

Proceedings of the 2nd Workshop on  
Logic and Practice of Programming (LPOP)

Held in conjunction with SPLASH 2020

November 15, 2020

## Table of Contents

Preface	3
Program	5
Organization	6
Invited talks	7
If You're Not Writing a Program, Don't Use a Programming Language <i>Leslie Lamport</i>	8
Logic, Probability, Knowledge, and Learning <i>Stuart Russell</i>	9
Tractable Boolean Circuits: Applications and Compilation Algorithms <i>Adnan Darwiche</i>	10
From CLP(R) to MiniZinc: There and Back Again <i>Peter Stuckey</i>	11
Papers	12
Deductive Synthesis of the Unification Algorithm: The Automation of Introspection <i>Richard Waldinger</i>	12
Logic in Program Analysis and Verification <i>Patrick Cousot</i>	22
Training Neural Networks to do Logic, with Logic <i>Paul Tarau</i>	26
Logical English <i>Robert Kowalski</i>	33
Flamingo, a Compiler and Runtime for Reactive ALM <i>Daniel Hines</i>	38

## Preface

Logic is fundamental to computer science. Since the development of logic programming in the 1960s, logic has seen a growing practical role. The purpose of this workshop is to be a bridge between different areas of computer science that use logic as a practical tool. We take advantage of the common language of formal logic to exchange ideas between these different areas. We have encouraged submissions from all areas of computer science that use formal logic in some aspect of the programming process. The goal is to encourage the transference of ideas and techniques among the areas.

LPOP 2020 is held in conjunction with SPLASH 2020, which was to be in Chicago, USA, but was moved online due to the pandemic. SPLASH has a long history of supporting workshops in many different areas of computer science, and LPOP is happy to be associated with it.

LPOP 2020 is a followup to the successful LPOP workshop held as part of the Federated Logic Conference in Oxford, UK in 2018. The earlier workshop focused on the integration of logic programming with imperative programming. LPOP 2020 broadens this goal to focus more generally on the practical use of logic as a crosscutting discipline through many areas of computer science. We seek synergies in these areas to help them progress faster by taking advantage of good ideas developed in other areas in order to improve the general practice of programming with logic.

There are invited talks by four distinguished researchers:

- Leslie Lamport (Microsoft Research) leads a discussion on his video (please view before the discussion at <https://youtu.be/wQiWwQcMKuw>), which describes how to use mathematical and logical notation to describe algorithms.
- Stuart Russell (UC Berkeley) explores learning through the use of probabilistic formalisms that draw on the expressive power of first-order logic.
- Adnan Darwiche (UCLA) describes tractable Boolean circuits and how they are used to attack problems beyond NP, learning from knowledge and data, and reasoning about machine learning systems.
- Peter Stuckey (U of Melbourne) discusses how seminal ideas in Constraint Logic Programming that were jettisoned in the transition to Constraint Problem Solving seem to be resurfacing as CSP attacks new, larger domains.

The program includes five presentations by authors of contributed position papers:

- Richard Waldinger describes the automatic derivation of the term unification algorithm by use of theorem proving techniques.
- Patrick Cousot discusses issues around the use of logic in abstract interpretation for program analysis and verification.

- Paul Tarau explores an approach to using deep learning techniques to learn to recognize theorems.
- Robert Kowalski describes "Logical English", a constrained natural language system for specifying programs in LPS, a logic language that includes imperative components.
- Daniel Hines describes Flamingo, an implementation for the ASP-based logic action language ALM, and how it can be applied to solve the RBAC challenge problem.

There are three discussion panels, made up of the presenters, to explore more deeply the issues raised in their talks on the practical use of logic in programming.

This organization, combining paper presentations, invited talks, and panels, is structured to encourage a deeper understanding of the various approaches and how they might mutually interact. We hope you enjoy the variety of talks and discussions!

We thank LPOP program committee members who provided timely helpful and insightful reviews. We thank Matthew Castellana for his excellent work for LPOP publicity matters. We thank Paul Fodor and Tuncay Tekle for their help hosting the LPOP zoom meeting. We thank SPLASH organizers, Jan Vitek, Hridesh Rajan, and Elmer van Chastelet in particular, for their extensive work on coordination in this difficult time.

November 2020

David Warren  
Peter Van Roy  
Y. Annie Liu

# Program

All times are CST time, Chicago, Nov. 15, 2020.

- 10:00-10:10 Opening and Introduction (Peter Van Roy)
- Session 1 Logic in Program Specification and Analysis** (Chair: Annie Liu)
- 10:10-10:50 Leslie Lamport (Invited Talk Q&A Session)  
*If You're Not Writing a Program, Don't Use a Programming Language*
- 10:50-11:05 Richard Waldinger  
*Deductive Synthesis of the Unification Algorithm: The Automation of Introspection*
- 11:05-11:20 Patrick Cousot  
*Logic in Program Analysis and Verification*
- 11:20-11:50 Panel: Leslie Lamport, Richard Waldinger, Patrick Cousot (Chair: David Warren)  
*Logic in Program Specification: Where has it Failed? How can we Fix it?*
- 11:50-12:00 Break + Ask Me Anything (Invited Guest/Host: Michael Leuschel/Jorge Lobo)
- Session 2 Logic in Artificial Intelligence and Machine Learning** (Chair: Peter Van Roy)
- 12:00-12:40 Stuart Russell (Invited Talk)  
*Logic, Probability, Knowledge, and Learning*
- 12:40-12:55 Paul Tarau  
*Training Neural Networks to Do Logic, with Logic*
- 12:55-13:05 Break + Ask Me Anything (Invited Guest/Host: Gopal Gupta/Joost Vennekens)
- Session 3 Logic and Implementation Tractability** (Chair: Annie Liu)
- 13:05-13:45 Adnan Darwiche (Invited Talk)  
*Tractable Boolean Circuits: Applications and Compilation Algorithms*
- 13:45-14:15 Panel: Stuart Russell, Paul Tarau, Adnan Darwiche (Chair: David Warren)  
*Logic in Artificial Intelligence: Don't Machine Learning and Neural Networks do it All?*
- 14:15-14:25 Break + Ask Me Anything (Invited Guest/Host: Manuel Hermenegildo/Martin Gebser)
- Session 4 Logic and Language Expressiveness** (Chair: Peter Van Roy)
- 14:25-15:05 Invited Talk: Peter Stuckey  
*From CLP(R) to MiniZinc: There and Back Again*
- 15:05-15:20 Robert Kowalski  
*Logical English*
- 15:20-15:30 Daniel Hines  
*Flamingo, a Compiler and Runtime for Reactive ALM Systems*
- 15:30-16:00 Panel: Peter Stuckey, Robert Kowalski, Daniel Hines (Chair: David Warren)  
*Since Logic Languages are so Good, Why Aren't They Pervasive?*
- 16:00 Closing (David Warren)

# Organization

## General Chairs

Y. Annie Liu, Stony Brook University, USA  
David Warren, Stony Brook University, USA

## Program Chair

Peter Van Roy, Université catholique de Louvain, Belgium

## Program Committee

Esra Erdem, Sabanci University, Turkey  
Martin Gebser, University of Klagenfurt, Austria  
Seyed Hossein Haeri, Université catholique de Louvain, Belgium  
Manuel Hermenegildo, IMDEA Software Institute and TU Madrid, Spain  
Y. Annie Liu, Stony Brook University, USA  
Jorge Lobo, ICREA and Universitat Pompeu Fabra, Spain  
Andreas Podelski, University of Freiburg, Germany  
Francesco Ricca, University of Calabria, Italy  
Marjan Sirjani, Malardalen University, Sweden  
Paul Tarau, University of North Texas, USA  
Peter Van Roy, Université catholique de Louvain, Belgium  
Joost Vennekens, Katholieke Universiteit Leuven, Belgium  
David Warren, Stony Brook University, USA

## Publicity Chair

Matthew Castellana, Stony Brook University, USA

## Invited talks

# If You're Not Writing a Program, Don't Use a Programming Language

*Leslie Lamport*

Algorithms are not programs. They can and should be written with math rather than programming languages or pseudo-languages. This applies to many more algorithms than the ones taught in algorithm courses.

This is a Q & A session for the following talk:

<https://youtu.be/wQiWwQcMKuw>

The talk itself is 50 minutes; the video includes Q & A and is longer. Participants are asked to watch the video before the talk.

# Logic, Probability, Knowledge, and Learning

*Stuart Russell*

One purpose of learning is to accumulate knowledge, which then becomes an input to enable further learning. I will examine this idea first in the context of logic and then in the context of probability. The idea becomes particularly powerful with probabilistic formalisms that draw on the expressive power of first-order logic, although there is still a long way to go before the potential of cumulative learning is fulfilled.

# Tractable Boolean Circuits: Applications and Compilation Algorithms

*Adnan Darwiche*

Tractable Boolean circuits have been playing an increasingly important role in AI and beyond, being also the basis for tractable probabilistic circuits. This includes (1) providing a systematic approach for tackling problems beyond NP, (2) allowing one to learn from certain combinations of knowledge and data, and (3) reasoning about the behavior of some machine learning systems. In this talk, I will review the basics and applications of tractable Boolean circuits, while also discussing the compilation of Boolean formula into tractable circuits: a critical process which can benefit from additional efforts by the broad computer science community.

## From CLP(R) to MiniZinc: There and Back Again

*Peter Stuckey*

Constraint logic programming (CLP) was a revolution in declarative programming showing how we could answer very interesting and complex questions by a combination of programmed search and constraint solving. But constraint programming (CP) moved away from its logic programming roots to concentrate on modelling, simply specifying a system of constraints, in the process losing the ability to do complex meta-search. MiniZinc is one of the leading constraint programming modelling languages. It was originally designed to tackle complex CP problems, typically small systems of complex constraints. But its uses have changed, often it is used to solve very large systems of simple constraints. This meant that many of the original assumptions in the design of MiniZinc are invalid. In this talk we will examine a new architecture for MiniZinc, which uses constraint solving for model optimization, and includes incremental solving and backtracking. In some sense the new architecture makes MiniZinc a CLP system, bringing us back to the roots of the field.

## Papers

## Deductive Synthesis of the Unification Algorithm: The Automation of Introspection

Richard Waldinger  
Artificial Intelligence Center  
SRI International  
waldinger@ai.sri.com

Unification is the problem of finding a common instance of two expressions. Introduced for mechanizing deduction in [Herbrand 30] and, later, in [Robinson 65], unification has also found application in such fields as logic programming, deductive databases, and natural language understanding.

Deductive program synthesis is an approach to automating the construction of a computer program by regarding programming as a problem in theorem proving. To construct a program that meets a given logical specification, a system proves the existence of an output entity that satisfies that specification. A program meeting the specification is then extracted from the proof.

The structure of the extracted program reflects the structure of the proof itself. A case analysis yields a test or conditional expression; a proof by induction yields recursion or other repetitive construct. The proof is restricted to be sufficiently constructive to allow program extraction to be carried out. For each specification there are many proofs and many corresponding programs.

The version of mathematical induction we employ is induction over a well-founded relation, i.e., a relation that admits no infinite decreasing sequences. For a given input entity, we prove the existence of an output entity that satisfies the specification, under the inductive assumption that the program we are trying to construct will meet the specification for all inputs that are less than the given input with respect to a well-founded relation. The relation is not known in advance; we establish its existence during the proof process and its definition is extracted from the proof via the same answer-extraction mechanism by which the program is constructed.

Program synthesis differs from program transformation in that our starting point is a specification that states the purpose of the desired program but may give no hint as to how that purpose is to be achieved. In particular, the specification may involve quantifiers or other non-executable constructs. In program transformation, we usually start with an executable program that aims for clarity and simplicity but may completely disregard such concerns as efficiency.

Although both program synthesis and logic programming begin with a statement in the notation of mathematical logic, for logic programming that statement is regarded as a program itself, and must be executable. In general, the form of that program is

restricted, often to Horn clauses. The result is not another program, but rather the output that is returned for a given input.

The unification algorithm was an early target for program synthesis efforts, but none of them have yielded a complete and automatic synthesis; proofs have been manual, partial, or interactive synthesis or verification (e.g., [Manna Waldinger 81], [Paulson 85], [Eriksson 84], [Nardi 89], [Armando et al. 97]). So an automatic synthesis is still a novel research effort. Our proof attempt follows the manual proof in [Manna Waldinger 81], but the theory has been reformulated to make it more machine oriented. The automatically constructed program was simpler than those previously obtained by hand.

## Program Synthesis

Program synthesis is a challenging application of theorem proving in that it involves proofs by mathematical induction as well as full-quantifier predicate-logic reasoning. Well-founded induction would be most naturally carried out in a higher-order-logic setting. Proofs typically involve case analysis, which adds to the complexity of the proof search.

The specification of the unification algorithm states that, for a given two expressions, we must find a most-general unifier; or, in the case in which the expressions are not unifiable, we must return a special *failure* symbol  $\perp$ . The proof is conducted in an axiomatic theory in which expressions and substitutions are defined and their properties are spelled out. Expressions comprise constants, variables, and composite terms.

The well-founded relation that serves as the basis for the inductive proof is constructed from two simpler well-founded relations: the *variable-reduction* relation, which holds if the variables of one expression are a proper subset of the other's, and the *proper-subexpression* relation, which holds if one expression is a proper subexpression of the other. The simpler relations are defined in the axiomatic theory and we expect the theorem prover to find a combination that enables the proof to go through.

To conduct a proof by mathematical induction, it is often necessary to generalize the theorem to obtain the benefit of a stronger induction hypotheses. (This observation is an instance of the [Polya 57] *Inventor's Paradox* and was famously exploited in the [Boyer Moore 79] theorem prover.) For the unification synthesis proof, we have found it necessary to strengthen the specification by requiring that the most-general unifier be *idempotent*; this means that applying it two (or more) times in succession has the same effect as applying it once. If the algorithm always returns an idempotent unifier, the arguments of the unification algorithm's recursive calls will be less than the given arguments with respect to the well-founded relation. We do not know how to discover this strengthening automatically, so we include idempotence in the initial specification.

For our experiments, the proof is performed by the theorem prover SNARK [Stickel et al. 2000], developed by the late Mark Stickel. SNARK is a first-order, resolution-based system implemented in Common Lisp. It has a *program (or answer) extraction* mechanism and a *constructive restriction* that allow us to obtain programs from proofs. We need to extract relations as well as programs, which would normally require a higher-order-logic theorem prover—some quantifiers range over relations. We do this in SNARK’s first-order setting by *reifying* the well-founded relation. In other words, we represent the relation by a first-order object  $w$ , and use a three-place relation  $<_w(e_1, e_2)$  to express that the well-founded relation  $w$  holds between pairs of arguments  $e_1$  and  $e_2$ . We can then quantify over  $w$ , just as if it ranged over other entities. (We shall refer to the relation as  $<_w$ )

## The Theory of Expressions and Substitutions

The proof is carried out in a theory of symbolic expressions (or terms), and substitutions [Baader Snyder 01]. We give a brisk introduction here. Our presentation is informal, but in the theory, all the properties are specified by axioms. We do not prove background properties that are proved in the literature. For simplicity, we do not consider general functional expressions  $f(e_1, e_2, \dots, e_n)$ , but restrict ourselves to a single *cons* function symbol  $\bullet$ , with composite expressions written  $(e_1 \bullet e_2)$ . The extension to arbitrary function symbols  $f$  doesn’t introduce any essential difficulty.

The atomic symbols in the theory comprise constants and variables; we distinguish variables with a prefix  $?$ ; thus  $?x$  is a variable and  $a$  is a constant. Non-atomic expressions are constructed by application of the *cons* function symbol  $\bullet$ ; thus  $((a \bullet ?x) \bullet ?y)$  is a non-atomic expression. The two components of a *cons* expression are called its *left* and *right* side, respectively. That is,  $left((e_1 \bullet e_2)) = e_1$  and  $right((e_1 \bullet e_2)) = e_2$ .

We define relations  $is-const(e)$ ,  $is-var(e)$ , and  $is-atom(e)$  to hold when  $e$  is a constant, a variable, and an atom, respectively.

The *proper-subexpression* relation  $e_1 \in e_2$  holds if the expression  $e_1$  occurs in the expression  $e_2$ , but they are not equal. Thus  $?x \in (a \bullet ?x)$  but not  $(a \bullet ?x) \in (a \bullet ?x)$ .

A *substitution* is a replacement of some (finite number) of the variables of an expression by corresponding terms. A substitution can be regarded as a set of replacements that are applied in parallel. For example, the set of replacements  $\{?x \leftarrow (a \bullet ?y), ?y \leftarrow b\}$  is a substitution. We denote by  $e \leftarrow \theta$  the result of applying the substitution  $\theta$  to the expression  $e$ ; for example,  $(b \bullet ?x) \leftarrow \{?x \leftarrow (a \bullet ?y), ?y \leftarrow b\}$  is  $(b \bullet (a \bullet ?y))$ . Because the replacements occur in parallel, their order is irrelevant.

The *empty* substitution  $\{\}$  makes no replacements; thus  $e \leftarrow \{\} = e$  for any expression  $e$ . We also introduce a special *failure* substitution  $\perp$ , which has the same result when applied to any expression  $e$ ; thus  $e_1 \leftarrow \perp = e_2 \leftarrow \perp$  for any expressions  $e_1$  and  $e_2$ . We do not regard  $\perp$  as a proper substitution. Regarding  $\perp$  as a substitution, however improper, is a novelty that allows us to simplify the specification and the final program.

We define the *composition*  $\theta_1 \blacklozenge \theta_2$  of two substitutions  $\theta_1$  and  $\theta_2$  to be the substitution that has the same effect as applying  $\theta_1$  followed by  $\theta_2$ . In other words,  $e \leftarrow (\theta_1 \blacklozenge \theta_2) = ((e \leftarrow \theta_1) \blacklozenge \theta_2)$  for all expressions  $e$ . For instance,  $\{?x \leftarrow (b \bullet ?x)\} \blacklozenge \{?x \leftarrow c, ?y \leftarrow d\} = \{?x \leftarrow (b \bullet c), ?y \leftarrow d\}$ . It is convenient for us to define  $(\theta \blacklozenge \perp) = (\perp \blacklozenge \theta) = \perp$ .

If  $e_1 \leftarrow \theta = e_2$  we say that  $e_2$  is an *instance* of  $e_1$ . If  $\theta_1 \blacklozenge \theta = \theta_2$  we say that  $\theta_1$  is *more general than*  $\theta_2$ . Note that thus a substitution is more general than itself, the empty substitution  $\{\}$  is more general than any substitution, and the failure substitution  $\perp$  is not more general than any proper substitution.

A substitution  $\theta$  is *idempotent* if applying it twice (or more) has the same effect as applying it once, that is, if  $\theta \blacklozenge \theta = \theta$ . For example,  $\{?x \leftarrow (a \bullet ?x)\}$  is not idempotent;  $\{?y \leftarrow ?z\}$  is. Both the empty substitution  $\{\}$  and the failing substitution  $\perp$  are idempotent. Idempotent substitutions do not have the same variable appearing on both the left side of an arrow and the right side of an arrow. Applying an idempotent substitution removes variables from an expression, though it may add others.

A *permutation substitution* is one that permutes its variables; for example,  $\{?x \leftarrow ?y, ?y \leftarrow ?x\}$  is a permutation substitution. The empty substitution is an idempotent permutation, but no other permutation is idempotent. Permutations are the only invertible substitutions: composing a permutation with its inverse substitution (obtained by reversing the arrows) yields the empty substitution  $\{\}$ .

A *unifier*  $\theta$  of two expressions  $e_1$  and  $e_2$  is a substitution that makes them identical; that is,  $e_1 \leftarrow \theta = e_2 \leftarrow \theta$ . For example,  $\{?x \leftarrow b, ?y \leftarrow a\}$  is a unifier of  $(a \bullet ?x)$  and  $(?y \bullet b)$ . Unifiers are not unique; for example,  $\{?x \leftarrow b, ?y \leftarrow a, ?z \leftarrow c\}$  is also unifier of  $(a \bullet ?x)$  and  $(?y \bullet b)$ .

In our nomenclature, the failing substitution  $\perp$  is a unifier of any two expressions. If two expressions are identical, any substitution is a unifier. If  $\theta$  is a unifier of two expressions, any instance of  $\theta$  is also a unifier. We say that two expressions are *unifiable* if they have a proper unifier, that is, any unifier other than  $\perp$ .

We say that a substitution  $\theta$  is a *most-general unifier* of two expressions  $e_1$  and  $e_2$  if  $\theta$  is a unifier of  $e_1$  and  $e_2$  and  $\theta$  is more general than any other unifier of  $e_1$  and  $e_2$ . For example,  $\{?x \leftarrow b, ?y \leftarrow a\}$  is a most-general unifier of  $(a \bullet ?x)$  and  $(?y \bullet b)$  but

$\{?x \leftarrow b, ?y \leftarrow a, ?z \leftarrow c\}$  is not most-general. Most-general unifiers (other than  $\perp$ ) are not unique, however. If we have one most-general unifier, we can construct another by composing it with a permutation substitution. Thus, if two expressions have a proper most-general unifier, they have an infinite number of them.

Our initial specification for a unification algorithm was that it should yield a most-general unifier for two given expressions  $e_1$  and  $e_2$ , i.e., a substitution that satisfies the condition

$$\begin{aligned} & mgu(\theta, e_1, e_2) \\ & \Leftrightarrow \\ & e_1 \triangleleft \theta = e_2 \triangleleft \theta \text{ and} \\ & (\forall \theta') \\ & \quad [if\ e_1 \triangleleft \theta' = e_2 \triangleleft \theta' \\ & \quad \text{then } (\exists \theta'')\ \theta' = \theta \blacklozenge \theta'']. \end{aligned}$$

In other words,  $\theta$  must be a unifier of the expressions  $e_1$  and  $e_2$  and must be more general than any other unifier of  $e_1$  and  $e_2$ .

As we remarked in the introduction, this specification was not strong enough for the induction proof to go through; we found it necessary to satisfy the simpler but stronger most-general idempotent unifier condition

$$\begin{aligned} & mgiu(\theta, e_1, e_2) \\ & \Leftrightarrow \\ & e_1 \triangleleft \theta = e_2 \triangleleft \theta \text{ and} \\ & (\forall \theta') \\ & \quad [if\ e_1 \triangleleft \theta' = e_2 \triangleleft \theta' \\ & \quad \text{then } \theta' = \theta \blacklozenge \theta']. \end{aligned}$$

This condition implies the earlier condition, as we can see by taking  $\theta''$  in that condition to be  $\theta$  itself. It also implies that  $\theta$  is idempotent, as we can see by taking  $\theta'$  in this condition to be  $\theta$ .

The idempotence condition also gives a unification algorithm more predictable behavior. While a nonidempotent unifier may involve any variable, it turns out that an idempotent unifier can involve only variables that occur in the unified expressions.

Idempotent most-general unifiers are not unique; for instance, both  $\{?x \leftarrow ?y\}$  and  $\{?y \leftarrow ?x\}$  are most-general idempotent unifiers of the two expressions  $?x$  and  $?y$ .

## The synthesis of a unification algorithm

Our specification is written as

$$\text{unify}(e_1, e_2) \Leftarrow \text{find } \theta \text{ such that } \text{mgiu}(\theta, e_1, e_2).$$

This requires us to prove the theorem that, for given expressions  $e_1$  and  $e_2$ ,  
 $(\exists \theta) \text{mgiu}(\theta, e_1, e_2)$ .

The proof is conducted under the induction hypothesis that the *unify* program being constructed will behave properly on all inputs  $e'_1$  and  $e'_2$  that are less than the original inputs  $e_1$  and  $e_2$  with respect to some well-founded relation  $<_w$ .

In other words, we conduct the proof with the benefit of the induction hypothesis

$$\begin{aligned} &\text{if } \langle e'_1, e'_2 \rangle <_w \langle e_1, e_2 \rangle \\ &\text{then } \text{mgiu}(\text{unify}(e'_1, e'_2), e'_1, e'_2). \end{aligned}$$

The relation  $<_w$  will not be given in advance—it will be discovered as byproduct of the proof discovery process. If the induction hypothesis is not used in the proof, the program extracted will not involve any recursive calls.

We will include in the theory axioms that define common well-founded relations for the theory of expressions and substitutions. These include the proper subexpression relation  $<_{\epsilon}$ , which holds if  $e'_1$  and  $e'_2$  are proper subexpressions of  $e_1$  and  $e_2$ , respectively, and the variable-reduction relation  $<_{\text{vars}}$ , which holds if the variables of  $e'_1$  and  $e'_2$  are a proper subset of the variables of  $e_1$  and  $e_2$ . These are defined (in the two-argument case) by

$$\begin{aligned} &\langle e'_1, e'_2 \rangle <_{\epsilon} \langle e_1, e_2 \rangle \text{ iff} \\ &e'_1 \in e_1 \text{ and } e'_2 \in e_2 \end{aligned}$$

and

$$\begin{aligned} &\langle e'_1, e'_2 \rangle <_{\text{vars}} \langle e_1, e_2 \rangle \text{ iff} \\ &\text{vars}(\langle e'_1, e'_2 \rangle) \subset \text{vars}(\langle e_1, e_2 \rangle). \end{aligned}$$

We shall also include axioms that allow us to combine well-founded relations to form more complex ones, such as their lexicographic combination  $<_{\text{lex}(w_1, w_2)}$ . But in our current experiments we provide the complex well-founded relation.

## Final Program

We skip ahead to the final program, obtained by SNARK from the above specification. We split between several cases; we hope ultimately to combine them in a single proof.

In the case in which  $e_1$  is given to be an atom, we have

```
unify( $e_1, e_2$ )  $\Leftarrow$ 
  if is-var( $e_1$ )
  then if  $e_1 = e_2$ 
    then {}
    else if  $e_1 \in e_2$ 
      then  $\perp$ 
      else  $e_1 \leftarrow e_2$ 
  else if is-var( $e_2$ )
    then  $e_2 \leftarrow e_1$ 
    else if  $e_1 = e_2$ 
      then {}
      else  $\perp$ .
```

Note that SNARK has invented the occurs-check  $e_1 \in e_2$ . The other base case, in which  $e_2$  is an atom but  $e_1$  is not, is similar.

For the remaining inductive case, in which neither  $e_1$  nor  $e_2$  is an atom, we obtain the program

```
unify( $e_1, e_2$ )  $\Leftarrow$ 
  unify(left( $e_1$ ), left( $e_2$ ))  $\blacklozenge$ 
  unify(right( $e_1$ )  $\blacktriangleleft$  unify(left( $e_1$ ), left( $e_2$ )),
    right( $e_2$ )  $\blacktriangleleft$  unify(left( $e_1$ ), left( $e_2$ ))).
```

In other words, we unify the left sides of the two arguments, obtaining a most-general unifier. We then apply that substitution to the right sides of the arguments and unify the results, obtaining a second most-general unifier. We then compose the two unifiers.

This program is simpler if not more efficient than the program obtained by hand (as published in [Manna Waldinger 81]):

```
unify( $e_1, e_2$ )  $\Leftarrow$ 
  let  $u_l$  be unify(left( $e_1$ ), left( $e_2$ )) in
  if  $u_l = \perp$  then  $\perp$ 
  else let  $u_r$  be unify(right( $e_1$ )  $\blacktriangleleft$   $u_l$ , right( $e_2$ )  $\blacktriangleleft$   $u_l$ ) in
    if  $u_r = \perp$  then  $\perp$ 
    else  $u_l \blacklozenge u_r$ .
```

The manually obtained program tested the two recursive unifications to see if they had failed; if so, it had returned  $\perp$  as the result. But SNARK observed that, since composing  $\perp$  with any substitution yields  $\perp$ , those two tests are redundant; they are already being performed during the composition operation.

For completeness, we include a [link](http://www.ai.sri.com/coffee/lpop-unify-proof-recursive-case.lisp)

<http://www.ai.sri.com/coffee/lpop-unify-proof-recursive-case.lisp>  
to the proof SNARK obtained in the inductive case, which was discovered in about 12 seconds. The proof consists of 45 rows, which include assertions (i.e., axioms, lemmas, and assumptions); some rows combine several inference steps. Note that we have given SNARK strong hints as to the final program, especially in providing the well-founded relation necessary to conduct the proof. We hope to back off from such hints in future work. Also, the base cases and the recursive case should be combined into a single proof.

The proof is expressed in SNARK's Lisp-like syntax, not the infix notation of this paper, and rows in the proof are in the conjunctive normal form employed by resolution theorem provers. Probably no one but SNARK has ever read the proof from top to bottom.

## Concluding Remarks

We argue that advances in theorem proving technology, such as better reasoning frameworks and heuristics, faster machines, and more parallelism, will allow it to serve a useful function in software development. For this to take place, it is necessary that automatic theorem provers support answer extraction with a constructive restriction and the ability to combine resolution-style proofs and mathematical induction.

It may seem that synthesizing a unification algorithm is a pointless exercise, since the theorem prover requires unification itself to carry out the derivation. It is possible, however, that in formulating a machine-oriented theory of expressions and substitutions we may enable a meta-logically equipped theorem prover to develop new special-purpose unification algorithms that enable fast reasoning in a particular theory, much as associative-commutative unification algorithms do. By providing axioms for a new function symbol, we could hope to synthesize a unification algorithm that exhibits good performance on expressions involving that symbol; that algorithm could then be incorporated into the theorem prover itself. We may even hope that a theorem prover may discover new theorem proving techniques. Since the unification algorithm is part of the theorem prover, reasoning about it can be regarded as a form of introspection, which may enable the system to improve itself.

Acknowledgements: The author would like to thank researchers at SRI International and the Kestrel Institute for discussion and comments, and the reviewers for helpful suggestions.

## References

[Armando et al 97]

Armando, A., Smaill, A., & Green, I. (1997). Automatic synthesis of recursive programs: the proof-planning paradigm. *Proceedings 12th IEEE International Conference Automated Software Engineering*, 2-9.

[Baader Snyder 01]

Franz Baader and Wayne Snyder (2001). "Unification Theory". In John Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, volume I, pages 447–533. Elsevier Science Publishers.

[Boyer Moore 79]

Boyer, R. S. and Moore, J S. (1979). *A Computational Logic*, Academic Press, New York.

[Eriksson 84]]

LH Eriksson, Synthesis of a unification algorithm in a logic programming calculus, *The Journal of Logic Programming* 1 (1), 3-18, 1984.

[Herbrand 30]

Jacques Herbrand (1930). *Recherches sur la théorie de la démonstration* (Ph.D. thesis). A. 1252. Université de Paris.

[Manna Waldinger 81]

Z. Manna and R. Waldinger, Deductive Synthesis of the Unification Algorithm, *Science of Computer Programming* 1 (1981), 5–48.

[Nardi 89]

Danielle Nardi, Formal synthesis of a unification algorithm by the deductive-tableau method, *The Journal of Logic Programming*, Volume 7, Issue 1, 1989, Pages 1-43,

[Paulson 85]

L. C. Paulson. Verifying the unification algorithm in LCF. *Sci. Computer Programming* 5 (1985), 143–170.

[Pólya 57]

Pólya, G. (1957). *How to solve it: A new aspect of mathematical method*. Princeton, N.J: Princeton University Press

[Robinson 65]

J. A. Robinson. 1965. A Machine-Oriented Logic Based on the Resolution Principle. *J. ACM* 12, 1 (Jan. 1965), 23–41.

[Stickel et al. 2000]

M. E. Stickel, R. J. Waldinger, and V. K. Chaudhri. A Guide to SNARK [www.ai.sri.com/snark/tutorial/tutorial.html](http://www.ai.sri.com/snark/tutorial/tutorial.html)

# Logic in program analysis and verification\*

Patrick Cousot

For program specification and verification, logic is a natural choice. However, for static analysis by abstract interpretation [10, 11], logic is rarely used, even as a user interface. We discuss the weaknesses of logic from this perspective.

## 1 Which logic for specification?

Choosing a decidable logic, such as Presburger arithmetic [20], to express program properties has the great advantage that user specified invariants can be mechanically checked. However, one can write a program that computes the multiplication  $*$  using iteration and addition  $+$  but the invariant is not expressible in Presburger arithmetic. Of course one can consider a richer specification logic by adding  $*$  to the logic, but then iteration of  $*$  yields exponentiation  $**$  which is not in the logic, *etc.* The common choice of first-order logic is limited by the lack of a recursion mechanism necessary to reason on programs [23]. Transitive closure is not expressible [19]. The next generalization is to use Datalog [22]. But by choosing more and more expressive logics for specification language, one loses the decidability and completeness of provers, has to resort to user-interaction, which is generally excluded in automatic static program analysis tools.

Finally, to discuss the soundness of static analyzers, second-order logic is required. For example [1], a Hoare triple,  $\{P\}C\{Q\}$  must satisfy  $\forall P, Q \in \wp(S) . \forall x \in P . \forall y \in S . x \xrightarrow{C} y \Rightarrow y \in Q$  where  $S$  is a set of states and  $\xrightarrow{C}$  the natural big-step semantics of command  $C$ . So called hyperproperties [3] would require  $P \in \wp(\wp(S))$ .

## 2 Which logic for property representation in a static analyzer?

Once a logic has been chosen for specification it, or a more expressive one, can be used as an internal representation of abstract properties.

An anonymous referee mentioned, with good reason, that “it is often convenient to use logic for describing, for instance how the domain of positive boolean functions can be used for analysis of the variable dependencies that arise within logic programs. By using global analysis via abstract interpretation a Prolog system like Ciao can reason with much richer information than, for example, traditional types as well as procedure-level properties such as determinacy, termination and non-failure. It is beneficial in this case that abstract interpretation algorithms are handled within the underlying logic programming preprocessor or and compiler” [18].

---

\*Position paper for LPOP 2020.

The great advantage of logical representations of program properties is obviously uniformity. Any program property can be represented by (the abstract syntax) of a formula in the logic. For example, the simple connectives ( $\vee$ ,  $\wedge$ ,  $\neg$ , *etc.*) are very simple tree manipulations.

However, uniformity has the disadvantage that many algorithms manipulating these properties are based on specific representations. For example, linear equalities or inequalities will use matrices to represent a set of constraints and a frame of higher-dimensional polyhedra [16]. The efficiency of the algorithms is highly dependant on an adequate representation of the properties. For example the disjunction  $\vee$  becomes an elaborated convex-hull algorithm. So, algorithmically, syntax-based representation uniformity is not tenable.

The concept of reduced product in abstract interpretation [11] was developed for expressing a conjunction of properties with non-uniform representations.

This problem was solved in the same way by SMT solvers, initially with Nelson-Oppen reduction [25, 13] and later with more general reductions when incorporating incomplete theories. Because the community has developed common interfaces for SMT solver competitions, such as SMTLIB2 [2], why not adopt them for specification? The obvious reason is that these formats were designed for machine interfaces, not memory representation, and are unreadable by the casual programmer.

### 3 Abstract domains

The algebraic concept of abstract domain in abstract interpretation and their combinations [11] were designed to handle non-uniform representations of program properties. Translating the concept in logic is difficult because the abstract properties have to be expressed in a concrete logic. If it is easy to express that  $x$  is in a given interval by  $a \leq x \wedge x \leq b$ , it is harder to express the concept of “to be a number between  $a$  and  $b$ ”. Abstract interpretation uses sets to express properties so this is  $x \in P$  or  $x \notin P$  with  $P \triangleq \{x \mid a \leq x \wedge x \leq b\}$  easily encoded as a pair of numbers  $[a, b]$ . Of course set theory is expressed in logic so the task is not impossible but the encoding would introduce a useless layer of difficulty. Operations in abstract domains are predictable algorithms. This is hardly the case for theorem proving in logic. Explaining why an SMT solver failed may be a titanic task [17].

An important idea about abstract domains is that of a hierarchy of abstractions, using Galois connections or concretization functions *e.g.* [5]. If a concretization can easily be used to relate logics at different levels of abstraction, this is much harder for Galois connections since in general logic formulas have no normal form ( $a \leq x \wedge x \leq b$  is also  $x \leq b \wedge a \leq x$  or  $a < x + 1 \wedge x - 1 < b$ , *etc.*). Reasoning on equivalence classes of formulæ with same interpretation to get a best abstraction is not impossible in logic but not entirely natural and algorithmically costly. Without Galois connections, calculational design [7] is impossible, completeness is much harder to prove, *etc.*

### 4 Induction

The central problem in program analysis is to infer inductive arguments in proofs. Of course one can ask the user to provide inductive arguments. This makes verification simpler than program analysis [14]. This is a common approach in the small but is infeasible in the large for program of several million lines [12].

An approximate induction or co-induction is formalized in abstract interpretation by widening, narrowing and their duals [6]. For example widening has a simple geometric interpretation:

extrapolated in the direction of growth.

This is very hard to translate in terms of logic since the changes in the syntax of the formula are not necessarily related to the changes in the semantics of the formula. The complexity of an object and its logical denotation may be completely unrelated. Short formulæ may be very difficult to prove because they encode very complex objects while huge formulæ may be very easy to prove. An example is Craig interpolation [21], which is a dual narrowing [6], but one for which the result is not unique, without best choice, which is problematic and multiplies the necessary attempts for finding the right induction. So for induction, even simple ones such as widening, narrowing, and their duals, the evolution of the formula during fixpoint iteration provides no clue for induction. Moreover, this evolution of the iterates must be monitored for induction [15] and it is hard to use a logic to reason on itself. The same problem arises in the even simpler context of typing where types are represented syntactically by terms. For example the hidden widening in Milner’s type inference system is syntactic identity of monotypes [24, 4] whereas liquid types [26] can handle only finitary abstract domains to avoid widening, which is known to be a severe expressivity limitation [9].

## 5 Conclusion

Logic reduces the representations of properties and formal reasonings to purely syntactic manipulations. Although logic is widely accepted for program verification and analysis, this is often limitative in static analysis by abstract interpretation. we prefer set theoretic and algebraic approaches to abstract interpretation for expressivity, efficiency, and scalability reasons. Of course this does not prevent us to widely use logic as a metalanguage, to study and prove the soundness or completeness of verification and static analysis methods [8].

## References

- [1] Ralph-Johan Back. Refinement calculus, lattices and higher order logic. In *NATO ASI PDC*, volume 118 of *NATO ASI Series*, pages 53–71. Springer, 1992.
- [2] Clark W. Barrett, Leonardo Mendonça de Moura, Silvio Ranise, Aaron Stump, and Cesare Tinelli. The SMT-LIB initiative and the rise of SMT - (HVC 2010 award talk). In *Haifa Verification Conference*, volume 6504 of *Lecture Notes in Computer Science*, page 3. Springer, 2010.
- [3] Michael R. Clarkson and Fred B. Schneider. Hyperproperties. *J. Comput. Secur.*, 18(6):1157–1210, 2010.
- [4] Patrick Cousot. Types as abstract interpretations. In *POPL*, pages 316–331. ACM Press.
- [5] Patrick Cousot. Constructive design of a hierarchy of semantics of a transition system by abstract interpretation. *Theor. Comput. Sci.*, 277(1-2):47–103, 2002.
- [6] Patrick Cousot. Abstracting induction by extrapolation and interpolation. In *VMCAI*, volume 8931 of *Lecture Notes in Computer Science*, pages 19–42. Springer, 2015.
- [7] Patrick Cousot. Calculational design of a regular model checker by abstract interpretation. In *ICTAC*, volume 11884 of *Lecture Notes in Computer Science*, pages 3–21. Springer, 2019.
- [8] Patrick Cousot. *Principles of Abstract Interpretation*. MIT Press, 2021 (to appear).

- [9] Patrick Cousot and Radhia Cousot. Comparing the Galois connection and widening/narrowing approaches to abstract interpretation. In *PLILP*, volume 631 of *Lecture Notes in Computer Science*, pages 269–295. Springer.
- [10] Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, pages 238–252. ACM, 1977.
- [11] Patrick Cousot and Radhia Cousot. Systematic design of program analysis frameworks. In *POPL*, pages 269–282. ACM Press, 1979.
- [12] Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, Antoine Miné, and Xavier Rival. Why does Astrée scale up? *Formal Methods in System Design*, 35(3):229–264, 2009.
- [13] Patrick Cousot, Radhia Cousot, and Laurent Mauborgne. Theories, solvers and static analysis by abstract interpretation. *J. ACM*, 59(6):31:1–31:56, 2012.
- [14] Patrick Cousot, Roberto Giacobazzi, and Francesco Ranzato. Program analysis is harder than verification: A computability perspective. In *CAV (2)*, volume 10982 of *Lecture Notes in Computer Science*, pages 75–95. Springer, 2018.
- [15] Patrick Cousot, Roberto Giacobazzi, and Francesco Ranzato. A<sup>2</sup>i: abstract<sup>2</sup> interpretation. *Proc. ACM Program. Lang.*, 3(POPL):42:1–42:31, 2019.
- [16] Patrick Cousot and Nicolas Halbwachs. Automatic discovery of linear restraints among variables of a program. In *POPL*, pages 84–96. ACM Press, 1978.
- [17] Leonardo Mendonça de Moura and Grant Olney Passmore. The strategy challenge in SMT solving. In *Automated Reasoning and Mathematics*, volume 7788 of *Lecture Notes in Computer Science*, pages 15–44. Springer, 2013.
- [18] Isabel Garcia-Contreras, José F. Morales, and Manuel V. Hermenegildo. Incremental analysis of logic programs with assertions and open predicates. In *LOPSTR*, volume 12042 of *Lecture Notes in Computer Science*, pages 36–56. Springer, 2019.
- [19] Erich Grädel. On transitive closure logic. In *CSL*, volume 626 of *Lecture Notes in Computer Science*, pages 149–163. Springer, 1991.
- [20] Christoph Haase. A survival guide to Presburger arithmetic. *ACM SIGLOG News*, 5(3):67–82, 2018.
- [21] Ranjit Jhala and Kenneth L. McMillan. Interpolant-based transition relation approximation. In *CAV*, volume 3576 of *Lecture Notes in Computer Science*, pages 39–51. Springer, 2005.
- [22] Uwe Keller. Some remarks on the definability of transitive closure in first-order logic and datalog. June 2004.
- [23] Tal Lev-Ami, Neil Immerman, Thomas W. Reps, Mooly Sagiv, Siddharth Srivastava, and Greta Yorsh. Simulating reachability using first-order logic with applications to verification of linked data structures. *Log. Methods Comput. Sci.*, 5(2), 2009.
- [24] Robin Milner. A theory of type polymorphism in programming. *J. Comput. Syst. Sci.*, 17(3):348–375.
- [25] Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.*, 1(2):245–257, 1979.
- [26] Patrick Maxim Rondon, Ming Kawaguchi, and Ranjit Jhala. Liquid types. In *PLDI*, pages 159–169. ACM, 2008.

# Training Neural Networks to Do Logic, with Logic

Paul Tarau

Department of Computer Science and Engineering  
University of North Texas  
*paul.tarau@unt.edu*

**Abstract.** We overview combinatorial generation algorithms, with focus on lambda terms and related type inference algorithms, all elegantly expressible in a logic programming language that supports backtracking and unification.

With help from these tools, we introduce methods to train neural networks as theorem provers. Our combinatorial generation algorithms provide pairs of lambda terms together with their inferred types. We make use of the Curry-Howard isomorphism between lambda terms and formulas in linear and intuitionistic logic corresponding to their types to train our neural networks on large, combinatorially generated datasets mapping formulas to their proof terms.

**Keywords:** logic programming tools for theorem proving, intuitionistic and linear logic, Curry-Howard isomorphism, neural theorem proving, neuro-symbolic computing

Renewed interest in neuro-symbolic computing [1, 2] and specifically experiments with training neural networks on theorem proving tasks [3] open the doors for applications of logic programming-based techniques and tools both as *explainable AI* interfaces and as *symbiotic logic and neural systems*, combining deductive strengths with pattern retrieval via machine learning.

SLD-resolution based logic programming languages (and in particular Prolog) provide a unified framework for implementing combinatorial generation algorithms, type-inference as well as search-intensive theorem provers.

Generating formulas of a given size in Prolog is quite easy as exemplified by the following code snippet, generating all implicational formulas of size  $N$ , counted by the well-known Catalan numbers.

```
gen_tree(N,Tree,Leaves):-gen_tree(Tree,N,0,Leaves, []).

gen_tree(V,N,N,[V|Vs],Vs).
gen_tree((A -> B),SN1,N3,Vs1,Vs3):-pred(SN1,N1),
    gen_tree(A,N1,N2,Vs1,Vs2),
    gen_tree(B,N2,N3,Vs2,Vs3).

pred(SN,N):-succ(N,SN).
```

The counts of generated trees match entry **A000108** in [4], representing the Catalan numbers [5], binary trees with  $N$  internal nodes.

All possible canonical labelings of variables can be derived from a set-partitioning algorithm, counted by the Bell numbers for the  $N + 1$  leaves of the trees with  $N$  internal nodes.

Similar generation algorithms can be built for several families of lambda terms (see [6]).

A theorem prover for intuitionistic propositional logic is derived directly from the sequent calculus formulas of Gentzen’s **LJ** calculus, modified by Roy Dyckhoff [7] with a rewriting rule for nested implications that avoids the need for loop checking.

In its simplest form, the Curry-Howard isomorphism [8] connects the implicational fragment of propositional intuitionistic logic with types in the *simply typed lambda calculus*. A low polynomial type inference algorithm associates a type (when it exists) to a lambda term. Harder (PSPACE-complete, see [9]) algorithms associate inhabitants to a given type expression with the resulting lambda term (typically in normal form) serving as a witness for the existence of a proof for the corresponding tautology.

*To train neural networks as theorem provers via the Curry-Howard isomorphism we need to efficiently generate all theorems up to a given size in the implicational fragment of propositional intuitionistic or linear logic, with help from their lambda term counterparts.*

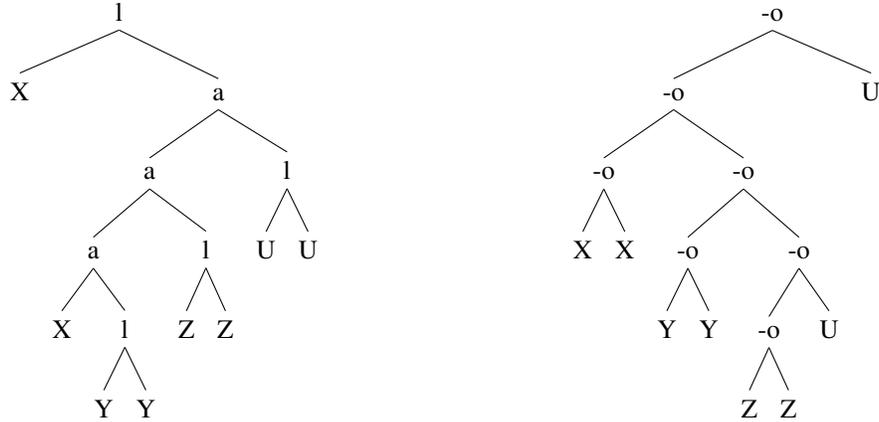
We will apply this mechanism to formulas inferred as types of simply typed lambda terms. After designing a compact encoding of the formula and proof term pairs we will discuss our experiments with training a **seq2seq** LSTM recurrent network [10] which will perform, in inference mode, with 92% accuracy as a theorem prover for the implicational fragment of **IPC** on unseen formulas from a test dataset.

*Linear Logic* [11], as a resource-control mechanism, constrains the use of formulas available as premises in a proof. We will reflect this constraint when generating lambda terms seen as proof terms of our formulas, on the other side of the Curry-Howard isomorphism. Linear Logic’s proof mechanism is known to be Turing-complete even in the propositional case. Thus, we will focus on generating all theorems of a given size for a restricted set of formulas, the Implicational fragment of Propositional Intuitionistic Linear Logic (**IPILL**), corresponding to principal types of lambda terms in normal form.

We start by filtering for linearity the proof terms associated by a Prolog-based theorem prover for Implicational Intuitionistic Logic. This works, but using for each formula a PSPACE-complete algorithm limits it to very small formulas. We take a few back and forth walks over the bridge between proof terms and theorems, provided by the Curry-Howard isomorphism, and derive step-by-step an efficient algorithm requiring a low polynomial effort per generated theorem. The resulting Prolog program runs in  $O(N)$  space for terms of size  $N$  and generates in a few hours **7,566,084,686** theorems in **IPILL**, together with their proof terms.

The figure below shows a lambda term normal form and its corresponding linear type (where we label lambda nodes with **l**, application nodes with **a**, linear implication nodes with **-o** and lambda and type variables with uppercase letters). Note the symmetries present in this “Goldilocks” special case between formulas and their proof terms, largely responsible for our interest in them.

$\lambda X.(((X \lambda Y.Y) \lambda Z.Z) \lambda U.U)$



As terms and clauses in logic programming languages have a tree (or equivalently, a directed acyclic graph) representation, easy to linearize to a canonical representation for both the theorems and their proofs.

Training the Neural Networks as Theorem Provers via the Curry-Howard Isomorphism proceeds as follows. Formulas/types and proofs/lambda terms are both trees, thus we can represent them as prefix strings. Note that for **IPILL** we can even find a *size definition* to give the same size on both sides:

- for lambda terms: leaves=0, lambda nodes=1, applications=1
- for -o formulas: leaves=0, lollipops = 1

We generate prefix encodings of formulas with lollipop=0, application=0, lambda=1, variables as uppercase letters, “:” as separator between formulas and proof terms.

**Example 1** *Provable formulas with their proof terms (for IPILL)*

OAA : 1AA  
 OA00ABB : 1A1BOBA  
 00ABOAB : 1A1BOAB  
 OA00AB00BCC : 1A1B1COC0BA  
 00000AAB00C0BD0CD00EEFF : 1A00A1B1C1D00CD0B1EE1FF

**Example 2** *Provable formulas with their proof terms and “?” if proof failed*

OA0B0000A0C0BODE0CODEFF : 1A1B1COC1D1E1F0000DAEBF  
 OA0B0000A0C0BODE0C0DFGH : ?  
 OA0B0000A0B0C0DE0DOCEFF : 1A1B1COC1D1E1F0000DABFE  
 OA0B0000A0B0C0DE0DOCFGG : ?

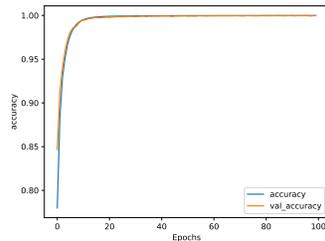
Similar formulas cover implicative intuitionistic propositional logic **IIPC**, also with normal forms in prefix notation.

*Recurrent Neural Networks* (RNN) keep track of dependencies within sequences. They work by using feedback from values at time  $t$  that is fed into computations at time

$t + 1$ . Among them, a *long short-term memory* (LSTM) network is a recurrent neural network architecture that can not only process single data points (such as images), but also entire sequences of data (such as text, speech or video). LSTM neural networks have feedback connections, that help them avoid vanishing or exploding gradient problems by also feeding *unchanged* values to the next layer. Our neural networks will learn the mapping between theorems and their proof-terms on arbitrarily large datasets generated by inferring the type of simply typed lambda terms in normal form, corresponding to formulas via the Curry-Howard correspondence. In particular, theorems in the linear fragment of IIPC will be obtained as types of linear lambda terms.

With trees as prefix string we will use “seq2seq” recurrent neural networks, in particular LSTM (long short term memory) networks that are good to handle long distance dependencies in the prefix terms, in a similar way to their original use for Natural Language processing applications. We plan to also experiment in the future with newer variants, possibly more interesting: **tree2tree**, **dag2dag** and several types of **graph neural networks** (e.g., convolutional, attention, spectral, torch geometric).

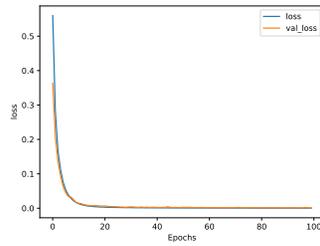
We have evaluated the performance of our neural networks working as theorem provers and our seq2seq LSTM recurrent neural network trained on encodings of theorems and their proof-terms performs *unusually well*. Besides the usual accuracy/loss curves, we have also tested them on unseen formulas, with **100/100** success on **IPLL** and **92/100** success on **IIPC** formulas.



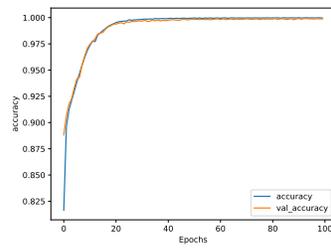
**Fig. 1.** Accuracy of the LSTM seq2seq neural network on our formula/proof term dataset for **IPLL**

We explain these unusually good results for **IPLL** in terms of a *size-preserving bijection between linear lambda terms in normal form and their principal types*. The bijection, first proven in [12] and given a geometric interpretation in [13] is based on a reversible transformation of oriented edges in the tree describing a linear lambda term in normal form, into corresponding oriented edges in the tree describing the linear implicational formula, acting as its principal type. Thus, by generating all typable linear lambda terms in normal form size  $N$ , we obtain, for free, a generator for all corresponding **IPLL** theorems of size  $N$ .

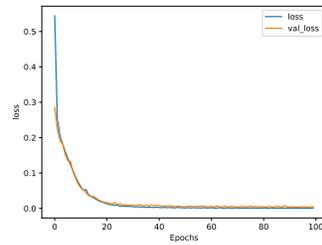
The dataset, containing generated theorems of several sizes and their proof-terms in postfix form, is available at <http://www.cse.unt.edu/~tarau/datasets/lltaut/> and can be used for correctness, performance and scalability testing for linear logic theorem provers. Our experiments with training Recurrent Neural Networks using our



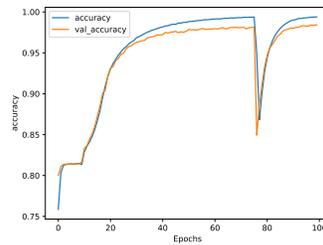
**Fig. 2.** Loss curve of the LSTM seq2seq neural network on our formula/proof term dataset for IPILL



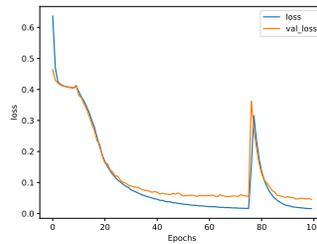
**Fig. 3.** Accuracy for IPILL + unprovable formulas



**Fig. 4.** Loss for IPILL + unprovable formulas



**Fig. 5.** Accuracy for IPC



**Fig. 6.** Loss for IIPC

implicational linear logic theorem dataset are available at: <https://github.com/ptarau/neuralgs>.

## Conclusion

There are basically two ways neural networks have been used to help with theorem proving. Theorem provers are computation-intensive search algorithms often Turing-complete (e.g., propositional linear logic, first order predicate calculus) or PSPACE-complete (e.g., IPC). Thus, a first approach is to help with fine-tuning the search, by selecting the right choices at choice points. Another way to integrate neural networks in symbolic computations has been via an interface to solve low-level “perception”-intensive tasks (e.g., working on learnable ground facts labeled with probabilities).

Our experiments hint to a third way that is more radical but also much simpler: *we replace the symbolic theorem prover with a neural network trained on a large enough training dataset.*

## References

1. Komendantskaya, E.: Unification neural networks: unification by error-correction learning. *Logic Journal of the IGPL* **19**(6) (2011) 821–847
2. Manhaeve, R., Dumancic, S., Kimmig, A., Demeester, T., De Raedt, L.: Deepproblog: Neural probabilistic logic programming. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R., eds.: *Advances in Neural Information Processing Systems* 31. Curran Associates, Inc. (2018) 3749–3759
3. Rocktäschel, T., Riedel, S.: End-to-end differentiable proving. In Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., eds.: *Advances in Neural Information Processing Systems* 30. Curran Associates, Inc. (2017) 3788–3800
4. Sloane, N.J.A.: *The On-Line Encyclopedia of Integer Sequences*. (2020) Published electronically at <https://oeis.org/>.
5. Stanley, R.P.: *Enumerative Combinatorics*. Wadsworth Publ. Co., Belmont, CA, USA (1986)
6. Tarau, P.: A Combinatorial Testing Framework for Intuitionistic Propositional Theorem Provers. In Alferes, J.J., Johansson, M., eds.: *Practical Aspects of Declarative Languages - 21th International Symposium, PADL 2019, Lisbon, Portugal, January 14-15, 2019, Proceedings*. Volume 11372 of *Lecture Notes in Computer Science*., Springer (2019) 115–132

7. Dyckhoff, R.: Contraction-free sequent calculi for intuitionistic logic. *Journal of Symbolic Logic* **57**(3) (1992) 795–807
8. Howard, W.: The Formulae-as-types Notion of Construction. In Seldin, J., Hindley, J., eds.: *To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*. Academic Press, London (1980) 479–490
9. Statman, R.: Intuitionistic Propositional Logic is Polynomial-Space Complete. *Theor. Comput. Sci.* **9** (1979) 67–72
10. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. *CoRR* **abs/1409.3215** (2014)
11. Girard, J.Y.: Linear logic. *Theoretical computer science* **50**(1) (1987) 1–101
12. Mints, G.E.: Closed categories and the theory of proofs. In: *Selected Papers in Proof Theory*, Bibliopolis (1992)
13. Zeilberger, N.: Balanced polymorphism and linear lambda calculus, talk at TYPES'15. (2015) <http://noamz.org/papers/linprin.pdf>.

# Logical English

Robert Kowalski  
Imperial College London

Logical English (LE) is a controlled natural language, in which English sentences are translated into a program in a logic programming (LP) language, such as Prolog, Datalog or Answer Set Programming (ASP). Similar controlled English languages that are also executed by translation to LP include ACE [Fuchs and Schwitter, 1996; Fuchs et al, 2008; Fuchs, 2013], PENG [Schwitter, 2002] and PENG<sup>ASP</sup> [Guy and Schwitter, 2017]. ACE and PENG are both intended for general-purpose knowledge representation and reasoning. In contrast, LE and PENG<sup>ASP</sup> are syntactic sugar for logic programs. PENG<sup>ASP</sup> provides syntactic sugar for ASP, and LE provides syntactic sugar for LPS [Kowalski and Sadri, 2015, 2016], which is an extension of pure Prolog, implemented in Prolog [Wielemaker et al, 2019].

The ultimate goal of LE is to serve as a general-purpose computer language, which can be understood by a reader without any training in computing, logic or mathematics. It is inspired in part by the language of law, which can be viewed as a programming language that is executed by humans rather than by computers [Kowalski, 1992]. LE can also be viewed as a domain-specific language for legal applications, similar to the English-like languages Blaux [Morris, 2020] and Lexon [Diedrich, 2020], both of which also have LP roots. Blaux is a combination of the LP language Flora-2 and the visual coding environment Blockly. Lexon on the other hand combines syntactic sugar for logic programs with higher-order logic, and compiles into Solidity, the programming language for the Ethereum blockchain.

LE is a declarative language, like the language of law. But LE also has an imperative character, inherited from LPS. Computation in both LE and LPS is similar to computation in conventional imperative languages, starting with an initial state, observing a potentially infinite stream of external events as input, and generating a potentially infinite stream of actions as output, while destructively updating a current state, with the goal of solving a problem or of simulating a real or imaginary world.

But unlike computation in conventional imperative languages, computation in LE also has a logical interpretation as generating a sequence of states and events, whose time-stamped history underlies a model that makes the program's goals true.

LE is a work in progress [Kowalski, 2019]. There have been three experimental implementations of variants of LE based on LPS or Prolog, focussed primarily on legal applications [Davila, 2017; Karadotchev, 2019; Fu, 2020]. The website <http://demo.logicalcontracts.com/> contains an example of the rock paper scissors game in a form of Logical English in the Fintech submenu of the examples menu. It also contains a number of LPS examples in the examples menu, which can be modified and executed online. Here are some examples based on this work:

- (1) If a player P1 plays a choice C1 and another player P2 plays a choice C2 and C1 beats C2 and it is not the case that the game is over then P1 receives the prize and it becomes the case that the game is over.
- (2) A transaction is governed by IsdaAgreement if a confirmation of the transaction states that the transaction is governed by IsdaAgreement and the transaction commences on a first day and IsdaAgreement is dated as of a second day and the first day is on or after the second day.
- (3) It becomes the case that a requirement is defaulted on a day when it is the end of the day and the requirement is potentially defaulted and the lender delivered a notice to the borrower on another day and the notice is that the requirement is potentially defaulted and the other day is 3 days before the day and it is not the case that the requirement is cured.

The first example translates into a reactive rule in LPS, written in the form *if antecedent then consequent*. Reactive rules in LPS represent goals that are made true by making their *consequents* true whenever their *antecedents* become true. *Consequents* can be made true either deliberately by performing actions or fortuitously by observing external events. The symbols P1, P2, C1 and C2 name variables. Their optional use in LE is similar to their use in legal texts. They provide names for variables, which look mathematical, but can be understood without mathematical training.

The second and third examples translate into ordinary LP clauses, which have the form *conclusion if conditions*. LP clauses define the models that can make goals true. In (2) and (3), as in ACE, PENG and LE more generally, variables can be represented by common nouns (such as *transaction* or *notice*) preceded by an article (*a*, *an* or *the*), as in *a transaction* and *the transaction*. The articles *a* and *an* are used for the first occurrence of the variable, and *the* is used for later occurrences of the same variable. Other variables of the same type in the same sentence can be introduced by preceding them with such adjectives as *first*, *second*, or *another* and *the other*, etc.

In (2) the condition *a confirmation of the transaction states that the transaction is governed by IsdaAgreement* illustrates the embedding of an object-level sentence inside a higher-order or meta-level sentence. This embedding is represented in Prolog by translating the phrase *is governed by* both into a predicate symbol, which is “used” at the object-level in the conclusion of the sentence, and into a function symbol, which is “mentioned” at the meta-level in a condition of the sentence.

Sentence (3) translates into an LP clause representing a causal relationship between an event *it is the end of a day* and a fluent *a requirement is defaulted on a day*, which is initiated by the event. In LPS, this relationship is implemented by adding the fluent to the current state if it is the end of a day and if the other conditions of the clause hold at

the end of the day. The representation uses the ontology of the event calculus [Kowalski and Sergot,1986], but an implementation involving destructive updates of a single current state. The frame axiom of the event calculus is an emergent property, which is true in any model that satisfies the goals, but it is not used for reasoning. The example could be expressed equally well in PENG<sup>ASP</sup>, which supports the writing of temporal specifications using an ASP-based adaptation of the event calculus.

To reduce ambiguity, LE has no pronouns, such as *he*, *she*, or *it*. To reduce the need for a dictionary, all nouns and verbs are expressed in the singular. The restriction to singular nouns means that LE does not use English quantifiers that require the use of plural nouns, such as *all* and *some*.

The use of articles in LE avoids the need for explicit quantification. In the case of a range-restricted LP clause (containing no variable in the conclusion that is not in the conditions), the natural reading of the articles in English conforms to the LP convention that all variables in the clause are universally quantified. But in the case of a non-range-restricted clause, the natural reading is that any variable in the conclusion that is not in the conditions is existentially quantified, For example:

- (4) An event of a person acquiring citizenship of the land of oz occurs on a day if the person is born in a place on the day and the place is in the land of oz.

Here the natural reading is that all variables are universally quantified except for the variable *an event*, which is existentially quantified. Moreover, although the scope of the universally quantified variables is limited to the clause, the existentially quantified variable has wider scope, as in the added clause:

- (5) A person celebrates the event if the person lives in the land of oz.

These readings of the English article are compatible with the interpretation of implicit quantifiers in existential (or  $\forall\exists$ ) rules, and with the elimination of existential quantifiers by skolemization [Baget et al, 2011].

All of the examples above are written in a basic form of LE, which is syntactic sugar for LPS. The plan is to develop LE as a series of extensions, starting from this basic form, introducing increasingly more natural syntaxes, while avoiding the introduction of ambiguity. For example, the LE sentence (2) above could be written in an extended form of LE as:

- (6) A transaction is governed by IsdaAgreement if the confirmation of the transaction states that the transaction is governed by IsdaAgreement and the transaction commences on a day that is on or after the day as of which IsdaAgreement is dated.

Here *a confirmation* is replaced by *the confirmation* to indicate that the relation between the confirmation and the transaction is “functional”, in the sense that there is only one confirmation for each transaction. The relative pronoun *that*, as well as the

preposition *as of* followed by *which*, introduces a restrictive relative clause, which inserts a logical condition into the text of another logical expression.

## LE and its relationship with other logics and other computer languages

LE and its logical underpinning LPS are based on a more general logic for abductive logic programming (ALP) [Kakas et al, 1992; Kowalski, 2011] in which logic programs are extended with abducible predicates (generalising actions in LPS) and with goals in first-order logic (generalising reactive rules and constraints). As in LPS, goals in this ALP logic are made true by a model determined by the logic program extended by facts expressed in the vocabulary of the abducible predicates.

In this ALP logic, goals of the form *if antecedent then consequent* are material implications, which can be satisfied *preventively* by making the *antecedent* false, or proactively by making the *consequent* true whether or not the *antecedent* ever becomes true. They can also be satisfied while performing unnecessary and irrelevant actions. In LPS and LE, goals of this form are reactive rules, which can be solved only *reactively* and *relevantly*, by generating actions to make *consequents* true whenever *antecedents* become true.

LPS is scaled down from ALP, losing some of the power of a problem-solving language, to compete more effectively with conventional computer languages for efficiency. However, the LE syntax for LPS introduces language features that are absent from both ALP and LPS, but which have been found to be useful in other computer languages. For example, even in the basic form of LE, the use of common nouns provides some of the features of a typed, object-oriented language. Other proposed extensions of LE provide some of the features of a functional language, as in *Monday is the day before the day before Wednesday*, which compiles into the LP relational form *Monday is the day before another day and the other day is before Wednesday*. The inclusion in LE of these and other features suggests that LE has the potential to compete with conventional computer languages not only for efficiency, but also for expressive power.

But no matter how LE compares with other computer languages today, the computer languages of the future need to be freed from such machine-oriented constructs as variables that name computer memory locations and assignment statements that manipulate the contents of computer memory. They also need to avoid the use of mathematical syntax in situations where it is unnecessary and where it makes the syntax unintelligible to the majority of readers.

Analogous developments are taking place in the legal domain, moving away from legal texts that can be understood only by legal professions, to texts written in plain language, so they can be understood by ordinary people [Williams, 2004]. We need to take similar steps in the domain of computing, accelerating the move away from languages that make people think like machines, to languages that make computers think more like people. Arguably, Logical English and its controlled natural language cousins provide a path, which is both logical and natural, for helping to reach this goal.

## References

- Baget, J.F., Leclère, M., Mugnier, M.L. and Salvat, E., 2011. On rules with existential variables: Walking the decidability line. *Artificial Intelligence*, 175(9-10), pp.1620-1654.
- Davila, J. 2017. Rock, Paper, Scissors. <http://demo.logicalcontracts.com/example/RockPaperScissorsBaseEN.pl>
- Diedrich, H. 2020. Lexon Bible: Hitchhiker's Guide to Digital Contracts. Wildfire Publishing.
- Fu, Z, 2020. Logical English (LE) for representing legal documents, MSc thesis. Imperial College London.
- Fuchs, N.E. 2013, Attempto Project. <http://attempto.ifi.uzh.ch/site/>
- Fuchs, N.E. and Schwitter, R., 1996. Attempto controlled english (ace). arXiv preprint [cmp-lg/9603003](https://arxiv.org/abs/cmp-lg/9603003).
- Fuchs, N.E., Kaljurand, K. and Kuhn, T., 2008. Attempto controlled english for knowledge representation. In *Reasoning Web* (pp. 104-124). Springer, Berlin, Heidelberg.
- Guy, S.C. and Schwitter, R., 2017. The PENG<sup>ASP</sup> system: architecture, language and authoring tool. *Language Resources and Evaluation*, 51(1), pp.67-92.
- Kakas, A.C., Kowalski, R.A. and Toni, F., 1992. Abductive logic programming. *Journal of logic and computation*, 2(6), pp.719-770.
- Karadotchev, V., 2019. First Steps Towards Logical English. MSc thesis. Imperial College London.
- Kowalski, R., 2011. *Computational logic and human thinking: how to be artificially intelligent*. Cambridge University Press.
- Kowalski, R., 1992. Legislation as Logic Programs In: *Logic Programming in Action* (eds. G. Comyn , N. E. Fuchs, M. J. Ratcliffe), Springer Verlag, pages 203-230.
- Kowalski, R. 2019. Logical English slides. <http://www.doc.ic.ac.uk/~rak/papers/Logical%20English.pdf>
- Kowalski, R. and Sadri, F., 2015. "Reactive Computing as Model Generation." *New Generation Computing*, Volume 33, Issue 1, pp 33-67
- Kowalski, R. and Sadri, F., 2016. Programming in logic without logic programming. *Theory and Practice of Logic Programming*, 16(03), pp.269-295.
- Kowalski, R. and Sergot, M., 1986, "A Logic-based Calculus of Events", in *New Generation Computing*, Vol. 4, No.1, February pp. 67-95.
- Morris, J., 2020. Blawx Alpha: User Friendly Rules as Code on the Web. <https://www.blawx.com/>
- Schwitter, R., 2002, English as a formal specification language. In *Proceedings. 13th International Workshop on Database and Expert Systems Applications* (pp. 228-232). IEEE.
- Schwitter, R., 2019. Augmenting an answer set based controlled natural language with temporal expressions. In *PRICAI 2019: Trends in Artificial Intelligence*, vol. 11670 of LNAI, 500–513. Springer.
- Wielemaker, J., Riguzzi, F., Kowalski, R. A., Lager, T., Sadri, F., & Calejo, M. 2019. "Using SWISH to realize interactive web-based tutorials for logic-based languages", *Theory and Practice of Logic Programming*, 19(2), 229-261.
- Williams, C., 2004. Legal English and plain language: An introduction. *ESP across Cultures*, 1(1), pp.111-124.

# Flamingo, a Compiler and Runtime for Reactive ALM

Daniel Hines

September 3rd, 2019

The Action Language with Modules,  $\mathcal{ALM}$ , is a declarative, relational language for describing state transition diagrams [3]. An extension of language  $\mathcal{AL}$  [2],  $\mathcal{ALM}$  is predicated on the notion that the key to success in industrial settings is not raw expressive power, but agility in the form of step-wise development, readability, and resilience to changing requirements. In this light,  $\mathcal{ALM}$  is carefully limited in its expressiveness so as to best support common-sense reasoning over real-world problems. Additionally,  $\mathcal{ALM}$  supports an inheritance-based module system for the support of reasoning libraries, enabling specific problems to be defined as special cases of more general ones.

Flamingo is a compiler and runtime for  $\mathcal{ALM}$  for Javascript. Flamingo currently targets the *reactive* case of  $\mathcal{ALM}$  systems, where all actions occur extrinsically and the system merely responds. Specifically, Flamingo is designed for user interface, and integrates with the React ecosystem<sup>1</sup>, allowing the rapid design and development of web applications.

Flamingo also enables the automatic verification of  $\mathcal{ALM}$  systems via  $\mathcal{ALM}$ 's support for negative state constraints. Inspired by systems like TLA+ [1] and Alloy [4], Flamingo builds on these by limiting the expressiveness of the modeling language to a much more beginner-friendly subset, while also supporting using the model directly in Javascript to power real web applications, without first translating by hand into an imperative language.

The Flamingo dialect of  $\mathcal{ALM}$  is an extension of Datalog that adds the notion of action and change. Indeed, Flamingo and Datalog share the same

---

<sup>1</sup><https://reactjs.org/>

complexity class (for the formal semantics of the Flamingo dialect, see <sup>2</sup>). Additionally, ALM adds a powerful sort and module system, enabling greater expressiveness, modularity, and error-checking than single-sorted Datalog. Flamingo specifications consist of two parts: the signature, which defines sorts and functions between them, and the axioms, the rules that govern the system. Functions have two varieties: *fluent*, which have an implicit time parameter and thus may change in value as a result of actions, and *static*, which do not change over time. Fluent functions are further divided into *basic fluents* which have inertia and retain their value from state to state unless directly changed by an action, and *defined*, which do not have inertia and are re-derived every state in terms of other fluents. Axioms also have three kinds: state constraints, which say “if [condition] is true in a given state, then [consequence] is true in that same state”; causal laws, which describe the consequence of a given action under some given condition; and executability conditions, which constrain which actions may occur under given conditions.

Flamingo can be used for “lightweight formal methods” in the spirit of Alloy. Sorts and fluent functions provide a modeling language to define possible states for your system. Axioms in the form of state constraints reduce the state space in a way that Flamingo can check automatically for contradiction. Executability conditions allow the user to assure Flamingo that a given action is impossible in certain circumstances.

Several important research and engineering problems remain for Flamingo. The first is how to model dynamic virtual systems, much like the RBAC challenge<sup>3</sup>, where entities are virtual and can be created *ex nihilo* as it were. We have had some success in this area using “table”-like entities, and created a very clean and compact implementation of the TodoMVC application<sup>4</sup> that you can find here on our website<sup>5</sup>. The second is how to extend Flamingo to infinite (or at least indefinite) domains and retain the ability to check them for correctness.

Our solution to the RBAC challenge is on Github as gist<sup>6</sup>. An introduction to Flamingo’s syntax and usage can be found on the documentation site<sup>7</sup>,

---

<sup>2</sup><https://flamingo-lang.org/docs/research/semantics>

<sup>3</sup><https://www3.cs.stonybrook.edu/liu/papers/RBACchallenge-LPOP18.pdf>

<sup>4</sup><https://todomvc.com>

<sup>5</sup><https://flamingo-lang.org/docs/todomvc/intro>

<sup>6</sup><https://gist.github.com/d4hines/136bc4b79d16f759081a1d62c1a3245f>

<sup>7</sup><https://flamingo-lang.org/docs/basics/basics>

along with a guide on integrating Flamingo with a React web application<sup>8</sup>.

## References

- [1] The TLA+ Home Page, <https://lamport.azurewebsites.net/tla/tla.html>.
- [2] Michael Gelfond and Yulia Kahl. *Knowledge Representation, Reasoning, and the Design of Intelligent Agents: The Answer-Set Programming Approach*. Cambridge University Press, 2014.
- [3] Daniela Incelezan and Michael Gelfond. Modular Action Language ALM. *Theory and Practice of Logic Programming*, 16(2):189–235, March 2016. arXiv: 1505.05022.
- [4] Daniel Jackson. *Software abstractions logic, language and analysis*. MIT Press, Cambridge, Mass., 2006.

---

<sup>8</sup><https://flamingo-lang.org/docs/todomvc/intro>