, “Jack eats eggs and Mary eats fish” cannot be reduced to “Jack and Mary eat eggs and fish,” otherwise one can no longer distinguish who eats what.

The problem for distributed networks is even more difficult. Multiple instantiation problems arise as soon as one node must be used by entities that have to be differentiated. If an n-ary predicate must be represented where either predicate arguments or arguments’ fillers need to use a common node, this node will have to be bound to different roles. However, since distributed architectures represent individual objects with many nodes, the problem will only occur if a significant number of shared nodes are assigned to different objects.

6.1.2 Relevance for Cognitive Science

Is it really necessary to solve the multiple instantiation problem? This issue was raised by Donald Norman (1986) in the context of the limits of Parallel Distributed Processing (PDP) models. He concluded that PDP limitations regarding the type-token problem could be a blessing in disguise because tasks involving multiple instantiation are precisely those tasks that cause difficulties for humans. Empirical evidence of these difficulties comes from:

Reasoning

The problem of multiple instantiation arises when one role must be bound to different objects or if a single object must be bound to different roles. Compare two equivalent questions: “If John is better than Dick and John is worse than Pete, then who is the best?” and “If John is better than Dick and Pete is better than John, then who is the best?”. These two sentences are

equivalent with regard to this definition of the multiple instantiation problem, since predicates like “Worse than” and “Better than” use the same arguments, i.e. “best-object” and “worst-object”. And in fact, no significant reaction-time and correctness differences between these two situations has been found (Clark, 1969; De Soto, London, & Handel, 1965; Sternberg, 1980). However, when the number of premises increases, the number of roles bound to one object can be increased. For instance, compare this situation where each object is bound to a maximum of two roles: “A is on the right of B, C is on the left of B, D is in front of C, E is in front of A” with the following situation where one object “A” is bound to three roles “A is on the right of B, C is on the left of A, D is in front of C, E is in front of A”. People are slower and make more mistakes in the second situation (Schaeken, Johnson-Laird, & d’Ydewalle, 1996; Carreiras, & Santamaria, 1997).

Studies of reasoning have not generally been designed to test multiple instantiation effects. However, Sougné, & French (1997) did manipulate the extent of multiple instantiation and demonstrated that increasing instantiations increases reaction time.

Similarity and Working Memory

The phonological similarity between letters or between words produces replacement errors or impaired immediate serial recall in working memory (WM) tasks (Baddeley 1966a, see chapter 3). However, semantic similarity does not produce as many errors. Visual similarity among items has also been found to impair WM (Longoni, Scalisi, 1994; Walker, Hitch, & Duroe, 1993). The relevance of these effects to the problem of multiple instantiation is based on an assumption of distributed representations — namely, the more similar stimuli are, the more they will share common features and common nodes. Recalling a set of similar stimuli should involve binding more shared nodes to different groups of nodes than recalling dissimilar stimuli. Therefore, recalling similar stimuli should be more difficult. Why is immediate serial recall not impaired for semantically similar words? The absence of a semantic similarity effect could be due to the type of WM tasks. We do not know the level at which semantic properties of stimuli are processed. For example, Sougné & French (1997) found in reasoning tasks that semantic similarity does, in fact, slow down the participants’ reaction time.

Similarity in perception

Perceptually similar objects should as well share something in the neurobiological level. Consequently, the perception of similar objects should require dealing with multiple instantiation. Treisman & Gormican (1988) showed that discriminating objects is more difficult if objects share common features. Ivry & Prinzmetal (1991) observed more illusory correlations between similar than dissimilar objects.

Repetition blindness

When a word is included twice inside a list of rapidly (less than 150 ms) presented words, humans tend to not report the second occurrence of the repeated word. Kanwisher (1987) called this repetition blindness. This effect has been observed with different types of stimuli (Kanwisher, Driver, & Machado, 1995; Hochhaus, & Johnston, 1996; Bavelier, & Potter, 1992), but the most intriguing effect has been shown by Harris & Morris (1999) and Morris & Harris (1997). They showed that repetition blindness can occur between different words if they share letters. The second occurrence of repeated letters can be omitted and the unreported letters of the second word can be combined to generate illusory words. For example, the stimulus “sleep creep azy” can be reported as “sleep crazy”. Repetition blindness has also been found to occur at the semantic level (MacKay, & Miller, 1994).

Animal studies

Terrace (1991) trained pigeons to reproduce a particular pecking sequence on 5 keys. The visual appearance of the keys was manipulated. After training, the time for the pigeons to reproduce the sequence was recorded. When the five keys are distinguishable by color, the time to reproduce the sequence is longer than when some keys are distinguishable by color and others by form. In the first case, the five keys share something in common. They are more similar than when 3 keys involve colors and two keys involve forms. Terrace (1991) manipulates the ease of chunking and shows that if there is no possibility of grouping colors on one side and forms on the other side it takes more time for the pigeon to reproduce the sequence. Terrace & Chen (1991) experiment 1, using the same materials, showed that it takes more time to train pigeons with homogeneous list of colors (increased similarity) than with a list involving a segregated list of colors and forms.

These data show that multiple instantiation can indeed cause problems for humans and animals. But they also show that the cognitive apparatus has developed the means of dealing with the problem. Confronted with multiple instantiation, people tend to be slower or to make more mistakes. Consequently, a good model should not only be able to deal with multiply instantiated concepts, but should also reflect the difficulties that humans have with them.

6.1.3 Connectionist Solutions

The solutions that connectionists have developed fall into 3 categories. The first type uses two systems, a LTKB and a working area into which copies of pieces of LTKB are loaded.

The second type of solution relies on several copies of the same concepts in the LTKB. The third uses different frequencies of oscillation.

Multiple copies loaded in a working area

This is the straightforward solution borrowed from symbolic models in which each additional instance of an entity that is required will be represented by an additional copy of this entity inside some sort of short term memory store (see Figure 6.1). Bookman and Alterman (1991) combine a localist semantic network that stores relationships between concepts with a distributed network of semantic features that determines which schema slots will get filled. Each combination of a concept instance and its associated role will lead to a new schema containing the information. Each additional instantiation of the concept will be stored in a new schema.


Another model, COMPOSIT (see Barnden, 1991, 1992, 1994; Barnden & Srinivas, 1991) uses two systems, a LTM and a WM, both of which are connectionist networks. In this model, WM is composed of several registers filled with activation patterns from LTM. Concept instances and their roles are stored in registers. Each additional instantiation is stored at a particular location in WM. Predicates, related roles and fillers are stored contiguously and thus can be linked to a particular role pertaining to another predicate. This allows the embedding of role binding required, for example, for recursive reasoning. This solution is probably the most powerful one and allows a broad range of high level inferencing capabilities. The processing of multiple instantiations is still unconstrained.

In conclusion, models that separate “LTM store” and “WM store” inadequately reflect difficulties people have when they perform multiple instantiation. For these models, even if WM store has a limited capacity, it is as easy to load one copy as to fill WM with copies of the same content from LTM, unless this is prevented in an “ad hoc” manner. Other solutions have been developed, however, in which WM store is the activated part of LTM.

Multiple copies of concepts inside the LTKB

ROBIN (Lange & Dyer, 1989a; Lange & Dyer, 1989b; Dyer, 1991; Lange, 1992) separates roles from concepts. Each concept has an associated node that outputs a particular constant value (called its signature). When a role node has the same activation as that of a concept signature, this concept is bound to the role. Lange (1992) describes a potential solution designed to handle multiple instantiation inside ROBIN “in which each frame will have more than one set of conceptual and binding units.” This solution involves each concept having more than one signature and each role more than one activation (see Figure 6.3).


Figure 6.3
Handling multiple
instantiation inside
ROBIN requires that
each concept have
more than one
signature and each
role more than one
activation.



SHRUTI (Shastri & Ajjanagadde, 1993; Mani, & Shastri, 1993) is a temporal synchrony variable binding connectionist model like INFERNET. It uses node-firing synchrony to bind objects to their roles. Multiple instantiation is achieved by the use of a bounded (usually 3) set of copies or banks of predicates and their argument slots. Activation is directed to an uninstantiated copy by means of a switch. Figure 6.4 shows schematically how SHRUTI deals with multiple instantiation. This model makes psychological predictions about both STM span and multiple instantiation abilities. SHRUTI predicts that the number of instantiations will be limited and that the time required for doing multiple instantiation will be proportional to the number of predicate banks. However, it is unclear how this system could handle a situation in which more instances are needed than copies are available. This number 3 is purely arbitrary, it does not follow from internal constraints within the system. Mani and Shastri (1993) have claimed that this solution is able to handle recursive statements, but unfortunately, they do not provide a detailed description of exactly how this is done. Their solution would seem to be difficult to implement in a purely distributed manner. Shastri and Ajjanagadde (1993) describe a version of SHRUTI in which each concept is represented by a set of nodes, but the sets cannot be overlapping (i.e. one concept cannot share a common node with another concept). Without this feature, producing similarity effects related to multiple instantiation, would seem problematic.

One difficulty with the solution involving multiple copies inside LTKB is that in SHRUTI the number of required nodes increases as the square of the number of necessary instantiations. Limiting the number of such copies does, indeed, lessen the severity of this problem but seems arbitrary since it does not follow from a natural constraint.

Figure 6.4
SHRUTI uses
multiple predicate
banks. To represent
"John loves Louise"
and "Louise loves
John" the system
needs two banks.
John is bound to the
role of lover in the
first predicate bank
and to the role of
lovee in the second
predicate bank,
while Louise is
bound to the role of
lovee in the first
predicate bank and
to the role of lover
in the second
predicate bank.



Period doubling

While some people may feel that no progress has been made on the problem of multiple instantiation since the early nineties, the following architecture may lead to renewed interest in this important problem. As we saw in chapter 2, INFERNET is a temporal synchrony variable binding distributed connectionist model. It achieves variable binding through temporal synchrony of node firing. In short, when one node fires in synchrony with another, they are temporarily bound together. It has a limited STM span and the content of STM is maintained by oscillations. In INFERNET, nodes pertaining to a doubly instantiated object will support two oscillations while those singly instantiated will support one oscillation. This makes doubly instantiated object nodes fire twice while singly instantiated fire once. This means that each new instance will occupy a new place in STM to avoid crosstalk. This solution has been compared, in chapter 2, to the neural phenomenon of bifurcation by period doubling. The number of instantiations that INFERNET can support is limited to 2 or 3 by STM storage capacity and by the range of oscillation frequencies that a node can support. One advantage of this constraint is that it is derived from hypothesized neural mechanisms instead of being arbitrarily chosen. The ability to treat a greater number of instantiations will be explained in section 6.3 One current limitation of this solution, is its inability to treat embedded roles required for recursive inferencing. This problem will be discussed in chapter 8. INFERNET uses distributed representations but does not require two separate systems, one for STM and the other for LTM. In addition, it does not make use of redundant copies of concepts inside LTM.

6.2 Double instantiations

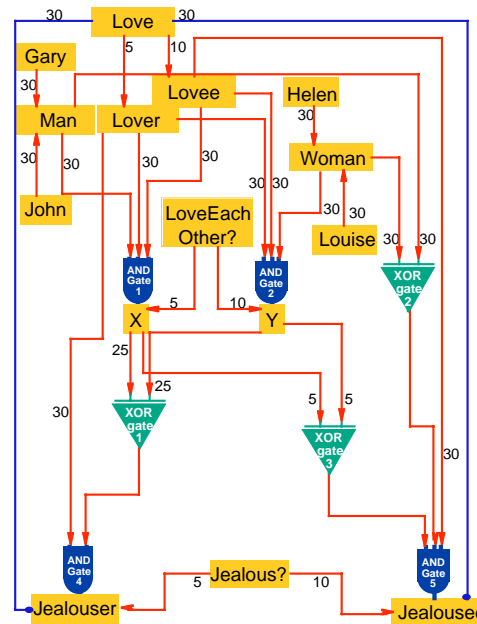
6.2.1 INFERNET's ability to handle double instantiations

Relational reasoning studies have explored how people reason with asymmetrical predicates like *bigger than*, *taller than*, *before*, *after* etc. This study will explore asymmetrical relations that, when doubly instantiated, can be symmetrical. For example: from “John loves Louise and Louise loves John” one can infer symmetry and assert that “John and Louise love each other”, from “John loves Louise and Louise loves Gary” one can infer asymmetry and assert that “John is jealous of Gary”. This kind of reasoning necessarily involves double instantiation. The capacity of INFERNET to do double instantiation will first be explored, then its performance will be compared to human data.

INFERNET LTM


INFERNET has a Long Term Knowledge Base that is used for encoding premises and answering queries. Figure 6.5 shows the knowledge necessary to make inferences about love and jealousy. Arrows represent connections; they are labeled with numbers that indicate the time required to propagate activation. Specifically, in this example, a delay of 30ms Δt_γ corresponds to the lag between two spikes of a node oscillating at 33Hz. This delay ensures that these concept-node spikes will synchronize after 30ms. INFERNET also implements AND-gates, which require all inputs to reach the target at the same time. Similarly, XOR-gates are only on when one of the inputs is active (see section 2.6).

Figure 6.5
The encoded
INFERNET
knowledge able to
deal with love
symmetry and
asymmetry.



The knowledge encoded, as shown in Figure 6.5, can correctly answer the query “Who loves each other?” and “Who is jealous of whom?” for all possible combination of two premises. In short, the connections represent the following facts: People’s love is reciprocated if one is a woman, who is both lovee and lover, and the other is a man who is also both lovee and lover; a “jealouser” is a lover whose love is not reciprocated and the person whom she/he loves, loves someone else; a “jealousee”, (the person of whom the “jealouser” is jealous) is the person who is the lovee of someone who is loved by someone else.

Figure 6.6
A. Idealized firing of nodes following “John loves Louise and Louise loves John”. B. Idealized firing of nodes following “John loves Louise and Louise loves Gary”.



Specifically and ideally, the connections depicted in Figure 6.5 will produce node firing as shown in Figure 6.6 A and B. The input “John loves Louise and Louise loves John” in Figure 6.6A causes the firing of “love”, “John”, “Louise”, “love”, “Louise” and “John” nodes in a single cycle defined by Δt_γ . All nodes fire twice in each cycle; they are instantiated twice. The nodes corresponding to “love” will enable nodes corresponding to “lover” and “lovee” to fire. The “John” nodes provoke the firing of “man” nodes and “Louise” nodes, those of “woman”. “John” nodes will fire successively in synchrony with “lover” and “lovee” nodes and connection strengths between them will increase. Shortly thereafter, the firing of “John” nodes will be sufficient to provoke the firing of both “lover” and “lovee” nodes. The same process will happen between “Louise” “lover” and “lovee” nodes. Thus the “lover” and “lovee” nodes will fire 4 times in a cycle. This is an idealized view since refractory period of nodes will prevent some of the nodes firing. From the synchrony of “lover”, “lovee” and “man” nodes firing, the AND-gate 1 will open and “X” nodes will fire in synchrony with “John” nodes and their connection strengths will increase. Similarly, from the synchrony of “lover”, “lovee” and “woman” nodes firing, the AND-gate 2 will open and “Y” nodes will fire in synchrony with “Louise” nodes and their connection strengths will increase.

The input “John loves Louise and Louise loves Gary” in Figure 6.6B provokes successively the firing of “love”, “John”, “Louise”, “love”, “Louise” and “Gary” nodes in a single cycle defined by Δt_γ . “Louise” and “love” nodes fire twice in a single cycle; they are instantiated twice. The “love” nodes will enable “lover” and “lovee” nodes to fire. The “John” and “Gary” nodes cause the firing of “man” nodes, and “Louise” nodes, those of “woman”. “John” nodes will fire in synchrony with “lover” nodes and connection strengths between them will increase. “Gary” nodes will fire in synchrony with “lovee” nodes and connection strengths between them will increase. “Louise” nodes will fire successively in synchrony with “lover” and “lovee” nodes and connection strengths between them will increase. The firing of “Louise” nodes will then be sufficient to cause the firing of both “lover” and “lovee” nodes. Thus “lover” and “lovee” nodes will fire 3 times in one cycle. From the synchrony of “lover”, “lovee” and “woman” nodes firing the AND-gate 2 will open and “Y” nodes will fire in synchrony with “Louise” nodes and their connection strengths will increase. However, AND-gate 1 will never open and “X” nodes will not fire. This state enables the opening of XOR-gate 1 and XOR-gate 3. From the combined effect of the opening of XOR-gate 1 and the firing of “lover” nodes, AND-gate 4 opens and “jealouser” nodes will fire in synchrony with “John” nodes. Their connection strengths increase. From the firing of “Man” nodes asynchronously with “woman” nodes the XOR-gate 2 opens. From the conjunction of XOR-gate 2 and 3 opening and “lovee” nodes firing, the AND-gate 5 opens and “jealousee” nodes begin to fire in synchrony with “Gary” nodes. Their connection strengths increase.

Hypothesis

Due to node firing “fatigue”, the “John”, Louise” and “Gary” nodes stop firing before “X”, “Y”, “jealouser” and “jealousee” nodes. The number of overlapping synchronies will determine how fast bindings are learned. Comparing Figure 6.6 A and B, one can see that there is less overlapping of “jealouser” and “jealousee” with their respecting instances than with “X” and “Y” and their instances. Subsequently, it will take more time for correct binding to occur in a jealousy context than in a reciprocal context.

Finding reciprocal love when the input contains symmetry (e.g. “John loves Louise and Louise loves John”) necessitates AND-gate 1 and 2. Finding who is jealous of whom when the input contains a jealous content (e.g. “John loves Louise and Louise loves Gary.”) requires for finding the jealous person, AND-gate 1 or 2, XOR-gate 1 and AND-gate 4, and for finding the “jealousee”, AND-gate 1 or 2, XOR-gate 2, XOR-gate 3 and AND-gate 5. Since the probability of error increases as the number of steps increases, finding “jealouser” and “jealousee” should be more difficult than finding reciprocal love. When there is no conclusion (e.g. finding who is jealous of whom when the input contains a symmetry or finding reciprocal love when the input contains a jealousy content), no gate is required so the response should be easy.

Design

The design of the INFERNET simulation has 1 between-subjects and 1 within-subjects variable. The between-subjects variable manipulates the symmetry. The within-subjects variable is the 2 questions asked to INFERNET (“Who loves each other?” and “Who is jealous of whom?”).


Material

Two statements were constructed. They differ according to their amount of symmetry or asymmetry. The first statement “*John loves Louise and Louise loves John*” is assigned to the “reciprocal context” group. The second statement, “*John loves Louise and Louise loves Gary*”, is the “jealousy context” group premise. Two questions per statement were constructed. The question “*Who loves each other?*” will be called the *Reciprocal question*. The question “*Who is jealous of whom?*” will be called the *Jealousy question*.

Procedure and parameters

Each symbol in the experiment is composed of 12 nodes connected each other with a delay of Δt_γ . In the learning phase, the statement is presented to INFERNET by making corresponding nodes fire in a particular order (Figure 6.7A, C). This input is repeated 10 times, which corresponds to the number of gamma cycles in a burst. The question is then presented to the system (Figure 6.7B, D) between each theta cycles. The question for reciprocal context will be the reciprocal question, for the jealousy context, it will be the jealousy question.

Figure 6.7
A. Binding learning
phase input for
reciprocal context.
B. Reciprocal
question phase
input. C. Binding
learning phase input
for jealousy context.
D. Jealousy
question phase
input.



The monitoring of the system response will be used to evaluate reaction time as explained in section 2.8. This process is repeated 10 times (theta wave). Every 1 ms interval, the nodes that are firing are recorded. The question phase follows. In this phase, the question is presented in input (Figure 6.7B, D). The reaction of the network is recorded in a file. The network reaction to the second question is recorded in the same manner. These data files will be used to calculate the proportion of correct INFERNET responses to the two questions.

Results

When the network has learned the bindings, it will be able to answer the questions correctly. Figure 6.8A shows object node firing time following the question “Who loves each other?”

for the “reciprocal context” group. “Love_each_other?” nodes fire as input and begin to oscillate. Activation is propagated to “X” and to “Y” which begin to oscillate in their own phase. “X” excites “John” whose nodes begin to fire in synchrony. “Y” excites “Louise” whose nodes fire in synchrony.

Figure 6.8
INFERNET reaction
to the questions A.
Love_each_other in
the reciprocal
context, B.
Jealous? in the
jealousy context.

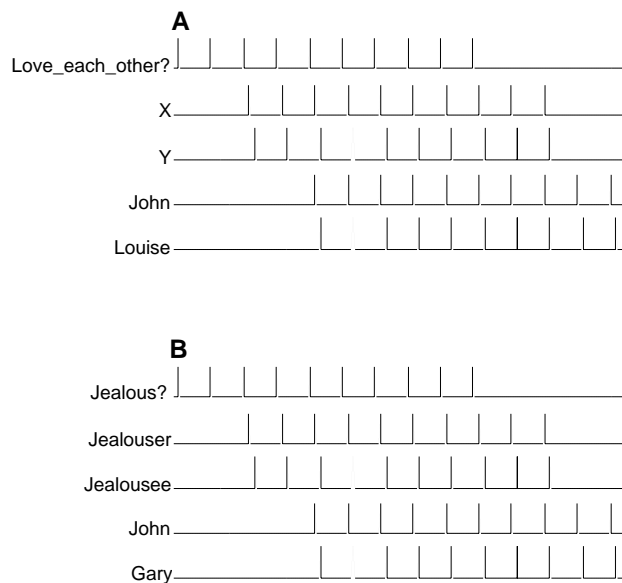


Figure 6.8B shows object node firing time following the question “Jealous?” for the “jealousy context” group. “Jealous?” nodes fire as input and begin to oscillate. Activation is propagated to “jealouser” and to “jealousee” which begin to oscillate in their own phase. “Jealouser” excites “John” whose nodes begin to fire in synchrony. “Jealousee” excites “Gary” whose nodes fire in synchrony.

The performance of INFERNET was measured as described in section 2.8. The frequency of correct inferences is reported in Table 6.1. INFERNET was run 30 times on each of the contexts (“reciprocal” and “jealousy”). The overall performance of INFERNET is quite good, except, as predicted, for the jealousy question in the jealousy context.

Table 6.1
INFERNET
frequencies of
correct responses

	Reciprocal context	Jealousy context
Reciprocal question	30	29
Jealousy question	28	16

According to a log-linear analysis, reported in Table 6.2, there is a significant difference between the two contexts (there are more correct inferences in the reciprocal context). There

is also a significant difference between the two questions (there are more correct inferences for the reciprocal question). Finally, there is no interaction effect.

Table 6.2
Log-linear analysis
of INFERNET
response
frequencies

Effect	DF	G ²	p
Independence	3	32.9947	.0000003
Context	1	14.68401	.0001271
Question	1	19.93984	.000008
Context*Question	1	.1578528	.6911438

The best fitting model is: *Context * Response + Question * Response* $G^2_{(2)}=1.944862$
 $p=.3781738$

Table 6.3 reports mean (and SD) reaction times for both groups on the reciprocal question only for the reciprocal context and the jealousy question only in the jealousy context. Here again, as expected, the jealousy question in the jealousy context is answered more slowly. The reaction times for the jealousy question in the reciprocal context and the reciprocal question in the jealousy context were not measured since INFERNET's response is nothing. Evaluating reaction times for these responses necessitates arbitrarily fixing a threshold which would indicate after how much silence will the correct answer be considered to be nothing.

Table 6.3
INFERNET reaction
times

	Reciprocal context		Jealousy context	
	Means	SD	means	SD
Reciprocal question	3692	429	-	-
Jealousy question	-	-	5283	2219

The difference between reaction times was tested by an ANOVA. Variables did not fit the normal distribution and variances were not homogeneous. No transformation allowed the assumptions to be met. The Welch procedure as defined in Howell (1987) was used. Table 6.4 shows the results.

Table 6.4
Welch procedure
test for the analysis
of INFERNET
reaction times

Effect	DF	F	p
context	1,31	14.86066	.0005458

The difference is significant. As expected, it takes more time to answer the jealousy question in the jealousy context than to answer the reciprocal question in the reciprocal context.

This experiment showed that INFERNET can deal with double instantiations. It also shows that INFERNET is sensitive to the length of the reasoning chain.

6.2.2 Experiment 4: Human abilities to treat double instantiation

The following experiment will test the same task on humans, it will provide a basis for evaluating INFERNET.

Participants and design

The experiment has a 1 between-subjects, 1 within-subjects design. Sixty participants received one of the two context premises. Both the question and premises were randomly chosen. The conditions differed according to symmetry of the expressed relationship. The 30 participants in each group were undergraduate psychology majors, 20 females and 10 males in each group, mean age: 21.73 and 21.6, (SD were 2.45, 2.19) respectively.

Material

Two relational statements were constructed for the two groups, manipulating relation symmetries. For the non symmetrical group “John loves Louise” and “Louise loves Gray” (jealousy context), for the “symmetrical group”: “John loves Louise” and “Louise loves John” (reciprocal context). Two questions were constructed: “Whose love is reciprocated?” and “Who is jealous of whom?”

Procedure

Each participant was seated approximately 50 cm in front of the monitor. The major premise appeared on the screen. Participants were asked to read it and to indicate when they had understood it. The major premise stayed on the screen during the entire experiment. Questions appeared on the screen, one at the time and in random order. Participants had to answer each question. The computer recorded the time required for them to respond. The experimenter recorded the response. Before presenting the experimental material, participants received training exercises with the same procedure, but with an arithmetic content. The refreshing frequency of the computer monitor was 75 Hz so the timing precision is about 13 ms. In the program written for the experiment, all processor interrupts were blocked to be sure that no other time-shared process interfered with timing measurements.

Results

Frequencies of correct responses are displayed in Table 6.5. As the reader can see, the less well answered question is the jealousy question in the jealousy context.

Table 6.5
Humans'
frequencies of
sound responses

	Reciprocal context	Jealousy context
Reciprocal question	30	26
Jealousy question	25	21

Are there any significant differences? Table 6.6 displays log-linear results. There is a significant difference between the two contexts and the two questions. There are more correct inferences in the reciprocal context and more correct inferences for the reciprocal question. Finally the interaction is not significant.

Table 6.6
Log-linear analysis
of humans'
response
frequencies

Effect	DF	G ²	p
Independence	3	14.20393	.0026439
Context	1	4.559318	.0327401
Question	1	7.11408	.0076481
Context*Question	1	2.778773	.0955311

The best fitting model is: *Context * Response + Question * Response*: $G^2(2)=3.027266$ $p=.2201246$.

Table 6.7 reports mean (and SD) reaction times for both groups on both questions. The reciprocal question in the reciprocal context is answered more rapidly than other questions.

Table 6.7
Humans' means and
SD reaction times

	Reciprocal context		Jealousy context	
	Means	SD	means	SD
Reciprocal question	3880	2216	7643	4692
Jealousy question	5607	2778	6557	3761

Data were analyzed by an ANOVA with 1 between-subjects variable (context) and 1 repeated measures variable (question) ANOVA. Variables did not fit the normal distribution and variances and covariances were not homogeneous. After transforming data into speed (inverse transformation), homogeneity of covariance assumption was met. Normality assumptions were also met. There are only two levels in the within-subjects factor, so the sphericity and compound symmetry assumptions do not apply. ANOVA results are displayed in Table 6.8.


Table 6.8
ANOVA on humans'
reaction times

Effect	DF	F	p
Context	1,58	11.06482	.0015307
Question	1,58	5.14168	.0271001
Context*Question	1,58	12.35669	.00086

The results shown in Table 6.8 indicate an effect of context, an effect of question and an interaction effect. Symmetrical relations induce faster responses. The reciprocal questions are answered faster. But the interaction indicates that, within the same context, questions matching the context are answered faster than questions that do not match the context and whose response is “nothing”.

6.2.3 Comparison between simulator and human data

Figure 6.9 compares the proportion of correct responses between INFERNET simulator and human data. These proportions appear to be similar but a statistical measure is necessary.



In Table 6.9, are displayed the statistical analysis concerning frequencies of correct responses. The Log-linear analysis was performed to test the differences between INFERNET and human on response correctness. The table shows that no effect contrasting INFERNET and human (group) is significant.

Table 6.9
Log-linear analysis
of humans vs.
INFERNET response
frequencies

Effect	DF	G ²	p
Independence	7	29.99081	.0000958
Group	1	.039261	.8429325
Group*Question	1	3.651286	.0560259
Group*Context	1	.753362	.3854138
Group*Context*Question	1	.0000018	.9989331

Figure 6.10 displays the comparison between INFERNET simulator and human reaction times. Reaction times appear to be longer for humans, especially for jealousy context group.

The reaction times were analyzed by an ANOVA with 2 between factors (context and contrast INFERNET-humans, called “group”). Variables did not fit the normal distribution and variances were not homogeneous across groups. Since no transformation met

assumptions, the Brown and Forsythe (1974) procedure for multi-way ANOVA was used. Results are displayed in Table 6.10.

Figure 6.10
Comparison
between Human and
Infernet reaction
times




Table 6.10
ANOVA on humans
vs. INFERNET
reaction times using
the Brown &
Forsythe procedure:
 $f=68.1559$

Effect	DF	DF _{corr}	F	p
Group	1,116	1,68	2.65476	.1078646
Context	1,116	1,68	22.61689	.0000107
Group*Context	1,116	1,68	1.46408	.2304716

As Table 6.10 shows, there is no significant difference between INFERNET and human data. There is no effect of group nor group*context interaction.

Discussion

INFERNET, like humans, can deal effectively with double instantiation. Performance decreases as the length of the reasoning chain increases. As a result, these empirical data showed, as INFERNET predicted, that dealing with symmetries is easier than dealing with asymmetries. Within the same context, questions matching the context are answered faster than questions that do not match the context and whose response is “nothing”. However, even if a non matching question and context takes more time, people still find a correct response.

6.3 More than 2 instantiations

6.3.1 INFERNET ability to handle more than two instantiations

The preceding experiment showed that INFERNET is able to handle double instantiations. This experiment will examine how INFERNET can deal with more than two instances. It will also show how the number of instantiations can impair INFERNET functioning.

To encode more than 2 instances, INFERNET takes instances 2 by 2, tries to find a chunk and take another pair of instances. For example, if there are 4 premises like: “John

In this experiment, the task will be to find who are the happy persons from four premises which involve 1, 2, 4 instances combined with 0, 1, 2 symmetries.

The INERNET LTM (Figure 6.11) is slightly modified compared to that one described in section 6.2.1 in order to answer the “Who are happy?” question. “Happy” nodes will fire whenever “X” and “Y” fires with a 15 ms delay (sign of reciprocal love) and if “jealouser” nodes do not fire. In that case, “happy” will be bound to the two persons involved in the statement. Connection weights will increase, thereafter, whenever “happy” fires, these two persons’ nodes will fire. The set of inhibitory connections from “happy” nodes prevents spurious representations from arising from the period doubling of people nodes which have been found to be happy. It also enables the connection linking man, woman, lover and lovee to the previous fillers to decrease in the recapitulating phase.

To provide an idea of how the system is functioning, 4 pairs of premises will be introduced into the system by hand. The next four figures will show how the system reacts and integrates these premises. The 4 pairs of premises will be A: “John loves Louise, Helen loves Gary”, B: “John loves Louise, Louise loves John”, C: “Peter loves Barbara, Barbara loves Allan” and D: “Gary loves Helen, Helen loves Gary”. B and D should enable “happy” to bind to the objects, C should enable “jealousy” and A should not modify anything.

Figure 6.12
A. Binding phase
encoding “John
loves Louise” and
“Helen loves Gary”.
B. Recapitulatory
phase following
binding phase,
nothing binds to
“happy”, “jealouser”
and “jealousee”.

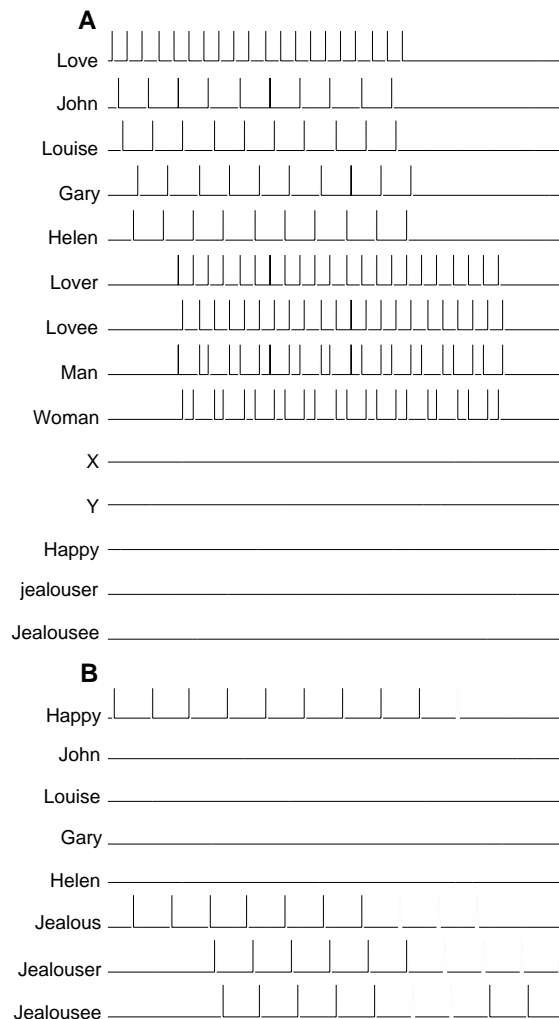


Figure 6.12A illustrates how INFERNET deals with the first premise “John loves Louise, Helen loves Gary”, “John” and “Helen” correctly bind to “lover” and “Louise” and “Gary” to “lovee”. Since none of them is both bound to “lover” and “lovee” the reasoning chain is blocked. Figure 6.12B shows how INFERNET recapitulates what has been encoded

in the binding learning phase. During this phase, forgetting occurs, the connection weights between the people nodes and “lovee” and “lover” nodes decrease to be ready for the next pair of premises.

Figure 6.13A illustrates how INFERNET deals with the second premise “John loves Louise” and “Louise loves John”, “John” and “Louise” both bind to “lover” and “lovee”. Since AND-gates 1 and 2 open, “X” and “Y” begin to oscillate which cause the opening of AND-gate 3. “Happy” begins to oscillate in synchrony with “John” and “Louise”. Figure 6.13B shows how INFERNET recapitulates what has been encoded in the binding learning phase. During this phase, forgetting occurs, the connection weights between the people nodes and “lovee” and “lover” nodes decrease to be ready for the next pair of premises. Since links between “happy” and “John” and “Louise” are used, they are maintained.

Figure 6.13
A. Binding phase
encoding “John
loves Louise” and
“Louise loves John”.
B. Recapitulatory
phase following
binding phase,
“John” and “Louise”
bind to “happy”.




Figure 6.14A illustrates how INFERNET deals with the third premise “Peter loves Barbara” and “Barbara loves Allan”, “Peter” binds to “lover”, “Allan” binds to “lovee” and “Barbara” binds to both “lover” and “lovee”. Only AND-gate 2 opens, and “Y” begins to

oscillate. Since “X” does not oscillate, XOR-gates 1 and 3 open, and since “woman” and “man” are not in synchrony XOR-gate 2 opens. AND-gate 4 will open and “lover” and “jealouser” will oscillate in synchrony with “Peter” while AND-gate 5 opens and “jealousee” nodes fires in synchrony with “Allan”. Figure 6.14B shows how INFERNET recapitulates what has been encoded in the binding-learning phase. Note that INFERNET did not forget what it learned in previous premises. “Happy” is still firing in synchrony with “John” and “Louise”. During this phase, forgetting occurs, the connection weights between the people nodes and “lovee” and “lover” nodes decrease to be ready for the next pair of premises. Since links from and to “happy”, from and to “jealouser” and from and to “jealousee” are used, they are maintained.

Figure 6.14
A. Binding phase
encoding “Peter
loves Barbara” and
“Barbara loves
Allan”. B.
Recapitulatory
phase following
binding phase,
“John” and “Louise”
still bind to “happy”,
“Peter” binds to
“jealouser” and
“Allan” binds to
“jealousee”.



Figure 6.15A illustrates how INFERNET deals with the last premise “Gary loves Helen” and “Helen loves Gary”, “Gary” and “Helen” bind both to “lover” and “lovee”. Since AND-gates 1 and 2 open, “X” and “Y” begin to oscillate, which causes the opening of AND-gate 3. “Happy” begins to oscillate in synchrony with “Gary” and “Helen”. Figure 6.15B shows how INFERNET recapitulates what has been encoded since the beginning. INFERNET did not forget what it learned in previous premises. “Happy” is still firing in synchrony with “John” and “Louise”, “Peter” is still firing in synchrony with “jealouser” and “Allan” with “jealousee”. Moreover, “Gary” and “Helen” fire in synchrony with “happy”. This constitutes INFERNET’s final answer.

Figure 6.15
A. Binding phase
encoding “Gary
loves Helen” and
“Helen loves Gary”.
B. Recapitulatory
phase following
binding phase,
“John” and “Louise”
still bind to “happy”,
“Helen”, “Gary” bind
to “happy”, “Peter”
still binds to
“jealouser” and
“Allan” still binds to
“jealousee”.



Hypothesis

In the previous section, the description is an idealized one. Some of the nodes will not fire because of the node's refractory period. consequently, the behavior of the system will be more erratic. Due to the number of gates required, finding happy people should be easier than finding jealous and jealous. Also the LTM has been encoded for increasing the chance of finding somebody happy, if nothing indicates that they are not happy (jealous). The system performance will also be a function of the number of combinations of premises. The greater the number of pairs of premises, the greater the chance of forgetting a conclusion. But also the more pairs, the more time it will take to arrive at the final conclusion. As a result, there should be an effect of the number of instances in the premises, and an effect of the number of symmetries. The greater the number of instances, the lower the performance, the fewer symmetries, the lower the performance.

Design

The design of INFERNET experiment has 2 between-subjects variables. The between-subjects variables are symmetry and number of instances. The dependent variable is the answer to the "Who is happy?" question.

Material

Six situations involving 4 statements were constructed. They differ according to the number of symmetries and the number of instances that are involved in them. The first situation involves 1 instance and no symmetry: "John loves Louise", "P (a, b)", "Q (c, d)" and "R (e, f)", which are arbitrary predicates of the same arity; the second situation involves 2 instances and no symmetry: "John loves Louise", "Louise loves Gary", "P (a, b)", "Q (c, d)"; the third situation includes 2 instances and 1 symmetry: "John loves Louise", "Louise loves John", "P (a, b)", "Q (c, d)"; the fourth situation involves 4 instances and no symmetry: "John loves Louise", "Louise loves Gary", "Gary loves Helen", "Helen loves John"; the fifth situation includes 4 instances and 1 symmetry: "John loves Louise", "Louise loves John", "Gary loves Helen", "Helen loves John"; the last situation involves 4 instances and 2 symmetries: "John loves Louise", "Louise loves John", "Gary loves Helen", "Helen loves Gary". The question "Who is happy?" is presented after each encoding of pairs of premises.

Procedure and parameters

Each symbol in the experiment is composed of 12 nodes connected to each other with a delay of Δt_γ . In the learning phase, a random pair of statements is presented to INFERNET. This input is repeated 10 times, which corresponds to the number of gamma cycles in a burst. The

question is presented to the system by making “happy” nodes and “jealous” node fire this is the recapitulatory phase. The following pairs of statements are successively presented in a random order alternating with a recapitulatory phase. For each situation, six pairs of combinations are presented in a random order. The evolution of responses during the binding-learning phase will be used to evaluate reaction time. The last recapitulatory phase will be used to evaluate the frequency of correct responses.

Results

The performance of INFERNET was measured as described in section 2.8. The frequency of correct inferences is an addition of different responses. In the various situations there is more than one correct response, so it was decided to consider a response correct if it was not incorrect.

In the first situation: “John loves Louise”, “P (a, b)”, “Q (c, d)” and “R (e, f)” since nothing is said about the love of “Louise”, it was equally considered correct stating that nobody, “Louise” or “John” were happy.

In the second situation: “John loves Louise”, “Louise loves Gary”, “P (a, b)”, “Q (c, d)”, correct responses are “nobody”, “Louise” and “Gary”.

In the third situation: “John loves Louise”, “Louise loves John”, “P (a, b)”, “Q (c, d)”, correct responses are “John and Louise”, “John”, and “Louise”.

In the fourth situation: “John loves Louise”, “Louise loves Gary”, “Gary loves Helen”, “Helen loves John”, the only correct response is “nobody”.

In the fifth situation: “John loves Louise”, “Louise loves John”, “Gary loves Helen”, “Helen loves John”, the correct responses are “John and Louise”, “John”, and “Louise”.

In the last situation: “John loves Louise”, “Louise loves John”, “Gary loves Helen”, “Helen loves Gary”, the correct responses are whatever combination of “John”, “Louise”, Gary”, and “Helen”.

The frequencies of correct responses are reported in Table 6.11. INFERNET was run 20 times on each of the situations. The overall performance of INFERNET is quite good.

Table 6.11
Count of sound
responses in
INFERNET

	1 instance	2 instances	4 instances
no Symmetry	20	16	20
1 Symmetry	-	19	20
2 Symmetries	-	-	20

Log-linear analysis was performed (see Table 6.12) with each cell increased by .00001 otherwise Log-linear cannot be computed (too many 0's). The G^2 are overestimated because of the large number of low frequency cells. Remember that in log-linear analysis all response categories are contrasted. In this study there are very few incorrect responses. Results with Yates' correction are provided. In this design there are 6 empty cells which have been counted as structural zeros. For the test of independence, the number of DF are: $(\text{number of cells} - 1) - (\text{number of modalities in symmetry} + \text{number of modalities in instances} + \text{number of modalities in response} - \text{number of variables}) - \text{number of empty cells} = (18 - 1) - (3 + 3 + 2 - 3) - 6 = 6$.

Table 6.12
Log linear analysis
on INFERNET
response
frequencies

	DF	G^2	p	Corrected G^2	p
Independence	6	13.61160	.0342889	8.058002	.2339216
Symmetry	2	2.184897	.3353943	1.514097	.4690489
Instance	2	10.78673	.0045467	6.296063	.0429366
Symmetry*Instance	4	.0000039	.9999999	.3059293	.9894282

Log-linear results are ambiguous since corrected G^2 does not lead to the same conclusion as normal G^2 does. However, since 4 cells are empirical 0's and the values of 2 other cells are very low (1 and 4), we are more prone to believe the corrected G^2 and to conclude the independence. So, in this case, the number of instances and the number of symmetries does not influence response correctness.

Reaction times were computed by examining the evolution of binding-learning during the six pairs of premise combinations. For each of the six combinations per situation, we looked for binding to "happy" "jealouser" and "jealousee" and obtained for each situation 6 reaction times with the procedure described in section 2.8. These 6 reaction times have been added.

Table 6.13 reports mean (and SD) reaction times for the six situations. Here, as expected, reaction times increase as the number of instances increases and as the number of symmetries decreases.

Table 6.13
INFERNET reaction
times means and
standard deviations

		1 instance	2 instances	4 instances
no Symmetry	Mean	3093	5442	13680
	SD	786	1865	1196
1 Symmetry	Mean	-	4610	10792
	SD	-	993	1178
2 Symmetries	Mean	-	-	7804
	SD	-	-	1640

Due to the incomplete design, effects of instances and symmetry were analyzed with type I sum of squares which express the gain of adding an effect to the model. Variables did not fit the normal distribution and variance were not homogeneous across group. Since no

transformation met this assumption, the Brown Forsythe procedure has been performed ($f=86.402977$).

The effect of interaction was computed by an analysis of variance in which groups “1 instance 0 symmetry” and “4 instances 2 symmetries” were removed. These group are removed since interaction could not be evaluated with these groups. RT did not fit the normal distribution across groups and variance and covariance were not homogeneous. Numerous transformations were tried but none of them makes the variance homogeneous. The Brown Forsythe procedure was performed ($f=59.1174$).

Table 6.14
Effects of the
number of instances
and symmetry on
INFERNET reaction
times

Effect	DF	DF corr	F	p
Conjunct effect	4,115	4,86	207.04	2.89 E-43
Instance	2,115	2,86	328.45	5.41 E-41
Symmetry gain	2,115	2,86	85.62	3.46 E-21
Instance*Symmetry	1,76	1,59	11.5948	.0011952

There is a significant effect of the number of instances and the number of symmetries. There is also a significant interaction effect: There is a greater decrease in reaction time when adding a symmetry in the 4-instance case than in the 2-instance case.

Post Hoc analysis were evaluated by planned comparisons which is a better procedure for incomplete designs (Tukey only applies to complete designs). Within the 2 instances cases, the presence of 1 symmetry reduces the reaction time with a low significant level: $F(1,114) = 3.926615$, $p = .0499345$. Within the 4 instances cases, the no-symmetry reaction times are longer than 1-symmetry $F(1,114) = 47.23826$, $p < .000001$ and longer than 2-symmetries $F(1,114) = 195.5868$, $p < .000001$ finally 1 symmetry reaction times are longer than 2 symmetries $F(1,114) = 50.58376$, $p < .000001$. Within the no-symmetry cases, 1-instance reaction times are faster than 2-instances $F(1,114) = 31.26194$, $p < .000001$, 1-instance reaction times are faster than 4-instances $F(1,114) = 635.0321$, $p < .000001$, and 2-instances are faster than 4-instances $F(1,114) = 384.4974$, $p < .000001$. Within the 1-symmetry cases, 2-instances reaction times are faster than 4-instances $F(1,114) = 216.595$, $p < .000001$.

In conclusion, INFERNET is sensitive to both the number of instances and the presence of symmetry.

6.3.2 Experiment 5: Human ability to handle more than two instantiations

The following experiment will test the same task on humans, it will provide a basis for evaluating INFERNET adequacy.

Participants and design

The experiment has a 2-way between-subjects design. 120 participants received one of the six sets of four randomly chosen premises and had to answer one question. The conditions differed according to the number of symmetries of the relationship involved and the number of instances. The 20 participants in each group were undergraduate psychology major, 83 females (13, 14, 12, 15, 14, 15 in the 6 groups), and 37 males (7, 6, 8, 5, 6, 5 in the 6 groups). Mean age was 21.86 (respectively 22.05 22.55 22.10 21.05 21.9 21.5) and SD was 2.22.

Material

Six sets of 4 relational statements were constructed for the six groups, manipulating the relational symmetries and instantiations. They are displayed in Table 6.15. One question was presented to the participants: “Who is happy about love?”

Table 6.15
Summary of the
design and material
of experiment 6

	1 instance	2 instances	4 instances
no symmetry	Peter is in love with Mary The cat ate the steak The car crashed into the wall The swallow came in the spring	Peter is in love with Mary The cat ate the steak Barbara is in love with Peter The car crashed into the wall	Peter is in love with Mary Barbara is in love with Peter Allan is in love with Barbara Mary is in love with Allan
1 symmetry		Peter is in love with Barbara The cat ate the steak Barbara is in love with Peter The car crashed into the wall	Peter is in love with Mary Barbara is in love with Allan Allan is in love with Mary Mary is in love with Peter
2 symmetries			Allan is in love with Mary Mary is in love with Allan Peter is in love with Barbara Barbara is in love with Peter

In French, each major premises involved the same number of words. To verify further that the six groups of stimuli were comparable, except for symmetries and instances, a lexical decision task were performed on the words used in the 6 groups (only differing words were tested). The participants were not the same, but also undergraduate psychology major (14 females, 6 males) mean age: 21.8 SD: 2.09. The mean (and SD) reaction times are displayed in Table 6.16.

Table 6.16
Means and SD
reaction time by
word present in each
stimulus

		1 instance	2 instances	4 instances
no Symmetry	Mean	751.8066	743.3882	720.3500
	SD	183.0917	153.3973	116.9653
1 Symmetry	Mean	-	744.1147	719.5781
	SD	-	156.6683	120.6780
2 Symmetries	Mean	-	-	720.3500
	SD	-	-	116.9653

A one-way within-subjects ANOVA (6 levels) on mean word latency for each set of premises was performed. Variances were homogeneous. Sphericity was not met so the Box procedure was performed, the correction factor is $\hat{\epsilon}=.242456$.

Table 6.17 shows the ANOVA results. As one can see, there is no significant difference between the six levels regarding for the mean reaction time to decide if strings were words or not. We are more confident now that the two groups stimuli are comparable.

Table 6.17
Anova on word
latency

Effect	DF	F	p
Repeated measure	1,23	1.101902	.304754

Procedure

Each participant was seated approximately 50 cm in front of the monitor. The major premises appeared on the screen. Participants were asked to read them and to indicate when they had understood it. The major premises stayed on the screen during the entire experiment. Questions appeared on the screen. Participants had to answer them. The computer recorded the time required for them to respond. The experimenter recorded the response. Before presenting the experimental material, participants received training exercises with the same procedure, but with an arithmetic content. The refreshing frequency of the computer monitor was 75 Hz so the timing precision was about 13 ms. In the program written for the experiment, we took care to block all processor interrupts to be sure that no other time shared process interfered with timing measurements.

Results

Frequencies of responses are displayed in Table 6.18. As the reader can see there is not much difference between group.

Table 6.18
Humans'
frequencies of
sound responses

	1 instance	2 instances	4 instances
no Symmetry	20	17	17
1 Symmetry	-	16	18
2 Symmetries	-	-	20

Are there any significant differences? Table 6.19 displays log-linear results. Since there were many low frequency cells a Yates' correction was performed. In this design there are 6 empty cells which have been counted as structural zeros. For the test of independence the number of DF are: $(\text{number of cells} - 1) - (\text{number of modalities in symmetry} + \text{number of modalities in instances} + \text{number of modalities in response} - \text{number of variables}) - \text{number of empty cells} = (18 - 1) - (3 + 3 + 2 - 3) - 6 = 6$. Even without the correction

independence is not rejected, which means that there is no significant difference between groups.

Table 6.19
Log-linear analysis
of human response
frequencies

	DF	G ²	p	G ² corrected	p
Independence	6	11.18376	.0828605	10.04552	.2618898

Table 6.20 reports mean (and SD) reaction times for the six situations. Here, as predicted by INFERNET, the reaction times increase as the number of instances increases and as the number of symmetries decreases.

Table 6.20
Humans reaction
times means and
standard deviations

		1 instance	2 instances	4 instances
no Symmetry	Mean	4708	5756	12759
	SD	3066	3303	11034
1 Symmetry	Mean	-	5484	9878
	SD	-	3597	6818
2 Symmetries	Mean	-	-	5044
	SD	-	-	3130

Due to the incomplete design, effects of instances and symmetry were analyzed with type I sum of squares which express the gain of adding the effect to the model. Variables did not fit the normal distribution and variances were not homogeneous. After inverse transformation (1/x), variances were homogeneous across groups. Variables were close to a normal distribution.

The effect of interaction was computed by an analysis of variance in which groups “1-instance no-symmetry” and “4-instances 2-symmetries” were removed. These group are removed since interaction could not be evaluated with these groups. After an inverse transformation (1/x), variances were homogeneous across groups.

Table 6.21
Effects of the
number of instances
and symmetry on
human reaction
times

Effect	DF	F	p
Conjunct effect	4,110	6.80	.0000618
Instance	2,110	6.13	.0028926
Symmetry gain	2,110	7.47	.0008748
Instance*Symmetry	1,72	.06453	.8001953

Similarly to INFERNET, there is a significant effect of the number of instances and the number of symmetry, but there is no interaction effect. Post Hoc analysis were evaluated by planned comparisons which is a better procedure for incomplete designs (Tukey only applies to complete designs). In the 4-instance cases, the no-symmetry reaction times are longer than for 2-symmetries $F(1,109) = 12.24739$, $p = .000676$ and 1-symmetry reaction times are

longer than 2-symmetries $F(1,109) = 8.10988$, $p = .005263$. In the no-symmetry cases, 1-instance reaction times are faster than 4-instances $F(1,109) = 15.78523$, $p = .0001277$ and 2-instances are faster than 4-instances $F(1,109) = 6.383726$, $p = .0129529$. In the 1-symmetry cases, 2-instances reaction times are faster than 4-instances $F(1,109) = 7.874046$, $p = .0059412$.

In conclusion, as INFERNET, people are sensitive to both the number of instances and the presence of symmetry.

6.3.3 Many instantiations: Comparison between simulator and humans

Figure 6.16 compares the proportion of correct responses between the INFERNET simulator and human data. These proportions are high in both group.




Table 6.22 displays the log-linear analysis. The Log-linear analysis was performed to test effect of the contrast INFERNET-human (called group) on response correctness.

For test of independence the number of DF is: $(\text{number of cells} - 1) - (\text{number of modalities in group} + \text{number of modalities in symmetry} + \text{number of modalities in instances} + \text{number of modalities in response} - \text{number of variables}) - \text{number of empty cells} = (36 - 1) - (2 + 3 + 3 + 2 - 4) - 12 = 17$.

The G^2 are overestimated because of numerous low frequency cells. G^2 with Yates' correction are provided.


As one can see, there is no significant difference between INFERNET simulator and human data.

Table 6.22
Log-linear analysis
of humans and
INFERNET response
frequencies

Effect	DF	G ²	p	G ² corrected	p
Independence	17	27.98497	.0451141	17.97059	.3907511
G*Inst	2	4.789335	.0912030	2.044148	.3598479
G*Sym	2	1.963081	.3747334	1.748373	.4172013
G*Sym*Inst	4	.0000262	.9999999	.5752323	.9657758

Figure 6.17 displays the comparison between INFERNET simulator and human reaction times. They are quite similar.

Figure 6.17
Comparison
between INFERNET
and human reaction
times



Comparisons of reaction times were analysed by an ANOVA with 2 between-subjects factors (Instances + symmetries, contrast INFERNET-humans). The grouping of Instance and symmetry has been done in order to avoid an incomplete design. The purpose of this analysis is to compare INFERNET and human reaction times. The present design is sufficient to provide an answer. Reaction times did not fit the normal distribution across groups and variances were not homogeneous across groups. No transformation met homogeneity of variances so the Brown Forsythe procedure was used $f=49.9810$. Table 6.23 provides ANOVA results. As Table 6.23 shows, there is no significant difference between INFERNET and humans.

Table 6.23
Contrast of the
effects of the
number of instances
and symmetry
between INFERNET
and humans.

Effect	DF	DF corr	F	p
Group	1,223	1,50	.29575	.5889760
Instance&Symmetry	5,223	5,50	28.43089	.0000000
Group*Inst&Symm	5,223	5,50	1.36165	.2546157

Discussion

This experiment shows that the more a task requires instantiations the greater the processing time. It also shows that people reduce the complexity of processing multiple

instantiation by a kind of chunking procedure. We explored the possibility of reducing the number of instances by detecting symmetries. The presence of symmetries reduces reaction times, indicating that somehow the task is facilitated by the presence of symmetries.

6.4 Similarity and multiple instantiation

For a distributed connectionist model, multiple instantiation will also affect related symbols. Symbols that share properties most likely share something in the neurobiological substrate. The effect of multiple instantiation should be observable when related symbols are used together. The following experiment tests this hypothesis.

6.4.1 Similarity and Multiple Instantiation INFERNET

One of the key features of distributed connectionist models is that a single symbol is represented by a large set of nodes, referred to here as a cell assembly. Moreover, a single node can participate in different cell assemblies. In INFERNET, a symbol is represented by a set of nodes firing in synchrony. The distributed nature of each symbol implies that closely related symbols have nodes in common. If two related symbols are needed simultaneously, and if they cannot belong to the same window of synchrony, the nodes that they share must be instantiated twice. In the present experiment, the number of closely related symbols was manipulated. The prediction was that if the number of instantiations of shared properties increased, a replacement process would be triggered. This replacement would take time and this would be reflected in the response times.


To test this prediction, we ran a conditional reasoning experiment comparing two situations. In the first situation, there is more similarity between concepts used in the major premise than in the second situation.

This task of conditional reasoning involves material implication. People's inferences use four types of logical rules: Modus Ponens $\frac{p \supset q, p}{q}$, Modus Tollens $\frac{p \supset q, \sim q}{\sim p}$ (correct inferences), Denying the antecedent $\frac{p \supset q, \sim p}{\sim q}$ and Affirming the consequent $\frac{p \supset q, q}{p}$ (sound only in material equivalence). For a more detailed introduction to conditional reasoning, see chapter 4.

INFERNET LTM

The INFERNET LTM that deals with conditional reasoning is the same as the one presented in chapter 4. Three predicates have been added to that LTM. These predicates are encoded as shown in Figure 6.18. Each of the two arguments are linked by the predicate.

Figure 6.18
Predicates used in
the experiment.



Hypothesis

The hypothesis that follows from this INFERNET structure is that the more the objects involved in the reasoning episode are similar or share nodes, the poorer the overall performance will be.

Design

The INFERNET experiment has a 1 between-subject variable and 1 within-subjects variable with 4 levels. The between-subjects variable manipulate the similarity of symbols. The within-subject variable was comprised of 4 questions put to INFERNET (corresponding to modus ponens, deny of the antecedent, affirmation of the consequent and modus tollens).

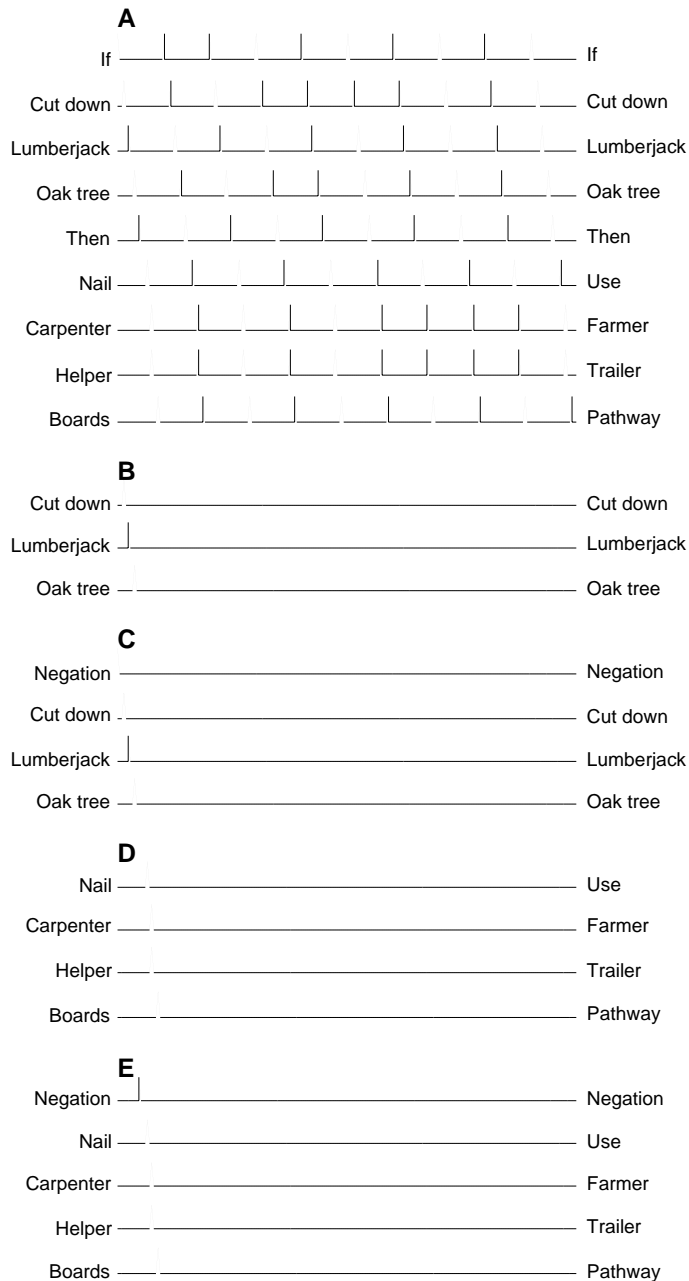
Material

Two major premises were constructed. They differ according to the similarity of symbols of which they are comprised. The first major premise: “If the lumberjack cuts down the oak tree, then the carpenter’s helper can nail the boards” involved a set of related symbols. The second premise: “If the lumberjack cuts down the oak tree, then the farmer’s tractor can use the pathway” uses fewer related symbols. Each symbol was represented by 12 nodes. In the first situation, more symbols making up the question shared nodes than in the second situation.

Procedure and parameters

Each object in the experiment is composed of 12 nodes connected to each other with a delay of Δt_γ . In the learning phase, the major premise is presented to INFERNET by causing the nodes corresponding to the objects in the premise to fire in a particular order (Figure 6.19A). This input is repeated 10 times which corresponds to the number of gamma cycles in a burst. On the left of the figure is indicated the object for the more similar group, on the right those of the less similar group. The four questions are then presented to the system (Figure 6.19B, C, D, E). The response of the system will be used to evaluate reaction time. This process is

Figure 6.19
A: Premise encoding
phase input; B:
Modus ponens
question input; C:
Deny of the
antecedent question
input; D: Affirmation
of the consequent
question input; E:
Modus tollens
question input.



Results

When the network has learned the bindings, it will be able to answer the questions related to *modus ponens*, denying of the antecedent, affirmation of the consequent and *modus tollens*.

The performance of INFERNET was measured as described in section 2.8. The frequency of inferences is reported in Table 6.24. INFERNET was run 20 times on each of the major premises (similar, less similar). The overall performance of INFERNET appear to be similar in both groups.

Table 6.24
INFERNET
frequencies of
stating each
inference

	MP	DA	AC	MT
Similar	18	18	18	18
Less similar	20	20	18	18

Is this difference significant? According to the log-linear analysis reported in Table 6.25, there is no significant difference between the two group's performance. There is no difference between the different questions. Finally there is no interaction effect. The G^2 displayed in Table 6.25 are underestimated because data were analysed with a between-subject design. The Rasch model could not be applied because of the between-within design. Statisticians have not yet studied the application of log-linear analysis on this type of design (Lindsey, personal communication). The G^2 values are probably correct because of many low frequency cells which overestimate G^2 . ANOVA between-within results are also displayed in Table 6.25. The sphericity assumption is met. Variances across groups for the four repeated measures are all homogeneous. ANOVA confirms the Log-linear analysis.

Table 6.25
Log-linear analysis
and ANOVA on
INFERNET response
frequencies

	DF	G^2	p	F	DF	p
Independence	7	7.223102	.4060516			
Inference	3	1.48052	.6867735	.6551724	3,114	.5813228
Similarity	1	1.48052	.2236932	.7600000	1,38	.3888035
Inference*Similarity	3	4.275273	.2332542	.6551724	3,114	.5813228

Table 6.26 reports mean (and SD) reaction times for both groups on the four questions. The "similar" group appears to be slower than the other group.

Table 6.26
INFERNET means
and SD reaction
times

		MP	DA	AC	MT
Similar	Mean	4527.000	4672.500	4393.500	4477.500
	SD	53.2225	215.6477	361.7687	362.0828
Less similar	Mean	4039.500	4614.000	4260.000	4431.000
	SD	493.2649	33.1504	476.5335	396.6956

Effects are computed by an analysis of variance with 1 between-subject variable (similarity) and 1 within-subjects variable (inference). The assumption of sphericity is a necessary and sufficient condition for the F to be valid. Since this assumption is violated the Box correction is used as described in appendix B $\epsilon=0.778328$. The analysis is displayed in Table 6.27.

Table 6.27
ANOVA on
INFERNET reaction
times

Effect	DF	DFcorr	F	p
Similarity	1,38	1,29	12.43043	.0014248
Inference	3,114	2,88	8.59370	.0003898
Similarity*Inference	3,114	2,88	3.55122	.0328704

Table 6.27 indicates a significant effect of similarity, a significant effect of the inference type and a small interaction effect. Similarity, which results in multiple instantiation, does affect INFERNET's performance, but it seems to affect Modus ponens more than other inferences. A Tukey post-hoc test indicates a significant difference $p=.00067$ between modus ponens reaction times in "similar" and "less similar" groups.

6.4.2 Experiment 6: Similarity and Multiple Instantiation for humans

The following experiment will test the same task on humans. It tests the effect of similarity on reasoning performance.

Participants and design

The experiment has a 1 between-subjects and 1 within-subjects design with 4 levels. Forty participants received one of the two randomly chosen major conditional premises and had to answer four questions selected in a random order. The 40 participants were undergraduate psychology majors, 26 females 14 males, (group 1, 8 males 12 females, median 3rd year; group 2, 6 males 14 females, median 3rd year;).

Material

Two conditional major premises were constructed. The first contained more related concepts than the second. The similar major premise was "If the lumberjack cuts down the oak tree, the carpenter's helper can nail the boards" the second involving less similar words was "If the lumberjack cuts down the oak tree, the farmer's tractor can use the pathway". These two major premises were respectively attributed to the "similar" and "less similar" group. Four minor premises or questions were constructed for each group each corresponding to the four types of inferences. They are reported in Table 6.28.

Table 6.28
The four questions
of each group

	Modus Ponens	Denying the Antecedent	Affirming the Consequent	Modus Tollens
Similar group	from the lumberjack cuts down the oak tree infer the carpenter's helper can nail the boards	from the lumberjack did not cut down the oak tree infer the carpenter's helper cannot nail the boards	from the carpenter's helper can nail the boards infer the lumberjack cuts down the oak tree	from the carpenter's helper did not nail the boards infer the lumberjack did not cut down the oak tree
Less similar group	from the lumberjack cuts down the oak tree infer the farmer's tractor can use the pathway	from the lumberjack did not cut down the oak tree infer the farmer's tractor cannot use the pathway	from the farmer's tractor can use the pathway infer the lumberjack cuts down the oak tree	from the farmer's tractor did not use the pathway infer the lumberjack did not cut down the oak tree

(In French, the two major premises involved the same number of words.) To verify further that the two group stimuli were comparable excepting for similarity, a lexical decision task was performed on the words used in the 2 groups (words that differed were tested). The participants were undergraduate psychology majors (14 females 6 males) mean age: 21.8 SD: 2.09. There was one missing value. The mean (and SD) reaction times are displayed in Table 6.29.

Table 6.29
Means and SD
reaction time by
word present in each
stimulus

	Similar	Less similar
Mean	775.5538	756.4038
SD	159.0371	248.5691

A one-way repeated measures ANOVA on mean word latency for each sentence was performed. Variances are homogeneous. There are only two levels in the within-subjects factor, so the sphericity and compound symmetry assumptions are not taken into account. Table 6.30 shows the ANOVA results. As one can see, there is no significant difference between the two groups for the mean reaction time to decide if strings are words or not. We can be confident now that the two groups' stimuli are comparable.

Table 6.30
Repeated measure
Anova on word
latency

Effect	DF	F	p
Repeated measure	1,19	.3303744	.5721826

Procedure

Each participant was seated approximately 50 cm in front of the monitor. The major premise appeared on the screen. Participants were asked to read it and to indicate when they had understood it. The major premise stayed on the screen during the entire experiment. Questions appeared on the screen, one at the time and in random order. Participants had to answer each question. The computer recorded the time required for them to respond. The experimenter recorded the response category. Before presenting the experimental material, participants received training exercises with the same procedure, but with an arithmetic content. The refreshing frequency of the computer monitor was 75 Hz so the timing

precision is about 13 ms. In the program made for the experiment, we took care to block all processor interrupts to be sure that no other time-shared process interfered with timing measurements.

Hypothesis

If INFERNET prediction is correct, it should be easier to perform the task when words involved in the major premise are less similar than when they are similar. Similar words should share something in their neurobiological representation. If these words are to be distinguished, their common parts must be bound to different things and must be multiply instantiated.

Results

The frequencies of responses are displayed in Table 6.31. As the reader can see, there is little difference between the two groups.

Table 6.31
Count of humans
inferences

	MP	DA	AC	MT
Similar	20	19	16	16
Less similar	20	19	19	17

Table 6.32 shows the statistical analysis concerning frequencies of inference. A Log-linear analysis was performed to test the effect of similarity on frequency of inference. The G^2 values are underestimated because data were analysed with a between-subjects design. The Rasch model could not be applied because of the between-within design. Statisticians have not yet studied the application of log-linear analysis on this type of design (Lindsey, personal communication). The G^2 are probably correct because of the presence of many low frequency cells which cause G^2 to be overestimated. Table 6.32 also displays ANOVA results on 1 between-subjects and 1 within-subjects factor with 4 levels. The sphericity assumption was violated so a Box correction ($\hat{\epsilon}=0.7810688$) on the number of degrees of freedom was performed. Variances for the 4 measures are all homogeneous across groups with the exception of AC. According to both analyses there is no effect of similarity or interaction and there is a significant effect of inference.

Table 6.32
Log-linear analysis
on humans
frequencies of
response

	DF	G^2	p	F	DF	p
Independence	7	14.18649	.0479962			
Inference	3	11.90358	.007721	4.897778	3,114	.0030725
Similarity	1	1.344195	.246296	.615385	1,38	.4376307
Inference*Similarity	3	1.014362	.797769	1.013333	3,114	.3895860

Table 6.33 shows mean (and SD) reaction times for the two groups and the four questions. The “similar” group appears to be slower and the reaction times increase from MP, DA, AC, MT.

Table 6.33
Humans' means and
SD reaction times

		MP	DA	AC	MT
Similar	Mean	4395.895	4965.000	5204.474	6521.737
	SD	1389.528	2894.142	2053.427	4529.169
Less similar	Mean	3468.053	3668.526	3843.210	4964.210
	SD	1274.845	1070.974	1486.863	2161.213

Effects are computed by an analysis of variance with 1 between-subjects variable and (similarity) 1 within-subject variable (inference) with 4 levels. The sphericity assumption is violated. After transformation of reaction times into speed ($1/x$), this assumption is no longer violated. Sphericity is a necessary and sufficient condition for the F to be valid. The ANOVA results are displayed in Table 6.34.

Table 6.34
ANOVA on humans'
reaction times

Effect	DF	F	p
Similarity	1,36	9.927828	.0032716
Inference	3,108	7.709543	.0001027
Similarity*Inference	3,108	.372422	.7730570

The results shown in Table 6.34 indicate a significant effect of similarity and inference type on the speed of processing. The more similarity, the more time it takes to draw an inference. This confirms INFERNET's hypothesis.

6.4.3 Similarity and Multiple Instantiation: Comparison between simulator and humans data

Figure 6.20 compares the proportion of inferences between INFERNET simulator and human data. These proportions are similar but a statistical measure is necessary.

Figure 6.20
Comparison
between human and
INFERNET
proportion of
inference




Table 6.35 displays log-linear results. The Log-linear analysis was performed to test differences between INFERNET and human inferences. The G^2 values are underestimated because the data were analysed with a between-subject design. The Rasch model could not be applied because of the between-within design. The G^2 are probably correct because of the numerous low frequency cells which cause the G^2 to be overestimated. Table 6.35 displays ANOVA results on 2 between-subjects and 1 within-subjects factors with 4 levels. The sphericity assumptions were violated, so a Box correction ($\epsilon=0.8696287$) on the number of degrees of freedom was performed. The variances for the 4 measures are all homogeneous across groups.


None of the differences is significant: there is no effect of the contrast between simulator and human (group). No interaction in which the variable group is involved is significant. ANOVA confirm these results.

Table 6.35
Log-linear analysis
and ANOVA on
frequencies of
inferences:
comparison between
INFERNET and
humans

	DF	G^2	p	F	DF	p
Independence	15	21.57718	.1194621			
Group	1	.173888	.6766793	.085011	1,66	.771533
Group*Inference	3	3.890072	.2735816	1.080963	2,198	.337628
Group*Similarity	1	.115133	.7343743	<.000001	1,66	>.99999
G*Inf*Sim	3	2.767412	.4289053	1.496718	2,198	.223187

Figure 6.21 displays the comparison between INFERNET simulator and human reaction times. The differences between “similar” and “less similar” groups are larger for humans than for INFERNET.

Figure 6.21
Comparison of
humans and
INFERNET reaction
times



Effects are computed by an ANOVA with 2 between-subjects variables (group and similarity) and 1 repeated measures variable (inference). The sphericity assumption is

violated. (Box correction $\hat{\epsilon}=.650408$ was used). These results must be taken with caution since the variances are not homogeneous. The results are displayed in Table 6.36.

Table 6.36
ANOVA on humans
and INFERNET
Reaction times

Effect	DF	F	p
Group	1,48	.462535	.499709
Group*Inference	2,144	8.12863	.000436
Group*Similarity	1,48	3.45526	.069188
Group*Similarity*Inference	2,144	.735571	.477671

The only significant effect is the interaction between group and inference. This means that INFERNET's reaction times do not increase monotonically from MP, DA, AC, MT, as humans' do.

6.4.4 Discussion

This experiment shows that, as INFERNET predicted, processing related symbols impairs processing. This similarity effect was already predicted for short term memory (Chapter 3). Here, multiple instantiation provides an explanation of the similarity effect on short term memory and reasoning. If objects to be stored in short term memory share common resource units and if they need to be simultaneously active and differentiated, they are multiply instantiated. A model like INFERNET is able to provide a probable explanation of this phenomenon based on rather simple neurobiological constraints.

6.5 General Discussion

This set of studies shows how a computational model can provide fruitful hypotheses about cognition. The initial constraint of having a short term memory as the activated part of the long term memory leads to questions about the mechanisms that make possible multiple instantiation. Period doubling has been hypothesized as a potential solution. The feasibility of this hypothesis and its conformity to human cognition was tested in experiment 4. The period doubling solution puts constraints on the treatment of multiply instantiated objects. The first prediction is that the difficulty of processing should be a function of the number of instances. The second prediction is: When there is too many instances, there must be a mechanism for reducing the number of instances, replacing the initial representation by an equivalent, but reduced one. Experiment 5 tested the first and the second prediction and showed that increasing the number of instances increases reaction times. This experiment also showed that when the number of instances can be reduced by the presence of symmetries, reaction times decrease. The third prediction follows from the nature of distributed representations. If the representations of objects or symbols are distributed, similar representations should share more nodes than very distinct ones. If similar objects

must be represented simultaneously but also must be differentiated, common nodes must be bound to different sets and therefore must be multiply instantiated. For this reason, similarity should impair performance. Experiment 6 confirmed the third prediction, i.e., processing similar objects increases reaction times.

Compared to other potential solutions to the multiple instantiation problem, INFERNET provides several advantages. Solutions which separate short term memory from long term memory do not provide constraints on multiple instantiation. These solutions should be rejected on the grounds that human performance is constrained by multiple instantiation. Solutions involving multiple copies of predicates and their arguments in long term memory are uneconomical and it cannot account for effects of similarity, since they tend to rely on localist representation schemes. Further, the constraints that they put on multiple instantiation are arbitrary. For example, in SHRUTI, a maximum of 3 copies of predicates are allowed. This number of copies does not emerge from an underlying mechanism and could just as well be set to 1000. By contrast, INFERNET's period-doubling solution is based on known neurobiological mechanisms and provides *emergent* constraints on multiple instantiation. INFERNET still cannot handle recursive predication, however. This will be discussed in the chapter 8.

Multiple instantiation effects would seem to be in contradiction with priming effects. Connectionist models explain priming by a pre-activation of shared units involved in the representation of successive objects. The difference between multiple instantiation and priming is the time scale. Multiple instantiation involves the *simultaneous* use of shared units while priming involves the *successive* use of shared units. INFERNET is sensitive to multiple instantiation and can still be sensitive to priming. Pre-activation of units should also facilitate successive use of them.


7 Exploring the INFERNET Simulator

In the previous chapters, INFERNET parameters were held constant. This chapter will explore INFERNET's reactions to variations in parameters.

A previous paper (Sougné, 1998b) discussed the fact that adding noise to the INFERNET simulator improved its performance. It was also discovered that INFERNET was rather resistant to the amount of overlapping distribution in the objects representation (see also sections 3.4.4 and 6.4). Figure 7.1 shows the simulation results of INFERNET. The task was to find “Whose love is reciprocated?” and “Who is jealous of whom?” when given “John loves Louise and Louise loves John” and “John loves Louise and Louise loves Gray.” The performance of the system was measured by the percentage of correct responses and by the time taken by the system to arrive at the correct bindings. In this experiment, each object was composed of 16 nodes. In the Non-overlapping condition, no object shared nodes with other objects, whereas in the overlapping condition, each object shared 4 nodes with two other objects which will never be bound to the same role. Noise was added at each time step. Experiments were composed of 20 trials for each of the four conditions.

Figure 7.1 shows that, in general, overlapping distributions reduce the percentage of correct answers, but when the response is correct, response time decreases (if there is no noise). In the task tested, most objects were doubly instantiated and the representational overlap means that more instantiations will occur. Consequently, certain nodes must be assigned to more than two windows of synchrony. For example, “Louise” nodes must be synchronized with “Lover” and “Lovee” nodes, but if “Louise” shares nodes with “Love”, these shared nodes must fire in additional windows of synchrony. Since these nodes are affected by a refractory period, some of the required spikes cannot occur and the proportion of correct answers thus decreases.

Figure 7.1
Effects of noise and
distribution. (A)
shows the
percentage of
correct responses,
(B) displays the time
taken by the system
to learn the correct
bindings.



Why does binding fixation convergence time decrease in the (noise-free) overlapping condition? Postsynaptic nodes require the conjunction of activation at a precise time to fire. If the conjunction involves input from different objects and if these objects share nodes in common, the increase in firing rate increases the chance of having a conjunction of activation that causes the firing of the postsynaptic node. On the other hand, this will also increase the number of inappropriate firings of these postsynaptic nodes. This increases response errors and the system rapidly reaches a local minimum. When noise is added, it provides a means of escaping from these local minima, thus improving the frequency of correct responses, but also increasing response times. Similar explanation for similar effect in simulated annealing has been provided by Kirkpatrick, Gelatt, & Vecchi (1983). Noise makes the system more erratic before reaching a stable point. It allows exploration of a bigger part of the space, which takes time, but improves the chance of finding the best answer.

The response of a non-linear system to weak periodic input can be enhanced by the presence of an optimized level of white noise. This phenomenon has been called stochastic resonance. Benzi, Sutera & Vulpiani (1981) were the first to describe this mechanism. They described a dynamical mechanism in which weak noise acting with a bistable model explained variations of the earth climate. The presence of white noise (i.e., purely random noise) is considered to be an important factor for the phenomenon to arise (Maddox, 1994). There are different domains in which stochastic resonance has been observed including climatology, technology of super-conductors, and biology (Wiesenfeld & Moss, 1995).

In neuroscience, various empirical studies describe stochastic resonance in sensory neuron firing. Douglass, Wilkens, Pantazelou & Moss (1993) showed that weak mechanical stimulation detection can be enhanced by an optimal noise intensity. Noise applied to a single mechanoreceptor cell of the crayfish enhances its responsiveness to weak stimuli. Levin & Miller (1996) showed similar facilitation of weak signal detection on mechanosensory system of the cricket. Bezrukov & Vodyanoy (1995) observed enhancement of transduction across ion channels by noise.

In the field of artificial neural networks, stochastic resonance has been observed by different authors (Collins, Chow & Imhoff, 1995; Stemmler, Usher & Niebur, 1995; Grossberg, & Grunewald, 1997). Stochastic resonance according to these studies enhances the detection of weak signals but not strong signals.

In the following sections, the effect of noise will be explored. Different kinds of noise will be explored. Noise will be related to different variables. Since stochastic resonance has been found to act on sub threshold signals, noise will be related to variation in INFERNET's node firing threshold.

7.1 Noise on connection weights

Noise was first applied to connection weights. These weights will have a chance to be increased or decreased.


7.1.1 Noise and Threshold

The first experiment recorded INFERNET's performance at different levels of noise for different values of the threshold Θ (see Chapter 2). For each activated connection, 2 random numbers are drawn between 0 and some maximum value. The first random number will be subtracted from the connection weight, the second will be added to the connection weight. The minimum and maximum of resulting random addition to weights were set to (0 0), (-10 20), (-20 40), (-30 60), (-40 80) and (-50 100). Remember that the minimal value of a connection weight is -127 and the maximum value is 127. When the minimum and maximum are set to 0, there is no noise on the connection weight. Table 7.1 displays the percentage of INFERNET's correct responses for a problem of finding reciprocal love from these premise: "John loves Louise and Louise loves John". The experiment was composed of 10 trials for each combination. Figure 7.2 displays a 3D plot of these data. Noise significantly improve the performance of the system which is very poor when there is no noise on connections. When noise amplitude is too high, however, there is less benefit for high values of threshold, and performance is reduced for low threshold values.

Table 7.1
Percentage of
INFERNET correct
response for
different levels of
noise and different
thresholds.

Threshold	Noise amplitude						
	0 0	-10 20	-20 40	-30 60	-40 80	-50 100	
280	30	99	93	68	30	3	
290	30	89	87	75	29	1	
300	10	99	99	79	36	7	
310	0	100	99	87	63	21	
320	0	80	89	88	75	1	
330	0	89	99	93	76	38	
340	0	80	98	97	82	33	
350	0	90	100	99	86	69	

Figure 7.2
3D plot of
percentage of
correct response for
different level of
noise on connection
weights and
different thresholds.



7.1.2 Increasing vs. decreasing connection strength with noise

When noise is added in a real neural system, it is done by producing a noisy stimulus e.g. air disturbance (Levin & Miller, 1996). This noise consists in adding something to the nervous system or increasing the overall activity of the system. This noise helps sub-threshold signals to cross the threshold (stochastic resonance). The following experiment tests the relative effect of decreasing the overall activity of the network compared to increasing the overall activity of the network. More precisely this second experiment explores the respective effects of noise which results in adding excitation to the network compared to noise which results in adding inhibition to the network (i. e. adding or subtracting a value on connection weights with noise). For each activated connection 2 random numbers are drawn between 0 and different maximum values. The first random number will be subtracted from the connection weight, the second will be added to the connection weight. The maximum of the random numbers that increase the connection weights was set to 0, 10, 20, 30, 40 and 50, and the minimum of the random numbers that decrease the connection weights was set to 0, -10, -20, -30, -40 and -50. Table 7.2, presents percentage of correct answers for each


combination of random numbers applied to addition and subtraction. Threshold was set low (280).

Table 7.2
Percentage of
INFERNET correct
response for
different levels of
noise applied to
addition or
subtraction

		Maximum adding value or amplitude					
		0	10	20	30	40	50
Minimum subtracting value	0	45	69	100	99	100	97
	-10	10	20	80	100	98	99
	-20	0	0	23	90	90	98
	-30	0	0	10	20	88	100
	-40	0	0	0	20	62	81
	-50	0	0	0	0	10	26

The problem was to find reciprocal love from these premise: “John loves Louise and Louise loves John”. The experiment was composed of 10 trials for each combination. Figure 7.3 displays a 3D plot of these data. Noise that raises weight values increases the performance of the system. Noise that lessens weight values reduces the performance unless the amplitude of adding noise is bigger.

Figure 7.3
3D plot of
percentage of
correct response for
different level of
noise applied to
addition or
subtraction on
connection weights



7.2 Does stochastic resonance depend on the problem?

7.2.1 Noise on connection and Threshold


The third experiment recorded INFERNET performance at different levels of noise for different values of the threshold Θ . This experiment is similar to the first experiment, but the problem is different. For each activated connection 2 random numbers are drawn between 0 and different maximum. The first random number will be subtracted from the connection weight, the second will be added to the connection weight. The minimum and maximum of

resulting random addition to weights were set to (0 0), (-10 20), (-20 40), (-30 60), (-40 80) and (-50 100). Table 7.3 displays the percentages of INFERNET correct response for a problem of finding jealousy from these premise: "John loves Louise and Louise loves Gary. The experiment was composed of 10 trials for each combination. Figure 7.4 displays a 3D plot of these data. Noise improves performance of the system which is very low when there is no noise on connection. When noise amplitude is too high, performance is reduced.

Table 7.3
Percentage of
INFERNET correct
response for
different levels of
noise and different
thresholds.

Threshold	Noise amplitude					
	0 0	-10 20	-20 40	-30 60	-40 80	-50 100
280	19	98	19	0	0	0
290	2	100	68	0	0	0
300	3	98	69	4	0	0
310	0	97	80	24	0	0
320	0	99	60	20	0	0
330	0	99	90	38	0	0
340	0	97	100	50	0	0
350	0	87	95	45	0	0

Figure 7.4
3D plot of
percentage of
correct response for
different level of
noise on connection
weights and
different thresholds.



Compared to the previous problem (experiment 1), there is more of a clear-cut optimal value for noise around 10.

7.2.2 Increasing vs. decreasing connection strength with noise

The fourth experiment explores the respective effects of adding or subtracting a value on connection weights with noise. For each activated connection, 2 random numbers are drawn between 0 and different maximum values. The first random number will be subtracted from the connection weight, the second will be added to the connection weight. The maximum of the random numbers that increase the connection weights was set to 0, 10, 20, 30, 40 and

50, and the minimum of the random numbers that decrease the connection weights was set to 0, -10, -20, -30, -40 and -50. Table 7.4, shows the percentage of correct answers for each combination of random numbers applied to addition and subtraction. The firing threshold was set low (280).


Table 7.4
Percentage of
INFERNET correct
response for
different levels of
noise applied to
addition or
subtraction

Minimum subtracting value		Maximum adding value or amplitude					
		0	10	20	30	40	50
0		11	74	98	88	90	100
-10		0	8	74	97	96	87
-20		0	0	3	73	95	99
-30		0	0	0	8	75	93
-40		0	0	0	0	9	74
-50		0	0	0	0	0	4

The problem was the same as in previous experiment. The experiment was composed of 10 trials for each combination. Figure 7.5 displays a 3D plot of these data. Noise raising weight values increases the performance of the system. Noise decreasing weight values reduces the performance, unless amplitude of increasing weight is bigger than the amplitude of decreasing weight.

Stochastic resonance seems to be unaffected by the problem. Similar effects were observed for the 2 different problems. Increasing weight with noise in both cases increases performance while decreasing weight randomly reduces the performance.

Figure 7.5
3D plot of
percentage of
correct response for
different level of
noise applied to
addition or
subtraction on
connection weights



7.3 Noise on spike generation

The following experiments will explore the effect of noise on node spiking. At any given time, each node has a low probability of firing, even in absence of activation. Conversely, for nodes whose activation crosses the firing threshold, there is the same low probability that it will not fire. This kind of noise is acting at a higher level than noise on connection weights.


7.3.1 Noise and Threshold

The fifth experiment explored different probabilities to force a spike or prevent a spike at different threshold levels. The experiment was composed of 3 trials for each combination. The problem was to find reciprocal love from these premises: “John loves Louise and Louise loves John”. Table 7.5 displays the percentage of correct answers. Figure 7.6 displays a 3D plot of these data.

Table 7.5
Percentage of
INFERNET correct
response for
different levels of
noise to force or
prevent a spike

Θ	Probability to force or prevent a spike												
	0	.001	.002	.003	.004	.005	.006	.007	.008	.009	.01	.011	.012
280	100	100	73	100	100	97	100	100	87	53	83	73	83
290	67	67	100	100	100	100	100	100	90	93	90	83	73
300	70	0	67	100	67	100	100	63	97	100	87	93	60
310	0	100	0	100	33	77	97	90	63	97	67	63	87
320	0	33	33	30	33	67	67	30	33	77	63	80	57
330	0	10	0	0	33	67	30	67	0	33	0	53	90
340	33	0	33	40	0	33	100	30	7	30	57	10	63
350	0	0	0	0	0	33	0	33	0	0	0	30	57

Figure 7.6
3D plot of
percentage of
correct response for
different level of
noise applied to
force or prevent a
spike



The effect of stochastic resonance is less pronounced than when noise is applied to connection weight. Noise increases performance when the threshold increases. At higher probabilities of noise, the performance decreases. In this case noise is a relatively rare and

discrete event which has an important effect. When noise is applied to connections, it has a less important effect but is applied much more widely.


7.3.2 Comparison between forcing and preventing a spike

The sixth experiment compared different probabilities of forcing a spike with various probabilities of preventing a spike. The experiment was composed of 10 trials for each combination and the threshold was set low (290). The problem was to find reciprocal love from these premises: “John loves Louise and Louise loves John”. Table 7.6 displays the percentage of correct answers. Figure 7.7 displays a 3D plot of these data.

Table 7.6
Percentage of
INFERNET correct
response for
different levels of
noise to force vs. to
prevent a spike

Probability to prevent a spike	Probability to force a spike											
	.001	.002	.003	.004	.005	.006	.007	.008	.009	.01	.011	.012
.001	90	84	100	100	99	89	98	85	91	72	74	73
.002	69	90	80	89	98	88	90	74	87	79	78	72
.003	70	90	71	88	99	90	87	74	88	84	74	78
.004	53	100	90	100	89	98	99	89	84	72	76	61
.005	60	90	100	100	89	87	95	78	90	86	70	64
.006	30	80	90	78	98	89	77	89	83	82	58	60
.007	67	100	60	70	99	86	88	97	96	84	81	69
.008	88	89	80	99	89	86	90	83	94	77	63	77
.009	80	85	70	99	100	90	95	81	93	82	86	78
.01	90	78	80	90	90	88	96	86	71	92	61	86
.011	79	60	99	88	89	79	95	82	83	74	75	67
.012	70	90	80	79	100	89	97	88	93	85	80	78

Figure 7.7
3D plot of
percentage of
correct response for
different level of
noise applied to
force vs. prevent a
spike



Performance of the system is quite erratic and stochastic resonance in this experiment is questionable. Including noise on such a high level event like a spike does not seem to improve the performance of the system. Stochastic resonance appears to require that noise be applied to a multitude of elements with a low perturbing modification. Stochastic resonance

facilitation seems to require a complex dynamical system: a system that is composed of a large number of interrelated variables.


7.4 Noise on Delay

Another type of noise was also explored. The following experiment tested the effect of noise on connection delays. The effect of noise on delays was already presented in section 3.4 to reduce discriminability of objects, thus reducing short term memory span. In this experiment various probabilities of applying noise on delays will be explored at different threshold values. The noise modifies the connection delay by adding to it a random value between -2 and 2 ms. The delay can be shortened or delayed by a maximum of 2 ms. Probabilities express the chance of adding this random value to the delay of an active connection. The experiment was composed of 10 trials for each combinaison of noise probability and threshold value. The problem was to find reciprocal love from these premises: "John loves Louise and Louise loves John". Table 7.7 displays the percentage of correct answers. Figure 7.8 displays a 3D plot of these data.

Table 7.7
Percentage of
correct answers for
different
probabilities of noise
on delay and
different threshold
values

Θ	Probability to modify connection delay																								
	0	.01	.02	.03	.04	.05	.06	.07	.08	.09	.1	.11	.12	.13	.14	.15	.16	.17	.18	.19	.2	.21	.22	.23	.24
280	62	51	31	50	13	35	37	58	39	20	18	0	12	0	16	9	17	9	0	0	0	0	0	0	3
290	37	10	10	20	12	20	30	17	11	18	10	20	0	10	11	14	10	0	0	0	0	0	0	0	0
300	19	6	26	4	10	10	0	10	0	20	0	0	9	8	0	0	0	0	0	0	0	0	0	0	0
310	10	0	20	20	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
320	0	11	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
330	0	0	10	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
340	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
350	0	0	0	0	7	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0


Figure 7.8
3D plot of
percentage of
correct response for
different threshold
and different level of
noise applied to
delays



Noise on delays does not enhance INFERNET's performance. The reason appears to be that a noisy delay does not cause node activation to cross the firing threshold. On the contrary, perturbation of delays tends to reduce the level of activation of a node. Since a node in INFERNET requires several excitatory signals arriving at the same time for it to fire, increasing noise on delays impairs the synchronous arrival of signals.

7.5 Noise and more realistic PSP curve

This last experiment will explore the effect of noise on connection weights using a Post Synaptic Potential curve which more closely resembles that of real neuron transmission. This new PSP function is plotted in Figure 7.8, together with the old PSP function. The new PSP curve gives more activation and this activation rises more slowly and decreases less sharply.



This experiment explores the respective effects of adding or subtracting a value on connection weights with noise together with the new PSP function. For each activated connection 2 random numbers are drawn between 0 and different maximum values. The first random number will be subtracted from the connection weight, the second will be added to the connection weight. The maximum of these random numbers were set to 0, 10, 20, 30, and 40, while the minimum of the random numbers that decrease the connection weights was set to 0, -10, -20, -30, and -40. The problem was to find reciprocal love from these premise: "John loves Louise and Louise loves John". The experiment was composed of 10 trials for each combination. The threshold was set low (280). Table 7.8, presents the percentage of correct answers for each combination of random numbers applied to addition and subtraction. Figure 7.10 displays a 3D plot of these data. Noise that raises weight values increases the performance of the system. Noise that decreases weight value reduce the performance unless adding weight value amplitude compensates.

Compare these data to those presented in sections 7.1.2 and 7.2.2. The difference of noise effect relies on the less pervasive effect of noise which decreases connection weights. Looking at the PSP curve in Figure 7.9 provides an explanation. The new PSP function makes a connection give more excitation and more inhibition, and this activation stays longer at a high value. If the noise adds too much to the connection weights, the associated nodes will receive more excitation and less inhibition. A lot of nodes will fire inappropriately.


Adding noise that decreases the value of excitatory weights and increases the effect of inhibitory connections will make the system reach an equilibrium. This fact does not contradict the idea that noise helps sub threshold states. It shows that transforming over threshold states into sub threshold states let stochastic resonance act.

Table 7.8
Percentage of
INFERNET correct
response for
different levels of
noise applied to
addition or
subtraction and with
the new PSP
function

	Maximum adding value or amplitude				
	0	10	20	30	40
Minimum	0	70	81	82	42
subtracting	-10	28	89	77	56
value	-20	9	16	81	50
	-30	0	0	44	74
	-40	0	0	0	20
					78

It has to be noted that this more neurally plausible PSP function does not improve the system's performance. It just requires different values of noise for reaching an optimum. The important factor in the PSP function is the decrease of activation (the leaky factor of integrate and fire neurons). This leak ensures that nodes will be sensitive to temporal factors like the precise time of convergence of excitation. This enables the precise functioning of temporal gates described in chapter 2.

Figure 7.10
3D plot of
percentage of
correct response for
different level of
noise applied to
addition or
subtraction on
connection weights
with the new PSP
function.



7.6 Discussion

This chapter provides evidence for stochastic resonance in the INFERNET simulator. Stochastic resonance seems to help sub threshold signals to cross the threshold. But we have

found that stochastic resonance also helps in reducing unwanted spurious activity. Stochastic resonance seems to appear if noise is randomly applied to a large number of variables like connections. In the above experiments there were 87552 connections. When noise is applied to a less numerous set of structures, like nodes, there is less facilitation by stochastic resonance. This shows that stochastic resonance effect is related to complex dynamical systems.

INFERNET not only preserves the connectionist quality of being resistant to perturbed activity but its performance has also been shown to be enhanced by small low-level perturbations. This means that INFERNET can process symbols while maintaining some of the advantages of distributed processing, such as resistance to perturbation and graceful degradation.

In exploring different PSP functions, our aim was to determine whether the simplification of neuron functioning in INFERNET were appropriate. The results are encouraging. A more realistic PSP function does not provide additional reasoning capabilities but significantly increases computational demands. Consequently, the level of simplification used in INFERNET would seem to be appropriate.

8 Conclusions

Based on neurobiological constraints, a cognitive model of binding has been constructed and implemented on a computer. This model is able to treat symbols and to display reasoning capabilities.

INFERNET is a connectionist model using integrate-and-fire nodes. In INFERNET, nodes can be in two different states: they can fire (“on”), or they can be at rest (“off”). A node fires at a precise moment and transmits activation to other connected nodes with some time course. When a node activation or potential reaches a threshold, it emits a spike. After firing, the potential is reset to some resting value. INFERNET solves the binding problem by synchrony. Each object is represented by a cluster of nodes firing in synchrony. If the firing distribution is tightly concentrated around the mean, the object is considered to be activated. Objects are bound to their roles by synchronous firing. This synchronous activity defines a window of synchrony: a delay within which the required nodes fire. This delay is constrained by the precision of synchrony (± 5 ms). In INFERNET, discrimination is achieved by successive windows of synchrony. Objects and bindings are maintained in memory by oscillation. Once a node is activated, it tends (but not necessarily) to begin oscillating at a γ (30-80 Hz) frequency range. The temporal gap between 2 spikes of a node is therefore from 14 to 33 ms.

The rhythmic activity and the precision of synchrony constrains the number of distinguishable entities that can be simultaneously maintained in short term memory. This constitutes the short term memory span of INFERNET. This span was studied in chapter 3 and corresponds well to human short term memory span. This constraint also applies to the treatment of predicates. In INFERNET, predicates and roles are linked by a specific temporal order. The activation of a predicate is always followed by the successive activations of its different roles, each of which is assigned to a particular window of synchrony. In chapter 5 the effect on the number of predicate arguments was found to impair reasoning both for INFERNET and for humans.

In INFERNET, the number of steps required for transmitting activation from one point to another is constrained by two factors. The first is the probability of error in transmission which increases as the number of steps increases. The second constraint concerns the lag between the starting point of the filler nodes' oscillation and the starting point of the role nodes' oscillation. The more steps that are required, the greater the lag between role and filler nodes' oscillation starting points. If the lag increases, the number of synchronizations between role and filler nodes firing decreases. In chapter 4 we showed that these constraints can explain effects of negation on reasoning.

The last constraint concerns multiple instantiation. INFERNET achieves multiple instantiation by means of period doubling. Nodes pertaining to a doubly instantiated concepts will sustain two oscillations. This means that these nodes will be able to synchronize with two different sets of nodes. Since nodes have a refractory period after firing, the lag between 2 successive spikes of a node cannot physically be reduced too far. Consequently, INFERNET's capacity to process multiple instantiations is limited. In chapter 6 we showed the limitations of INFERNET and humans in the processing of multiple instantiation.

8.1 Comparisons with other work

The advantages of INFERNET over related works involve both its neural and psychological plausibility.

In chapter 1, it was argued that the only neurally plausible mechanisms for binding were either firing rate, firing latency, phase, or firing synchrony. Among these possibilities connectionist models have only explored binding by synchrony. Other implemented connectionist solutions do not have neurobiological justifications. Among models of binding by synchrony, SHRUTI (Shastri & Ajjanagadde, 1993), LISA, (Hummel & Holyoak, 1997) are the main competitors of INFERNET. INFERNET is the only integrate-and-fire neurally based model. Only INFERNET uses a Post-Synaptic Potential function with a refractory period. As a result, the nodes in INFERNET are neurally more plausible than those in SHRUTI and LISA. Moreover, only INFERNET and LISA use distributed representation, while SHRUTI does not.

In terms of symbolic processing, symbolic models are too powerful, while classical connectionist models are not powerful enough to account for human cognition. INFERNET has a constrained human-like symbolic processing capability. It has been shown to be comparable to human in a relatively wide range of reasoning tasks. For example, INFERNET displays human-like effects in short term memory (STM) capacity, in serial effects on STM, in similarity effects on STM, in double dissociation between STM and long

term memory, in the negation effects on conditional reasoning, in the treatment of predicates, and in effects related to multiple instantiation.

No competing connectionist model has such a large scope of human cognition simulation abilities. However, some models display abilities that are not yet possible with INFERNET. LISA, for example, can account for human-like analogy making which necessitates predicate embedding. The work of Jensen & Lisman (1998) can account for the Sternberg effect (Sternberg, 1966) which necessitates matching abilities. SHRUTI has abilities of parsing and planning, but its performance was never compared carefully with human data.

8.2 Problems and their potential solutions

8.2.1 Learning

One of the main limitation of INFERNET is its lack of ability to build a Long Term Knowledge base. How could INFERNET learn to build the interconnected successive gates as that described in Figure 4.3? Some solutions based on Hebbian learning have been tried, but none of them has succeeded in learning long chains of knowledge like the successively connected AND-gates.


What would be needed is an algorithm that has power similar to the backpropagation algorithm for feedforward networks. If, by chance, the network correctly made node A fire at the right moment, one could strengthen connections that made the node A fire. But how to reinforce the preceding connections, connections responsible for the firing of nodes that in turn provoked the firing of node A? Do they all need to be strengthened? There is a distinct possibility of ending up with a network that make the node A fire whatever the input is and this, of course, is undesirable.

The problem is similar to the learning of syn-fire chains (Abeles et al., 1993) which were introduced in section 2.6. In recording spike timing of different cortical cells, Abeles et al. (1993) observed that when a neuron A fired, neuron B would fire 151ms later while neuron C would fire 289ms later with a precision across trials of 1 ms! Such long delays, would require dozens of transmission delays from presynaptic to postsynaptic neuron. How could such a long chain of successive neuron firing could be learned? Hertz & Prügel-Bennet (1996), have studied learning of syn-fire chains but this learning is at each step (layer by layer) controlled. An external input produces a syn-fire chain and the Hebbian learning tends to reproduce it independently.

What we need is an algorithm that does not monitor each step of the chain. It would require a means of backpropagating the error. If a node fires at the right time, the weights between it and its presynaptic nodes must be reinforced, but the weights arriving at these presynaptic nodes must *also* be reinforced, although to a much lesser extent.

Some recent neurobiological discoveries could give an insight. Engert (1998) recorded Long term depression (LTD) at different synapses around a neuron. He found that the LTD was spreading around the synapse that was actually responsible for LTD, with the particular pattern displayed in Figure 8.1. LTD happens when the presynaptic cell A signal arrives at the postsynaptic cell B membrane after the spike of this cell B. In this case, the efficiency of the cell A-B connecting synapse is reduced. Engert (1998) found that synapses connecting other presynaptic cells to cell B were also depressed. He also observed that synapses connecting the presynaptic cell A to other cells than B were also depressed. Moreover, synapses connecting cells to cell A were as well depressed. This observation shows that synapse modification is backtracked along the chain of successive connections that produced the incorrect stimulation of the postsynaptic cell B. Research of this kind is essential for the future study of learning.

Figure 8.1
Backward spreading
Long Term
Depression nearby
the synapse
originating it.




8.2.2 Embedding of predicates

This problem was mentioned in chapters 5 and 6. Embedded predicates are predicates whose arguments are predicates. For example, observing in a game of billiards that “the blue

ball went into the pocket because the man shot the cue ball, which collided with the blue ball” requires embedding predicates. Figure 8.2 represents schematically how predicates are embedded. Get (blue-ball, pocket) is the consequent of the antecedent Collide (cue-ball, blue-ball) which itself is a consequent of the antecedent Shot (man, cue-ball).


A good solution to this problem should reflect the difficulties that people have dealing with embedded predicates. In chapter 5, evidence was provided to show that people have considerable difficulty nesting 3 predicate levels and almost never nest 4 levels.



A possible solution would be to make predicates bind to an argument of another predicate by synchrony. In Figure 8.3, the predicate “goes into” fires in synchrony with “consequent”, followed by the firing of its own arguments: “blue-ball” and “pocket”. The predicate “collide” synchronises with “antecedent” followed by the firing of its own arguments: “cue-ball” and “blue-ball”. This solution would work with two levels of embedding. When more levels of embedding are needed, one possibility would be to alternate two pairs of embeddings. The first pair has just been described, the next pair would be represented by the synchronous firing of “consequent” with the predicate “collide” followed by its arguments firing “cue-ball” and “blue-ball”, as well as the synchronous firing of “antecedent” with the predicate “shot” followed by its own arguments firing “man” and “cue-ball”.

This solution would be constrained by the possibility of interference between different roles that predicates bind to. For example, “collide” is successively bound to “antecedent” then to “consequent” in Figure 8.3. This is a theoretical solution that has yet to be implemented and tested.

Figure 8.3
Potential solution for
the representation
of embedded
predicates



8.2.3 Representation

Predicates and rules in INFERNET have a particular fixed representations. Role nodes fires after predicate nodes. For example, the expression “love (lover, lovee)” is represented by making “love” nodes fire 5 ms before “lover” nodes and 10 ms before “lovee” nodes. How this structure is achieved was not the purpose of this work, but it raises some questions. A system cannot rely solely on temporal or spatial structure to produce an unambiguous representation. There are too many different temporal or spatial structures that could be considered as equivalent. “John loves Mary” in English has a temporal fixed structure of the type Subject-Verb-Complement. “Mary is loved by John” has a virtually identical meaning but nonetheless, does not share the same temporal structure. Does every predicate and its arguments have two representations: active and passive? This would hardly be economical. What about languages that do not assign predicate roles by the means of their location in the sentence? In Latin, for example, the order of subject verb and complements is of no importance; the roles are assigned by a particular suffix. The temporal structure of vision is even more complex. If one see a cat on a car’s roof, this should lead to the equivalent representation, whether the cat was seen first or the car was seen first. Moreover, the speed of perception should not influence the resulting temporal structure of the representation.

A system must be able to tune perception into a fixed and unambiguous representation that will be used for reasoning. This system requires top-down control, a means of controlling the coherence of representations based on indices like spatial position (in the cat on the roof example). It would require nodes sensitive to a particular organizational structure. These nodes should have the ability to trigger the firing of other nodes at a particular moment. This could be done by means of connections that act on other connections (amplification of excitation, amplification of inhibition, inhibition of excitation and inhibition

of inhibition, etc., see chapter 2). It would also need a means of mapping the representations at hand, incorporating those representations into the Long Term Knowledge base, and, finally, being sensitive to inconsistencies.

8.2.4 The control of attention and consciousness

This problem was already mentioned in sections 1.3.2 and 3.5. People seem to have the ability to control their cognition. This requires the ability to perceive thoughts and to modify them.

It seems that there is no brain center of consciousness. Consciousness must emerge from many active memory representations located in different areas. There is evidence for just such an interaction. Understanding language needs both hemispheres (Gazzaniga, Le Doux & Wilson, 1977). Inter-hemispheric disconnection remove awareness by one hemisphere of processes originating in the other. Each hemisphere seems to be conscious of what happens inside it but unconscious of what happens in the other hemisphere (Mark, 1995). So, there must be different consciousnesses and the brain should has to make them congruent. Crick & Koch (1990, 1992, 1998) suggested that the mechanism for linking different consciousnesses in the different parts of the brain was temporal synchronization of neuron spikes. There could also be other possible neural correlates to consciousness: firing rate and phases. According to Frith, Perry & Lumer (1999) numerous studies have explored neural correlates of consciousness but they have encountered methodological problems and no clear answer is yet available.

Finding neural correlates of consciousness is not the only problem. We must understand also how consciousness controls perception and action. This will require examining the relations between perception and attention and between intention and action. According to Frith, Perry & Lumer (1999), we are still far from a clear answer.

8.3 Tractability of the solution

To understand the problem of tractability, imagine a computer program solving the “travelling salesman problem”. This problem involves a series of cities in which the salesman has clients. The salesman has to visit all cities. The problem is to find the optimal succession of cities in the trip (the trip has to be as short as possible). If there are 3 cities to visit, there are 3! (9) different trips to compare. If there are 6 cities to visit, there are 6! (240) different trips to compare. As the number of cities increases, the problem quickly leads to combinatory explosion, becomes “intractable”, i.e., too huge to be solved in any reasonable amount of time. For example, if there are 1000 cities, there will be 1000! trips to evaluate! This is a so

large number that no computer could find a solution. Even if the computer is able to evaluate 1 trip every ms, it would require more than 10^{2554} milleniums to complete the task!

The theory of complexity examines problems of this type and their potential solutions. The order of the complexity of a process is an expression of how the time or memory taken to perform a task increases as the number of data elements n increases. In the above example the order of complexity is factorial $O(n!)$. A problem is tractable if its solution can be attained with a finite and reasonable temporal and spatial computational resources, otherwise the problem is said to be untractable (see Garey & Johnson, 1979).

The order of complexity depends on the problem, the algorithm and the architecture of the computer. It is clear that the INFERNET is untractable on a serial computer for certain problems. In other words, as the number of objects increases, the computational resources increase by a more than polynomial factor. However, the brain is not a serial architecture. Implemented on a brain-like computer, the INFERNET algorithm would render problems tractable with a constant time order of complexity $O(1)$. In this case, every node could calculate their input in parallel and every connection could propagate their potential in parallel.

The problem of tractability is important because if a model fails to make problems tractable on a brain-like machine, this model has to be rejected. Every model that sequentially scan Long Term Knowledge base, for example, has to be rejected, since the more facts that are in this Long Term Knowledge base, the longer it will take to sequentially scan it. This is important because, for humans, the retrieval time from long term memory does not depend on the size of it. There is the exception of the “fan effect” (Anderson, 1983) but this only applies to material that has been learned recently.

8.4 Epilogue

In order to build a computer program that models human cognition, we have to describe precisely the mechanisms that will lead to a behavior as close as possible to human behavior. The originality of INFERNET is to attempt to justify these mechanisms by neurobiological processes. As a result, all properties of INFERNET are at some level mechanistic. But no claim has been made about INFERNET’s completeness in accounting for cognition. A materialist approach may ultimately prove to be insufficient to explain all aspects of human cognition, but attempts to go as far as we can with this approach will help us better understanding the basic mechanisms underlying human cognition, in general and human reasoning in particular.

Appendices

A INFERNET algorithm description with a pseudo-code

Here follows the pseudo-code description of INFERNET algorithm. The sub-routine *a* executes a theta cycle: reset a series of variables, executes a series of gamma cycles (sub-routine *b*) and reduces weights of unused connections (sub-routine *c*). The sub-routine *b* executes a gamma cycle which is Δt_γ ms long. It executes sub-routine *d* and *e*. The sub-routine *d* executes at each ms a series of computations, it integrates inputs, reduce resistance, looks for nodes that fire (sub-routine *f*), and for them, it resets their activation, assign their resistance, execute binding learning (sub-routine *g*) and propagates their activation. The sub-routine *f* takes inputs and looks for nodes which potential cross the threshold.

For each theta cycle from 0 to max-theta

a Reset the activation of each node to 0

Reset the list of active connections

Reset node firing counter

For each gamma cycle from 0 to max-gamma

b Do from $ms = 0$ to $\Delta t\gamma$

d For each node integrate activation taking care of PSP curve

For each node, reduce the resistance

For each node

f If activation > noisy threshold or if node belongs to the input list
unless that node firing counter is > 10.

i Put that node reference in the firing list of nodes

Write the node reference in a file

For each node firing assign maximum resistance to input

For each node firing reset the activation to 0

For each node firing in this ms, If binding learning is on

g Increase each other plastic connection

For each firing node

h Propagate activation with appropriate noisy weight and noisy delay

For each node that has fired in this cycle

e Increment firing counter

For each post-synaptic connection

c Reduce weight of unused connection

B Statistical analysis

The empirical data reported here are analyzed in a non standard manner compared to what is found in the psychological literature. The motivation was to avoid certain flaws that occur in the analysis of frequency tables and the violation of assumptions in the analysis of variance. This section provides a description and justification of these methods. Two types of measures were collected: frequencies of responses and reaction times. Frequencies have been analyzed by the Log-linear Analysis, and reaction times by Analysis of variance (ANOVA).

B.1 Frequency tables

Suppose that a researcher collects responses from subjects to a question from two groups and two different contexts. Responses are coded into two categories. We obtain a $2 \times 2 \times 2$ frequency table. Table B.1 provides an example of data.

Table B.1
Example of multi-
way frequency table

	Group 1		Group 2	
	Response α	Response β	Response α	Response β
Context A	20	0	19	1
Context B	15	5	17	3

Many researchers use ANOVA to analyze this type of data, assigning, for example, 0 to responses α and 1 to responses β . This metric transformation of categorical variable is not appropriate (see Bakeman & Robinson, 1994). This is even more obvious when the dependent categorical variable has more than two possible values. Coding Oriental=1, White=2 and Black=3, for example, implies that, at the very least, the interval separating Orientals and Whites is the same as that separating Whites and Blacks and is half that separating Orientals and Blacks. This assumption is, of course, nonsensical and is the reason that an ANOVA is inappropriate here. In general, when you wish to explain a categorical dependent variable by a set of categorical variable, the Log-linear analysis is the appropriate technique.

Log-linear analysis provides a way to test the different factors (e.g. “context” and “group” in table B.1) and their interaction for statistical significance. The principle of Log-linear analysis is to compute expected frequencies on the basis of marginal frequencies. These expected frequencies are those which would be expected if the factors were unrelated. Significant deviations of observed frequencies from expected frequencies reflect a relationship between variables.

Testing an hypothesis with a Log-linear analysis tests a particular model to see if it fits the data. If there is a significant difference between observed frequencies and expected frequencies the model is rejected. One can also compare the differences between two models to estimate the effect of a particular variable or interaction. The log-linear analysis uses the maximum likelihood ratio chi-square G^2 .

For the interested reader, here follows a more detailed description. In our example, we are interested in the effect of the “Context” and the “Group” and their interaction on the response. These effects are tested by dropping interactions between the variable and the response from a more complete model. The interaction effect can be computed by testing the saturated model ($C+G+R+C*G+C*R+G*R+C*G*R$) and the model without the interaction ($C+G+R+C*G+C*R+G*R$). The difference of the two G^2 reflects the effect of the interaction. The effect of the Group on the response is computed by the difference between the model with all 2-way interactions ($C+G+R+C*G+C*R+G*R$) and the model without the interaction between Group and Response ($C+G+R+C*G+C*R$). Similarly, the effect of context is evaluated by the difference between these models ($C+G+R+C*G+C*R+G*R$) ($C+G+R+C*G+G*R$). The independence is evaluated by a model that exclude relationship between Context and/or Group and the response ($C+G+R+C*G$) in our example, since $p < .05$ independence is rejected. Table B.2 displays the log-linear analysis of the data provided in table B.1.

We conclude by stating that the context influences the response.

Table B.2 Example of Log-linear analysis	Model	DF	G^2	ΔG^2	ΔDF	p	Effect
	$C+G+R+C*G+C*R+G*R+C*G*R$	0	0	-	-	-	
	$C+G+R+C*G+C*R+G*R$	1	1.906	1.906	1	.167	$C*G*R$
	$C+G+R+C*G+C*R$	2	2.042	.136	1	.712	$G*R$
	$C+G+R+C*G+G*R$	2	8.806	6.9	1	.009	$C*R$
	$C+G+R+C*G$	3	8.931			.03	

The G^2 is sometimes over evaluated when there are many low frequency cells (less than 10). Some authors (e.g. Goodman, 1971) proposed adding a small constant 1/2 to each cell. This is similar to Yates correction. Agresti (1990), argues that apart from the saturated model, this correction smooths too much the data and makes the test too conservative. In the present study, this correction has been used, keeping in mind that the true G^2 is somewhere between the two estimations.

The Log-linear analysis provides also a way to deal with structural zeros. Structural zeros are zeros that reflect a non-existence (e.g. a pregnant male). This possibility will be useful for analyzing incomplete designs.

For repeated measures, Rasch model could be used. A supplementary variable is added. This categorical variable is obtained by assigning to each pattern of response a category. For example, if three questions are asked to each subject, and if people responses fall into three categories: a, b, c, there will be $3^3 = 27$ categories of count patterns. More details on that procedure can be found in Conaway (1989), Lindsay, Clogg & Grego (1989). According to Lindsey (personal communication), it is neither reasonable nor necessary when there are only two repetitions per individual to apply this technique. It would give the same results. The Rasch model could not be applied for between-within design. Statisticians have not yet studied the application of log-linear analysis on this type of design (Lindsey, personal communication).

B.2 Reaction times

Reaction times have been analyzed by the well-known ANOVA. Special emphasis has been put on remedies to the violation of assumptions.

In the analysis of variance it is assumed that the dependent variable is normally distributed within groups in the population. The Fisher-Snedecor F statistic is quite robust to deviation from normality. Moreover, the normality assumption is about the population and not samples. In the present study we will not worry too much about departure from normality.

One major assumption is the homogeneity of variance. It is assumed that variances in different groups are approximately equal. When violated, this assumption is quite serious. The first remedy when variances have been found to be unequal was to try to transform the data in a natural way. For reaction time in ms, one can use $1/x$ transformation and deal with speed instead of time. When this transformation did not make variances homogeneous, different remedies were applied depending on the design.

For one-way ANOVA, one can think of non-parametric alternatives. However, there are parametric techniques that provide correction for the violation of homogeneity assumption. The first is the Box procedure (see Box 1954a), the second is the Welch procedure (see Welch, 1951). According to Kohr & Games (1974), the Welch procedure seems to be better than Box one, specially when the number of data points is unequal across groups. We chose the Welch procedure for one-way ANOVA. Description of the Welch procedure can be found in Howell (1987).

For multi-way ANOVA, these latter techniques are not appropriate. Box (1954b) describes a correction for two-way ANOVA, but this solution is limited to the cases where variances across one group variable are unequal. This solution seems also limited to cases where factors $k > 2$. Finally, Box (1954b) does not provides a method for testing interaction effects. An appropriate technique was provided by Brown & Forsythe (1974). Throughout

the chapter it will be referred as the Brown & Forsythe procedure. This procedure is derived from the Welch procedure for one way ANOVA.

In more details, and for the interested reader, here follows the Brown & Forsythe procedure description. All F are computed as usual but are distributed approximately with usual DF effect and f DF for their denominator where f is defined as:

$$f = \frac{1}{\sum_i \sum_j \frac{c_{ij}^2}{n-1}} \quad (B.1)$$

where

$$c_{ij} = \frac{s_{ij}^2}{\sum_i \sum_j s_{ij}^2} \quad (B.2)$$

and

$$s_{ij}^2 = \frac{\sum (x_{ij} - \bar{x}_{ij})^2}{n_{ij} - 1} \quad (B.3)$$

In repeated-measures designs where factors have more than 2 levels, there are two additional assumptions: compound symmetry and sphericity. Compound symmetry requires equality of variances and covariances of the different repeated measures. This is a *sufficient* condition for the F to be valid. The sphericity assumption requires that the different levels in the repeated measure factor are not correlated across subjects. This suggests that a particular participant modify her response as successive trials are performed. This is a *necessary* and *sufficient* condition for the F to be valid. Sphericity is tested by the Mauchly's Test. If the test is significant, the assumption is violated. When this assumption is violated, a correction can be performed: the Box adjustment to degrees of freedom. Details of this correction can be found in various textbooks (e.g. Howell, 1987; Hays, 1994). Suffice to say that a correction $\hat{\epsilon}$ is computed. The numerator degrees of freedom is calculated by multiplying $\hat{\epsilon}$ with original numerator degrees of freedom. The denominator degrees of freedom is calculated by multiplying $\hat{\epsilon}$ with original denominator degrees of freedom.

C. INFERNET code and simulator

The code was written in common lisp, (MCL 4.2). It is available at the following address:

<http://www.fapse.ulg.ac.be/Lab/cogsci/jsougne/INFERNET.lisp.zip>

A standalone application is also available for Macintosh computers at this address:

<http://www.fapse.ulg.ac.be/Lab/cogsci/jsougne/INFERNET.sea>

Glossary of symbols

t	The time in consideration
i	post-synaptic node
j	presynaptic node
k	node that modify a synapse
$t_j^{(f)}$	Firing time of node j
$t_i^{(f)}$	Firing time of node i
$t_k^{(f)}$	Firing time of node k
F_j	Set of firing time of node j
F_i	Set of firing time of node i
F_k	Set of firing times of node k
$V_i(t)$	potential of a node i at time t
Γ_i	Set of presynaptic j node to node i
κ_{ij}	Set of presynaptic nodes to ij synapse

$\varepsilon()$	Post-Synaptic Potential function
$\eta()$	Refractory period function
w_{ij}	Connection strength from pre-synaptic node j to post-synaptic node i
\hat{w}_{ij}	Noisy connection strength
$\hat{w}_{k \rightarrow ij}$	Noisy the connections strength that modify ij connection
d	Delay of a connection
\hat{d}	Noisy delay on connection
x	difference between the time in consideration, the time of the presynaptic node firing and the noisy delay on the connection: $x = t - t_j^{(f)} - \hat{d}$
$\mathcal{H}()$	A stepwise function
Θ	Firing threshold
$\hat{t}_i^{(f)}$	Time of the last spike of the node i
u	The difference between between the time in consideration and the time of the last spike of the node: $u = t - \hat{t}_i^{(f)}$
Δw_{ij}^+	The increase of connection strength between node i and node j
c	Learning constant
Δw_{ij}^-	The decrease of connection strength between node i and node j
e	Forgetting constant

Δt_γ	Delay defined by the gamma wave frequency. If the oscillation has a frequency of 40Hz, $\Delta t_\gamma = 25$
\mathcal{F}_i^θ	The set of firing times of node i in a theta cycle
\mathcal{F}_j^θ	The set of firing times of node j in a theta cycle
$\vec{F}i_\gamma$	The firing vector of node i for a particular gamma cycle
\vec{X}_i	Sum of $\vec{F}i_\gamma$ over gamma and theta cycles $= \sum_{p=1}^{n_\theta} \sum_{q=1}^{n_\gamma} \vec{F}i_\gamma$ for node i
r	correlation
S	The set of binding pairs that are required to consider that a response belongs to a category.
ρ_S	The normalized signed square correlation (or signed coefficient of determination) between a set of paired cell assemblies S
$\frac{\sum y_t}{\sum z_t}$	The ratio of the two sets of nodes spike density.
n_A	The number of node in A cell assembly
n_B	The number of node in B cell assembly
Δt_θ	The delay in ms defined by a theta wave
n_θ	The number of theta cycles
n_γ	The number of gamma cycles
\supset	Material implication
\equiv	Material equivalence

$\&$	Conjunction
V	Disjunction
\sim	Negation
$\hat{\epsilon}$	Box correction on the number of degrees of freedom for violation of the sphericity assumption. Corrected DF are obtained by multiplying original DF with $\hat{\epsilon}$ for both numerator and denominator DF.
f	Brown & Forsythe correction for violation of homogeneity of variances. f is the denominator DF

References

- Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. New-York: Cambridge University Press.
- Abeles, M., Prut, Y., Bergman, H., Vaadia, E., & Aertsen, A. (1993). Integration, Synchronicity and Periodicity. In A. Aertsen (Ed.), *Brain Theory: Spatio-Temporal Aspects of Brain Function* (pp. 149-181). Amsterdam: Elsevier.
- Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann Machines. *Cognitive Science*, 9, 147-169.
- Adrian, E. D. (1926). The impulses produced by sensory nerve endings. *Journal of Physiology*, 61, 42-72.
- Aertsen, A., Vaadia, E., Abeles, M., Ahissar, E., & Bergman, H. (1992). Neural interactions in the frontal cortex of a behaving monkey: Signs of dependance on stimulus context and behavioral states. *Journal für Hirnforschung*, 32, 735-743.
- Agresti, A. (1990). *Categorical data analysis*. New-York: John Wiley & Sons.
- Ajjanagadde, V. (1994). Unclear Distinctions lead to Unnecessary Shortcomings: Examining the rule vs fact, role vs filler, and type vs predicate distinctions from a connectionist representation and reasoning perspective. *Proceedings of the AAAI-94* (pp. 846-851). AAAI Press.
- Anderson, J. R. (1983). *The Architecture of Cognition*. Cambridge: Harvard University Press.
- Ans, B., & Rousset, S. (1997). Avoiding catastrophic forgetting by coupling two reverberating neural networks. *Académie des Sciences: Sciences de la vie*, 320, 989-997.
- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its controlled process. In K. W. Spence, & J. T. Spence (Eds.) *The psychology of learning and motivation* (vol. 2) (pp. 89-195). New-York: Academic Press.
- Baddeley, A. D. (1966a). Short term memory for word sequences as a function of acoustic, semantic and formal similarity. *Quarterly Journal of Experimental Psychology*, 18, 362-365.
- Baddeley, A. D. (1966b). The influence of acoustic and semantic similarity on long term memory for word sequence. *Quarterly Journal of Experimental Psychology*, 18, 302-309.
- Baddeley, A. D. (1986). *Working Memory*. Oxford: Oxford University Press.
- Baddeley, A. D., & Warrington, E. K. (1970). Amnesia and the distinction between long and short-term memory. *Journal of Verbal Learning and Verbal Behavior*, 9, 176-189.
- Baddeley, A. D., Thomson, N., & Buchanan, M. (1975). Word length and the structure of short-term memory. *Journal of Verbal Learning and Verbal Behavior*, 14, 575-589.
- Baddeley, A. D., & Hitch, G. J. (1977). Recency re-examined. In S. Dornic (Ed.), *Attention and Performance VI*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Baird, B. (1996). A cortical network model of cognitive attentional streams, rhythmic expectation, and auditory stream segregation. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 726-726). Mahwah, NJ: Lawrence Erlbaum Associates.
- Baird, B. (1997). A cortical model of cognitive 40 Hz attentional streams, rhythmic expectation, and auditory stream segregation. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 25-30). Mahwah, NJ: Lawrence Erlbaum Associates.

- Bakeman, R., & Robinson, B. F. (1994). *Understanding Log-Linear Analysis with ILOG: An interactive approach*. Hillsdale, NJ: Lawrence Erlbaum Ass.
- Barlow, H. B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology? *Perception, 1*, 371-394.
- Barnden, J. A. (1991). Encoding complex symbolic data structures with some unusual connectionist techniques. In J. A. Barnden, & J. B. Pollack (Eds.) *Advances in Connectionist and Neural Computation Theory Volume 1: High-level connectionist models* (pp. 180-240). Norwood, NJ: Ablex Publishing Corporation.
- Barnden, J. A. (1992). Overcoming rule-based rigidity and connectionist limitations through massively-parallel case-based reasoning. *International Journal of Man-Machine Studies, 36*, 221-246.
- Barnden, J. A. (1994). On the connectionist implementation of analogy and working memory matching. In J. A. Barnden, & K. J. Holyoak (Eds.) *Advances in Connectionist and Neural Computation Theory Volume 3: Analogy, metaphor, and reminding* (pp. 327-374). Norwood, NJ: Ablex Publishing Corporation.
- Barnden, J. A., & Pollack, J. B. (1991). Introduction: Problems for high-level connectionism. In J. A. Barnden, & J. B. Pollack (Eds.), *Advances in Connectionist and Neural Computation Theory Volume 1: High-level connectionist models* (pp. 1-16). Norwood, NJ: Ablex Publishing Corporation.
- Barnden, J. A., & Srinivas, K. (1991). Encoding techniques for complex information structures in connectionist systems. *Connection Science, 3*, 269-315.
- Bavelier, D., & Potter, M. C. (1992). Visual and phonological codes in repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance, 18*, 134-147.
- Bechtel, W., & Abrahamsen, A. (1991). *Connectionism and the mind*. Cambridge, MA: Basil Blackwell.
- Benzi, R., Sutera, A., & Vulpiani, A. (1981). The mechanism of stochastic resonance. *Journal of Physics A: Mathematics and General, 14*, L453-L457.
- Berkeley, I. S. N., Dawson, M. R. W., Medler, D. A., Schopflocher, D. P., & Hornsby, L. (1995). Density plots of hidden value unit activations reveal interpretable bands. *Connection Science, 7*, 167-186.
- Beth, E. W., & Piaget, J. (1961). *Epistémologie mathématique et psychologie*. Paris: PUF.
- Beth, E. W., Grize, J. B., Martin, R., Matalon, B., Naess, A., & Piaget, J. (1962). *Implication, formalisation et logique naturelle*. Paris: PUF.
- Bezrukov, S. M., & Vodyanov, I. (1995). Noise induced enhancement of signal transduction across voltage-dependent ion channels. *Nature, 378*, 362-364.
- Bookman, L. A., & Alterman, R. (1991). Schema recognition for text understanding: An analog semantic feature approach. In J. A. Barnden, & J. B. Pollack, (Eds.), *Advances in Connectionist and Neural Computation Theory Volume 1: High-level connectionist models* (pp. 87-122). Norwood, NJ: Ablex Publishing Corporation.
- Box, G. E. P. (1954a). Some theorems on the quadratic forms applied in the study of analysis of variance problems, I. Effect of inequality of variance in the one-way classification. *Annals of Mathematical Statistics, 25*, 290-302.
- Box, G. E. P. (1954b). Some theorems on the quadratic forms applied in the study of analysis of variance problems, II. Effect of inequality of variance and of correlation between errors in the two-way classification. *Annals of Mathematical Statistics, 25*, 484-498.
- Bragin, A., Jando, G., Nadasdy, Z., Hetke, J., Wise, K., & Busaki, G. (1994). Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat. *Journal of Neuroscience, 15*, 47-60.
- Braine, M. D. S. (1978). On the Relation between the Natural Logic of Reasoning and Standard Logic. *Psychological Review, 85*, 1-21.
- Braine, M. D. S., & O'Brien, D. P. (1991). A Theory of If: A Lexical Entry, Reasoning Program, and Pragmatic Principles. *Psychological Review, 98*, 182-203.
- Broadbent, D. E. (1975). The magic number seven after fifteen years. In A. Kennedy, & A. Wilkes (Eds.), *Studies in long-term memory* (pp. 3-18). New-York: John Wiley & Sons.
- Brody, C. D. (1998). Slow Covariations in Neuronal Resting Potentials Can Lead to Artefactually Fast Cross-Correlations in their Spike Trains. *Journal of Neurophysiology, 80*, 3345-3351.
- Brody, C. D. (1999a). Correlation without Synchrony. *submitted to Neural Computation*
- Brody, C. D. (1999b). Disambiguating Different Covariation Types. *submitted to Neural Computation*.

- Brown, M. B., & Forsythe, A. B. (1974). The ANOVA and Multiple Comparisons for Data with Heterogeneous Variances. *Biometrics*, 30, 719-724.
- Calvin, W. H. (1983). A stone's throw and its launch window: timing precision and its implications. *Journal of Theoretical Biology*, 104, 121-135.
- Canavier, C. C., Clark, J. W., & Byrne, J. H. (1990). Routes to chaos in a model of a bursting neuron. *Biophysical Journal*, 57, 1245-1251.
- Carreiras, M., & Santamaria, C. (1997). Reasoning about Relations: Spatial and Nonspatial Problems. *Thinking and Reasoning*, 3, 191-208.
- Chase, W. G., & Simon, H. A. (1973). Perception in Chess. *Cognitive Psychology*, 4, 55-81.
- Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, 3, 57-65.
- Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic Reasoning Schemas. *Cognitive Psychology*, 17, 391-416.
- Cheng, P. W., & Holyoak, K. J. (1989). On the Natural Selection of Reasoning Theories. *Cognition*, 33, 285-313.
- Clark, H. H. (1969). Linguistic process in deductive reasoning. *Psychological Review*, 76, 387-404.
- Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequence. *Journal of Experimental Psychology: General*, 120, 235-253.
- Cleeremans, A., & French, R. M. (1996). From chicken squawking to cognition: Levels of description and the computational approach in psychology. *Psychologica Belgica*, 36, 5-29.
- Clossman, G. (1988). *A Model of Categorization and Learning in a Connectionist Broadcast System*. Unpublished doctoral dissertation. Indiana University, Bloomington.
- Collins, J. J., Chow, C. C., & Imhoff, T. T. (1995). Stochastic resonance without tuning. *Nature*, 376, 236-238.
- Conaway, M. R. (1989). Analysis of repeated categorical measurements with conditional likelihood methods. *Journal of the American Statistical Association*, 84, 53-62.
- Cosmides, L. (1989). The Logic of Social Exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, 31, 187-276.
- Cowan, N. (1998). Visual and auditory working memory capacity. *Trends in Cognitive Sciences*, 2, 77-78.
- Crick, F. H. C., & Koch, C. (1990). Toward a neurobiological theory of consciousness. *Seminars in neurosciences*, 2, 263-275.
- Crick, F. H. C., & Koch, C. (1992). The problem of consciousness. *Scientific American*, 267, 153-159.
- Crick, F. H. C., & Koch, C. (1998). Consciousness and neuroscience. *Cerebral Cortex*, 8, 97-107.
- D'Andrade, R. (1982). *Reason versus logic*. Paper presented at the symposium on the Ecology of Cognition, Greensboro, NC.
- De Soto, C. B., London, M., & Handel, S. (1965). Social reasoning and spatial paralogic. *Journal of Personality and Social Psychology*, 2, 513-521.
- Deese, J., & Kaufman, R. A. (1957). Serial effects in recall of unorganized and sequentially organized verbal material. *Journal of Experimental Psychology*, 54, 180-187.
- Defays, D., French, R. M. and Sougné, J. (1999). L'apport de l'intelligence artificielle à la psychologie. In J. A. Rondal (Ed.), *Introduction aux Sciences Psychologiques* (pp. 379-415). Labor : Bruxelles.
- Dellarosa-Cummins, D. (1996). Evidence of deontic reasoning in 3- and 4-years old children. *Memory & Cognition*, 24, 823-829.
- Douglas, J. K., Wilkens, L., Pantazelou, E., & Moss, F. (1993). Noise enhancement of information transfer in crayfish mechanoreceptors by stochastic resonance. *Nature*, 365, 337-340.
- Dyer, M. G. (1991). Symbolic NeuroEngineering for Natural Language Processing: A Multilevel Research Approach. In J. A. Barnden, & J. B. Pollack, (Eds.), *Advances in Connectionist and Neural Computation Theory Volume 1: High-level connectionist models* (pp. 32-86). Norwood, NJ: Ablex Publishing Corporation.
- Eckhorn, R., Bauer, R., Jorden, W., Brosch, M., Kruse, W., Munk, M., & Reitboeck, H. J. (1988). Coherent oscillations: A mechanism for feature linking in the visual cortex? *Biological Cybernetics*, 60, 121-130.

- Engel, A. K., König, P., Gray, C. M., & Singer, W. (1990). Stimulus-dependent neuronal oscillations in cat visual cortex: Inter-columnar interaction as determined by cross-correlation analysis. *European Journal of Neuroscience*, 2, 588-606.
- Engel, A. K., König, P., & Singer, W. (1991). Direct physiological evidence for scene segmentation by temporal coding. *Proceedings of the National Academy of Sciences, U. S. A.*, 88, 9136-9140.
- Engel, A. K., König, P., Kreiter, A. K., & Singer, W. (1991). Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex, *Science*, 252: 1177-1179.
- Engel, A. K., Kreiter, A. K., König, P., & Singer, W. (1991). Synchronisation of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. *Proceedings of the National Academy of Sciences, U. S. A.*, 88, 6048-6052.
- Engert, F. (1998). *Pairing-induced LTP in hippocampal slice cultures is not strictly input-specific*. Talk presented at the winterschool Networks with spiking neurons and synaptic plasticity, Berlin december 10-12.
- Evans, J. St. B. T. (1977). Linguistic factors in reasoning. *Quarterly Journal of Experimental Psychology*, 29, 297-306.
- Evans, J. St. B. T. (1989). *Bias in Human Reasoning: Causes and Consequences*. London: Lawrence-Erlbaum Ass.
- Evans, J. St. B. T. (1998). Matching Bias in Conditional Reasoning: Do we understand it after 25 years? *Thinking and Reasoning*, 4, 45-82.
- Evans, J. St. B. T., Newstead, S. E., & Byrne, R. M. J. (1993). *Human Reasoning: The psychology of deduction*. Hove: Lawrence Erlbaum Associates.
- Evans, J. St. B. T., Clibbens, J., & Rood, B. (1995). Bias in conditional inference: Implications for mental models and mental logic. *Quarterly Journal of Experimental Psychology*, 48A, 644-670.
- Evans, J. St. B. T., Clibbens, J., & Rood, B. (1996). The role of implicit and explicit negation in conditional reasoning bias. *Journal of Memory and Language*, 35, 392-409.
- Feldman, J. A. (1982). Dynamic Connections in Neural Networks. *Biological Cybernetics*, 46, 27-39.
- Fisher, D. L. (1984). Central capacity limits in consistent mapping, visual search tasks: four channels or more? *Cognitive Psychology*, 16, 449-484.
- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3-71.
- French, R. M. (1995). *The Subtlety of Sameness: A theory and computer model of analogy-making*. Cambridge, MA: MIT Press.
- French, R. M. (1997a). Selective memory loss in aphasics: An insight from pseudo-recurrent connectionist network. In J. A. Bullinaria, D. W. Glasspool, & G. Houghton (Eds.), *Proceedings of the Fourth Neural Computation and Psychology Workshop: Connectionist Representation* (pp. 183-195). London: Springer-Verlag.
- French, R. M. (1997b). Pseudo-recurrent connectionist networks: An approach to the 'sensitivity-stability' dilemma. *Connection Science*, 9, 353-379.
- French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3, 128-135.
- French, R. M., & Mareschal, D. (1998). Could Category-Specific Semantic Deficits Reflect Differences in the Distributions of Features Within a Unified Semantic Memory? In *Proceedings of the Twentieth Annual Cognitive Science Society Conference* (pp. 368-373). Mahwah, NJ: Lawrence Erlbaum Ass.
- Frith, C., Perry, R., & Lumer, E. (1999). The neural correlates of conscious experience: An experimental framework. *Trends in Cognitive Sciences*, 3, 105-114.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and Intractability*. New York: Freeman.
- Gawne, T. J., Kjaer, T. W., & Richmond, B. J. (1996). Latency: Another potential code for feature binding in striate cortex. *Journal of Neurophysiology*, 76, 1356-1360.
- Gazzaniga, M. S., Le Doux, J. E., & Wilson, D. H. (1977). Language, praxis and the right hemisphere: Clues to some mechanisms of consciousness. *Neurology*, 27, 1144-1147.
- Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10, 227-300.
- Gerstner, W. (1999). Spiking Neurons. In W. Maass & C. M. Bishop (Eds.) *Pulsed Neural Networks*. Cambridge, MA: MIT Press.

- Gerstner, W., Kreiter, A. K., Markram, H., & Herz, A. V. M. (1997). Neural codes: Firing rates and beyond. *Proceedings of the National Academy of Sciences, U. S. A.*, 94, 12740-12741.
- Gigerenzer, G., & Hug, K. (1992). Domain-specific reasoning: Social contracts, cheating, and perspective change. *Cognition*, 43, 127-171.
- Giroto, V., Mazzocco, A., & Cherubini, P. (1992). Judgements of deontic relevance in reasoning: A reply to Jackson and Griggs. *Quarterly Journal of Experimental Psychology*, 45A, 547-574.
- Glass, L., & Mackay, M. C. (1988). *From Clocks to Chaos*. Princeton University Press.
- Golomb, D., Wang, X.-J., & Rinzel, J. (1994). Synchronization Properties of Spindle Oscillations in a Thalamic Reticular Nucleus Model. *Journal of Neurophysiology*, 72, 1109-1126.
- Goodman, L. A. (1971). The analysis of multidimensional contingency tables: Stepwise procedures and direct estimation methods for building models for multiple classifications. *Technometrics*, 13, 33-61.
- Goswami, U., & Brown, A. L. (1989). Melting Chocolate and Melting Snowmen: Analogical reasoning and causal relations. *Cognition*, 35, 69-95.
- Gray, C. M., Engel, A. K., König, A., & Singer, W. (1992). Synchronization of oscillatory neuronal responses in cat striate cortex: Temporal properties. *Visual Neuroscience*, 8, 337-347.
- Gray, C. M., König, A., Engel, A. K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338, 334-337.
- Green, D. W., Over, D. E., & Pyne, R. A. (1997). Probability and choice in the selection task. *Thinking and Reasoning*, 3, 209-235.
- Griggs, R. A., & Cox, J. R. (1982). The Elusive Thematic-Materials Effect in Wason's Selection Task. *British Journal of Psychology*, 73, 407-420.
- Griggs, R. A., & Cox, J. R. (1993). Permission schemas and the selection task. *Quarterly Journal of Experimental Psychology*, 46A, 637-651.
- Grossberg, S., & Somers, D. (1992). Synchronized oscillations for binding spatially distributed feature codes into coherent spatial patterns. In G. Carpenter, & S. Grossberg (Eds.), *Neural networks for vision and image processing* (pp. 385-405). Cambridge, MA: MIT Press.
- Grossberg, S., & Grunewald, A. (1997). Cortical synchronization and perceptual framing. *Journal of Cognitive Neuroscience*, 9, 117-132.
- Grossberg, S., Boardman, I., & Cohen, M. (1997). Neural dynamics of variable-rate speech categorization. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 235-253.
- Hadley, R. F. (1996). Connectionism, Systematicity and Nomic Necessity. *Proceedings of the Eighteen conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Ass.
- Halford, G. S., & Wilson, W. H. (1980). A Category Theory Approach to Cognitive Development. *Cognitive Psychology*, 12, 356-411.
- Halford, G. S., Bain, J. D., & Mayberry, M. T. (1984). Does a Concurrent Memory Load Interfere with Reasoning? *Current Psychological Research and Reviews*, 3, 14-23.
- Halford, G. S., Mayberry, M. T., & Bain, J. D. (1986). Capacity Limitations in Children's Reasoning: A dual task approach. *Child Development*, 57, 616-627.
- Halford, G. S., Wilson, W. H., Guo, J., Gayler, R. W., Wiles, J., & Stewart, J. E. M. (1994). Connectionist Implications for Processing Capacity Limitations in Analogies. In K. J. Holyoak, & J. A. Barnden (Eds.), *Advances in Connectionist and Neural Computation Theory, Volume 2: Analogical Connections* (pp. 363-415). Norwood, NJ: Ablex.
- Harris, C. L., & Morris, A. L. (1999). Orthographic repetition blindness. *Quarterly Journal of Experimental Psychology* (in press).
- Hartley, T., & Houghton, G. (1996). A linguistically constrained model of short-term memory for nonwords. *Journal of Memory and Language*, 35, 1-31.
- Hawkins, R. D., Kandel, E. R., & Siegelbaum, S. A. (1993). Learning to modulate transmitter release: Themes and variations in synaptic plasticity. *Annual Review of Neuroscience*, 16, 625-665.
- Hays, W. L. (1994). *Statistics (5th ed.)*. Forth-Worth: Harcourt Brace College publishers.
- Hebb, D. O. (1949). *The Organization of Behavior*. New York: Wiley.
- Henderson, J. B. (1994). *Description based parsing in a connectionist network*. Unpublished doctoral dissertation. University of Pennsylvania, Philadelphia.

- Hertz, J., & Prügel-Bennet, A. (1996). Learning synfire chains: Turning noise into signal. *International Journal of Neural Systems*, 7, 445-450.
- Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. in D. E. Rumelhart, & J. L. McClelland (Eds.), *Parallel Distributed Processing Volume 1* (pp. 77-109). Cambridge, MA: MIT Press.
- Hinton, G. E., & Plaut, D. C. (1987). Using fast weights to deblur old memories. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Hochhaus, L., & Johnston, J. C. (1996). Perceptual repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 355-366.
- Hofstadter, D. R. (1984). *The Copycat project: An experiment in nondeterminism and creative analogies*. AI Memo No. 755, Massachusetts Institute of Technology.
- Holyoak, K. J., & Spellman, B. A. (1993). Thinking. *Annual Review of Psychology*, 44, 265-315.
- Holyoak, K. J., & Cheng, P. W. (1995). Pragmatic reasoning with a point of view. *Thinking and Reasoning*, 1, 289-313.
- Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational properties. *Proceedings of the National Academy of Sciences, U. S. A.*, 79, 2554-2558.
- Howell, D. C. (1987). *Statistical methods for psychology*. Boston: Pws Publishers.
- Hulme, C., Maughan, S., & Brown, G. D. A. (1991). Memory for familiar and unfamiliar words: Evidence for a long-term memory contribution to short-term memory span. *Journal of Memory and Language*, 30, 685-701.
- Hummel, J. E., & Holyoak, K. J. (1997). Distributed representation of structure: A theory of analogical access and mapping. *Psychological Review*, 104, 427-466.
- Ingram, D. (1985). The psychological reality of children's grammars and its relation to grammatical theory. *Lingua*, 66, 79-103.
- Inhelder, B., & Piaget, J. (1955). *De la logique de l'enfant à la logique de l'adolescent*. Paris: PUF.
- Ivry, R. B., & Prinzmetal, W. (1991). Effect of feature similarity on illusory conjunctions. *Perception Psychophysics*, 49, 105-116.
- Jackson, S. L., & Griggs, R. A. (1990). The elusive pragmatic reasoning schemas effect. *Quarterly Journal of Experimental Psychology*, 42A, 353-373.
- Jensen, O., & Lisman, J. E. (1996). Novel lists of 7±2 known items can be reliably stored in an oscillatory short-term memory network: Interaction with long term memory. *Learning & Memory*, 3, 257-263.
- Jensen, O., & Lisman, J. E. (1998). An Oscillatory Short-Term Memory Buffer Model Can Account for Data on the Sternberg Task. *The Journal of Neuroscience*, 18, 10688-10699.
- Johnson-Laird, P. N. (1983). *Mental Models: Towards a cognitive science of language, inference, and consciousness*. Cambridge: Cambridge University Press.
- Johnson-Laird, P. N., & Tagart, J. (1969). How implication is understood. *American Journal of Psychology*, 2, 367-373.
- Johnson-Laird, P. N., Legrenzi, P., & Legrenzi, M. S. (1972). Reasoning and a sense of reality. *British Journal of Psychology*, 63, 395-400.
- Johnson-Laird, P. N., & Byrne, R. M. J. (1991). *Deduction*. London: Lawrence-Erlbaum Ass.
- Johnson-Laird, P. N., & Byrne, R. M. J. (1992). Modal reasoning, models and Manktelow and Over. *Cognition*, 43, 173-182.
- Johnson-Laird, P. N., Byrne, R. M. J., & Schaeken, W. (1992). Propositional reasoning by model. *Psychological Review*, 99, 418-439.
- Johnson-Laird, P. N., & Byrne, R. M. J. (1993). Precis of deduction. *Behavioral and Brain Sciences*, 16, 323-380.
- Kanerva, P. (1990). *Sparse distributed memory*. Cambridge, MA, MIT-Press.
- Kanwisher, N. G. (1987). Repetition Blindness: Type recognition without token individuation. *Cognition*, 27, 117-143.
- Kanwisher, N. G., Driver, J., & Machado, L. (1995). Spatial repetition blindness is modulated by selective attention to color or shape. *Cognitive Psychology*, 29, 303-337.

- Keppel, G., & Underwood, B. J. (1962). Proactive inhibition in short-term retention of single items. *Journal of Verbal Learning and Verbal Behavior*, 1, 153-161.
- Kirby, K. N. (1994). Probabilities and utilities of fictional outcomes in Wason's four-card selection task. *Cognition*, 51, 1-28.
- Kirpatrick, S., Gelatt, C. D. J., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science*, 220, 671-680.
- Kohr, R. L., & Games, P. A. (1974). Robustness of the analysis of variance, the Welch procedure and a Box procedure to heterogeneous variances. *The Journal of Experimental Education*, 43, 61-69.
- Konnerth, A., Tsien, R. Y., Mikoshiba, K., & Altman, J. (1996). *Coincidence detection in the nervous system*. Strasbourg: Human Frontier Science Program.
- Kroger, J. K., Cheng, P. W., & Holyoak, K. J. (1993). Evoking the permission schema: The impact of explicit negation and violation-checking context. *Quarterly Journal of Experimental Psychology*, 46A, 615-635.
- Kroger, J. K., & Holyoak, K. J. (1997). What it means to be "the same": The impact of relational complexity on processing efficiency. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*. (pp. 406-411). Mahwah, NJ: Lawrence Erlbaum Associates.
- Lange, T. E. (1992). Lexical and pragmatic disambiguation and re-interpretation in connectionist networks. *International Journal of Man-Machine Studies*, 36, 191-220.
- Lange, T. E., & Dyer, M. G. (1989a). High-level inferencing in a connectionist network. *Connection Science*, 1, 181-217.
- Lange, T. E., & Dyer, M. G. (1989b). Frame selection in a connectionist model of high-level inferencing. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society* (pp. 706-713). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Lashley, K. (1950). In search of the engram. In *Society of Experimental Biology Symposium, 4: Psychological Mechanisms in Animal Behavior*. (pp. 454-480). Cambridge: Cambridge University Press.
- Levin, J. E., & Miller, J. P. (1996). Broadband neural encoding in the cricket cercal sensory system enhanced by stochastic resonance. *Nature*, 380, 165-168.
- Levy, J. P., & Bairaktaris, D. (1995). Connectionist dual-weight architectures. *Language and Cognitive Processes*, 10, 265-283.
- Liberman, N., & Klar, Y. (1996). Hypothesis testing in Wason's selection task: Social exchange cheating detection or task understanding. *Cognition*, 58, 127-156.
- Lindsay, B., Clogg, C. C., & Grego, J. (1991). Semiparametric estimation in the Rasch model and related exponential response models, including a simple latent class model for item analysis. *Journal of the American Statistical Association*, 86, 96-107.
- Lisman, J. E., & Idiart, M. A. P. (1995). Storage of 7 ± 2 Short-Term Memories in Oscillatory Subcycles. *Science*, 267, 1512-1515.
- Logie, R. H. (1986). Visuo-spatial processing in working memory. *Quarterly Journal of Experimental Psychology*, 38A, 229-247.
- Logie, R. H. (1996). The seven ages of working memory. In Richardson, J. T. E., Engle, R. W., Hasher, L., Logie, R. H., Stoltzfus, E. R., & Zachs, R. T. (Eds.), *Working memory and human cognition* (pp. 31-65). New York: Oxford University Press.
- Longoni, A. M., & Scalisi, T. G. (1994). Developmental aspects of phonemic and visual similarity effects: Further evidence in Italian children. *International Journal of Behavioral Development*, 17, 57-71.
- Löwel, S., & Singer, W. (1992). Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science*, 255, 209-212.
- Luck, S. J., & Vogel, E. K. (1998). Response from Luck and Vogel. *Trends in Cognitive Sciences*, 2, 78-80.
- Maass, W. (1996a). *Networks of spiking neurons: The third generation of neural network models*. Techreport. Technische Universität Graz.
- Maass, W. (1996b). *On the computational power of noisy spiking neurons*. Neurocolt techreport NC-TR-96-016. University of London.
- Maass, W. (1996c). *Analog computations on networks of spiking neurons*. Neurocolt techreport NC-TR-96-014. University of London.
- Maass, W., & Bishop, C. M. (1999). *Pulsed Neural Networks*. Cambridge, MA: MIT Press.

- MacKay, D. G., & Miller, M. D. (1994). Semantic blindness: Repeated concepts are difficult to encode and recall under time pressure. *Psychological Science*, 5, 52-55.
- MacKay, W. A. (1997). Synchronized neuronal oscillations and their role in motor processes. *Trends in Cognitive Sciences*, 1, 176-183.
- Maddox, J. (1994). Bringing more order out of noisiness. *Nature*, 369, 271-271.
- Mani, D. R., & Shastri, L. (1993). Reflexive Reasoning with Multiple Instantiation in a Connectionist Reasoning System with a Type Hierarchy. *Connection Science*, 5, 205-242.
- Manktelow, K. I., & Evans, J. S. B. T. (1979). Facilitation of Reasoning by Realism: Effect or non-effect? *British Journal of Psychology*, 70, 477-488.
- Manktelow, K. I., & Over, D. E. (1991). Social roles and utilities in reasoning with deontic conditionals. *Cognition*, 39, 85-105.
- Manktelow, K. I., & Over, D. E. (1992). Utility and deontic reasoning: Some comments on Johnson-Laird and Byrne. *Cognition*, 43, 183-188.
- Manktelow, K. I., Sutherland, E. J., & Over, D. E. (1995). Probabilistic factors in deontic reasoning. *Thinking and Reasoning*, 1, 201-219.
- Marcus, S. L., & Rips, L. J. (1979). Conditional Reasoning. *Journal of Verbal Learning and Verbal Behavior*, 18, 199-233.
- Mark, V. W. (1995). Diagnostic communicative behavior in a split brain subject. In S. R. Hameroff, A. W. Kaszniak, & A. C. Scott (Eds.), *Toward a science of consciousness: Contributions from the 1994 Tucson conference on nature and mind*. Cambridge, MA: MIT-Press.
- Markram, H., Wang, Y., & Tsodyks, M. (1998). Differential signaling via the same axon of neocortical pyramidal neurons. *Proceedings of the National Academy of Sciences, U. S. A.*, 95, 5323-5328.
- Markram, H., Pikus, D., Gupta, A., & Tsodyks, M. (1998). Potential for multiple mechanisms, phenomena and algorithms for synaptic plasticity at single synapses. *Neuropharmacology*, 37, 489-500.
- Marsh, G. R. (1975). Age differences in evoked potential correlates of a memory scanning process. *Experimental Aging Research*, 1, 3-16.
- Maybery, M. T., Bain, J. D., & Halford, G. S. (1986). Information-Processing Demands of Transitive Inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 600-613.
- McCarthy, J. (1988). Epistemological Challenges for Connectionism. *Behavioral and Brain Science*, 11, 44.
- McClelland, J. L., (1986). The programmable blackboard model of reading. In J. L. McClelland, & D. E. Rumelhart (Eds.), *Parallel Distributed Processing Volume 2* (pp.122-169). Cambridge, MA: MIT Press.
- McClelland, J. L., McNaughton, B., & O'Reilly, R. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102, 419-457.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 115-133.
- Miller, G. A. (1956). The magical number seven, plus or minus two. *Psychological Review*, 63, 81-97.
- Moore, G. P., Segundo, J. P., Perkel, D. H., & Levitan, H. (1970). Statistical Signs of Synaptic Interaction in Neurons. *Biophysical Journal*, 10, 876-900.
- Morris, A. L., & Harris, C. L. (1997). A sublexical locus for repetition blindness: Evidence from illusory words. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 519-524). Mahwah: Lawrence Erlbaum Associates.
- Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, 64, 482-488.
- Murthy, V. N., & Fetz, E. E. (1992). Coherent 25- to 35- Hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proceedings of the National Academy of Sciences, U. S. A.*, 89, 5670-5674.
- Nakamura, K., Mikami, A., & Kubota, K. (1992). Oscillatory neuronal activity related to visual short-term memory in monkey temporal pole. *Neuroreport*, 3, 117-120.
- Nelson, J. I., Salin, P. A., Munk, M. H. J., Arzi, M., & Bullier, J. (1992). Spatial and temporal coherence in cortico-cortical connections: A cross-correlation study in area 17 and 18 in the cat. *Visual Neuroscience*, 9, 21-38.
- Newstead, S. E., Ellis, M. C., Evans, J. S. B. T., & Dennis, I. (1997). Conditional reasoning with realistic material. *Thinking and Reasoning*, 3, 49-76.

- Norman, D. A. (1986). Reflections on Cognition and Parallel Distributed Processing. In J. L. McClelland, & D. E. Rumelhart (Eds.), *Parallel Distributed Processing Vol. 2* (pp. 531-546). Cambridge, MA: MIT Press.
- Noveck, I. A., & O'Brien, D. P. (1996). To what extent do pragmatic reasoning schemas affect performance on Wason's selection task? *Quarterly Journal of Experimental Psychology*, 49A, 463-489.
- O'Keefe, J., & Recce, M. (1993). Phase relationship between hippocampal place units and the hippocampal theta rhythm. *Hippocampus*, 3, 317-330.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101, 608-631.
- Oaksford, M., & Chater, N. (1995). Theories of Reasoning and the Computational Explanation of Everyday Inference. *Thinking and Reasoning*, 1, 121-152.
- Oaksford, M., & Chater, N. (1995b). Two and three stage models of deontic reasoning. *Thinking and Reasoning*, 1, 350-357.
- Oaksford, M., & Stenning, K. (1992). Reasoning with Conditionals Containing Negated Constituents. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 835-854.
- Parkin, A. J. (1998). The central executive does not exist. *Journal of the International Neuropsychological Society*, 4, 518-522.
- Peterson, L. R., & Peterson, M. J. (1959). Short-term retention of individual verbal items. *Journal of Experimental Psychology*, 58, 193-198.
- Pitts, W., & McCulloch, W. S. (1947). How we know universals: the perception of auditory and visual forms. *Bulletin of Mathematical Biophysics*, 9, 127-147.
- Platt, R. D., & Griggs, R. A. (1993). Darwinian algorithms and the Wason selection task: A factorial analysis of social contract selection task problems. *Cognition*, 48, 163-192.
- Pollard, P., & Evans, J. St. B. T. (1980). The influence of logic on conditional reasoning performance. *Quarterly Journal of Experimental Psychology*, 32, 605-624.
- Postman, L., & Phillips, L.W. (1965). Short-term temporal changes in free recall. *Quarterly Journal of Experimental Psychology*, 17, 132-138.
- Premack, D. (1983). The Code of Man and Beasts. *The Behavioral and Brain Sciences*, 6, 125-167.
- Riehle, A., Grün, S., Diesmann, M., & Aertsen, A. (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278, 1950-1953.
- Rips, L. J. (1983). Cognitive Processes in Propositional Reasoning. *Psychological Review*, 90, 38-71.
- Rips, L. J. (1990). Paralogical reasoning: Evans, Johnson-Laird, and Byrne on liar and truth-teller puzzles. *Cognition*, 36, 291-314.
- Rips, L. J. (1994). *The psychology of proof: Deductive reasoning in human thinking*. Cambridge, MA: MIT-Press.
- Roberge, J. J. (1976a). The effect of negation on adult's disjunctive reasoning abilities. *Journal of General Psychology*, 91, 23-28.
- Roberge, J. J. (1976b). Reasoning with exclusive disjunctive arguments. *Quarterly Journal of Experimental Psychology*, 28, 419-427.
- Roberge, J. J. (1977). Effects of content on inclusive disjunction reasoning. *Quarterly Journal of Experimental Psychology*, 29, 669-676.
- Roberge, J. J. (1978). Linguistic and psychometric factors in propositional reasoning. *Quarterly Journal of Experimental Psychology*, 30, 705-716.
- Roelfsema, P. R., Engel, A. K., König, P., & Singer, W. (1996). The role of neuronal synchronization in response selection: A biologically plausible theory of structured representations in the visual cortex. *Journal of Cognitive Neuroscience*, 8, 603-625.
- Roelfsema, P. R., Engel, A. K., König, P., & Singer, W. (1997). Visuomotor integration is associated with zero time-lag synchronization among cortical areas. *Nature*, 385, 157-161.
- Rosenblatt, F. (1958). The Perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65, 386-408.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart, & J. L. McClelland (Eds.), *Parallel Distributed Processing Volume 1*. (pp. 318-362). Cambridge, MA: MIT Press.

- Schaeken, W., Johnson-Laird, P. N., & d'Ydewalle, G. (1996). Tense, Aspect, and Temporal Reasoning. *Thinking and Reasoning*, 2, 309-327.
- Shallice, T., & Warrington, E. K. (1970). Independent functioning of verbal memory stores: A neuropsychological study. *Quarterly Journal of Experimental Psychology*, 22, 261-273.
- Shastri, L. (1999). Advances in SHRUTI—A neurally motivated model of knowledge representation and rapid inference using temporal synchrony. *Journal of Applied Artificial Intelligence* (to appear).
- Shastri, L., & Ajjanagadde, V. (1993). From Simple Associations to Systematic Reasoning: A connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Science*, 16, 417-494.
- Shastri, L., Grannes, D. J., Narayanan, S., & Feldman, J. A. (1998). *A connectionist encoding of schemas and reactive plans*. ICSI techreport. University of California, Berkeley.
- Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annual Review of Physiology*, 55, 349-74.
- Singer, W. (1995). Synchronization of neuronal responses as a putative binding mechanism. In A. Arbib (ed.), *The Handbook of Brain Theory and Neural Networks* (pp. 960-964). Cambridge, MA: MIT Press.
- Singer, W., & Gray, C. M. (1995). Visual Feature Integration and the Temporal Correlation Hypothesis. *Annual Review of Neuroscience*, 18, 555-586.
- Singer, W., Engel, A. K., Kreiter, A. K., Munk, M. H. J., Neuenschwander, S., & Roelfsema, P. R. (1997). Neuronal assemblies: necessity, signature and detectability. *Trends in Cognitive Sciences*, 1, 252-261.
- Smolensky, P. (1987). *On variable binding and the representation of symbolic structures in connectionist systems*. Technical report CU-CS-355-87, University of Colorado Boulder.
- Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46, 159-216.
- Sougné, J. (1996). A Connectionist Model of Reflective Reasoning Using Temporal Properties of Node Firing. In *Proceedings of the Eighteenth Annual Conference of the Cognitive Science* (pp. 666-671). Mahwah, NJ: Lawrence Erlbaum Associates.
- Sougné, J. (1997). Variable oscillation frequencies for solving the problem of multiple instantiation. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 1058-1058). Mahwah, NJ: Lawrence Erlbaum Associates.
- Sougné, J. (1998a). Connectionism and the problem of multiple instantiation. *Trends in Cognitive Sciences*, 2, 183-189.
- Sougné, J. (1998b). Period Doubling as a Means of Representing Multiply Instantiated Entities. In *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 1007-1012). Mahwah, NJ: Lawrence Erlbaum Associates.
- Sougné, J., & French, R. M. (1997). A Neurobiologically Inspired Model of Working Memory Based on Neuronal Synchrony and Rythmicity. In J. A. Bullinaria, D. W. Glasspool, and G. Houghton (Eds.), *Proceedings of the Fourth Neural Computation and Psychology Workshop: Connectionist Representations* (pp. 155-167). London: Springer-Verlag.
- Sperber, D., Cara, F., & Girotto, V. (1995). Relevance Theory Explains the Selection Task. *Cognition*, 57, 31-95.
- Stanovich, K. E., & West, R. F. (1998). Cognitive ability and variation in selection task performance. *Thinking and Reasoning*, 4, 193-230.
- Stemmler, M., Usher, M., & Niebur, E. (1995). Lateral interactions in primary visual cortex: A model bridging physiology and psychophysics. *Science*, 269, 1877-1880.
- Sternberg, R. J. (1980). Representation and process in linear syllogistic reasoning. *Journal of Experimental Psychology: General*, 109, 119-159.
- Sternberg, S. (1966). High Speed Scanning in Human Memory. *Science*, 153, 652-654.
- Sun, R. (1992). On variable binding in connectionist networks. *Connection Science*, 4, 93-124.
- Sun, R. (1995). Robust reasoning: integrating rule-based and similarity-based reasoning. *Artificial Intelligence*, 75, 241-295.
- Taplin, J. E. (1971). Reasoning with conditional sentences. *Journal of Verbal Learning and Verbal Behavior*, 10, 219-225.
- Terrace, H. S. (1991). Chunking during serial learning by a pigeon: I. Basic evidence. *Journal of Experimental Psychology: Animal Behavior Processes*, 17, 81-93.

- Terrace, H. S., & Chen, S. (1991). Chunking during serial learning by a pigeon: III. What are the necessary conditions for establishing a chunk? *Journal of Experimental Psychology: Animal Behavior Processes*, 17, 107-118.
- Thomas, E., Van Hulle, M. M., & Vogels, R. (1999). A study of the neuronal encoding of categorization with the use of a Kohonen network. *Proceedings of the Fifth Neural Computation and Psychology Workshop: Connectionist Models in Cognitive Neuroscience*. London: Springer-Verlag.
- Touretsky, D. S. (1995). Connectionist and symbolic representations. In A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks* (pp. 243-247). Cambridge, MA: MIT Press.
- Touretsky, D. S., & Hinton, G. E. (1988). A distributed connectionist production system. *Cognitive Science*, 12, 423-466.
- Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology*, 14, 107-141.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, 95, 15-48.
- Turvey, M. T., Brick, P., & Osborn, J. (1970). Proactive interference in short-term memory as a function of prior item retention interval. *Quarterly Journal of Experimental Psychology*, 22, 142-147.
- Vaadia, E., Ahissar, E., Bergmann, H., & Lavner, Y. (1991). Correlated activity of neurons: A neural code for higher brain functions? In J. Kruger (Ed.), *Neuronal cooperativity* (pp. 249-279). Berlin: Springer-Verlag.
- von der Malsburg, C. (1981). The correlation theory of brain function. Internal report 81-2, MPI für Biophysikalische Chemie, Göttingen, Reprinted In K. Schulten, & H.J. Van Hemmen, (Eds.) (1994) *Models of neural networks*, 2nd ed. Berlin: Springer.
- Walker, P., Hitch, G. J., & Duroe, S. (1993). The effect of visual similarity on short-term memory for spatial location: Implications for the capacity of visual short-term memory. *Acta Psychologica*, 83, 203-224.
- Wang, X., & Rinzal, J. (1995). Oscillatory and Bursting Properties of Neurons. In A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks* (pp. 686-691). Cambridge, MA: MIT Press.
- Wason, P. C. (1966). Reasoning. In B. M. Foss (Ed.), *New Horizons in Psychology I* (pp.135-151). Harmandsworth: Penguin.
- Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, 20, 273-281.
- Wason, P. C. (1969). Regression in reasoning? *British Journal of Psychology*, 60, 471-480.
- Wason, P. C., & Shapiro, D. (1971). Natural and Contrived Experience in a Reasoning Problem. *Quarterly Journal of Experimental Psychology*, 23, 63-71.
- Watkins, O. C., & Watkins, M. J. (1977). Serial recall and the modality effect: Effects of word frequency. *Journal of Experimental Psychology: Human Learning and Memory*, 3, 712-718.
- Wehr, M., & Laurent, G. (1996). Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature*, 384, 162-166.
- Welch, B. L. (1951). On the Comparison of Several Mean Values: An alternative approach. *Biometrika*, 38, 330-336.
- Wickens, D. D., Born, D. G., & Allen, C. K. (1963). Proactive inhibition and item similarity in short-term memory. *Journal of Verbal Learning and Verbal Behavior*, 2, 440-445.
- Wiesenfeld, K., & Moss, F. (1995). Stochastic resonance and the benefits of noise: from ice ages to crayfish and SQUIDS. *Nature*, 373, 33-36.
- Wildman, T. M., & Fletcher, H. J. (1977). Developmental increases and decreases in solutions of conditional syllogism problems. *Developmental Psychology*, 13, 630-636.
- Wilson, M., & Shepherd, G. M. (1995). Olfactory Cortex. In A. Arbib (Ed.) *The Handbook of Brain Theory and Neural Networks* (pp. 669-673). Cambridge, MA: MIT Press.
- Yantis, S. (1992). Multi-element visual tracking: attention and perceptual organization. *Cognitive Psychology*, 24, 295-340.

Author index

A

Abeles, M. 22, 30, 31, 35, 39, 48, 72, 171, 189
Abrahamsen, A. 70, 190
Ackley, D. H. 12, 189
Adrian, E. D. 9, 189
Aertsen, A. 22, 30, 31, 48, 189, 197
Agresti, A. 180, 189
Ahissar, E. 22, 189, 199
Ajjanagadde, V. 19, 20, 36, 49, 93, 116, 189, 198
Allen, C. K. 47, 199
Alterman, R. 115, 190
Altman, J. 39, 195
Anderson, J. R. 176, 189
Ans, B. 57, 189
Arzi, M. 22, 34, 196
Atkinson, R. C. 55, 189

B

Baddeley, A. D. 45, 47, 48, 51, 53, 55, 113, 189
Bain, J. D. 92, 193, 196
Bairaktaris, D. 55, 56, 195
Baird, B. 55, 189
Bakeman, R. 179, 190
Barlow, H. B. 8, 190
Barnden, J. A. 18, 56, 109, 115, 190
Bauer, R. 22, 31, 191
Bavelier, D. 114, 190
Bechtel, W. 70, 190
Benzi, R. 156, 190
Bergman, H. 22, 30, 31, 48, 189, 199
Berkeley, I. S. N. 70, 190
Beth, E. W. 63, 190
Bezrukov, S. M. 157, 190
Bishop, C. M. 12, 195
Boardman, I. 56, 193
Bookman, L. A. 115, 190
Born, D. G. 47, 199

Box, G. E. P. 181, 182, 190
Bragin, A. 22, 190
Braine, M. D. S. 63, 64, 190
Brick, P. 47, 199
Broadbent, D. E. 46, 190
Brody, C. D. 41, 190
Brosch, M. 22, 31, 191
Brown, A. L. 92, 193
Brown, G. D. A. 55, 194
Brown, M. B. 103, 128, 137, 142, 181, 182, 191
Buchanan, M. 51, 189
Bullier, J. 22, 34, 196
Busaki, G. 22, 190
Byrne, J. H. 40, 191
Byrne, R. M. J. 64, 69, 93, 192, 194

C

Calvin, W. H. 39, 191
Canavier, C. C. 40, 191
Cara, F. 70, 198
Carreiras, M. 113, 191
Chase, W. G. 55, 191
Chater, N. 68, 69, 191, 197
Chen, S. 114, 199
Cheng, P. W. 64, 65, 66, 67, 68, 191, 194, 195
Cherubini, P. 67, 193
Chow, C. C. 157, 191
Clark, H. H. 113, 191
Clark, J. W. 40, 191
Cleeremans, A. 2, 56, 191
Clibbens, J. 70, 192
Clogg, C. C. 181, 195
Clossman, G. 19, 191
Cohen, M. 56, 193
Collins, J. J. 157, 191
Conaway, M. R. 181, 191
Cosmides, L. 66, 191
Cowan, N. 45, 48, 191
Cox, J. R. 64, 65, 66, 67, 193
Crick, F. H. C. 55, 175, 191

D

d'Ydewalle, G. 113, 198
 Dawson, M. R. W. 70, 190
 De Soto, C. B. 113, 191
 Deese, J. 47, 191
 Defays, D. 2, 191
 Dellarosa-Cummins, D. 69, 191
 Dennis, I. 69, 196
 Diesmann, M. 22, 197
 Douglass, J. K. 157, 191
 Driver, J. 114, 194
 Duroe, S. 48, 113, 199
 Dyer, M. G. 17, 18, 39, 109, 112, 115, 191, 195

E

Eckhorn, R. 22, 31, 34, 191
 Ellis, M. C. 69, 196
 Engel, A. K. 11, 22, 23, 31, 33, 34, 35, 41, 192, 193, 197, 198
 Engert, F. 172, 192
 Evans, J. St. B. T. 63, 65, 69, 70, 72, 192, 196, 197

F

Feldman, J. A. 15, 20, 39, 192, 198
 Fetz, E. E. 22, 34, 196
 Fisher, D. L. 46, 192
 Fletcher, H. J. 63, 69, 199
 Fodor, J. A. 16, 192
 Forsythe, A. B. 103, 128, 137, 142, 181, 182, 191
 French, R. M. 2, 8, 19, 27, 36, 49, 56, 57, 113, 191, 192, 198
 Frith, C. 175, 192

G

Games, P. A. 181, 195
 Garey, M. R. 176, 192
 Gawne, T. J. 10, 192
 Gayler, R. W. 17, 92, 193
 Gazzaniga, M. S. 175, 192
 Gelatt, C. D. J. 156, 195
 Gentner, D. 92, 192
 Gerstner, W. 19, 192, 193
 Gigerenzer, G. 67, 193
 Girotto, V. 67, 70, 193, 198
 Glass, L. 40, 193
 Golomb, D. 41, 193
 Goodman, L. A. 180, 193
 Gormican, S. 15, 113, 199
 Goswami, U. 92, 193
 Grannes, D. J. 20, 198
 Gray, C. M. 22, 34, 35, 48, 192, 193, 198

Green, D. W. 68, 193
 Grego, J. 181, 195
 Griggs, R. A. 64, 65, 66, 67, 193, 194, 197
 Grize, J. B. 63, 190
 Grossberg, S. 19, 56, 157, 193
 Grün, S. 22, 197
 Grunewald, A. 157, 193
 Guo, J. 17, 92, 193
 Gupta, A. 56, 196

H

Hadley, R. F. 16, 193
 Halford, G. S. 17, 89, 90, 92, 93, 104, 193, 196
 Handel, S. 113, 191
 Harris, C. L. 114, 193, 196
 Hartley, T. 56, 193
 Hawkins, R. D. 193
 Hays, W. L. 182, 193
 Hebb, D. O. 1, 8, 32, 56, 193
 Henderson, J. B. 20, 193
 Hertz, J. 171, 194
 Herz, A. V. M. 19, 193
 Hetke, J. 22, 190
 Hinton, G. E. 12, 16, 56, 189, 194, 197, 199
 Hitch, G. J. 47, 48, 113, 189, 199
 Hochhaus, L. 114, 194
 Hofstadter, D. R. 27, 194
 Holyoak, K. J. 14, 19, 21, 64, 65, 66, 67, 68, 86, 92, 93, 170, 191, 194, 195
 Hopfield, J. 12, 194
 Hornsby, L. 70, 190
 Houghton, G. 56, 193
 Howell, D. C. 124, 181, 182, 194
 Hug, K. 67, 193
 Hulme, C. 55
 Hummel, J. E. 19, 21, 93, 170, 194

I

Idiart, M. A. P. 36, 49, 195
 Imhoff, T. T. 157, 191
 Ingram, D. 16, 194
 Inheller, B. 63, 194
 Ivry, R. B. 15, 113, 194

J

Jackson, S. L. 67, 194
 Jando, G. 22, 190
 Jensen, O. 36, 49, 56, 171, 194
 Johnson, D. S. 176, 192
 Johnson-Laird, P. N. 63, 64, 65, 93, 113, 194, 198
 Johnston, J. C. 114, 194
 Jorden, W. 22, 31, 191

K

- Kandel, E. R. 193
 Kanerva, P. 45, 194
 Kanwisher, N. G. 114, 194
 Kaufman, R. A. 47, 191
 Keppel, G. 47, 195
 Kirby, K. N. 68, 195
 Kirpatrick, S. 156, 195
 Kjaer, T. W. 10, 192
 Klar, Y. 67, 195
 Koch, C. 55, 175, 191
 Kohr, R. L. 181, 195
 König, P. 11, 22, 23, 33, 34, 35, 192, 193, 197
 Konnerth, A. 39, 195
 Kreiter, A. K. 11, 19, 22, 23, 31, 33, 34, 41, 192, 193, 198
 Kroger, J. K. 67, 92, 195
 Kruse, W. 22, 31, 191
 Kubota, K. 31, 49, 196

L

- Lange, T. E. 17, 18, 115, 195
 Lashley, K. 8, 195
 Laurent, G. 11, 199
 Lavner, Y. 22, 199
 Le Doux, J. E. 175, 192
 Legrenzi, M. S. 65, 194
 Legrenzi, P. 65, 194
 Levin, J. E. 157, 158, 195
 Levitan, H. 41, 196
 Levy, J. P. 55, 56, 195
 Liberman, N. 67, 195
 Lindsay, B. 181, 195
 Lisman, J. E. 36, 49, 56, 171, 194, 195
 Logie, R. H. 55, 195
 London, M. 113, 191
 Longoni, A. M. 48, 113, 195
 Löwel, S. 23, 195
 Luck, S. J. 36, 45, 46, 49, 50, 195
 Lumer, E. 175, 192

M

- Maass, W. 12, 14, 195
 Machado, L. 114, 194
 MacKay, D. G. 114, 196
 MacKay, W. A. 35, 196
 Mackay, M. C. 40, 193
 Maddox, J. 156, 196
 Mani, D. R. 109, 116, 196
 Manktelow, K. I. 65, 67, 68, 196
 Marcus, S. L. 63, 196
 Mareschal, D. 57, 192
 Mark, V. W. 175, 196
 Markram, H. 19, 56, 193, 196
 Marsh, G. R. 196
 Martin, R. 63, 190

- Matalon, B. 63, 190
 Maughan, S. 55, 194
 Mayberry, M. T. 92, 193
 Mazzocco, A. 67, 193
 McCarthy, J. 15, 196
 McClelland, J. L. 16, 56, 191, 194, 196
 McCulloch, W. S. 12, 196, 197
 McNaughton, B. 57, 196
 Medler, D. A. 70, 190
 Mikami, A. 31, 49, 196
 Mikoshiba, K. 39, 195
 Miller, G. A. 46, 89, 196
 Miller, J. P. 157, 158, 195
 Miller, M. D. 114, 196
 Moore, G. P. 41, 196
 Morris, A. L. 114, 193, 196
 Moss, F. 156, 157, 191, 199
 Munk, M. H. J. 22, 31, 34, 41, 191, 196, 198
 Murdock, B. B. 47, 52, 196
 Murthy, V. N. 22, 34, 196

N

- Nadasdy, Z. 22, 190
 Naess, A. 63, 190
 Nakamura, K. 31, 49, 196
 Narayanan, S. 20, 198
 Nelson, J. I. 22, 34, 196
 Neuenschwander, S. 22, 31, 41, 198
 Newstead, S. E. 69, 192, 196
 Niebur, E. 157, 198
 Norman, D. A. 109, 112, 197
 Noveck, I. A. 67, 197

O

- O'Brien, D. P. 63, 67, 190, 197
 O'Keefe, J. 10, 197
 O'Reilly, R. 57, 196
 Oaksford, M. 68, 69, 70, 191, 197
 Osborn, J. 47, 199
 Over, D. E. 67, 68, 193, 196

P

- Pantazidou, E. 157, 191
 Parkin, A. J. 55, 197
 Perkel, D. H. 41, 196
 Perry, R. 175, 192
 Peterson, L. R. 46, 197
 Peterson, M. J. 46, 197
 Phillips, L. W. 47, 197
 Piaget, J. 63, 64, 190, 194
 Pikus, D. 56, 196
 Pitts, W. 12, 196, 197
 Platt, R. D. 67, 197
 Plaut, D. C. 56, 194
 Pollack, J. B. 109, 190
 Pollard, P. 65, 69, 72, 197
 Postman, L. 47, 197

Potter, M. C. 114, 190
 Premack, D. 92, 197
 Prinzmetal, W. 15, 113, 194
 Prügel-Bennet, A. 171, 194
 Prut, Y. 30, 31, 48, 189
 Pylyshyn, Z. W. 16, 192
 Pyne, R. A. 68, 193

R

Recce, M. 10, 197
 Reitboeck, H. J. 22, 31, 191
 Richmond, B. J. 10, 192
 Riehle, A. 22, 197
 Rinzel, J. 31, 41, 48, 193, 199
 Rips, L. J. 63, 64, 196, 197
 Roberge, J. J. 93, 197
 Robinson, B. F. 179, 190
 Roelfsema, P. R. 22, 31, 35, 41, 197, 198
 Rood, B. 70, 192
 Rosenblatt, F. 12, 197
 Rousset, S. 57, 189
 Rumelhart, D. E. 12, 16, 194, 197

S

Salin, P. A. 22, 34, 196
 Santamaria, C. 113, 191
 Scalisi, T. G. 48, 113, 195
 Schaeken, W. 93, 113, 194, 198
 Schmidt, H. 15, 199
 Schopflocher, D. P. 70, 190
 Segundo, J. P. 41, 196
 Sejnowski, T. J. 12, 189
 Shallice, T. 48, 198
 Shapiro, D. 65, 199
 Shastri, L. 19, 20, 36, 49, 56, 93, 109, 116, 170, 196, 198
 Shepherd, G. M. 31, 35, 48, 199
 Shiffrin, R. M. 55, 189
 Siegelbaum, S. A. 193
 Simon, H. A. 55, 191
 Singer, W. 11, 22, 23, 31, 33, 34, 35, 41, 48, 192, 193, 195, 197, 198
 Smolensky, P. 17, 198
 Somers, D. 19, 193
 Sougné, J. 2, 19, 23, 36, 42, 49, 56, 113, 155, 191, 198
 Spellman, B. A. 14, 86, 194
 Sperber, D. 70, 198
 Srinivas, K. 18, 56, 115, 190
 Stanovich, K. E. 69, 198
 Stemmler, M. 157, 198
 Stenning, K. 70, 197
 Sternberg, R. J. 113, 198
 Sternberg, S. 56, 171, 198
 Stewart, J. E. M. 17, 92, 193

Sun, R. 18, 198
 Sutura, A. 156, 190
 Sutherland, E. J. 68, 196

T

Tagart, J. 63, 194
 Taplin, J. E. 63, 198
 Terrace, H. S. 114, 198, 199
 Thomas, E. 9, 199
 Thomson, N. 51, 189
 Toupin, C. 92, 192
 Touretsky, D. S. 16, 17, 199
 Treisman, A. 15, 113, 199
 Tsien, R. Y. 39, 195
 Tsodyks, M. 56, 196
 Turvey, M. T. 47, 199

U

Underwood, B. J. 47, 195
 Usher, M. 157, 198

V

Vaadia, E. 22, 30, 31, 48, 189, 199
 Van Hulle, M. M. 9, 199
 Vecchi, M. P. 156, 195
 Vodyanoy, I. 157, 190
 Vogel, E. K. 36, 45, 46, 49, 50, 195
 Vogels, R. 9, 199
 von der Malsburg, C. 33, 199
 Vulpiani, A. 156, 190

W

Walker, P. 48, 113, 199
 Wang, X. 31, 41, 48, 193, 199
 Wang, Y. 56, 196
 Warrington, E. K. 48, 51, 189, 198
 Wason, P. C. 62, 63, 65, 66, 67, 68, 69, 199
 Watkins, M. J. 55, 199
 Watkins, O. C. 55, 199
 Wehr, M. 11, 199
 Welch, B. L. 124, 181, 182, 199
 West, R. F. 69, 198
 Wickens, D. D. 47, 199
 Wiesenfeld, K. 156, 199
 Wildman, T. M. 63, 69, 199
 Wiles, J. 17, 92, 193
 Wilkens, L. 157, 191
 Williams, R. J. 12, 197
 Wilson, D. H. 175, 192
 Wilson, M. 31, 35, 48, 199
 Wilson, W. H. 17, 89, 90, 92, 193
 Wise, K. 22, 190

Y

Yantis, S. 46, 199

Subject index

A

Absolute refractory period 5, 13, 29
Acoustic cortex 22
Acoustic similarity 53
Action potential viii, 4, 13, 28, 30
Affirmation of the consequent (AC) 62, 63, 144-146
AMPA 4, 26
AMPA channels 4
Analog computations 14
Analogical inference 17, 21
Analogy 20, 22, 90-92, 171
AND-gate 37, 38, 72, 76, 77, 94, 97, 118, 120, 121, 131-133, 171

Arity 134
ART 19
Attention 48, 49, 175
Axon 2, 3, 5, 27
Axon terminal 3, 5, 27
Axonal tree 3
Axonal hillock 5

B

Backpropagation 12, 171
Bayesian approach of reasoning 68
Behaviorism 1
Bifurcation 40, 41, 117
Binding vii, viii, ix, 5, 14-24, 33-36, 39, 41, 43, 44, 49, 56, 59, 62, 70, 76, 87, 93, 94, 96, 103, 104, 110, 112, 113, 115-117, 121, 122, 130-133, 136, 146, 155, 156, 169, 170, 177, 187
Binding fixation 87, 156

Binding learning 49, 50-52, 54, 96, 122, 131, 132, 135, 136, 177
Binding-learning rate 54
Binding problem vii, viii, 14-24, 70, 110, 169

Binocular vision 23
Boltzmann machines 12
Boolean function 12, 14
Box procedure 181, 182, 188
Brittleness 14
Broadcast-net 19
Brown & Forsythe correction 128, 137, 142, 181, 182, 188
Brown-Peterson technique 46
Burst 7, 49, 73, 96, 122, 134, 144

C

Ca²⁺ 4, 27
cAMP 27
Catastrophic forgetting 8, 57
Catastrophic interference 8, 57
Category specific neurons 9
Ceiling effect 69
Cell assemblies 8, 9, 32, 43, 143, 187
Chunking 46, 48, 49, 89, 104-108, 114, 128, 129, 143

Cl- 4, 27
Cl- channels 4, 27
Coarse coding 16, 17
Coincidence detection 37, 39
Combinatorial explosion 16, 175
Compartmental model 14
Complex dynamic system 23, 164, 167
COMPOSIT 18, 56, 115
Compound symmetry assumption 100, 101, 105, 126, 148, 182
Cognitive modeling 1-2
Conditional reasoning 62-65, 70-72, 80, 84, 104, 143, 171
Confabulations 20

- Conjunction (&) 188
 Connection delays 36, 164
 Consciousness 55, 175
 CONSYDERR 18, 19
 Content addressable memory 45
 Content effect 65
 Contingency table 78, 82
 Corpus callosum 23
 Correlation 11, 15, 41-44, 50, 113, 187
 Correlation normalized 42
 Correlogram 11
 Cortex 9, 10, 22, 31
 Cortical areas vii, 5, 9, 10, 14, 22, 34
 Cortical cells 39, 171
 Cortical neuron viii
 Cortico-cortical connection 23
 Cost-benefit rule 67
 Cross-correlation 11, 41
 Cross-correlogram 11
 Crosstalk 15, 39, 40, 117
- D**
- Day-off problem 67
 Decay 47, 51, 52
 Deductive reasoning 62, 93
 Defective truth table 63
 Delayed synchronization 39, 41
 Dendrite 2, 3, 14, 27
 Dendritic delays 5, 26
 Dendritic integration 5
 Denial inference 93
 Denial of the antecedent (DA) 62, 72, 104, 105, 143-146, 148
 Denial of the consequent 72
 Deontic reasoning 65-69
 Discrimination viii, xii, 15, 24, 31, 33, 34, 169
 Disjunction (V) 93, 94, 108, 188
 Disjunctive reasoning 93, 99, 108
 Disjunctive syllogism 93, 94
 Distributed representation 8, 18, 22, 32, 40, 110, 112, 113, 115, 117, 152, 167, 170
 Double dissociation 48, 53, 54, 170
 Double instantiation 40, 73, 108, 117, 118, 124, 125, 128, 155, 170
 Double negation 21, 24, 59, 70, 72, 77-87
 Drinking age rule 65, 68
 Dual component architecture 55
 Dual weight architecture 56
 Dual trace architecture 56
 Dynamic binding 15
- E**
- Effect of acoustic similarity 53
 Effect of double negation 77
 Effect of noise 155-168
 Effect of similarity 53, 147, 149, 150
 Embedded oscillations 8
 Embedded predicates 92, 172-174
 Embedded roles 115, 117
 Enabler nodes 19, 20
 Excitation 8, 27, 52, 158, 165, 166, 174
 Excitatory 3, 4, 6, 20, 26, 27, 30, 37, 49, 165, 166
 Excitatory post-synaptic potential (EPSP) 6, 26, 27
 Expected utility 68
 Experts 55
- F**
- Fan effect 176
 Fatigue 121
 Feedforward networks 12, 171
 Filler vii, viii, ix, 15, 17, 18, 36, 40, 59, 61, 62, 112, 115, 129, 170
 Filler nodes ix, 18, 36, 59, 61, 62, 170
 Filler signature 18
 Firing fatigue 121
 Firing latency 170
 Firing rate 11, 12, 19, 156, 170, 175
 Firing rate binding 19
 Firing threshold 157, 161, 162, 165, 186
 Forgetting 8, 45-47, 51, 131, 132, 134, 186
 Frequency of oscillation viii, 6, 31, 35, 48, 93, 115, 117
 Frequency table analysis 179
 Frontal cortex 22
 Fully distributed representation 8, 25, 32, 37
- G**
- GABA channels 26
 GABA neurotransmitter 4
 Gamma wave 6-8, 11, 23, 31, 32, 35, 36, 41-44, 52, 56, 73, 89, 94, 96, 122, 134, 144, 177, 187
 General-purpose rule theory 64
 Generalization ix, 8, 14, 20
 Graceful degradation 167
 Grandmother cell 8, 16, 20
- H**
- Hebbian learning 36, 37, 171
 Hemispheres 22, 23, 34, 175
 Hidden layer 12
 High-level cognition 2, 14
 Hippocampus 10, 11, 22
 Hodgkin-Huxley 13

- Homogeneity of variance assumption 101, 126, 142, 181, 188
- Hopfield networks 12
- I**
- Illusory conjunction 15
- Illusory correlations 15, 113
- Impaired LTM 48, 54
- Impaired STM 48, 54
- Implicit negation 67, 69
- Implicit reasoning 20
- Incomplete design 136, 137, 140, 142, 180
- Induction 21
- Inference paradigm 62, 63, 69
- Inferior temporal cortex 9
- INFERNET 25-44
- Inhibition 27, 47, 76, 158, 165, 174, 175
- Inhibitory 4, 6, 20, 26, 27, 37, 49, 129, 166
- Inhibitory post-synaptic potential (IPSP) 6, 26, 27
- Inhibitory transmitters 4
- Integrate-and-fire neurons viii, 12, 13, 23, 166, 169, 170
- Inter-hemispheric disconnection 175
- Intercellular space 3
- Interference 8, 20, 45, 47, 51, 52, 57, 173
- Ion channels 3, 4, 14, 26, 27, 157
- K**
- K⁺ 4, 27
- K⁺ channels 4, 14
- L**
- Latency code 10, 170
- Leakage factor 13, 21, 26, 166
- Leaky integrator 13
- Learning 12, 16, 21, 22, 36, 37, 42, 44, 50, 57, 62, 73, 86, 96, 122, 131, 134, 144, 171, 172, 177, 186
- Lexical decision task 60, 99, 105, 138, 148
- Linear output function 12
- LISA 19, 21, 22, 93, 103, 170, 171
- Localist representation 8, 17-19, 40, 112, 115, 153
- Log-linear analysis 78, 79, 82, 84, 85, 98, 100, 102, 107, 123, 124, 126, 127, 136, 139, 140-142, 146, 149, 151, 179-181
- Logic 62-64
- Logical gates 37-39
- Logical reasoning 63
- Long term depression (LTD) 172
- Long term knowledge base (LTKB) 37, 39, 71, 109-118, 171, 175, 176
- Long term memory (LTM) ix, 18, 21, 24, 45-48, 52-57, 71, 94, 115, 117, 118, 129, 134, 143, 152, 153, 176
- M**
- Matching bias 69
- Material equivalence (\equiv) 104, 143, 187
- Material implication (\supset) 62, 66, 70, 71, 104, 143, 187
- Mauchly test 182
- Max likelihood chi-square G^2 180
- McCulloch & Pitts neuron 12
- Mechanoreceptor cell 157
- Mechanosensory system 157
- Membrane potential 4, 5
- Membrane pumps 4
- Membrane resistance 5, 6, 13, 28
- Memory 45-58
- Memory scan effect 56
- Memory span ix, 55, 164, 169
- Mental logic 63, 64
- Mental models 63, 64
- Modus ponens 62, 72, 104, 105, 143-148
- Modus tollens 62, 63, 76, 104, 105, 143-148
- Multiple instantiation ix, 23, 24, 39-41, 56, 59, 109-154, 170, 171
- Multiple store model of memory 46-48, 55
- Myelinated axon 5
- N**
- N-ary predicates 15, 24, 40, 112
- Na⁺ 4, 14, 43, 187
- Na⁺ channels 4
- Negation effect 67, 69
- Negative conclusion bias 69, 70, 72, 77-79, 81, 82, 84, 86, 87
- Neuron 2-8
- Neuron model 1, 14
- Neurotransmitter 3, 4, 5, 13, 27
- Neurotransmitter release 3, 5, 13, 27
- NMDA 4, 26
- NMDA channels 4, 26
- Noise ix, 9, 13, 14, 23, 24, 26, 30, 35, 50, 54, 155-167
- Noisy delay 26, 27, 165, 186

- Non-linear system 156
 Non-symmetric inferences 111
 Normality assumption 101, 126, 181
 Normalized correlation 42, 44, 50
 Novices 55
 Nucleus 2
- O**
- Obligation schema 68
 Occam's razor principle 55
 Olfactory system 11
 OR-gate 37, 38
 Order of complexity 176
 Oscillation viii, ix, xii, xiii, 6-11, 23, 31, 32, 35, 36, 40, 42, 48, 55, 59, 61, 72, 93, 115, 117, 169, 170, 187
 Oscillation frequency viii, 6, 35, 48, 93, 117
 Outer product 17
- P**
- Parallel distributed processing 112
 Parametric test 181
 Parietal cortex 22
 Parsing 20, 171
 Perception vii, 9, 45, 92, 113, 174, 175
 Perceptron 12
 Peri-stimulus time histogram 9, 10
 Period doubling 40, 117, 129, 152, 153, 170
 Permission schema 66
 Phase 2, 10, 11, 24, 31, 35, 36, 40, 49, 50-52, 56, 72, 73, 76, 90, 93-104, 106, 108, 122, 123, 129-135, 144, 145, 170, 175
 Phase-based models 93
 Phase code 10, 170
 Phonological similarity 48, 55, 113
 Phonological store 55
 Postal rule 67
 Post-hoc Tukey 80, 83, 84, 107, 137, 140, 147
 Post-synaptic excitation 27
 Post-synaptic inhibition 27
 Post-synaptic potential (PSP) 6, 13, 21, 22, 27, 165-167, 170, 186
 Potentiation 37
 Pragmatic reasoning schemas 63-65, 68
 Precision of synchrony 35, 41, 48, 93, 169
 Predicate ix, 15, 16, 18-24, 33-35, 39-41, 59, 60, 89-118, 134, 143, 144, 153, 169, 171-174
- Predicates & dimensions of variation 89
 Predicates & cognitive complexity 89
 Pre-synaptic amplification 27
 Pre-synaptic facilitation 27
 Pre-synaptic inhibition 27
 Primacy effect ix, 47, 52
 Priming 153
 Proactive inhibition 47
 Probabilistic model of reasoning 68
 PropAI 64
 Pseudo-memories 39, 112
 PSYCOP 63
 Pulse code 9, 11
 Pulsed networks 12, 19, 21-25
 Purkinje cell 3
 Pyramidal cell 3
- R**
- Rasch model 78, 82, 85, 107, 146, 149, 151, 181
 Rate code 9, 11, 19, 170, 175
 Reasoning ix, 14, 20, 24, 59, 61-72, 80, 84, 86, 93, 95, 99, 103, 104, 108, 112, 113, 115, 118, 124, 128, 130, 143, 144, 147, 152, 167, 169, 170, 171, 174, 176
 Recapitulatory phase 129-135
 Recency effect ix, 47, 52
 Receptor molecules 3
 Recurrent sigmoidal neural networks 12
 Recursion 23, 111, 115-117, 153
 Refractoriness 13
 Refractory period ix, 5, 6, 13, 22, 28-30, 41, 59, 120, 134, 155, 170, 186
 Relational reasoning 118, 138
 Relative refractory period 5, 29
 Repetition blindness 114
 Representation problem 174
 Reset (post-spike) viii, 13, 24, 28, 169
 Resting period 31, 49
 Resting potential 4
 Rhythmic bursts 7
 Right and duty conditional regulation 68
 ROBIN 17-19, 115, 116

- Role binding 35, 39, 104, 112, 115
- S**
- Scan effect 56, 171
- Selection task paradigm 69
- Semantic network 115
- Semantic similarity 48, 53, 113
- Sensory memory 45
- Sensory neuron 157
- Serial effect 47
- Serotonin 27
- Short term memory (STM) ix, 24, 31, 35, 36, 39, 45-56, 59, 89, 111, 112, 115, 117, 152, 153, 164, 169, 170
- SHRUTI 19-22, 93, 116, 117, 153, 170, 171
- Sigmoid function 12, 46, 51
- Signature 17, 18, 115, 116
- Similarity & multiple instantiation 143-152
- Similarity effect 48, 53, 113, 152
- Simulated annealing 156
- Social contract theory 66, 67
- Soma 2
- Somatosensory cortex 22
- Sparse coding 8
- Sparse distribution 32
- Speech rate effect 56
- Sphericity assumption 78, 79, 82, 83, 85, 86, 100, 101, 105, 107, 126, 139, 146-151, 182, 188
- Spiking neural network 14
- STAR 93, 103
- Stepwise function 44, 186
- Sternberg memory scan effect 56, 171
- Stochastic resonance 156-167
- Strabismus 23
- Striate cortex 10
- Structural zeros 136, 139, 180
- Sub-threshold states 158
- Super-conductors 156
- Syllogism 93, 94
- Symbolic models ix, 2, 14, 21, 59, 64, 103, 109, 110, 115, 170
- Symmetries 125, 128, 129, 134-143, 152
- Syn-fire chains 39, 72, 171
- Synapse viii, xi, xii, 3, 4, 27, 29, 30, 39, 56, 172, 185
- Synaptic delay 5
- Synchrony viii, ix, 9, 11, 19-24, 30-42, 48-52, 55, 59, 62, 71-73, 76, 77, 89, 93-95, 97, 103, 106, 116-118, 120, 123, 131-133, 143, 155, 165, 169, 170, 173, 175
- Systematicity 16, 92
- T**
- Temporal cortex 9
- Temporal gates 37-39, 72, 166
- Tensor product 17, 93, 103
- Terminals 3, 5, 27
- Theorem proving 14
- Theta wave 7, 31, 36, 37, 44, 73, 96, 122, 145, 187
- Threshold function 29, 46
- Threshold (noisy) 158, 160
- Tractability 16, 175, 176
- Transitive inferences 92, 111
- Transmission delays 39, 171
- Travelling salesman problem 175
- Trigram 47
- Truth table task 62, 63, 69
- Tukey test 80, 83, 84, 107, 137, 140, 147
- Type-token problem 109, 112
- U**
- Unary predicates 15
- Unmyelinated axon 5
- V**
- Verification bias 63
- Vesicles 3
- Voltage 13, 27
- W**
- Wason selection task 62, 63, 65-69
- Welch correction for inhomogeneity 124, 181, 182
- White noise 156
- Window of synchrony viii, ix, 34, 35, 48-51, 59, 60, 73, 77, 89, 93, 95, 106, 143, 155, 169
- Word-length effect 56
- Working memory (WM) 18, 21, 45, 46, 55, 113, 115
- X**
- XOR-gate 37, 38, 72, 118, 120, 121, 132
- Y**
- Yates correction 136, 139, 141, 180
- Z**
- Zeros (structural) 180