In L.R. Gleitman, & A.K. Joshi. Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society 918-923. 2000.

Simulating Conditional Reasoning Containing Negations: A Computer Model and Human Data

Jacques Sougné (J.Sougne@ulg.ac.be)
University of Liège
Department of Psychology
Bât B32, Sart Tilman
4000 Liège Belgium

Abstract

Moddeling human conditional reasoning of the type "if p then q" containing negations poses a challenge for connectionism. A network of spiking neurons (INFERNET) was used to model this type of conditional reasoning. This model also provides insights on certain human limitations. The model is compared to empirical data, and classical explanations. Statistical analysis shows that the model's performance not only surpasses classical explanations but also provides a very good overall fit to empirical data. INFERNET simulator results are also compared to human performance. The simulations compare well with both human performance and limitations.

Introduction

INFERNET (Sougné, 1996, 1998a, 1998b, Sougné & French, 1997) 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 Working Memory (WM) span and the content of WM is maintained by oscillations. Once a node is activated, it tends to fire rhythmically at a particular frequency. This technique is used to represent n-ary predicates (Sougné, 1996), relational reasoning with multiple instantiation (Sougné, 1998a; Sougné, 1998b), working memory (Sougné & French, 1997) and conditional reasoning (Sougné, 1996). This paper shows how the model handles negated conditionals.

Many psychological studies in the area of deductive reasoning have focused on conditional reasoning of the type *"if p then q."* Of course, some logicians would deny that material implication is really what humans mean by "if…then". Nonetheless, here are transcribed rules related to material implication: modus ponens (MP) *If p then q; p; infer q* and modus tollens (MT) *If p then q; ~q; infer ~p* (~ stands for not). While most humans follow modus ponens, it is different for modus tollens. People also use two inappropriate rules related to material equivalence: Denial of the antecedent (DA) *If p then q; ~p; infer ~q,* and Affirmation of the consequent (AC) *If p then q; q; infer p.* Throughout this paper the "if p then q" form will be called the "major premise", *p* the antecedent, *q* the consequent.

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 1 shows these four forms and the inferences resulting from the application of the four rules (MT, DA, AC, MT).

Table 1: 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*
If p then q	p	q	not p	not q	q	p	not q	not p
If p then not q	p	not q	not p	q	not q	p	q	not p
If not p then q	not p	q	p	not q	q	not p	not q	p
If not p then not q	not p	not q	p	q	not q	not p	q	p

Empirical studies reveal that negations do modify the frequencies of rule application (Evans, 1977, Wildman & Fletcher, 1977, Pollard & Evans, 1980). Pollard & Evans (1980) explain these data with what they call "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 that "the letter is not an X" would have a higher probability (25/26) than concluding that "the letter is an X" (1/26). Oaksford & Chater (1994) provide a similar explanation. There is also an interpretation of negation effect in terms of a "Matching bias": a tendency to verify cases that are stated in the major premise. However, this bias concerns only certain procedures like the "Wason Selection Task", the "Truth Table Task" or the "Evans construction task" in which participants have to test or verify a major premise instead of applying it. Moreover, matching bias is closely related to implicit negation (Evans, 1998). The present study focuses on explicit negation. 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 or Evans, Clibbens & Rood, 1995).

In this paper, the INFERNET simulator's performance will be compared with human data. INFERNET suggests hypothesis related to the difficulty of removing double negations. An experiment was also done in order to collect reaction time data in a production task which were not available in previous studies.

INFERNET

INFERNET is a network of spiking neurons (Maass & Bishop, 1999). In INFERNET, nodes can be in two different states: they can fire (be on), or they can be at rest (be 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 $V_i^{(t)}$ reaches a threshold, it emits a spike. After firing, the potential is reset to some resting value V_r. Inputs increase the node potential, but some part of the node potential is lost at each time step. Spiking neuron models use a post synaptic potential function. Integration of input in INFERNET is a variation of standard input integration. In INFERNET there are two main types of connections: either they act on nodes (synaptic link) or on synapses (presynaptic link). Unlike most links, these latter links act on *connections* rather than nodes (French, 1995). Moreover each of these connections can be excitatory or inhibitory. There are six types of connections: synaptic excitation, synaptic inhibition, presynaptic amplification of an excitation, presynaptic inhibition of an excitation, presynaptic inhibition of an inhibition and presynaptic amplification of an inhibition (figure 1). In addition to the weight of a connection, there is a delay parameter associated with each connection. A delay of 10 means that the effect of the presynaptic node firing on the postsynaptic node will take 10 units of time. A unit of time has been taken to simulate 1 ms. In addition, connection weights are modified by a random factor that injects white noise into the signal propagation.

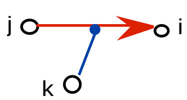

Figure 1: Example of synaptic and presynaptic connection in INFERNET. The node k inhibits the exitatory connection from j to i

The potential of node i at time t, $V_i^{(t)}$ is:

$$V_i^{(t)} = \sum_{j \in \Gamma_i} \sum_{t_j^{(f)} \in F_j} \left[w_{ij} + \sum_{k \in \mathrm{K}_{ij}} \sum_{t_k^{(f)} \in F_k} w_{k \to ij}\, \varepsilon_{k \to ij}(x) \right] \varepsilon_{ij}(x) - \eta_i(u) \qquad (1)$$

The potential of node i: $V_i^{(t)}$ is affected by connection weights coming from presynaptic node j: w_{ij} but also by the connection weights that modify this connection $w_{k \to ij}$. The set of presynaptic to node i is $\Gamma_i = \{ j | j \text{ is presynaptic to } i \}$. F_j is the set of all firing times of presynaptic nodes j: $t_j^{(f)}$. The set of presynaptic to synapse ij is $\mathrm{K}_{ij} = \{ k | k \text{ is presynaptic to } ij \text{ synapse} \}$. F_k is the set of all firing times of k nodes: $t_k^{(f)}$. These are the nodes from which start a connection acting on the connection ij. The connection weight linking node k to synapse ij is designed by $w_{k \to ij}$. The equations $\varepsilon_{ij}(x)$ and $\varepsilon_{k \to ij}(x)$ express the postsynaptic potential function. A value $\eta_i(u)$ associated with the refractory state of nodes is substracted. When $V_i^{(t)}$ reaches the threshold Θ, node i fires and V_i is reset to a resting value V_r,

Representation in INFERNET

How does the brain represent the world? Two contrasting hypotheses are often presented in neuroscience: the code used by neurons is either a rate code or a pulse code. INFERNET relies on a pulse code, specifically, phase and synchrony. In INFERNET, a symbol is represented by a cluster of nodes and is activated if its nodes fire in synchrony (the firing distribution is tightly concentrated around the mean: figure 2). Different symbols share nodes, so representations are distributed (see Sougné, 1998b), or more accurately, semi-distributed.

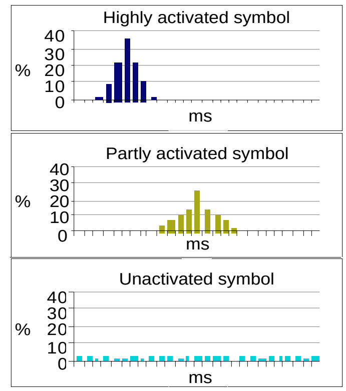

Figure 2: Symbols are represented as a set of nodes firing in synchrony.

There is considerable neurobiological evidence for considering synchrony as a possible binding mechanism in the brain (Roelfsema, Engel, König & Singer, 1996, Singer, 1993, Singer & Gray, 1995). In INFERNET, attributes are bound to an object and objects are bound to their roles by synchronous firing. For example, to represent "the red rose on the green lawn", the attribute "red" must fire in synchrony with the object "rose" and they must fire synchronously with nodes belonging to the role "supported object" (Figure 3).

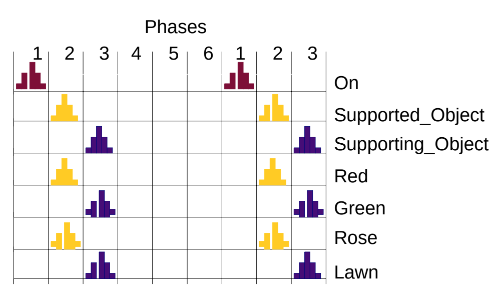

Figure 3: The "red rose on the green lawn" requires binding of symbols with their roles.

Discrimination is achieved by successive synchronies, for example, to discriminate a red rose on a green lawn. The nodes belonging to "red", "rose" and "supported object" must fire in synchrony and those corresponding to "green", "lawn" and "supporting object" must also fire in synchrony. Further, these two sets of nodes must fire asynchronously in different phases for "the red rose on the green lawn" to be perceived. Engel, Kreiter, König, & Singer (1991) provide evidence to show that if several objects are present in a scene, several groups of cells are clustered in distinct windows of synchrony.

A number of neurobiological parameters are involved in representations that rely on clusters of nodes firing simultaneously. The first is the frequency of oscillation. Certain specific oscillatory activities seem to facilitate synchronization (Roelfsema et al., 1996, Singer, 1993). In INFERNET once a node is activated, it tends to begin oscillating at a γ frequency range, whose lower limit is 30Hz and upper limit varies, according to various authors, from 70Hz (Abeles, Prut, Bergman, Vaadia & Aertsen, 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 γ waves have been observed to be associated with attention (Wang & Rinzel, 1995) and with associative memory (Wilson & Shepherd, 1995) and therefore seem to be a primary 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 and al. (1993), it is about 5 ms, sometimes less, and depends on the oscillation frequency. This allows us to approximate the number of windows of synchrony that can be differentiated, i.e., approximately 25/5 = 5, based on a typical frequency of 40Hz. If we assume that a window of synchrony corresponds to an item, a word, an idea, an object in a scene, or a chunk in working memory (WM), this puts WM span at approximately 5, with a small amount of variance since precision is proportional to oscillation frequency. This corresponds to current estimates of human WM span (see Cowan, 1998). The more the system needs to discriminate objects in WM, the more precise the synchrony should be. Since this parameter is bounded, it can lead to WM overload where windows of synchrony can no longer be distinguished. Therefore, the number of distinct items and the number of predicate arguments (Sougné, 1996) in WM is limited. Finally, following Lisman and Idiart (1995), the representation is maintained in WM by bursts of γ waves. Similar explanations for the brain's ability to store short-term memory items can be found in the literature (Hummel & Holyoak, 1997; Jensen & Lisman, 1998; Lisman and Idiart 1995; Shastri & Ajjanagadde, 1993).

Inference in INFERNET

INFERNET implements logical gates sensitive to input timing. AND-gates require all inputs to reach the target at the same time. This is achieved by a set of excitatory and inhibitory links combined with presynaptic inhibition and facilitation (see Hawkins, Kandel, and Siegelbaum, 1993, for neurobiological evidence of this mechanism). Similarly, XOR-gates are only on when one of the inputs is active. These gates are related to the phenomenon of *coincidence detection* (Konnerth, Tsien, Mikoshiba, & Altman, 1996, Singer, Engel, Kreiter, Munk, Neuenschwander, & Roelfsema, 1997).

INFERNET has a Long Term Knowledge Base that is used for encoding premises and answering queries. Figure 4 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, can correctly answer queries related to material implication.

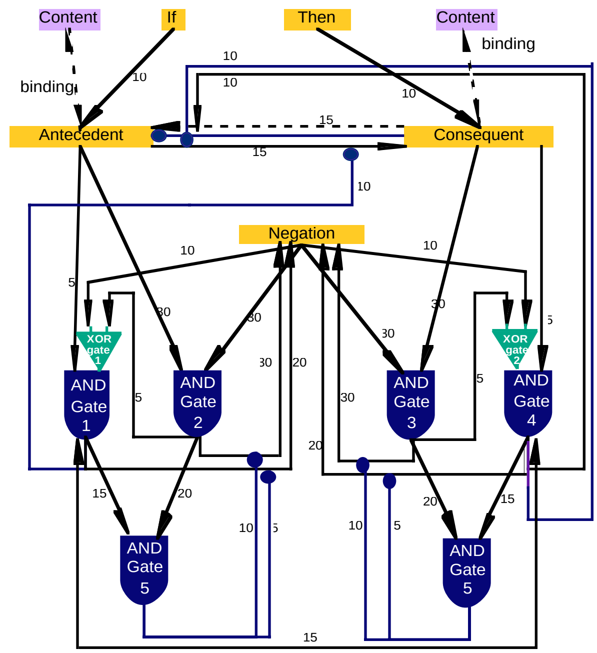

Figure 4: The encoded knowledge necessary to deal with negated conditionals

The first capacity that INFERNET must have is the ability to distinguish 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: *If ~p then q,* the connection between the *AND-gate 2* and *p* will be strengthened as well as connections between *p* and *Antecedent*. After this phase, the firing of *p* nodes will be sufficient to induce the synchronous firing of 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 it occured 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. *If ~p then ~q*), and if the minor premise is in the affirmative form (e.g. *p*), *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. 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.

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 whenever *AND-gate 5* (see figure 4) is needed, a decrease in performance should occur. This effect is due to an increase of the number of steps required to propagate the activation and to this gate's role of blocking the oscillation of *negation* nodes. *AND-gate 5* is required to treat double negations. Therefore this hypothesis predicts a decrease in DA errors for major premises *If p then ~q* and *If ~p then ~q* and a poorer MT performance for major premises *If ~p then q* and *If ~p then ~q*.

In order to contrast classical and INFERNET hypotheses frequencies of inference and reaction times will be used.

INFERNET Simulation Results

Normalized correlation between obtained data and different possible answers was computed for the 40 trials. This is a correlation between data observed and data for perfect answers. The proportion of correct responses was obtained by combining the correlations obtained on different trials, taking care to ensure that correlations are not additive (see Sougné, 1999 for computation details). INFERNET simulator results are reported in figure 5.

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.

Response times for the simulator are measured by monitoring the encoding phase. After each γ wave burst, the questions are presented and responses are collected. Since the INFERNET simulator has a resolution of 1ms, the response time is determined by the time (in ms) for the normalized correlation to reach a threshold. INFERNET simulator mean reaction times are reported in figure 6. The reaction times show that MP responses are faster than others and that a double negation results in slower reaction times.

An experiment was conducted to provide data that could be compared with INFERNET. Normally, data about negation effects on conditional reasoning do not provide reaction times and are collected with forced choice responses. The comparison between machine and human data will allow us to test INFERNET.

Experiment and comparison with INFERNET simulator

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 21.3 and SD was 2.1.

Material

Four major premises were constructed, alternating positive and negative antecedents and consequents. 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 ~X*, Negative antecedent, positive consequent: *If the number is ~3 then the letter is X*, and Negative antecedent, negative consequent: *If the number is ~3 then the letter is ~X*. Each major premise presentation was followed by four questions: *The number is 3, what do you conclude?*, *The number is ~3, what do you conclude?*, *The letter is X, what do you conclude?*, *The letter is ~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 understood it. The major premise stayed on the screen when the subsequent questions were displayed. Questions then appeared on the screen, one at the time and in random order. Participants had to answer each question. The computer recorded the time required 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.

Results

Frequencies of stating each inference are shown on figure 5. According to the "Negative Conclusion Bias" hypothesis, there should be more DA type inferences for major premises *If 3 then X* and *If ~3 then X*, more AC type inferences for major premises *If ~3 then X* and *If ~3 then ~X*, more MT type inference for major premises *If 3 then X* and *If 3 then ~X*, 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 *If 3 then ~X* and *If ~3 then ~X*, and fewer MT type inferences for major premises *If ~3 then X* and *If ~3 then ~X.*

Data were analyzed by a Loglinear analysis which provides a means to analyze multi-way frequency table. Loglinear analysis evaluates the effect of each variable and of their interaction[1]. Moreover, loglinear analysis evaluates each model that could explain the data, this gives us a way to compare INFERNET and the classical "Negative Conclusion Bias" model.

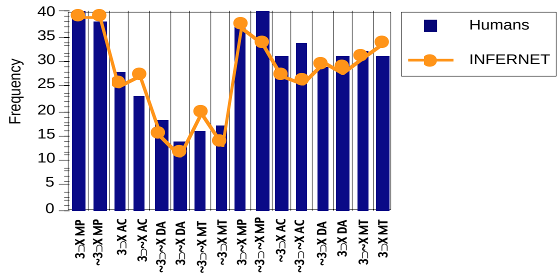

Figure 5: Graph of comparison between human and INFERNET simulator frequencies of inference.

In addition to the effect of the conclusion sign (i.e. with or without "not" in the conclusion) being significant (194 positive and 265 negative conclusions, $G^2_{(1)}=44.135$, $p<.0001$, 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), $G^2_{(1)}=59.358$, $p<.0001$. Forward inferences (MP+DA) are more often done (247 inferences) than backward inferences (AC+MT) (212 inferences), $G^2_{(1)}=11.092$, $p<.001$. 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), $G^2_{(1)}=4.893$, $p<.03$. 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, $G^2_{(1)}=30.226$, $p<.0001$. The INFERNET model is the best fitting model $G^2_{(20)}=12.88$, $p=.88$, while Negative Conclusion bias with the exception of MP cases provides a poor fit: $G^2_{(22)}=21.92$, $p=.46$. The difference between these two models is significant: $G^2_{(2)}=9.04$, $p<.01$.

The INFERNET data are not significantly different from these results. The comparison with human data can be done by adding one group factor to the analysis (Human vs INFERNET simulator). The effect of group is not significant, none of the interactions are significant.

Figure 6 shows mean reaction times for the 4 major premises and the four questions. The two hypotheses to compare are the same as above. The use of ANOVA[2] with 4 within-subject factors reveals a significant effect of the variables "Positive or Negative Conclusion": $F(1,9)=11.02$ $p<.01$ (negative conclusion bias). However, a Post Hoc Tukey test reveals that cases in which the conclusion is negative are only faster than those involving double negation. The double negation effect is significant: $F(1,9)=12.79$ $p<.01$. A post hoc Tukey reveals that reaction times for cases of "Double Negation" are significantly longer than others. INFERNET reaction times are faster than those of humans, but INFERNET does not account for the time of reading the question and producing an utterance.

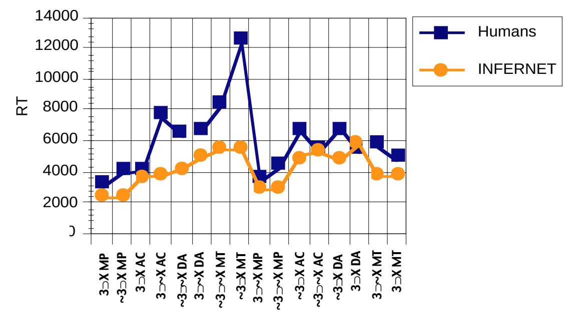

Figure 6: Graph of comparisons between human and INFERNET simulator mean reaction times.

Conclusions

Connectionist moddeling of human reasoning is a difficult challenge. Even though Holyoak & Spellman (1993) have described human reasoning in terms of constraint satisfaction, few connectionist systems has been designed for moddeling 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 paper 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.

Acknowledgments

This research was supported by the Belgian PAI Grant p4/19 Special thanks to Robert French for his assistance in the work presented here.

---

1 All the following G^2 are underestimated because data were analysed with a between subjects design. A method for analysing within designs exists but in this case, it would require a 2^{16} table to analyse. However, this would not be feasible. Note, however, that a within-subjects ANOVA gave the same results.

2 In this analysis, degrees of freedoms have been corrected because of violation of the sphericity assumption: Box correction $\hat{\varepsilon}=.23$

References

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.* Amsterdam: Elsevier.

Cowan, N. (1998). Visual and auditory working memory capacity. *Trends in Cognitive Sciences, 2,* 77-78.

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. *Proc. Natl. Acad. Sci. U. S. A., 88,* 6048-6052.

Evans, J. St. B. T. (1977). Linguistic factors in reasoning. *Quarterly Journal of Experimental Psychology, 29,* 297-306.

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.

French, R. M. (1995). *The Subtlety of Sameness: A theory and computer model of analogy-making.* Cambridge, MA: MIT Press.

Hawkins, R.D, Kandel, E. R. and Siegelbaum, S. A. (1993). Learning to modulate transmitter release: Themes and variations in synaptic plasticity. *Annu. Rev. Neurosci., 16,* 625-665.

Holyoak, K. J., & Spellman, B. A. (1993). Thinking. *Annual Review of Psychology, 44,* 265-315.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representation of structure: A theory of analogical access and mapping. *Psychological Review, 104,* 427-466.

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.

Konnerth, A., Tsien, R.Y., Mikoshiba, K. and Altman, J. (1996). *Coincidence detection in the nervous system.* Strasbourg: Human Frontier Science Program.

Lisman, J. E., & Idiart, M. A. P. (1995). Storage of 7 ± 2 Short-Term Memories in Oscillatory Subcycles. *Science, 267,* 1512-1515.

Maass, W. & Bishop, C. M. (1999). *Pulsed Neural Networks.* Cambridge, MA: MIT Press.

MacKay, W. A. (1997). Synchronized Neuronal Oscillations and their Role in Motor Process. *Trends in Cognitive Sciences, 1,* 176-183.

Oaksford, M., & Stenning, K. (1992). Reasoning with Conditionals Containing Negated Constituents. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 835-854.

Oaksford, M., & Chater, N. (1994). A Rational Explanation of the Selection Task. *Psychological Review, 101,* 608-631.

Pollard, P. & Evans, J. St. B. T. (1980). The influence of logic on conditional reasoning performance. *Quarterly Journal of Experimental Psychology, 32,* 605-624.

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.

Shastri, L. and 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.

Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annu. Rev. Physiol., 55,* 349-74.

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.

Sougné, J. (1996). A Connectionist Model of Reflective Reasoning Using Temporal Properties of Node Firing. *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society.* Mahwah, NJ: Lawrence Erlbaum Ass.

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. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society.* Mahwah, NJ: Lawrence Erlbaum Ass.

Sougné, J. (1999). *INFERNET: A neurocomputational model of binding and inference.* Unpublished doctoral dissertation, Université de Liège.

Sougné, J. and 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.* London: Springer-Verlag.

Sperber, D., Cara, F., & Girotto, V. (1995). Relevance Theory Explains the Selection Task. *Cognition, 57,* 31-95.

Wang, X. & Rinzel, J. (1995). Oscillatory and Bursting Properties of Neurons. In A. Arbib (Ed.) *The Handbook of Brain Theory and Neural Networks.* Cambridge, MA: MIT Press.

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.* Cambridge, MA: MIT Press.