Predicting Individual Spatial Reasoners: A Comparison of Five Cognitive Computational Theories

Marco Ragni and **Paulina Friemann** and **Enver Bakija** and **Novian Habibie**
Yannick Leinhos and **Dennis Pohnke** and **Yvan Satyawan** and **Maya Schöchlin** and **Rabea Turon**
Cognitive Computation Lab, Technical Faculty
Georges-Köhler-Allee, University of Freiburg, 79110 Freiburg, Germany

Abstract

While there is a plethora of cognitive models for spatial relational reasoning, only few of them have been implemented and less have been compared to each other. Additionally, a quantitative benchmark consisting of core spatial relational reasoning problems is missing. And, if empirical data is available it reports aggregated response patterns only, and not the responses of each individual human reasoner. Accordingly, most cognitive models do only aim to explain or reproduce these aggregated response patterns. This paper approaches these issues from a cognitive computational perspective: (1) To establish a first benchmark, we conducted an experiment on reasoning with cardinal direction relations, (2) where necessary, we reimplemented existing cognitive models for spatial relational reasoning, ranging from connectionist approaches to symbolic theories and analyze these theories based on diagnostic criterias, and (3) we evaluated the cognitive models on the benchmark data and extended them where necessary to give predictions for individual reasoners. We discuss implications for theories of spatial reasoning.

Keywords: Spatial Cognition; Reasoning; Cognitive Models; Cardinal Direction; Model Comparison

Introduction

Spatial reasoning is ubiquitous. When we travel, navigate, or communicate about spatial information, we process spatial information and draw inferences. Consider the following problem about cardinal directions:

(A1) The tower is north of the city.
The city is north-west of the mountain.
Where is the mountain in relation to the tower?

A human reasoner may quickly conclude that the answer to this problem is 'The mountain is south-east of the tower'. However, a cognitive model, as you will see later, would only predict about 64% of the responses of individual reasoners correctly, as inter-individual differences are present. Predictions can become even more difficult:

(A2) The train station is north-west of the library.
The library is south-east of the church.
Where is the church in relation to the train station?

This spatial description is ambiguous, or *indeterminate*, i.e., there are several arrangements possible. Hence, it has no straightforward answer – all relations are possible. When looking at the aggregated data, a correct prediction is almost impossible, since the answers are, with a few exceptions, nearly uniformly distributed. Therefore, to better understand the cognitive processes behind the integration of spatial information, we chose to compare various cognitive models of spatial relational reasoning on their performance on modeling the responses of individual participants.

The remainder of the paper is structured as follows: In the next section we briefly report the conducted experiment on reasoning with cardinal directions and the framework which was used for the comparative quantitative evaluation on individual empirical data. We will then introduce the cognitive models we (re-)implemented and extended for the experimental task. Lastly, implications for the domain of spatial reasoning are discussed.

Human Reasoning with Cardinal Directions

As a benchmark for spatial relational reasoning, we conducted an experiment about transitive inferences with cardinal direction relations in the line of Ragni and Becker (2010). As already mentioned, spatial descriptions can be determinate or indeterminate. In the case of oppositional directions (as in problem (A1)), this is easy to see. However, also other descriptions can be indeterminate, depending on the interpretation. The description 'A is south-west of B. B is north of C.' can lead to various representations including the following, depending on the assumed interpretation of the distances conveyed by the relation:

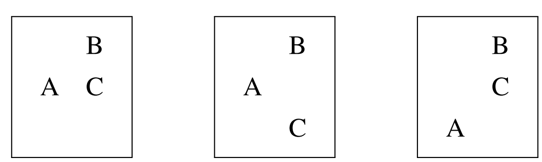

According to various studies (e.g., Knauff, Rauh, & Schlieder, 1995; Ragni & Knauff, 2013), humans do not simply construct all possible models, but have often a preference for one specific representation - the *preferred mental model*. Only if prompted to search for other models, for instance by the experimental task, will humans consider the other possibilities. Preferences for reasoning with cardinal directions have been investigated in Schultheis, Bertel, and Barkowsky (2014) and Ragni and Becker (2010). A preference for arrangements obeying a prototypical triangle shape were found (Schultheis et al., 2014) as well as a preference for main cardinal directions, e.g., 'north' over 'north-west' (Ragni & Becker, 2010).

	NW	N	NE	E	SE	S	SW	W
NW	[SE] NE	[SE] SE	[SE,S,SW] S 0,763	[SE,S,SW] S 0.737	[ALL] E, W 0.195	[NE,E,SE] E 0.903	[NE,E,SE] E 0.844	[SE] SE
N	[SE] SE	[S] S	[SW] SW	[SW] SW	[NW,W,SW] W 0.806	[N,S] S 0.529	[NE,E,S] E 0.857	[SE] SE
NE	[SE,S,SW] S 0.829	[SW] SW	[SW] SW	[SW] SW	[NW,W,SW] W 0.871	[NW,W,SW] W 0.935	ALL E, NW	[SE,S,SW] S 0.758
E	[SE,S,SW] S 0.81	[SW] SW	[SW] SW	[W] W	[NW] NW	[NW] NW	[NE,N,NW] N 0.829	[E,W] E, W 0.5
SE	[ALL] E, SE, W 0.915	[NW,W,SW] W 0.929	[NW,W,SW] W 0.818	[NW] NW	[NW] NW	[NW] NW	[NE,N,NW] N 0.765	[NE,N,NW] N 0.841
S	[NE,E,SE] E 0.920	[N,S] S 0.553	[NW,W,SW] W 0.844	[NW] NW	[NW] NW	[N] N	[NE] NE	[NE] NE
SW	[NE,E,SE] E 0.858	[NE,E,SE] E 0.853	[ALL] E 0.366	[NE,N,NW] N 0.833	[NE,N,NW] N 0.770	[NE] NE	[NE] NE	[NE] NE
W	[SE] SE	[SE] SE	[SE,S,SW] S 0.818	[E,W] E 0.541	[NE,N,NW] N 0.816	[NE] NE	[NE] NE	[E] E

Figure 1: Response preferences for the 64 problems of the cardinal direction experiment. The row represents the relation in the first premise (e.g., A is NW of B), the column the respective relation in the second premise (e.g., B is NE of C). In each cell the first row depicts the logically valid relations, the second row the most frequently chosen answer. In the indeterminate case, the third row contains the relative frequency of the preferred relation.

Method

We tested 49 participants in a web experiment on Amazon's Mechanical Turk. In the main part, participants were presented with 64 spatial reasoning problems with cardinal directions. All problems were of the form 'A r_1 B. B r_2 C.' with each r_1 and r_2 being one of the 8 cardinal direction relations *north*, *north-east*, *east*, *south-east*, *south*, *south-west*, *west*, and *north-west*. Instead of A, B, and C different buildings based on their frequency in the English language were used. The task of the participants was to give a relation that holds between C and A. The premises were presented sequentially in a self-paced procedure. The order of the problems was randomized separately for every participant. Participants responded by pressing the respective key/s (e.g., "nw" for north-west).

Based on previously defined exclusion criteria — more than two fast guessing responses (RTs < 0.2s), more than two wrong responses to standard problems (i.e., North-North, South-South, West-West, East-East) and telling that they wrote down the premises or drew pictures on paper — eight participants were excluded. Thus, the final sample size was $N = 41$ participants.

Results

Overall, 84.0 % of the problems were solved correctly, i.e., participants gave a valid answer. In the indeterminate cases preferences can be observed. These are depicted in Figure 1.

Evaluating Models on the Individual Reasoner

The Cognitive Computation for Behavioral Reasoning Analysis (CCOBRA) framework[1] is a benchmarking tool implemented in Python. Its goal is to test models and how good these simulate the experimental procedure of individual participants. The models are presented with the same task in the same sequence with the same response options. By providing precise responses to individual tasks, models are evaluated based on their predictive accuracy.

Models are allowed to train on a data set consisting of tasks and related human responses of individuals not present in the evaluation data. In the test phase, the models are presented with novel empirical data on which they are to give a prediction regarding the conclusion drawn by the current participant. Additionally, after predicting the response to a task, they are presented with the true response and thus allowed to adapt to an individual participant. Hence, CCOBRA extends the traditional cognitive modeling problem by moving beyond the level of aggregates. As a result, the modeling problem gets harder, but the outcomes can be interpreted more intuitively. Higher predictive scores correspond directly to a better grasp of the processes underlying an individual human reasoner's cognitive system.

We divided the gathered empirical data into a training and a test set: One third of the participants were randomly assigned to the training set, and the other participants were assigned to the test set.

[1] `https://github.com/CognitiveComputationLab/ccobra`

Five Cognitive Models for Spatial Reasoning

Model Selection Criteria

Models were selected with respect to the following selection criteria, of which every model reported here fulfills at least two: (i) the cognitive model is developed, or easily extendable, for human reasoning with cardinal directions, (ii) the model already has an implemented version or is easily implementable, (iii) the model makes a prediction concerning complexity of task, (iv) the model is a stand-alone implementations, (v) the model offers explanations for basal principles of spatial reasoning. We identified six cognitive models for spatial reasoning existent in the literature and contacted the respective authors. The six models were: The Spatial Probabilistic Model (Ragni & Becker, 2010), Verbal Spatial Reasoning Model (Krumnack, Bucher, Nejasmic, & Knauff, 2010), the spatial architecture CASIMIR (Schultheis & Barkowsky, 2011), the Spatial Artificial Neural Network (Ragni & Klein, 2012), PRISM (Ragni & Knauff, 2013), the Dynamic Field Theory (DFT) (Kounatidou, Richter, & Schöner, 2018). The spatial architecture CASIMIR (Schultheis & Barkowsky, 2011) was not available and due to its size and dependence between long-term memory and reasoning processes it was not possible to reimplement it. In the following, we briefly report the models.

The Spatial Probabilistic Model (Ragni & Becker, 2010)

We reimplemented the spatial probabilistic model developed by Ragni and Becker (2010).

The Unit Layout Model. This model is used as a heuristic for calculating detours by computing the conditional probability of relations between different locations (Ragni & Becker, 2010). It is represented as a lookup table that contains every possible direction relation between R_1 and R_2. Example of one unit layout lookup table can be seen in Figure 2.

Gains. For representing some cognitive phenomena, gains were added to certain probabilities. E.g., the given data shows that participants prefer the direction *east* over the direction *west* if they have the choice between them. In this example the model adds a certain value (usually the value is optimized in the pre-train-function) to the probability $p(\text{“}east\text{”})$ and normalizes all probabilities (here for all directions) afterwards.

Implementation Details. Let B' be the set of cardinal relations, where each of them represents an applicable relation (in this case, direction). Given three locations a, b, c, the relations between them are stated as $R_1, R_2, R_3 \in B'$ which are applied as aR_1b, bR_1c, and aR_3c. The relative frequency of R_3 for R_1, R_2 (called $f_{R_1,R_2}(R_3)$) is parametrized in probability distribution $P_{R_1,R_2}(R_3)$. The preferred relation is then:

$$M(R_1,R_2) = \operatorname*{argmax}_{R_3 \in B} P_{R_1,R_2}(R_3) \tag{1}$$

SE-NW	S-NW	SW-NW	SW-N	SW-NE
E-NW	*a*	W-NW	W-N	W-NE
NE-NW	N-NW	NW-NW	NW-N	NW-NE
NE-W	N-W	NW-W	*c*	NW-E
NE-SE	N-SE	NW-SE	NW-S	NW-SE

Figure 2: The unit layout for R_3 = NW. Field a is to the northwest of c. All other field are uniquely labeled with relations $R_1 - R_2$. It holds for each of them that field a is R_1-wards of it and it is R_2-wards of c (Ragni & Becker, 2010).

and by using Bayes Rule in equation (Ragni & Becker, 2010), it becomes:

$$P_{R_1,R_2}(R_3) := P(R_3|R_1,R_2) = \frac{P(R_1,R_2|R_3)P(R_3)}{P(R_1,R_2)} \tag{2}$$

where $P(R_1,R_2)$ is assumed to have a unit distribution and $P(R_3)$ has a unit distribution with a probability gain for the main cardinals and gain towards the east. These gains are motivated by the given data and are added to the respective probabilities of the directions. After adding the gains, it is necessary to normalize the probabilities. As mentioned, calculation of $P(R_1,R_2|R_3)$ is done using the unit layout's lookup table that contains every possible direction relation between R_1 and R_2.

$$P(R_1,R_2|R_3) = \frac{c^{-1}_{R_1,R_1}}{\sum_{R'_1,R'_2 \in} c^{-1}_{R'_1,R'_2}} \tag{3}$$

with the cost function (Ragni & Becker, 2010):

$$c^{R_3}_{R_1,R_2} := \frac{d([a]^{R_3},[R_1,R_2]^{R_3}) + d([R_1,R_2]^{R_3}),[c]^{R_3}}{d([a]^{R_3},[c]^{R_3})} \tag{4}$$

Verbal Spatial Reasoning Model (Krumnack et al., 2010)

Verbal reasoning is based on the assumption that the human mind constructs a verbal representation of a problem, and the reasoning process then uses mechanisms similar to those of language processing to draw or validate a conclusion as proposed by Polk and Newell (1995).

The parameter-free verbal model (Krumnack et al., 2010; Krumnack, Bucher, Nejasmic, Nebel, & Knauff, 2011) suggests that individuals construct a queue of object terms in their mind that can be read like a sentence. A mental cost metric determines where a new object is inserted. It assumes that breaking links between objects costs more than creating new links, and searching for an object is more efficient in the direction of the queue. This direction is determined upon insertion of the first relation and depends on a cultural left-right preference (Maass & Russo, 2003).

Extension of the Model. The model by Krumnack et al. has been developed for one-dimensional spatial relational problems only. Hence, we expanded the model for cardinal directions, while keeping the structure of the queue. This is done by adding a direction encoding for the vertical and horizontal plane to each link, with positive values for "north" and "east", and negative ones for "south" and "west". If the angle between the direction of the new relation and the queue direction is more than 90°, the new object is inserted before the reference object, otherwise at the end of the queue. Problem (A1) generates the following queue:

$$\begin{array}{lccccc} & tower & \rightarrow & city & \rightarrow & mountain \\ vertical & -1 & & -1 & & 0 \\ horizontal & 0 & & 1 & & 0 \end{array} \tag{5}$$

To predict a response the model sums up all the direction encodings between the two objects in the queue and decodes them into cardinal directions. E.g., in the queue above, the model would sum up all the direction instructions from "tower" to "mountain", receiving a negative total in the vertical and a positive one in the horizontal plane, which means that one must go "south" and "east" to reach the "mountain" starting from the "tower". This results in "the tower is northwest of the mountain". Individual adaption of the queue direction was implemented to account for (cultural) preference.

The Spatial Artificial Neural Network (Ragni & Klein, 2012)

We adapted the implementation of an Artificial Neural Network (ANN) (e.g., Zurada, 1992) for spatial relational reasoning with cardinal directions (Ragni & Klein, 2012) to work with the CCOBRA framework to predict individual subjects' responses to spatial relational problems.

Implementation Details. As in Ragni and Klein (2012), we used one hidden layer and a full connectivity between the layers, and trained the network with backpropagation (Rumelhart, Hinton, & Williams, 1986). The network is based on calculations on point algebra, and treats x- and y-directions independently. The network is hence tested twice on each premise-pair. First, the x-direction (west-east) is calculated, and second, the y-direction (north-south). Consequently, the network consists of 2 input nodes. Three output nodes semantically describe the spatial relation between the first and the last object on the tested axis.

All parameters were tuned manually to maximize correctness on the limited data provided. Learning rate and momentum factor were tuned in 0.1 steps in a range of 0 to 1. A value of 0.1 for both parameters yielded the best results.

To train the network, we performed 10 iterations on the training set with each of the eight possible response choices respectively.

Lastly, to perform a prediction on individual participants in the test set, the given task is given to the network, again separately for x- and y-direction. The highest valued response is returned as prediction.

PRISM (Ragni & Knauff, 2013)

We re-implemented the PRISM model, an implementation of the preferred mental model theory (Ragni & Knauff, 2013). It simulates the construction of preferred mental models and can vary this preferred model to find alternative conclusions. A spatial working memory structure is operationalized by a spatial array. In short, it consists of a mental array and a spatial focus which inserts tokens into the array, inspects the array to find new spatial relations, and relocates tokens in the array to generate alternative models of the problem description, if necessary. The focus also introduces a general measure of difficulty based on the number of necessary focus operations (rather than the number of models).

Dynamic Field Theory (DFT) (Kounatidou et al., 2018)

Kounatidou et al. (2018) proposed a cognitive model to solve the preference effect for relations right or left based on the Dynamic Field Theory (DFT) (Schöner, Spencer, & the DFT Research Group, 2015). The architecture can be divided into five functional parts. The first part involves discrete conceptual nodes for the input premises whose activation is translated into continuous activation in later fields. The second functional part is the attention part which forms peaks of activation for objects that are currently attended. The third part is the scene representation in which the spatial scene as well as the color of the objects in the scene is stored. The fourth part is concerned with spatial transformations that put the objects in the correct relation according to the given premise. And the final part is concerned with the organization of all the involved processes, including starting processes, checking if processes are finished and resetting activation to their resting state after all processes of a premise are completed.

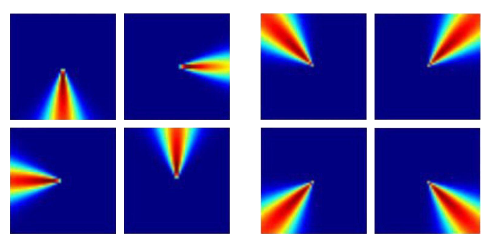

Figure 3: Cardinal direction spatial activation templates can be seen on the left, along with the new extended inter-cardinal templates on the right.

Extension of the Model. The original architecture of Kounatidou et al. (2018) can only create two-dimensional scene representations with the four cardinal relations (north, south, east, and west) between the objects. However, the benchmark data of the Cardinal Direction Experimentin-

Table 1: Overview of the evaluated models for spatial relational reasoning.

Cognitive Model	Cognitive complexity measure	Assumptions about WM representation	Predictions of phenomena	Generalizable to other relations	Accuracy predicting individuals
Verbal Reasoning	Yes	Minimal mental model	Yes	Yes	64%
Bayesian	No	None	No	No	64%
PRISM	Yes	Minimal mental model	Yes	Yes	63%
ANN	No	None	No	Yes	63%
DFT	No	Open	Yes	Yes	62%

Note. *Cognitive complexity measure* refers to whether the model gives an explicit account of the difficulty of a reasoning problem. *Assumptions about WM representation* refers to whether the models make any assumptions about the human working memory. *Predictions of empirical phenomena* refers to whether the models make new empirically testable predictions. *Generalizable* means whether the model can be extended to other spatial relations. *Accuracy predicting individuals* denotes the percentage of correctly predicted answers for the individual participants.

cludes inter-cardinal directions. Therefore, we extended the architecture with new concept nodes and corresponding spatial activation templates for the inter-cardinal directions. For these, we took the existing spatial relation templates and performed a component multiplication and normalization operation on them. For example, to get north-east, we took the product of a component multiplication between the north and east spatial activation templates and normalized it, such that the peaks were equivalent to those found in the cardinal directions. The resulting templates can be seen in Figure 3.

The implementation of this architecture was done within the CEDAR framework (Lomp, Zibner, Richter, Rano, & Schöner, 2013), which provides a way to create models based on dynamic neural fields. However, it was not possible to connect the framework to CCOBRA's evaluation function. Therefore, evaluation was performed by hand, which was possible because the model is deterministic, i.e., generates the same output for each participant.

Results and Discussion

If we just consider the accuracy to predict each individual reasoner, then we see that we reach about 64% of the predictions (see Table 1). The probabilistic approach and the Verbal Reasoning model performed the best. However, overall accuracy was very comparable for all models. Considering each single participant (see Figure 4), the different models reach a prediction accuracy of up to 90%. So a first result is: Though the models have been developed for predicting the most frequent answer, the prediction rate for many individual participants is high. It seems that the aggregate responses do capture general cognitive processes.

But the models differ in some theoretical aspects: Some do make predictions about the difficulty of problems and are *process models*, e.g., the Verbal Reasoning model and PRISM, and they do predict which symbolic mental representation is built. While this is not necessarily reflected in the accuracy, it allows to make *predictions on a phenomenological level*, i.e., the model can generate predictions about new phenomena that can be tested. This is specifically a limitation of the current version of the Spatial ANN and the Bayesian approach. They can fit the data, but testable novel predictions cannot be drawn. A further point is that models may not be too restricted to some specific spatial relations. Here extensions of the specific Bayesian approach is not straightforward.

Limitations of the approaches. The Verbal Reasoning model performed relatively well in the task it was built for: linear orderings in one dimension. However, as of yet it is not able to predict instances where a reasoner gives a logically incorrect answer. In the future, the assumptions of this cognitive model should be tested more rigorously. It could be possible that the introduction of an individual mental cost threshold would solve the problem of giving incorrect solutions. Possibilities for individual adaptation have to be explored further, since the paths used here did not improve performance. The implementation of the Neural Dynamic Field Architecture brings some limitation with it. These include the inability to rearrange existing objects in its spatial memory, to place a new premise if the reference point does not exist in its spatial memory yet, even if the target does exist, the inability to adapt to new information and the limited size of the spatial memory. The model is only able to append objects at the end of each cardinal direction (e.g., left-most position with regards to west or top-most position with regards to north). It is unable to insert objects in between two existing objects. If the architecture produces a response that is incorrect to the information or to a specific individual, it is unable to adapt or be trained specifically to respond differently. Moreover, all parameters are hard-coded and must be manually tuned.

Conclusion

The current state of the art demonstrates that it can fit about 64% of the data. The models vary to a great extend, but are very similar in their predictions. More research is necessary to understand the mechanics of human reasoning for such a simple task as transitive inferences in cardinal directions.

References

Knauff, M., Rauh, R., & Schlieder, C. (1995). Preferred mental models in qualitative spatial reasoning: A cognitive

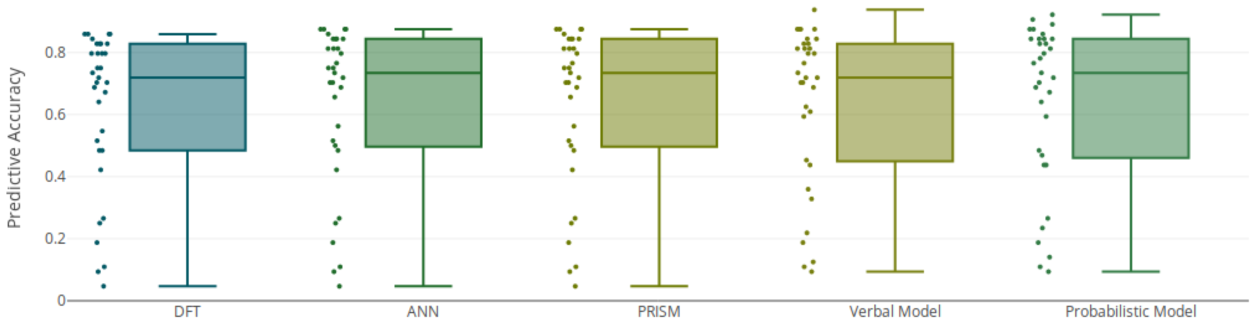

Figure 4: Percentage of correct predictions for each individual participant (dots) and for the population (box plots).

assessment of Allen's calculus. In J. Moore & J. Lehman (Eds.), *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 200–205).

Kounatidou, P., Richter, M., & Schöner, G. (2018). A neural dynamic architecture that autonomously builds mental models. In C. Kalish, M. A. Rau, X. J. Zhu, & T. T. Rogers (Eds.), *Proceedings of the 40th Annual Meeting of the Cognitive Science Society.*

Krumnack, A., Bucher, L., Nejasmic, J., & Knauff, M. (2010). Spatial reasoning as verbal reasoning. In S. Ohlsson (Ed.), *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society.*

Krumnack, A., Bucher, L., Nejasmic, J., Nebel, B., & Knauff, M. (2011). A model for relational reasoning as verbal reasoning. *Cognitive Systems Research*, *12*, 377–392.

Lomp, O., Zibner, S., Richter, M., Rano, I., & Schöner, G. (2013). A Software Framework for Cognition, Embodiment, Dynamics, and Autonomy in Robotics: Cedar. In *Artificial Neural Networks and Machine Learning – ICANN 2013* (Vol. 8131, p. 475-482). Springer Berlin Heidelberg.

Maass, A., & Russo, A. (2003). Directional bias in the mental representation of spatial events: Nature or culture? *Psychological Science*, *14*, 296–301.

Polk, T. A., & Newell, A. (1995). Deduction as verbal reasoning. *Psychological Review*, *102*, 533.

Ragni, M., & Becker, B. (2010). Preferences in cardinal direction. In S. Ohlsson (Ed.), *Proceedings of the 32nd annual meeting of the cognitive science society.*

Ragni, M., & Klein, A. (2012). Deductive reasoning - using artificial neural networks to simulate preferential reasoning. In J. Filipe & J. Kacprzyk (Eds.), *Proceedings of the 4th International Joint Conference on Computational Intelligence* (pp. 635–638).

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

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, *323*(6088), 533–536. doi: 10.1038/323533a0

Schöner, G., Spencer, J. P., & the DFT Research Group. (2015). *Dynamic thinking: A primer on dynamic field theory*. Oxford University Press.

Schultheis, H., & Barkowsky, T. (2011). Casimir: An Architecture for Mental Spatial Knowledge Processing. *Topics in Cognitive Science*, *3*(4), 778–795.

Schultheis, H., Bertel, S., & Barkowsky, T. (2014). Modeling Mental Spatial Reasoning About Cardinal Directions. *Cognitive Science*, *38*(8), 1521–1561.

Zurada, J. M. (1992). *Introduction to Artificial Neural Systems*. West Publishing Company St. Paul.