

Learning Propositional Logic From Scratch

Abdul Rahim Nizamani¹ and Claes Strannegård^{2,3}

¹Department of Applied Information Technology, University of Gothenburg.

abdulrahim.nizamani@gu.se

²Department of Applied Information Technology, Chalmers University of Technology;

³Department of Philosophy, Linguistics and Theory of Science, University of Gothenburg.

claes.strannegard@gu.se

Abstract

We present a computational model for developing intelligent agents that are able to reason in multiple symbolic domains. The agents have deductive and inductive reasoning abilities. The deductive mechanism is based on a simple cognitive model with bounded cognitive resources. The main learning mechanism is a formalization of Occam's razor. Agents constructed in our model can learn generalized knowledge from concrete examples. For example, an agent can learn elementary arithmetic and propositional logic, and then compute correct answers to previously unseen questions such as "what is $27 * 9$?" and "is $P \vee (P \rightarrow Q)$ a tautology?". We illustrate the learning process in the case of propositional logic, where an agent first learns the syntax and then some basic logical laws. The performance of this agent beats average human scores in terms of accuracy when tested on tautologies used in an experiment reported in [16].

Keywords: multi-domain agents, logical reasoning, propositional logic.

1 Introduction

Artificial Intelligence (AI) is concerned with mathematical modeling of intelligence. Cognitive science, on the other hand, focuses on constructing formal models of human cognition and reasoning. One of the methods used in cognitive science is cognitive modeling that develops computational models of mental processes and

structures. It is a method for understanding human intelligence and can be a valuable tool for AI as well:

"Cognitive modelers attempting to explain human intelligence share a puzzle with artificial intelligence researchers aiming to create computers that exhibit human-level intelligence: how can a system composed of relatively unintelligent parts (such as neurons or transistors) behave intelligently?" [3]

As argued in [3], cognitive modeling and artificial intelligence can help each other to solve their shared puzzle. This has been experimented in some recent AI designs that incorporated simple cognitive models to achieve human level performance in their relevant domains (e.g., [13, 16]).

Cognitive models can be constructed as standalone models, or as components within cognitive architectures that are general frameworks to simulate and study specific mental processes. Many such cognitive architectures exist, such as ACT-R [1], Soar [7], and CHREST [5].

ACT-R is a general cognitive architecture containing multiple modules that represent different mental structures (such as memory, vision and speech). The basic assumption in ACT-R is that memory is divided in declarative and procedural knowledge.

Soar is a general cognitive architecture that is intended for developing intelligent systems. Though it is similar to ACT-R, it does not require the cognitive models to be plausible with respect to human cognition.

Both ACT-R and Soar are based on production systems and use explicit production rules that are roughly of the form “if *condition* then *action*”.

The working memory (WM) is one of the core components in cognitive models of reasoning. The WM model was first developed by Baddeley and Hitch [2]. WM is a mental workspace and a precious cognitive resource with severe restrictions on the number of information chunks it can hold at any point in time. Earlier findings suggest its capacity to be between 5 and 9 chunks [10]. WM capacity is highly correlated with general intelligence (Spearman’s *g* factor) [6].

Other cognitive resources used in cognitive modeling include declarative memory, procedural memory, visual short-term memory, and attention.

Artificial general intelligence (AGI) focuses on the “general” intelligence as opposed to traditional AI that aims at more specific applications of intelligence [18]. Cognitive modeling is a commonly used method in AGI systems. Research in AGI designs has utilized many cognitive architectures, e.g., ACT-R [11], CHREST [8] and Sigma [12]. New cognitive architectures have also been proposed [17].

Occam’s razor is a simple scientific principle that expresses a preference for the simplest solution to a problem. Mathematical formulations of Occam’s razor include Solomonoff complexity, Kolmogorov complexity, and also Levin complexity [9].

The work presented in this paper builds upon our earlier models of deductive reasoning [13, 16, 4] and inductive reasoning [15]. Our objective is to develop a unified design for inductive and deductive reasoning in symbolic domains. We use the same formal design as in [14] (which uses the example of arithmetic), and apply it to the domain of propositional logic.

In this paper, we present a formal model for intelligent agents that can be trained to compete with humans in propositional reasoning. An agent is presented that learns the syntax and standard axioms of propositional logic from scratch. It can then prove tautologies and perform above average human score.

Section 2 presents the computational model, section 3 shows the training and evaluation of an agent, section 4 discusses the results, and section 5 concludes the paper.

2 Computational model

Our computational model is inspired by cognitive modeling and employs a simple cognitive model with bounded cognitive resources. An explicit representation of working memory enables the agent to perform deductive computations, and its long-term memory stores symbolic concepts and axioms. Learning is achieved by a simple formulation of Occam’s razor together with a small set of strategies.

Definition 1 (Tag) A tag is a string of unicode characters with no punctuation marks such as period, comma or colon.

Definition 2 (Variable) A variable is a string of the form $\sigma : \tau$, where τ is a tag and $\sigma \in \{x, y, z\}$.

For example, $x:\text{Formula}$ and $y:\text{Digit}$ are variables.

Definition 3 (Term) A term is a finite binary tree whose nodes contain tags or variables, with the restriction that variables can only appear in the leaf nodes.

Hereafter, we will write terms in conventional linearized form using infix notation, such as $2 + 3$ and $P \wedge Q$. For the purpose of implementing this model as a computer program, we use a standard parser and linearizer to read and write terms in text format.

Definition 4 (Axiom) An axiom has the form $(\tau, a \iff b)$ or $(\tau, a \mapsto b)$, where τ is a tag, and a and b are terms.

Example 1 (Axioms) Following are some axioms.

$$(\text{Taut}, P \vee Q \mapsto P) \quad (1)$$

$$(\text{Taut}, \perp \vee x : \text{Formula} \iff x : \text{Formula}) \quad (2)$$

$$(\text{Arith}, 3 * 2 \iff 6) \quad (3)$$

$$(\text{Arith}, x : \text{Number} + 0 \iff x : \text{Number}) \quad (4)$$

Definition 5 (Theory) A theory is a finite set of axioms.

Definition 6 (Chunk) A chunk is an item of the form (τ, t) where τ is a tag and t is a term.

Definition 7 (Assignment) An assignment is a partial function from variables to terms.

An instance of an assignment is the following.

$$\sigma = \{(x : \text{Formula}, \top), (y : \text{Formula}, P)\}$$

By extension, assignments are defined from terms to terms. E.g.,

$$\sigma(x : \text{Formula} \rightarrow y : \text{Formula}) = \top \rightarrow Q$$

For a variable-free term t , $\sigma(t) = t$.

Definition 8 (Shallow rewrite) A shallow rewrite replaces a term t' with a term t'' if there is an axiom $(\tau, t_1 \mapsto t_2)$ in the theory and there is an assignment σ such that $\sigma(t_1) = t'$ and $\sigma(t_2) = t''$.

$$\frac{t'}{t''} (\tau, t_1 \mapsto t_2)$$

Definition 9 (Deep rewrite) A deep rewrite replaces a single occurrence of a subterm t' in a term t (which can be t itself) with another term t'' , if there is an axiom $(\tau, t_1 \Longrightarrow t_2)$ in the theory and there exists an assignment σ such that $\sigma(t_1) = t'$ and $\sigma(t_2) = t''$.

$$\frac{t(t')}{t(t'')} (\tau, t_1 \Longrightarrow t_2)$$

By extension, we shall call an axiom *deep* if it has the form $(\tau, t' \Longrightarrow t'')$ and *shallow* if it is like $(\tau, t' \mapsto t'')$.

Example 2 Below is an example of a shallow rewrite.

$$\frac{P \vee Q}{P} (\text{Taut}, P \vee Q \mapsto P)$$

Example 3 An instance of a deep rewrite is the following.

$$\frac{(\perp \vee P) \rightarrow Q}{P \rightarrow Q} (2)$$

Note that a shallow rewrite operates only on the term itself, while a deep rewrite can also apply on a subterm within the term. Thus, a deep rewrite encompasses the shallow rewrite, but not vice versa.

Semantically, deep axioms are intended to preserve equivalence and shallow axioms to preserve sufficiency. Thus, a shallow axiom such as $(\text{Taut}, P \vee Q \mapsto P)$ means that, in order to prove the logical truth of $P \vee Q$, it is sufficient to prove that P is true. Although deep axioms preserve equivalence, the deep rewrite is unidirectional. For example, the axiom $(\text{Taut}, \neg\neg P \Longrightarrow P)$ rewrites $\neg\neg P$ to P , but not vice versa.

Definition 10 (Computation) A computation is a sequence of distinct terms (t_0, \dots, t_n) such that t_{i+1} is computed from t_i by applying deep or shallow rewrite.

Note that computations are acyclic. Hereafter, computations will be written vertically from top to bottom, with every transition marked with the relevant axiom used by the rewrite rules.

Example 4 Below is an example of a computation in the domain of propositional logic.

$$\frac{\neg\top \rightarrow Q}{\perp \rightarrow Q} (\text{Taut}, \neg\top \Longrightarrow \perp)$$

$$\frac{\perp \rightarrow Q}{\top} (\text{Taut}, \perp \rightarrow x : \text{Formula} \Longrightarrow \top)$$

Example 5 Following is a computation in the domain of Arithmetic.

$$\frac{(3+0)*2}{3*2} (4)$$

$$\frac{3*2}{6} (3)$$

Definition 11 (Term size) The size of a term t , denoted $s(t)$, is the number of nodes in t .

The definition of *size* is extended to axioms as, for an axiom $a = (\tau, t_1 \Longrightarrow t_2)$, its size is defined as $s(a) = s(t_1) + s(t_2)$. The size of a theory T is $s(T) = \sum\{s(a) : a \in T\}$.

Definition 12 (Agent) An agent is a tuple (T, C, W, L, D) , where

T is a theory, representing beliefs in the declarative memory;

C is a set of chunks, that represents concepts in the declarative memory;

W is a natural number that represents working memory capacity;

L is a natural number that limits the maximum length of computations (representing attention span);

D is a natural number that bounds the maximum size of theories.

Definition 13 (Bounded computation) Given an agent $A = (T, C, W, L, D)$, a bounded computation is a sequence of distinct terms (t_0, \dots, t_n) such that:

1. $n \leq L$;
2. for each term t_i , $s(t_i) \leq W$;
3. each transition from t_i to t_{i+1} uses an axiom from T ;

4. for $1 \leq i < j \leq n$, t_i does not appear as a subterm in t_j .

Condition 4 above is a weak restriction on increasing sizes of terms in a computation.

Definition 14 (Induction Item) An Induction Item is a tuple (τ, t_1, t_2, u) , where τ is a tag, t_1 and t_2 are terms, and u is the utility (an integer). In addition, t_1 and t_2 do not contain any variables.

Definition 15 (Induction Problem) An Induction Problem (IP) is a finite set of induction items.

Example 6 Below is an example of an IP.

$$(\text{Taut}, P \vee T, T, 1) \quad (5)$$

$$(\text{Taut}, Q \vee T, T, 1) \quad (6)$$

$$(\text{Taut}, P \vee Q, Q, -1) \quad (7)$$

Definition 16 (Item computation) Agent A computes an induction item (τ, t_1, t_2, u) if there is a bounded computation from t_1 to t_2 satisfying the following condition.

For any term t_i in the computation and any axiom $(\tau, w_1 \Longrightarrow w_2)$ or $(\tau, w_1 \longmapsto w_2)$, if w_1 contains a variable $(x: \text{Tag})$ that matches with a subterm t' in t_i , then t' must be of type Tag , i.e., there must be a bounded computation (for type-checking) from t' to Tag using only axioms with tag Lang .

Example 7 Example 4 contains a bounded computation of induction item $(\text{Taut}, \neg T \rightarrow Q, T, 1)$, in addition to the following type-checking computation.

$$\frac{\frac{Q}{\text{Boolean}} (\text{Lang}, Q \Longrightarrow \text{Boolean})}{\text{Formula}} (\text{Lang}, \text{Boolean} \Longrightarrow \text{Formula})$$

Item computations are goal-driven bounded computations. We use the standard A^* search algorithm to find bounded computations of minimum length.

Definition 17 (Performance) The performance p of agent A on induction problem IP is computed as $p = \sum \{u : (\tau, t_1, t_2, u) \in IP \text{ and } A \text{ computes } (\tau, t_1, t_2, u)\}$. We use the convention $\sum \emptyset = 0$ to ensure that p is always defined.

Induction items with $u < 0$, e.g., (7), are negative examples. Such an induction item, if solved, will reduce the performance of the agent. Thus, to achieve maximum performance, an agent must solve all items with $u > 0$ and must not solve any item with $u < 0$.

As we shall see now, if an agent cannot compute an induction item with $u > 0$, it tries to extend its theory in a way that enables it to compute that item.

Definition 18 (Strategy) A strategy is a method to form axioms for constructing new theories.

Our model includes the following strategies.

Strategy 1 (Abstraction) Axiom $(\tau, t' \Longrightarrow w')$ is formed from the induction item (τ, t, w, u) by abstraction, where t' is obtained by replacing one or more subterms of t with variables, and w' is obtained similarly from w .

Abstraction is a straightforward way of generalizing observations. Below is an example.

Example 8 Here are some of the axioms constructed by abstraction from item $(\text{Taut}, P \wedge P, P, 1)$.

$$(\text{Taut}, x : \text{Formula} \Longrightarrow P)$$

$$(\text{Taut}, x : \text{Formula} \wedge x : \text{Formula} \Longrightarrow P)$$

$$(\text{Taut}, x : \text{Formula} \wedge x : \text{Formula} \Longrightarrow x : \text{Formula})$$

Strategy 2 (Crossover) Axioms $(\tau, t' \Longrightarrow w')$ and $(\tau, t' \longmapsto w')$ are formed from chunks (τ, t) and (τ, w) by crossover, where t' is obtained by replacing a subterm of t with a subterm of w (w' is obtained similarly).

Crossover is a powerful but computationally expensive method of constructing new theories. All the chunks in the agent's memory and also a small number of variables are used to generate axioms by crossover. Simple heuristics are applied to reduce the number of generated theories.

Example 9 Following are some of the axioms formed by crossover from chunks $(\text{Taut}, P \wedge Q)$ and $(\text{Taut}, T \vee \perp)$.

$$(\text{Taut}, P \Longrightarrow T)$$

$$(\text{Taut}, T \Longrightarrow P \wedge Q)$$

$$(\text{Taut}, P \vee Q \Longrightarrow \perp)$$

$$(\text{Taut}, Q \vee (P \wedge Q) \Longrightarrow \perp)$$

Strategy 3 (Memorization) The axiom $(\tau, t_1 \Longrightarrow t_2)$ is formed by memorization from the induction item (τ, t_1, t_2, u) if $u > 0$.

Memorization is the strategy to remember the examples literally. It will be used if no other strategy from the previous strategies works.

Example 10 Axiom $(\text{Formula}, \neg \perp \implies \top)$ is formed by memorization from item $(\text{Formula}, \neg \perp, \top, 1)$.

The set of strategies can easily be extended if required. We also have a strategy to find recursive definitions of functions; for example, for solving number sequence problems such as 8, 11, 14, ?. It is not described here however, being irrelevant for propositional logic.

Definition 19 (Update function) The update function takes as input agent $A = (T, C, W, L, D)$ and an IP, and outputs agent $A' = (T \cup T', C \cup C', W, L, D)$.

C' is obtained from IP by taking all the subterms of the terms in the items of IP.

$T' = \emptyset$ if all items in IP with positive utility have bounded computations; otherwise T' is obtained by the following procedure.

First, construct a set Θ of new theories whose axioms are generated by the strategies mentioned earlier (strategies 1 to 3). Each theory T^* in Θ is subject to a set of conditions, such as the following:

- number of axioms in $T^* \leq 3$;
- no axiom of T^* appears in T ;
- $s(T^*) \leq D$ if T^* is formed by Crossover.

Theory T' is selected from Θ by applying the following selection criteria in the listed order.

1. Agent A' performs optimally on the IP.
2. $s(T')$ is as small as possible.
3. T' contains the maximum number of deep axioms.
4. T' has the maximum number of variable tokens.
5. T' has the minimum number of variable types.
6. T' is lexicographically minimal.

Condition 2 above is a formal application of Occam's razor and selects an explanation (i.e., theory) of the smallest size that can account for the observations (IP).

Strategies 1 to 3 are used one by one, such that if strategy n leads to a suitable T' , the Update function halts, otherwise it proceeds to strategy $n + 1$.

3 Results

3.1 Implementation

We implemented our model in the functional programming language Haskell. The code consists of approximately 2,500 lines. The program, called OccamStar, takes an agent and an IP as input, and outputs a new

agent with updated memory (that replaces the input agent). It also produces bounded computations of the items in IP.

3.2 Learning process

To illustrate the learning process, we show here a set of examples that allows the agent to acquire the elementary axioms of propositional logic.

Let the agent A be initiated with $T = \emptyset$, $C = \emptyset$, $W = 8$, $L = 10$, and $D = 4$.

Example 11 Let IP consist of the following items.

$$(\text{Lang}, \top, \text{Boolean}, 1) \quad (8)$$

$$(\text{Lang}, \perp, \text{Boolean}, 1) \quad (9)$$

OccamStar takes agent A and the above IP, and updates A by adding the following axioms to the theory.

$$(\text{Lang}, \top \implies \text{Boolean}) \quad (10)$$

$$(\text{Lang}, \perp \implies \text{Boolean}) \quad (11)$$

Item 8 is computed with this theory as follows.

$$\frac{\top}{\text{Boolean}} \quad (10)$$

With a similar IP, agent A learns the following axioms.

$$(\text{Lang}, P \implies \text{Boolean}) \quad (12)$$

$$(\text{Lang}, Q \implies \text{Boolean}) \quad (13)$$

$$(\text{Lang}, R \implies \text{Boolean}) \quad (14)$$

Example 12 Now, continuing the learning process, let IP be given as follows.

$$(\text{Lang}, \top, \text{Formula}, 1) \quad (15)$$

$$(\text{Lang}, P, \text{Formula}, 1) \quad (16)$$

Then, agent A learns the following axiom.

$$(\text{Lang}, \text{Boolean} \implies \text{Formula}) \quad (17)$$

Intuitively, agent A knows that a boolean symbol is a formula. The bounded computation for item (15) is as follows.

$$\frac{\top}{\text{Boolean}} \quad (10)$$

$$\frac{\text{Boolean}}{\text{Formula}} \quad (17)$$

Example 13 Let IP contain the following items.

$$(\text{Lang}, \top \wedge \perp, \text{Formula}, 1) \quad (18)$$

$$(\text{Lang}, P \wedge Q, \text{Formula}, 1) \quad (19)$$

Then, the following axiom is added to the theory.

$$(\text{Lang}, \text{Formula} \wedge \text{Formula} \implies \text{Formula}) \quad (20)$$

When the IP contains the tag Lang , the agent prefers the most general tag. From (17), it knows that the tag Formula is more general than Boolean . Thus, it selects (20) instead of the following axiom that is equally valid.

$$(\text{Lang}, \text{Boolean} \wedge \text{Boolean} \implies \text{Formula}) \quad (21)$$

With similar IP s, agent A adds the following axioms.

$$(\text{Lang}, \neg \text{Formula} \implies \text{Formula}) \quad (22)$$

$$(\text{Lang}, \text{Formula} \vee \text{Formula} \implies \text{Formula}) \quad (23)$$

$$(\text{Lang}, \text{Formula} \rightarrow \text{Formula} \implies \text{Formula}) \quad (24)$$

$$(\text{Lang}, \text{Formula} \leftrightarrow \text{Formula} \implies \text{Formula}) \quad (25)$$

Example 14 Next, let IP consist of the following items.

$$(\text{Taut}, \top \vee P, \top, 1) \quad (26)$$

$$(\text{Taut}, \top \vee (\neg Q), \top, 1) \quad (27)$$

$$(\text{Taut}, P \vee Q, P, -1) \quad (28)$$

Then, agent A acquires the following axiom.

$$(\text{Taut}, \top \vee x : \text{Formula} \implies \top) \quad (29)$$

Agent A is now able to solve new problems such as the following.

$$\frac{\top \vee (\neg R)}{\top} \quad (29)$$

The above computation requires an additional type-checking computation to verify that $\neg R$ is a Formula .

$$\frac{\neg R}{\neg \text{Boolean}} \quad (14)$$

$$\frac{\neg \text{Boolean}}{\neg \text{Formula}} \quad (17)$$

$$\frac{\neg \text{Formula}}{\text{Formula}} \quad (22)$$

The inclusion of the negative item (28) in example 14 invalidates the following erroneous axiom, which would have been learned by the agent otherwise.

$$(\text{Taut}, x : \text{Formula} \vee y : \text{Formula} \implies x : \text{Formula}) \quad (30)$$

Example 15 Let IP have the following items.

$$(\text{Taut}, \perp \vee P, P, 1) \quad (31)$$

$$(\text{Taut}, \perp \vee \neg Q, \neg Q, 1) \quad (32)$$

$$(\text{Taut}, P \vee Q, Q, -1) \quad (33)$$

Then, agent A learns the following.

$$(\text{Taut}, \perp \vee x : \text{Formula} \implies x : \text{Formula}) \quad (34)$$

The agent can also learn shallow axioms when a negative item is included in the IP to invalidate a similar deep axiom. The following example illustrates this.

Example 16 Let IP contain the following items.

$$(\text{Taut}, P \vee Q, P, 1) \quad (35)$$

$$(\text{Taut}, P \wedge (P \vee Q), P \wedge P, -1) \quad (36)$$

Then, agent A learns the following.

$$(\text{Taut}, x : \text{Formula} \vee y : \text{Formula} \implies x : \text{Formula}) \quad (37)$$

3.3 Evaluation

An experiment reported in [16] presented 40 tautologies and 40 non-tautologies, mixed randomly, to participants who were asked to decide which ones were tautologies. The subjects were university students of computer science who had received tuition in propositional logic. The mean accuracy of participants on tautologies was 32.1, with the maximum score of 35 (out of 40). All tautologies except one (trial 10) were solved by at least one of the participants.

The computer program *OccamStar* was used to train an agent LA with the basic axioms of propositional logic, a subset of the axioms used by [16] in their computational model. Parameters were set as follows: $T = \emptyset$, $C = \emptyset$, $W = 8$, $L = 10$ and $D = 4$. The training proceeded in the same fashion as illustrated with examples previously, starting with an empty theory and resulting in the theory that is given in Appendix A.

Agent LA was able to solve 34 tautologies out of the 40. It computed the solutions within a few minutes on a standard personal computer. Sample computations of tautologies produced by LA are presented below.

Computation of trial 3:

$$\frac{(P \vee P) \leftrightarrow \neg \neg P}{P \leftrightarrow \neg \neg P} \quad (A43)$$

$$\frac{P \leftrightarrow \neg \neg P}{P \leftrightarrow P} \quad (A57)$$

$$\frac{P \leftrightarrow P}{\top} \quad (A18)$$

Computation of trial 17:

$$\frac{P \vee (P \vee \neg(Q \wedge P))}{P \vee \neg(Q \wedge P)} \quad (\text{A83})$$

$$\frac{P \vee \neg(Q \wedge P)}{\neg(Q \wedge P) \vee P} \quad (\text{A45})$$

$$\frac{\neg(Q \wedge P) \vee P}{(Q \wedge P) \rightarrow P} \quad (\text{A60})$$

$$\frac{(Q \wedge P) \rightarrow P}{\top} \quad (\text{A22})$$

Computation of trial 33:

$$\frac{(P \leftrightarrow (Q \leftrightarrow (P \vee P))) \rightarrow Q}{(P \leftrightarrow (Q \leftrightarrow P)) \rightarrow Q} \quad (\text{A43})$$

$$\frac{(P \leftrightarrow (Q \leftrightarrow P)) \rightarrow Q}{(P \leftrightarrow (P \leftrightarrow Q)) \rightarrow Q} \quad (\text{A46})$$

$$\frac{(P \leftrightarrow (P \leftrightarrow Q)) \rightarrow Q}{((P \leftrightarrow P) \leftrightarrow Q) \rightarrow Q} \quad (\text{A52})$$

$$\frac{((P \leftrightarrow P) \leftrightarrow Q) \rightarrow Q}{(\top \leftrightarrow Q) \rightarrow Q} \quad (\text{A18})$$

$$\frac{(\top \leftrightarrow Q) \rightarrow Q}{Q \rightarrow Q} \quad (\text{A39})$$

$$\frac{Q \rightarrow Q}{\top} \quad (\text{A17})$$

Computation of trial 37:

$$\frac{(P \wedge Q) \rightarrow (Q \vee R)}{\neg(P \wedge Q) \vee (Q \vee R)} \quad (\text{A61})$$

$$\frac{\neg(P \wedge Q) \vee (Q \vee R)}{(\neg(P \wedge Q) \vee Q) \vee R} \quad (\text{A50})$$

$$\frac{(\neg(P \wedge Q) \vee Q) \vee R}{(\neg(P \wedge Q) \vee Q)} \quad (\text{A82})$$

$$\frac{(\neg(P \wedge Q) \vee Q)}{(P \wedge Q) \rightarrow Q} \quad (\text{A60})$$

$$\frac{(P \wedge Q) \rightarrow Q}{\top} \quad (\text{A22})$$

4 Discussion

The model described here incorporates some elements of cognitive modeling to reduce the combinatorial explosion of the search space, and achieves human level performance in multiple domains. We have demonstrated this in the case of propositional logic, where an agent trained in logic was able to compete with human scores in an experiment on tautologies. Nothing about our agent was geared specifically to logic, however. Indeed, the agent can equally well be trained in arithmetic as was done in [14]. When trained in arithmetic, the agent was able to solve previously unseen deductive problems (e.g., what is $65 * 7$?) and inductive problems (e.g., what number comes next in 8, 11, 14?). A similar model developed earlier was able to solve number sequence problems not seen before from real IQ tests and it scored above average human performance [15].

Cognitive modeling for the purposes of AI does not need to be psychologically verified. AI aims to achieve human level performance (or beyond) on intelligence

problems, therefore the methods are not required to be cognitively grounded. Thus, AI does not need psychologically plausible cognitive models, rather only particular elements from the field of cognitive modeling that can help reduce the computational complexity of the solution. This can be compared with the biologically inspired computing methods such as neural networks, ant colony optimization algorithms or genetic algorithms. Such algorithms do not necessarily imitate any biological systems, rather are designed with inspiration from certain properties of natural systems.

The bounds over cognitive resources, such as working memory capacity (W), can be manipulated to vary the agent performance and learning complexity. For instance, with $W = 10$, agent LA correctly solved 36 tautologies out of 40. Learning additional axioms helps it score even higher. It is still an open question whether the agent can acquire a complete theory of propositional logic, and this requires further work.

5 Conclusion

The results presented here suggest that resource bounded cognitive models can be helpful in AI for constructing intelligent agents who perform at or above human level. For the purposes of AI, such cognitive models need not be psychologically plausible, as long as they help solve the problems. We believe that a more refined cognitive model would strengthen our model further however, for example in acquiring larger theories and reducing the search complexity.

We have presented an agent architecture that can acquire knowledge in multiple symbolic domains. Thus, it represents a restricted form of artificial general intelligence. Potential application areas of such agents include automated theorem proving, inductive program synthesis, intelligent pedagogical systems, and computational linguistic systems.

Acknowledgement

This research was supported by The Swedish Research Council (grant 421-2012-1000).

A Agent theory

Following is the theory learned by agent LA. For readability, $x:\text{Formula}$ is shortened to x , $y:\text{Formula}$ is short-

ened to y , and z : Formula is written z .

Language syntax:

- A1. (Lang, $\top \implies \text{Boolean}$)
- A2. (Lang, $\perp \implies \text{Boolean}$)
- A3. (Lang, $P \implies \text{Boolean}$)
- A4. (Lang, $Q \implies \text{Boolean}$)
- A5. (Lang, $R \implies \text{Boolean}$)
- A6. (Lang, $\text{Boolean} \implies \text{Formula}$)
- A7. (Lang, $\neg \text{Formula} \implies \text{Formula}$)
- A8. (Lang, $\text{Formula} \vee \text{Formula} \implies \text{Formula}$)
- A9. (Lang, $\text{Formula} \wedge \text{Formula} \implies \text{Formula}$)
- A10. (Lang, $\text{Formula} \rightarrow \text{Formula} \implies \text{Formula}$)
- A11. (Lang, $\text{Formula} \leftrightarrow \text{Formula} \implies \text{Formula}$)

Tautologies:

- A12. (Taut, $\neg \perp \implies \top$)
- A13. (Taut, $x \vee \top \implies \top$)
- A14. (Taut, $\top \vee x \implies \top$)
- A15. (Taut, $x \rightarrow \top \implies \top$)
- A16. (Taut, $\perp \rightarrow x \implies \top$)
- A17. (Taut, $x \rightarrow x \implies \top$)
- A18. (Taut, $x \leftrightarrow x \implies \top$)
- A19. (Taut, $x \vee \neg x \implies \top$)
- A20. (Taut, $\neg x \vee x \implies \top$)
- A21. (Taut, $(x \wedge y) \rightarrow x \implies \top$)
- A22. (Taut, $(x \wedge y) \rightarrow y \implies \top$)
- A23. (Taut, $x \rightarrow (x \vee y) \implies \top$)
- A24. (Taut, $y \rightarrow (x \vee y) \implies \top$)

Contradictions:

- A25. (Taut, $\neg \top \implies \perp$)
- A26. (Taut, $x \wedge \perp \implies \perp$)
- A27. (Taut, $\perp \wedge x \implies \perp$)
- A28. (Taut, $x \wedge \neg x \implies \perp$)
- A29. (Taut, $\neg x \wedge x \implies \perp$)
- A30. (Taut, $x \leftrightarrow \neg x \implies \perp$)
- A31. (Taut, $\neg x \leftrightarrow x \implies \perp$)

Identity:

- A32. (Taut, $x \vee \perp \implies x$)
- A33. (Taut, $\perp \vee x \implies x$)
- A34. (Taut, $x \wedge \top \implies x$)
- A35. (Taut, $\top \wedge x \implies x$)
- A36. (Taut, $\top \rightarrow x \implies x$)
- A37. (Taut, $x \rightarrow \perp \implies \neg x$)

- A38. (Taut, $x \leftrightarrow \top \implies x$)
- A39. (Taut, $\top \leftrightarrow x \implies x$)
- A40. (Taut, $x \leftrightarrow \perp \implies \neg x$)
- A41. (Taut, $\perp \leftrightarrow x \implies \neg x$)

Idempotence:

- A42. (Taut, $x \wedge x \implies x$)
- A43. (Taut, $x \vee x \implies x$)

Commutativity:

- A44. (Taut, $x \wedge y \implies y \wedge x$)
- A45. (Taut, $x \vee y \implies y \vee x$)
- A46. (Taut, $x \leftrightarrow y \implies y \leftrightarrow x$)

Associativity:

- A47. (Taut, $(x \wedge y) \wedge z \implies x \wedge (y \wedge z)$)
- A48. (Taut, $x \wedge (y \wedge z) \implies (x \wedge y) \wedge z$)
- A49. (Taut, $(x \vee y) \vee z \implies x \vee (y \vee z)$)
- A50. (Taut, $x \vee (y \vee z) \implies (x \vee y) \vee z$)
- A51. (Taut, $(x \leftrightarrow y) \leftrightarrow z \implies x \leftrightarrow (y \leftrightarrow z)$)
- A52. (Taut, $x \leftrightarrow (y \leftrightarrow z) \implies (x \leftrightarrow y) \leftrightarrow z$)

Distributivity:

- A53. (Taut, $x \wedge (y \vee z) \implies (x \wedge y) \vee (x \wedge z)$)
- A54. (Taut, $x \vee (y \wedge z) \implies (x \vee y) \wedge (x \vee z)$)
- A55. (Taut, $(x \wedge y) \vee (x \wedge z) \implies x \wedge (y \vee z)$)
- A56. (Taut, $(x \vee y) \wedge (x \vee z) \implies x \vee (y \wedge z)$)

Negation:

- A57. (Taut, $\neg \neg x \implies x$)
- A58. (Taut, $x \rightarrow \neg x \implies \neg x$)
- A59. (Taut, $\neg x \rightarrow x \implies x$)
- A60. (Taut, $\neg x \vee y \implies x \rightarrow y$)
- A61. (Taut, $x \rightarrow y \implies \neg x \vee y$)
- A62. (Taut, $\neg(x \rightarrow y) \implies x \wedge \neg y$)
- A63. (Taut, $x \wedge \neg y \implies \neg(x \rightarrow y)$)
- A64. (Taut, $\neg y \rightarrow \neg x \implies x \rightarrow y$)
- A65. (Taut, $x \rightarrow y \implies \neg y \rightarrow \neg x$)

De Morgan:

- A66. (Taut, $\neg x \wedge \neg y \implies \neg(x \vee y)$)
- A67. (Taut, $\neg(x \vee y) \implies \neg x \wedge \neg y$)
- A68. (Taut, $\neg(\neg x \wedge \neg y) \implies x \vee y$)
- A69. (Taut, $x \vee y \implies \neg(\neg x \wedge \neg y)$)
- A70. (Taut, $\neg(x \wedge \neg y) \implies \neg x \vee y$)
- A71. (Taut, $\neg x \vee y \implies \neg(x \wedge \neg y)$)

- A72. (Taut, $\neg(\neg x \wedge y) \iff x \vee \neg y$)
A73. (Taut, $x \vee \neg y \iff \neg(\neg x \wedge y)$)
A74. (Taut, $\neg x \vee \neg y \iff \neg(x \wedge y)$)
A75. (Taut, $\neg(x \wedge y) \iff \neg x \vee \neg y$)
A76. (Taut, $\neg(\neg x \vee \neg y) \iff x \wedge y$)
A77. (Taut, $x \wedge y \iff \neg(\neg x \vee \neg y)$)
A78. (Taut, $\neg(x \vee \neg y) \iff \neg x \wedge y$)
A79. (Taut, $\neg x \wedge y \iff \neg(x \vee \neg y)$)
A80. (Taut, $\neg(\neg x \vee y) \iff x \wedge \neg y$)
A81. (Taut, $x \wedge \neg y \iff \neg(\neg x \vee y)$)

Axioms for shallow rewrite:

- A82. (Taut, $x \vee y \iff x$)
A83. (Taut, $x \vee y \iff y$)
A84. (Taut, $x \rightarrow y \iff y$)

B Tautologies

Below is the list of 40 tautologies that were used for testing agent LA. The same test items were used for testing human performance [16]. Column L shows the minimum length of bounded computations generated by LA. The symbol * indicates that agent LA failed to find a computation. LA successfully identified the tautologies in 34 of 40 cases. Accuracy and latency measures of these tautologies are listed in [16]. Trial numbers are non-consecutive as non-tautologies have been omitted from the list.

Trial	Formula	L
1	$P \rightarrow ((\neg\neg P \rightarrow P) \vee Q)$	4
2	$(\neg P \leftrightarrow P) \rightarrow (Q \leftrightarrow Q)$	2
3	$(P \vee P) \rightarrow \neg\neg P$	3
5	$(\neg P \wedge P) \rightarrow Q$	2
7	$P \rightarrow (Q \rightarrow P)$	2
9	$\neg(\neg P \leftrightarrow P)$	2
10	$(\neg(P \rightarrow Q) \wedge R) \rightarrow \neg Q$	*
11	$\neg\neg\neg(P \wedge \neg P)$	3
12	$(P \rightarrow P) \vee Q$	2
15	$(P \wedge P) \vee \neg P$	2
17	$P \vee (P \vee (\neg(Q \wedge P)))$	4
18	$(P \leftrightarrow \neg P) \rightarrow (Q \leftrightarrow P)$	2
21	$\neg((P \rightarrow \neg Q) \vee P) \rightarrow Q$	*
22	$((P \vee \neg Q) \vee Q) \vee P$	4
28	$(P \wedge P) \rightarrow P$	1
29	$(P \wedge P) \leftrightarrow P$	2

32	$\neg\neg(P \rightarrow Q) \vee \neg Q$	3
33	$(P \leftrightarrow P) \vee (P \wedge \neg Q)$	6
35	$(P \leftrightarrow P) \vee (P \wedge \neg Q)$	2
37	$(P \wedge Q) \rightarrow (Q \vee R)$	5
40	$\neg P \leftrightarrow \neg P$	1
42	$(\neg P \vee Q) \vee P$	3
43	$\neg(P \rightarrow P) \rightarrow Q$	3
44	$P \leftrightarrow \neg\neg P$	2
45	$P \rightarrow ((P \leftrightarrow Q) \vee (Q \rightarrow R))$	*
50	$(P \rightarrow Q) \rightarrow (P \rightarrow P)$	2
52	$\neg(P \wedge P) \vee (Q \rightarrow Q)$	2
56	$P \vee (P \rightarrow \neg Q)$	4
57	$P \vee ((P \rightarrow Q) \vee P)$	5
62	$(P \wedge ((Q \rightarrow P) \leftrightarrow P)) \leftrightarrow P$	*
63	$P \vee (Q \vee ((Q \wedge R) \leftrightarrow P))$	*
64	$P \vee ((Q \wedge Q) \rightarrow Q)$	2
66	$P \vee (P \rightarrow Q)$	4
67	$P \rightarrow (P \vee Q)$	1
69	$\neg(P \rightarrow Q) \vee (Q \rightarrow Q)$	2
70	$\neg\neg((P \rightarrow Q) \wedge Q) \wedge Q$	*
71	$P \vee (P \rightarrow \neg\neg(Q \leftrightarrow Q))$	4
72	$(P \rightarrow \neg Q) \vee P$	4
76	$(P \rightarrow \neg Q) \vee Q$	4
77	$\neg P \vee (Q \rightarrow (P \vee P))$	4

References

- [1] John R Anderson, Daniel Bothell, Michael D Byrne, Scott Douglass, Christian Lebiere, and Yulin Qin. An integrated theory of the mind. *Psychological review*, 111(4):1036, 2004.
- [2] Alan Baddeley. *Working Memory, Thought, and Action*. Oxford University Press, 2007.
- [3] Nicholas L Cassimatis. Artificial Intelligence and Cognitive Modeling Have the Same Problem. In *Theoretical Foundations of Artificial General Intelligence*, pages 11–24. Springer, 2012.
- [4] Fredrik Engström, Abdul Rahim Nizamani, and Claes Strannegård. Understandable Explanations in Description Logic. submitted.
- [5] Fernand Gobet and Peter CR Lane. The CHREST Architecture of Cognition: The Role of Perception in General Intelligence. In *Procs 3rd Conf on Artificial General Intelligence*. Atlantis Press, 2010.
- [6] Patrick C Kyllonen. Is Working Memory Capacity Spearman’s g? *Human abilities: Their nature and measurement*, pages 49–75, 1996.

- [7] John Laird. *The Soar Cognitive Architecture*. MIT Press, 2012.
- [8] Peter Lane and Fernand Gobet. CHREST Models of Implicit Learning and Board Game Interpretation. In *Artificial General Intelligence*, pages 148–157. Springer, 2012.
- [9] Ming Li and Paul MB Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, 2009.
- [10] George A Miller. The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. *Psychological review*, 63(2):81, 1956.
- [11] Alessandro Oltramari and Christian Lebiere. Pursuing Artificial General Intelligence by Leveraging the Knowledge Capabilities of ACT-R. In *Artificial General Intelligence*, pages 199–208. Springer, 2012.
- [12] David V Pynadath, Paul S Rosenbloom, Stacy C Marsella, and Lingshan Li. Modeling Two-Player Games in the Sigma Graphical Cognitive Architecture. In *Artificial General Intelligence*, pages 98–108. Springer, 2013.
- [13] Claes Strannegård, Fredrik Engström, Abdul Rahim Nizamani, and Lance Rips. Reasoning About Truth in First-Order Logic. *Journal of Logic, Language and Information*, 22(1):115–137, 2013.
- [14] Claes Strannegård, Abdul Rahim Nizamani, and Ulf Persson. Learning Symbolic Domains From Scratch. In press.
- [15] Claes Strannegård, Abdul Rahim Nizamani, Anders Sjöberg, and Fredrik Engström. Bounded Kolmogorov Complexity Based on Cognitive Models. In *Proceedings of the 6th conference on Artificial General Intelligence*, pages 130–139. Springer, 2013.
- [16] Claes Strannegård, Simon Ulfsbäcker, David Hedqvist, and Tommy Gärling. Reasoning Processes in Propositional Logic. *Journal of Logic, Language and Information*, 19(3):283–314, 2010.
- [17] Claes Strannegård, Rickard von Haugwitz, Johan Wessberg, and Christian Balkenius. A Cognitive Architecture Based on Dual Process Theory. In *Artificial General Intelligence*, pages 140–149. Springer, 2013.
- [18] Pei Wang and Ben Goertzel. Introduction: Aspects of Artificial General Intelligence. In *Proceedings of the 2007 conference on Advances in Artificial General Intelligence: Concepts, Architectures and Algorithms: Proceedings of the AGI Workshop 2006*, pages 1–16. IOS Press, 2007.