Towards a unified reasoning theory:
An evaluation of the Human Reasoning Module in Spatial Reasoning

Matthias Frorath, Rebecca Albrecht, and Marco Ragni
{frorath, albrechr, ragni}@cs.uni-freiburg.de
Center for Cognitive Science
University of Freiburg, Germany

Abstract

Lately two modules which aim to improve the modeling of human reasoning have been introduced into the ACT-R cognitive architecture – the Human Reasoning Module (HRM) and the Pre-Attentive and Attentive Vision Module (PAAV). This is a first attempt to create a domain-specific infrastructure as an extension of a cognitive architecture. The HRM defines the basic functionality of two different theories in human reasoning – the mental model theory and the mental logic theory. Thus, it is a step towards a "unified theory of human reasoning". In this article we use a model of the continuity effect in spatial reasoning to evaluate this approach. The results show that the HRM is a clearly more convenient way to define ACT-R models for reasoning domains. However, as a unified theory of reasoning improvements are necessary before quantitative measures of all aspects in a reasoning task can be predicted.

Keywords: Spatial Reasoning; Human Reasoning Module; Cognitive modelling; ACT-R

Introduction

While navigating in a new environment, assembling furniture, or setting the table we always process spatial relational information. Consider the following abstract information:

> Object A is to the left of Object B.
> Object C is to the right of Object B.

The task of a reasoner is to infer a relation between the objects A and C (a *generation task*) or to test if a specific relation holds (a *verification task*). If only one arrangement (which we call in the following a *model*) can be built from the given information, we call it a *determinate problem*, otherwise an *indeterminate problem*. Indeterminate problems are empirically often more difficult than determinate problems (Johnson-Laird & Byrne, 1991). The way relational information is presented, e.g., if the information is presented *continuously*, e.g., A is to the left of B, B is to the left of C, C is to the left of D, *semi-continuously* e.g., B is to the left of C, A is to the left of B, C is to the left of D, or *discontinuously* e.g., A is to the left of B, C is to the left of D, B is to the left of C, has an impact on reasoning difficulty (e.g., Knauff, Rauh, Schlieder, & Strube, 1998; Ragni & Knauff, 2013). We call each sentence that contains relational information a *premise*. Please note that the relational information is identical, and only the order of the premises changes.

In order to explain human spatial reasoning processes various psychological theories have been proposed. Two major theories are the *mental model theory* (MMT, Johnson-Laird, 1983), that assumes that humans construct, inspect, and vary mental models; and the theory of mental logic, a *rule-based reasoning theory* (e.g., Rips, 1994). The mental model theory was recently extended by the *preferred mental model theory* and its computational model: PRISM (Ragni & Knauff, 2013). To evaluate the theories their predictions must be tested against empirical data. Consequently, we have implemented both theories (Brüssow, Ragni, Frorath, Konieczny, & Fangmeier, 2013) in the cognitive architecture ACT-R (Anderson, 2007). While predictions of the MMT can explain the data, this support is missing for the rule-based theory (Ragni, 2008).

ACT-R provides principally enough "structure", i.e., modules to model human spatial reasoning problems. Findings from psychometrics (e.g., the Block-Tapping task, Vandierendonck, Kemps, Fastame, & Szmalec, 2004) to neuroscience (Prado, Chadha, & Booth, 2011; Knauff, 2013; Ragni, Franzmeier, Wenczel, & Maier, 2014) support that humans use for spatial reasoning a specialized mental structure, e.g., an amodal or multimodal representation of objects as proposed by the MMT. Hence, researchers have proposed an extension, a specialized spatial module for ACT-R (e.g., Gunzelmann & Lyon, 2006; Lyon, Gunzelmann, & Gluck, 2008; Douglass, 2007).

Another recently introduced approach assumes two specialized ACT-R modules to model spatial relational reasoning, the *Pre-Attentive and Attentive Vision Module* (PAAV, Nyamsuren and Taatgen (2013)) and the *Human Reasoning Module* (HRM, Nyamsuren and Taatgen (2014)). The HRM was introduced with the goal "to create a unified theory of human reasoning". Its functionality includes to build models and to use inference rules (e.g., transitivity rules). The PAAV is an alternative for the default ACT-R vision module, and gives access to the newly introduced *Visual Short Term Memory* (VSTM), where analogical representations, i.e., mental models could be built.

The HRM has already been evaluated with regard to its ability to model the qualitative difference between determinate and indeterminate problems in spatial relational reasoning (e.g., the problems from Byrne & Johnson-Laird, 1989) in Nyamsuren and Taatgen (2014). Another aspect, as outlined above, is to test how the HRM deals with the way relational information is presented. Hence we created a model to analyze if performance differences induced by the continuity effect can be reproduced and which implications the continuity effect has on the processing of spatial reasoning.

The paper is structured as follows: Firstly, we summarize the core functionality provided by the PAAV and the HRM.

In a second step we introduce an empirical investigation of the continuity effect. Based on a model by Nyamsuren and Taatgen we created an extended cognitive model. This model was extended to solve problems from the empirical investigation of the continuity effect. The model's results are then compared with the empirical data. A discussion of additional properties or limitations of the modules for improving the predictions concludes the paper.

ACT-R Modules PAAV and HRM

The authors describe their intention for providing the *Human Reasoning Module* (*HRM*) as "a single system that can express different facets of reasoning", including deduction and induction, deterministic and probabilistic inference, using rules and mental models. As such, the HRM differs from most other ACT-R modules in two aspects. Firstly, the HRM includes a set of production rules instead of a mathematical function which models subsymbolic mechanisms. With these production rules the HRM is able to request other modules like the declarative module. Secondly, the HRM has direct access to the *Pre-Attentive and Attentive Vision Module's* (*PAAV*) *Visual Short Term Memory* (*VSTM*) for requesting task-specific information.

In the first step of the reasoning process the HRM relies on the VSTM's functionality to store objects. The VSTM is represented by a two-dimensional array and stores two types of information: firstly all information that was processed by the default visual buffer is placed automatically into the VSTM. Secondly, by an explicit request task objects are stored such that premise information are represented. The VSTM offers a limited capacity of objects (default: 4 objects) which are stored for a limited amount of time (default: 10 seconds). When the capacity is exceeded the oldest object is cleared from the VSTM. By requesting the visual-memory buffer a model can access an object in the VSTM. Accessing an object in the VSTM resets the associated time stamp.

The HRM uses a specific type of chunks, the *assertion* chunks, with three slots: a *property*, a *subject* and an *object*. For spatial reasoning, chunks like `(p1 ISA assertion property left-of subject A object B)` are used to represent the fact that 'object A is to the left of object B'.

During the reasoning process all premises are stored as assertion chunks in the declarative memory. As mentioned before, the HRM offers a set of production rules for two types of reasoning, a set for *forward reasoning* and a set for *backward reasoning*. Each reasoning process starts by sending a request to the *reasoner buffer*. Forward reasoning is used for generation tasks where a rather unspecified request is sent to the reasoner buffer. The task consists of infering some information that follows given the premises in the declarative memory. Backward reasoning is used for verification tasks where a mostly or completely specified assertion is sent to the reasoner buffer, e.g., specifying a subject and an object and the missing relation has to be inferred, or a completely specified assertion with property, subject and object is given that has to be verified or rejected. In our experiment we used a verification task, thus we concentrate on backward reasoning.

Backward reasoning is a three step process, the *backward reasoning pipeline*. In case one of the steps succeeds proving the assertion, the backward reasoning pipeline is stopped. In the first step (*bottom-up reasoning*) valid assertions are requested from the VSTM based on the request. One of the valid assertions is then randomly chosen. This step does not involve using production rule requests to the VSTM, but direct access to the VSTM. Thus accessing the VSTM does not cost time. In the second step (*declarative retrieval*) the assertion is requested from the declarative module. In the third step (*top-down reasoning*) inference rules (like transitivity, opposite rules) are used to prove the assertion. Inference rules and assertions are retrieved from the declarative memory and for checking the applicability of a inference rule the backward reasoning pipeline is requested recursively. Therefore, this step may include several requests to the VSTM and to the declarative memory.

Experiment

In the behavioral experiment semi-continuous (SC) and discontinuous (DC) spatial reasoning problems were tested (cp. Table 1). Semi-continuous problems consist of premises that do not have terms that appear in both, the second and the third premise. For integrating the third premise information from the first premise is necessary. Discontinuous problems, however, consist of premises that do not have terms that appear in both, the first and the second premise. As a result, at the time where the second premise is presented, no information is available how the new information can be integrated with the first premise. This leads to different complexities between SC and DC problems that are known as the premise order effect (Ragni & Knauff, 2013) or the continuity effect (Ehrlich & Johnson-Laird, 1982), respectively. The goal of this experiment is to allow for an in-depth evaluation of the presented ACT-R model and therefore the respective modules.

In the following sections we call terms or objects that appear in multiple premises *common terms*. Using the SC (right) condition from Table 1, the common term in the first and the second premise would be the term *C*.

Method

Participants. We tested fourty-five students from the University of Freiburg (27 female, mean age: 22.86 years) who received course credit or a nominal fee for their participation.

Design and Materials. In total each participant received 64 four-term series conclusion verification problems in randomized order. Each problem consisted of three premises and a conclusion. No indeterminate problems were tested, i.e., in all problems only one correct model could be constructed. Half of the problems were semi-continuous (SC), the other half were discontinuous (DC) problems. The 32 problems in each category included 16 problems with a cor-

Table 1: Empirically tested problem categories. Two problems have a semi-continuous premise order (SC), and the other two problems have a discontinuous premise order (DC). The first problem in each category demands to add terms in the model construction to the right of already included terms (right), the second problems to add terms to the left (left).

Category	Premise 1	Premise 2	Premise 3
SC (right)	B left of C	C left of D	A left of B
SC (left)	B left of C	A left of B	C left of D
DC (right)	A left of B	C left of D	B left of C
DC (left)	C left of D	A left of B	B left of C

rect and 16 problems with an incorrect conclusion. The 16 correct/incorrect problems in one problem category varied in the direction in which new terms need to be integrated in a partial model (to the left or to the right) (cf. Table 1).

Procedure. Participants were tested individually in a quite room using a computer that administered the experiment. Preceding the experiment participants received three practice trials with feedback. Each premise and conclusion was presented subsequently in a self-paced manner (indicated by a key-press). As premise terms fruit names were used, e.g., "The apple is to the left of the orange.". Participants were asked if a conclusion holds given the previously presented premises. The answer was given by pressing a key corresponding to a "yes" or "no" answer. Processing times for each premise and the conclusion were recorded as well as the given answer.

Results and Discussion

In the empirical analysis we used linear mixed-effect models. We examined a possible influence of the premise relations on the processing times, also known as the figural effect (Knauff, Rauh, & Schlieder, 1995). The investigated categories can be seen in Table 2.

Table 2: Analyzed problem categories for discontinuous problems. All four problems have a discontinuous premise order (DC). The first relation names the direction where terms of the second premise are inserted. The second relation (with -of) names the relation inside each premise.

Category	Premise 1	Premise 2	Premise 3
Right, Left-of	A left of B	C left of D	B left of C
Left, Left-of	C left of D	A left of B	B left of C
Right, Right-of	B right of A	D right of C	C right of B
Left, Right-of	D right of C	B right of A	C right of B

In each premise phase we did not find a significant difference between the four conditions. In the following data analysis we collapsed all four conditions to one single DC condition (premise 1: F(3,38) = 1.09, p = 0.36; premise 2: F(3,38), p = .11; premise 3: F(3,39) = 0.23, p = .88). This was conducted analogously for the SC condition.

We analyzed differences in the processing time in each premise and the conclusion for semi-continuous (SC) and discontinuous (DC) problems. In the first premise (F(1, 40) = .45; p = .51), the second premise (F(1, 42) = 2.81; p = .1) and the conclusion (F(1, 41) = .26; p = .61) no significant differences in processing time between semi-continuous and discontinuous problems were found. In the third premise the processing time in semi-continuous problems was significantly shorter than in discontinuous problems (F(1, 42) = 22.56; $p < .0001$). Figure 1 shows a graphical overview of the results. These results support the assumption that in discontinuous problems additional processes, especially in the third premise, are necessary to process the respective premise and to integrate information into the mental model.

Additional to the differences in semi-continuous and discontinuous problems we investigated which strategy could be used in discontinuous problems. In these problems the first and second premises do not share a common term. As a result, no integrated representation can be constructed after the second premise was presented. Two different strategies have been proposed in the literature to handle this situation; (1) both premises are integrated into one mental model (e.g. adding the terms of the second premise to the right of the terms of the first premise), the information from the third premise then may result in a belief revision (Nejasmic, Krumnack, Bucher, & Knauff, 2011); (2) For both, the first and second, premises separate models are constructed, when the third premise is presented those models are integrated. In order to determine which strategy was used by the participants, we analyzed the response time in the third premise. In strategy (1) it was proposed that terms in the second premise are integrated into the mental model. If this information is inconsistent given the third premise a belief revision mechanism is used. As a result, the response time in the third premise should be significantly higher if belief revision is necessary (e.g., DC (right) vs. DC (left) in Table 1). This effect was not found (F(1, 41) = 0; p = .98).

ACT-R Model

Model Description

Original model by Nyamsuren and Taatgen (2014). They created a cognitive model to demonstrate the functionality of the HRM and the PAAV's VSTM in spatial reasoning tasks. For their demonstration Nyamsuren and Taatgen chose determinate and indeterminate five-term problems in a generation task setting with two-dimensional relations as introduced in Byrne and Johnson-Laird (1989). The model clearly focused on mental model construction in the VSTM and the reasoning processes in the HRM. Therefore, premises were placed directly in ACT-R's imaginal buffer. In order to represent all terms in one problem in the VSTM the capacity was increased to five objects.

For evaluating the quality of their model and, therefore, their modules they compared correctness in determinate and indeterminate problems with existing experiments (e.g., Byrne and Johnson-Laird (1989)). The model predicts 100% correctness in determinate problems. In indeterminate problems, due to several possible valid mental models, errors only occur during mental model modification.

An important aspect of Nyamsuren and Taatgen's model is the integration of new information into the VSTM. When the second premise is presented the terms of the first premise have already been included into the VSTM. In order to integrate the terms of the second premise the model checks whether the object of the corresponding assertion is already in the VSTM. If this is the case, the subject is integrated into the VSTM with regard to the relation represented by the property. If the object is NOT yet in the VSTM the *opposite* relation is requested from the HRM. As a result, the model makes some interesting predictions about the time to integrate premises into the VSTM.

Extension for the Continuity Effect. In order to model the continuity effect with the HRM and PAAV, we extended the cognitive model by adding the functionality necessary to solve discontinuous problems.

Instead of placing premises into the imaginal buffer directly, we used the ACT-R vision module to read premises and conclusions from the screen. Premises and conclusions are presented in the form of ”A L B” representing the sentence ”A is to the left of B”. Similarly, in order to include the self-paced character of our experiment the model pressed keys (using the manual buffer) to view the next premise or the conclusion. As in the experiment processing times and answers were recorded.

We presented all premise information visually, so all chunks from the visual buffer were inserted into the VSTM automatically. This had an impact on the model's predictions. Including the manually added term information a maximal capacity of 13 objects was necessary. In order to counter this effect we extended the model by a rehearsal mechanism. Each time premise information is integrated in the VSTM and before the model views the next premise all manually added objects into the VSTM are rehearsed. As a result, only the visually presented premise information is removed from the VSTM and terms from the mental model are held in the VSTM. With this mechanism still a maximum capacity of 7 objects is necessary to hold the complete mental model in the VSTM. Note that the rehearsal mechanism has no significant impact on differences in the evaluated problem categories.

In order to also allow for the processing of discontinuous problems additional changes to the model were necessary. As explained in the Experiment Section there are two possible strategies how discontinuous information can be integrated. Empirical data suggests that premise information are only integrated after they can be linked by a common term, i.e., when the third premise has been presented. However, different implementations of this strategy are possible. We chose the strategy to keep the first premise inside the VSTM and to store the information from the second premise in an assertion chunk in the declarative memory. Once the third premise is presented and integrated into the VSTM the assertion chunk is recalled from the declarative memory. The model checks if there are now common terms in the assertion and the VSTM and when a common term is found, the information from this assertion is integrated into the VSTM.

Model Evaluation

Premise processing times. Figure 1 shows a comparison between the model predictions and the empirical data for the processing time of the three premises in the semi- and discontinuous problems.

For evaluating our model for the continuity effect we compared model predictions with the overall response time and the processing time for all premises. In the premise processing phases the overall response times could not be predicted. The integration of all the premises is too fast compared to the human reaction times. For this reason we concentrated on comparing the qualitative reaction time differences. We found that for the first and the third premise the reaction time trends could be modeled. While the first premise has no significant difference for both continuities no differences could be found in the model predictions. For the third premise a significant difference between semi-continuous and discontinuous problems could be found. The model predicts higher processing times for discontinuous problems due to the integration of the postponed integration of the second premise.

For the second premise we did not found a significant difference between semi-continuous and discontinuous problems. The model, however, does predict a lower processing time for discontinuous problems than for semi-continuous problems. After noticing that no common term can be found in the VSTM the model stores the second premise in the declarative memory without integrating terms into the VSTM.

Correctness and Conclusion Answer Time. In order to evaluate our choice concerning the maximum capacity of the VSTM we used correctness as a measure. Table 3 shows the correctness for correct and incorrect conclusions of the empirical data (H) and for capacities of 4 (default) to 7. For incorrect conclusions the model predicts a correctness of 100% for each capacity. For correct conclusions the correctness drops from 100% to under 50% when the capacity is not sufficient to store all terms in a problem. There is no capacity which predicts the human correctness. The reason for this rapid drop in correctness is the switch between the mental model approach and the rule-based reasoning approach in the top-down reasoning process. Inference rules (e.g., transitivity) are used to validate a conclusion. This process includes several requests to the declarative module and, thus, is highly error prone. It can also be noted that with the current module implementation the correctness for incorrect conclusions cannot be lower than 100%. The reason is that there are no mecha-

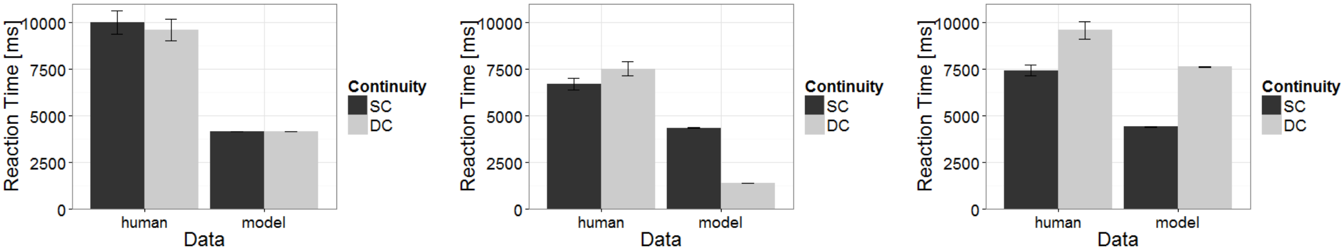

(a) First premise.

(b) Second premise.

(c) Third premise

Figure 1: Processing time for the first, second, and third premise comparing human data and model predictions for semi-continuous problems (SC) and discontinuous problems (DC).

nisms to accept an incorrect conclusion in the top-down reasoning mechanism.

A comparison between the empirical data and the model predictions for a VSTM capacity of 7 shows that the time to reject an incorrect conclusion is significantly higher than to accept a correct one (cp. Fig. 2). In the latter case only the VSTM needs to be checked. In contrast for an incorrect conclusion all three steps including the top-down reasoning mechanism is necessary before a rejection.

Table 3: Proportions of correct answered problems in empirical data (H) and model predictions for VSTM capacities of 4 to 7 comparing correct and incorrect conclusions.

	H	4	5	6	7
Correct	0.74	0.24	0.27	0.47	1.0
Incorrect	0.84	1.0	1.0	1.0	1.0

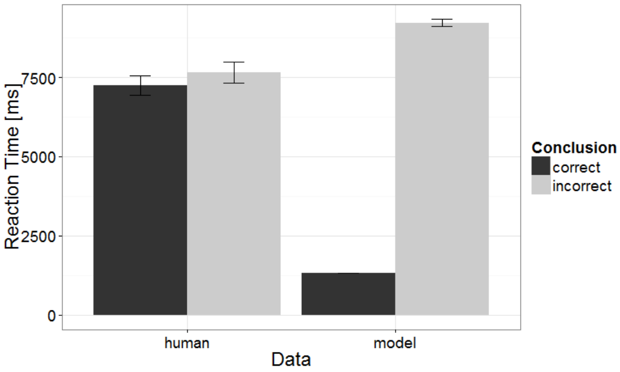

Figure 2: Time to process a conclusion and give an answer for correct and incorrect conclusions.

Conclusion

The Human-Reasoning Module (HRM) aims “to create a unified theory of human reasoning” (Nyamsuren & Taatgen, 2014). The module achieves this mainly by introducing two major changes to the ACT-R philosophy. (1) The HRM includes a set of production rules instead of a mathematical function which models subsymbolic mechanisms; (2) the HRM has direct access to other modules, especially the PAAV. (1) introduces a certain flow of information and control into the ACT-R system which cannot be influenced by the modeler. This introduction of limitations to the ACT-R architecture restricts what can be explained and adds effects that cannot be explained. Of course, introduced restrictions to a cognitive architecture must be theoretically sound and empirically validated. The aim of this paper is to evaluate these introducted restrictions by analyzing a model for the continuity effect.

The presented model is able to predict empirical data for the processing of the first and third premise, but not for the second premise. In discontinuous problems, the second premise cannot be integrated into the existing mental model. Our model stores this premise in declarative memory to recall it later. Other approaches, like the spatial buffer of Douglass (2007) or in the Spatial and Visual System of the cognitive architecture SOAR (Wintermute, 2009) use hierarchical spatial objects. An extension by such a mechanism should be evaluated.

Especially in the construction phase of the mental model overall response times could not be predicted. The model is too fast in the integration of all the premises. The HRM does not include any assumptions on this phase in spatial reasoning. The switch from a mental model-based approach to a rule-based approach only occurs when a conclusion must be validated. Thus, additional process assumptions should be considered, e.g., focus operations defined in the PRISM model (Ragni & Knauff, 2013).

The correctness could not be predicted as well for several evaluated capacities associated with the mental model representation. For incorrect conclusions the correctness is in each case 100%. In case of correct conclusions the correctness drops from 100% to under 50% as soon as the capacity is not sufficient to hold all terms in a problem. Orthogonally, the response time for a correct conclusion is significantly lower than in the empirical data if all information is accessible in the VSTM. In order to address these issues a decay mechanism instead of a fixed capacity should be considered to limit the mental representation.

To conclude, the HRM and PAAV are interesting approaches to allow for a more convenient model definition and to introduce restrictions to ACT-R. Additional improvements are possible to better predict empirical effects in higher-level cognition using ACT-R.

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* New York: Oxford University Press.

Brüssow, S., Ragni, M., Frorath, M., Konieczny, L., & Fangmeier, T. (2013). Premise annotation in mental model construction: An ACT-R approach to processing indeterminacy in spatial relational reasoning. *Cognitive Systems Research*, *24*(1), 52–61.

Byrne, R. M., & Johnson-Laird, P. (1989). Spatial reasoning. *Journal of memory and language*, *28*(5), 564–575.

Douglass, S. A. (2007). *A computational model of situated action*. dissertation, Carnegie Mellon University.

Ehrlich, K., & Johnson-Laird, P. N. (1982). Spatial descriptions and referential continuity. *Journal of verbal learning and verbal behavior*, *21*(3), 296–306.

Gunzelmann, G., & Lyon, D. R. (2006). Qualitative and quantitative reasoning and instance-based learning in spatial orientation. In *Proceedings of the 28th annual conference of the cognitive science society* (pp. 303–308). Vancouver, British Columbia, Canada.

Johnson-Laird, P. N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, MA, USA: Harvard University Press.

Johnson-Laird, P. N., & Byrne, R. M. J. (1991). *Deduction*. Hillsdale, NJ: Erlbaum.

Knauff, M. (2013). *Space to Reason: A Spatial Theory of Human Thought*. MIT Press.

Knauff, M., Rauh, R., & Schlieder, C. (1995). Preferred mental models in qualitative spatial reasoning: A cognitive assessment of allen's calculus. In *Proceedings of the seventeenth annual conference of the cognitive science society* (pp. 200–205).

Knauff, M., Rauh, R., Schlieder, C., & Strube, G. (1998). Continuity effect and figural bias in spatial relational inference. In *Proceedings of the twentieth annual conference of the cognitive science society* (pp. 573–578). Mahwah, NJ: Erlbaum.

Lyon, D. R., Gunzelmann, G., & Gluck, K. A. (2008). A computational model of spatial visualization capacity. *Cognitive Psychology*, *57*, 122–152.

Nejasmic, J., Krumnack, A., Bucher, L., & Knauff, M. (2011). Cognitive processes underlying the continuity effect in spatial reasoning. In *Proceedings of the 33rd annual conference of the cognitive science society* (pp. 1127–1132).

Nyamsuren, E., & Taatgen, N. A. (2013). Pre-attentive and attentive vision module. *Cognitive Systems Research*, *24*, 62–71.

Nyamsuren, E., & Taatgen, N. A. (2014). Human reasoning module. *Biologically Inspired Cognitive Architectures*, *8*, 1–18.

Prado, J., Chadha, A., & Booth, J. R. (2011). The brain network for deductive reasoning: a quantitative meta-analysis of 28 neuroimaging studies. *Journal of Cognitive Neuroscience*, *23*(11), 3483–3497.

Ragni, M. (2008). Human logic in spatial reasoning. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 933–939). Austin, TX: Cognitive Science Society.

Ragni, M., Franzmeier, I., Wenczel, F., & Maier, S. (2014). The role of the posterior parietal cortex in relational reasoning. *Cognitive Processing*.

Ragni, M., & Knauff, M. (2013). A theory and a computational model of spatial reasoning with preferred mental models. *Psychological Review*, *120*(3), 561–588.

Rips, L. J. (1994). *The psychology of proof: Deductive reasoning in human thinking*. Cambridge, MA: The MIT Press.

Vandierendonck, A., Kemps, E., Fastame, M. C., & Szmalec, A. (2004). Working memory components of the corsi blocks task. *British Journal of Psychology*, *95*(1), 57–79.

Wintermute, S. (2009). *An overview of spatial processing in soar/svs* (Tech. Rep.). Technical Report CCA-TR-2009–01, Center for Cognitive Architecture, University of Michigan.