The Thirty-Third International
FLAIRS Conference (FLAIRS-33)

An Adaptive Model for Cognitive Reasoning

Jonas Bischofberger, Marco Ragni*
Cognitive Computation Lab, University of Freiburg, 79110 Freiburg, Germany
JonasBischofberger@web.de, ragni@cs.uni-freiburg.de

Abstract

How humans reason about syllogistic statements is a problem that currently lacks a comprehensive, universally accepted explanatory theory. The goal of this article is twofold: First, it sheds light on the actual predictive quality of existing theories by providing a standardized implementation of a subset of them. To that end, the theories are algorithmically formalized, including their capabilities for adaptation to an individual reasoner. The implementations are modular with regard to mental operations defined by the cognitive theories. Based on such operations, a novel composite approach is devised, resulting in a prediction model for predicting an individual reasoner before she draws the inference. It uses sequences of operations, selected from possibly different theories, to form its predictions. Among the basic models, our implementations of PHM, mReasoner and Verbal Models make the best predictions. The composite model is able to significantly surpass it by exploiting synergies between different models. Therefore, the composite approach is a promising tool to model and study syllogistic reasoning and possibly other reasoning tasks as well.

Introduction

Reasoning is one of the most distinguishing human characteristics compared to other living creatures. One particular domain in the study of reasoning concerns itself with syllogistic inferences. Such an inference may look like this:

No optician is a camper.
Some campers are barkeepers.

Therefore: Some barkeepers are not opticians.

In experiments, reasoners often demonstrate systematic deviations from the classical logical conclusion to syllogisms (see for example Begg and Denny 1969). In the last century, multiple theories have been developed to explain patterns in human responses to syllogistic reasoning tasks. None of these theories seems to be fully satisfactory, as discussed in a meta-analysis by Khemlani and Johnson-Laird (2012). However, there is no comprehensive comparison of syllogistic theories yet that takes their adaptive potential towards individual reasoners into account. The goal of this article is to present explicit and modular implementations with respect to mental operations for a number of syllogistic theories. The set of all operations across different theories is used to create a novel composite model of reasoning. Our evaluation shows that this model exceeds the predictive performance of its constituent models by finding new and appropriate mental processes.

Background

A syllogism is a kind of logical argument made up of categorical propositions. The four types of propositions are: *All X are Y*, *Some X are Y*, *No X are Y* and *Some X are not Y*. A syllogism consists of two categorical premises which share exactly one term, often denoted as *B*. When terms like *optician* and *camper* are replaced by generic terms, there are 64 different kinds of syllogisms, like the following one.

All A are B.
Some B are C.

A well-formed conclusion to a syllogism relates the two terms that are not shared. These terms are often denoted with *A* and *C*. With four types of propositions and two ways to order *A* and *C*, there are eight possible conclusions. A special response in many experiments is *No valid conclusion* (NVC), indicating that nothing follows.

There are at least twelve theories that seek to explain and model reasoning on syllogisms. They can roughly be categorized into three domains (Khemlani and Johnson-Laird 2012). First, *heuristic theories* proposing that conclusions are drawn quickly and intuitively, based on apparent features of the syllogism. Second, *rule-based theories*, which propose that inference takes place deliberately by applying formal inference rules to mental representations of propositions, similar to logical deduction. Third, there are theories using *sets, diagrams or models* rather than propositions as mental representation. These theories suggest that reasoning happens by encoding, manipulating and drawing conclusions from such representations.

*The project has been partially supported by DFG grants RA1934/3-1 and RA1934/4-1

Basic Models

The reasoning task to be modeled is defined by its possible input items and responses. The set of input items C_{Item} contains the 64 types of syllogisms. $C_{Responses}$ contains all subsets of the eight possible conclusions plus NVC. A model needs to implement a parameterized prediction function $f_p : C_{Item} \rightarrow C_{Responses}$. The internal parameters p can then be fitted to seen pairs of item and response. A final prediction is obtained by choosing uniformly among the responses predicted by f_p.

From the twelve theories of syllogistic reasoning the following seven have been formalized and implemented: Atmosphere, Matching, Illicit Conversion and PHM as heuristic theories, PSYCOP as rule-based theory and Mental Models (both mReasoner and a classical version without heuristics) and Verbal Models as model-based theories.

Exemplarically, Verbal Models (Polk and Newell 1995) proposes three different operations: Encoding, concluding and reencoding. Initially, the premises are encoded into a mental model. A mental model is a representation of a situation in which the presented propositions are true. For example, the syllogism *All A are B. Some B are not C.* quantifies three sets and if we consider individual elements the following model could be mentally formed:

```
a'  b'  -c   (one individual)
a'  b'       (another individual)
```

VM draws conclusions from a mental model according to certain pattern-matching rules. If no conclusion is found, the problem is reencoded, leading to a new mental model. For reencoding, VM chooses a term in the model and tries to extract additional information about it from the premises. This process is repeated until a mental model is found which entails a conclusion or all available additional information has been considered.

A simpler theory than VM is the Atmosphere theory (Woodworth and Sells 1935; Revlis 1975). Its core operation is a heuristic that determines the type of an accepted conclusion from the types of premises in a specific way: The quantifier of the conclusion is particular if at least one of the premises is particularly quantified, otherwise the quantifier is universal. A similar rule holds for the quality of the quantifier. The quantifier of the conclusion is negative if at least one of the premise quantifiers is negative, otherwise the quantifier is universal. In that sense, particular and negative quantifiers are dominant. The Atmosphere heuristic naturally leads to an algorithmic prediction model by choosing those conclusions for prediction which are of the type that follows from the heuristic.

A Composite Approach to Modeling Reasoning

The modular algorithmic implementations of different reasoning models form the basis for an abstract reasoning model that replaces concrete operations with classes of operations and concrete mental representations with a set of unified mental representations. That way, operations from different theories can operate on a common representation,

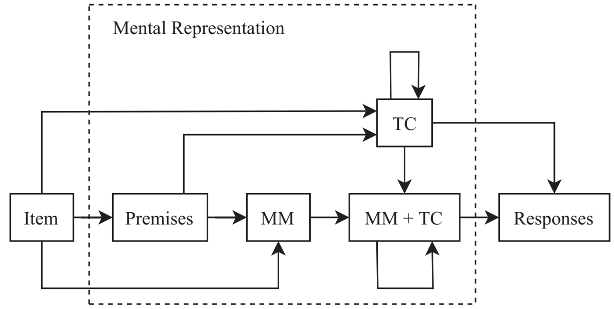

Figure 1: Abstract reasoning model for syllogistic reasoning.

opening up the possibility of combining these operations into new mixed-model reasoning paths.

An abstract reasoning model is defined as a graph with a set of states S and a set of directed edges E. For every $s \in S$, a set of possible content C_s is defined. The set C_{Item} contains the possible input items and $C_{Responses}$ all subsets of possible responses. The pair $(C_{Item}, C_{Responses})$ defines the reasoning task.

Considering the implemented syllogistic theories, we use three different types of mental representations: First, tentative conclusions $C_{TC} = C_{Responses}$ as produced by heuristics or derived from a mental model, including NVC. Second, a mental model representation C_{MM} that unifies mental and verbal models. And third, sets of two to four premises $C_{Premises}$ as produced by Illicit Conversion. Additionally, there can be pairs of mental models and tentative conclusions $C_{MM+TC} = C_{MM} \times C_{TC}$ that are kept in mind simultaneously. The according abstract reasoning model is shown in Figure 1. Nodes in the abstract reasoning model correspond to classes of content, edges correspond to classes of operations. By considering operations as being related to edges in the abstract reasoning model, we can use sequences of operations that lead from the initial state $Item \in S$ to the target state $Responses \in S$ to make predictions. The set of prediction sequences can be obtained by expanding a tree where each node $N = (s, c, v)$ corresponds to a state $s \in S$ in the abstract reasoning model and contains some specific content $c \in C_s$ associated with this state. Additionally, every node holds a particular variable assignment $v \in V$ that is used to define applicability constraints of an operation, for example to restrict the length of prediction sequences.

An operation $\Omega = (O, f_{var}, \phi)$ consists of several elements: First, the actual content transformation operation $O : C_{pre} \rightarrow C_{post}$ that is defined on the edge $(pre, post) \in E$. Implicitly, O defines a state transition from the state $pre \in S \setminus \{Responses\}$ to the state $post \in S$. Additionally, the application of Ω may alter the current state of variables via a variable transition function $f_{var} : V \rightarrow V$, for instance by incrementing a counter. The function $\phi : S \times V \rightarrow \{True, False\}$ defines the applicability of Ω, depending on the variables and state of a node, but not on its content. An operation Ω is *applicable* in a node $N = (s, c, v)$ iff $\phi(s, v)$ evaluates to true. A check for $s = pre$ is by default included in ϕ. An applicable operation

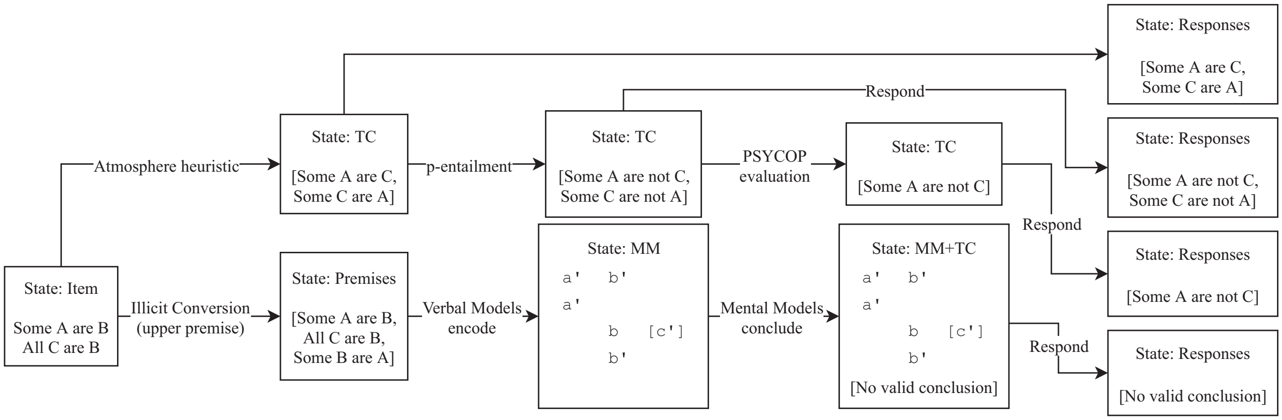

Figure 2: Exemplary subtree of T_i for i = *Some A are B. All C are B* with four prediction sequences. For every node, state and content are depicted. Values of variables are not shown.

$\Omega = (O, f_{var}, \phi)$ can be applied to a node $N = (pre, c, v)$ to yield a new node $\Omega(N) = (post, O(c), f_{var}(v))$. Now, for every $i \in C_{Item}$, a prediction tree T_i is expanded, starting with the initial node $(Item, i, v_{init})$. Child nodes are added to a node by application of an operation until no more operations are applicable to any node. A subgraph of such a prediction tree is shown in Figure 2.

A prediction sequence in T_i is a path from the root node to some leaf node $L = (Responses, c, v)$ that corresponds to the target state $Responses$. Since the applicability of operations depends only on the state and the variable assignment, but not on the content of a node, T_i looks the same for each $i \in C_{Item}$ except for the content of its nodes. In particular, they have the same set of prediction sequences, which makes them comparable across different items.

Evaluation and Discussion

All models are evaluated using the CCOBRA framework[1]. For training and testing, the data provided in the CCOBRA repository is used. The `Veser 2018` dataset with 2058 items over 33 participants is used as training data and the `Ragni 2016` dataset including 8896 participants over 139 subjects is used for evaluation. For comparison, four benchmark models are added. The models' overall predictive accuracies are shown in Figure 3.

Basic models Every implemented model is better than guessing and thus captures some amount of structure in the data. The simple heuristic models Atmosphere and Matching show the lowest performance, predicting between 22 and 24 percent of responses correctly. In contrast to these two models, the remaining six basic models all come with the ability to predict NVC and, as expected, perform above the NVC benchmark. The non-heuristic Mental Models version, PSYCOP and Illicit Conversion constitute a mediocre block, predicting around 35 percent of the responses correctly.

The strongest basic models are Verbal Models, mReasoner and PHM, which almost reach the MFA benchmark. Both Verbal Models (Polk and Newell 1995) and mReasoner (Khemlani and Johnson-Laird 2016) have previously been used successfully to explain individual differences in syllogistic reasoning. This evaluation underlines the strength of these models in adaption to individual reasoners. The strongest basic model is PHM, which is the only theory in the top group that is not model-based. The strong performance of PHM provides evidence that not only model-based approaches but also a purely heuristic theory can perform close to the MFA benchmark in syllogistic reasoning.

Composite model The composite model outperforms the best basic models, predicting about 44 percent of the responses correctly. It also slightly exceeds the MFA benchmark, which strongly suggests that it successfully adapts to individual reasoning patterns. For comparison, a version of the composite model was implemented which allows only prediction sequences containing operations from one and the same basic model. This version gets only around 42% of its predictions correct. Thus, the improvement that the composite model provides over its constituent models seems to have a significant synergistic component that comes from allowing prediction sequences that combine operations not only from the same, but also from different models. The composite model using only single-model-sequences does not significantly outperform the best basic models, even though it has access to the bulk of their operations. Though, it cannot use their entire adaption potential. The internal parameters of the operations used by the composite model need to be fixed, so the operations themselves cannot be adapted to individual reasoners. In effect, part of the improvement gained by the adaptability on the level of prediction sequences is eaten up by a lack of adaptability on the level of single operations. To judge the predictive quality of the composite model in more detail, it would be interesting to have an upper bound specifically defined for the evaluation set, revealing which performance can in principle be achieved. Such an upper bound is likely not far beyond the MFA benchmark (Riesterer, Brand, and Ragni 2019). Minimally outperforming this benchmark seems to place the composite model

[1]https://github.com/CognitiveComputationLab/CCOBRA

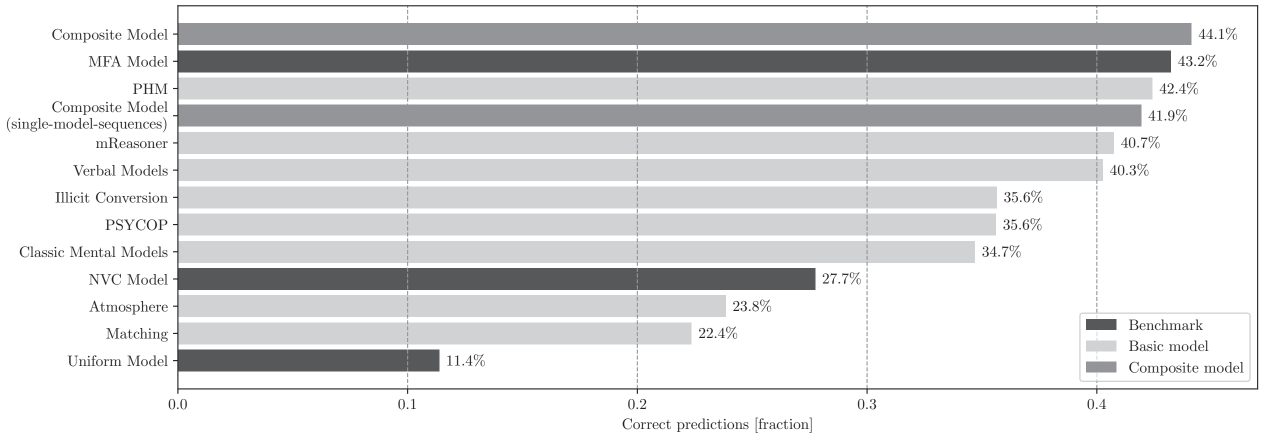

Figure 3: Model evaluation results. *Uniform Model*: Randomly predict one of the nine possible responses. *NVC Model*: Always predict NVC. *MFA Model*: Always predict the most frequent response to the respective syllogism in the training data (breaking ties by randomly choosing among the most frequent responses).

not far away from the theoretical upper limit, while still leaving room for improvement.

Conclusion and Outlook

A set of seven different theories of syllogistic reasoning has been systematically converted into parametric prediction models. The evaluation of these models sheds light on the theories' potentials for individual adaption, which has previously lain largely dormant in the context of model evaluation. It has been shown that PHM, mReasoner and Verbal Models show the best performances, with PSYCOP in midfield and simple heuristic theories like Atmosphere and Matching below. These results make a case for model-based and heuristic models. Also, a novel approach to combining basic reasoning models has been proposed and applied to the syllogistic single-response task. While even the best basic models still perform below the MFA benchmark, the composite model slightly but significantly outperforms both its best constituent models and the MFA benchmark. A synergistic component of this improvement, resulting from combining operations from different models, has been exposed.

To garner further insights, a methodically sound analysis of the performance of the composite model is required. One could examine, which operations are particularly important or negligible to achieve its successful adaption and gain information about how existing models might be improved. For the future, it might also be interesting to apply a composite approach either to other syllogistic tasks like conclusion verification or to other reasoning domains. The presented approach can also be used to create models with a carefully designed set of reasoning sequences. The composite approach itself can further be improved as well to accommodate to use stochastic operations. In contrast to classical reasoning theories which are monolithic, identifying and recombining operations in existing theories is a next level of models that can be adaptive for humans.

References

Begg, I., and Denny, J. P. 1969. Empirical reconciliation of atmosphere and conversion interpretations of syllogistic reasoning errors. *Journal of Experimental Psychology* 81(2):351–354.

Khemlani, S., and Johnson-Laird, P. N. 2012. Theories of the syllogism: A meta-analysis. *Psychological Bulletin* 138(3):427–457.

Khemlani, S., and Johnson-Laird, P. N. 2016. How people differ in syllogistic reasoning. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, 2165–2170. Austin, TX: Cognitive Science Society.

Polk, T. A., and Newell, A. 1995. Deduction as verbal reasoning. *Psychological Review* 102(3):533–566.

Revlis, R. 1975. Two models of syllogistic reasoning: Feature selection and conversion. *Journal of Verbal Learning and Verbal Behavior* 14(2):180–195.

Riesterer, N.; Brand, D.; and Ragni, M. 2019. Predictive modeling of individual human cognition: Upper bounds and a new perspective on performance. In *Proceedings of the 17th International Conference on Cognitive Modeling*, 178–183. Waterloo, Canada: University of Waterloo.

Woodworth, R. S., and Sells, S. B. 1935. An atmosphere effect in formal syllogistic reasoning. *Journal of Experimental Psychology* 18(4):451–460.