



Inductive logic programming at 30

Andrew Cropper¹ · Sebastijan Dumančić² · Richard Evans³ · Stephen H. Muggleton³

Received: 1 June 2021 / Revised: 6 August 2021 / Accepted: 27 September 2021 /
Published online: 9 November 2021
© The Author(s) 2021

Abstract

Inductive logic programming (ILP) is a form of logic-based machine learning. The goal is to induce a hypothesis (a logic program) that generalises given training examples and background knowledge. As ILP turns 30, we review the last decade of research. We focus on (i) new meta-level search methods, (ii) techniques for learning recursive programs, (iii) new approaches for predicate invention, and (iv) the use of different technologies. We conclude by discussing current limitations of ILP and directions for future research.

Keywords Inductive logic programming · Relational learning · Program synthesis · Program induction

1 Introduction

Inductive logic programming (ILP) (Muggleton, 1991; Muggleton & De Raedt, 1994) is a form of machine learning (ML). As with other forms of ML, the goal is to induce a hypothesis that generalises training examples. However, whereas most forms of ML use vectors/tensors to represent data, ILP uses logic programs (sets of logical rules). Moreover, whereas most forms of ML learn functions, ILP learns relations.

Editors: Nikos Katzouris, Alexander Artikis, Luc De Raedt, Artur d’Avila Garcez, Ute Schmid, Jay Pujara.

✉ Andrew Cropper
andrew.cropper@cs.ox.ac.uk

Sebastijan Dumančić
sebastijan.dumancic@cs.kuleuven.be

Richard Evans
richardevans@google.com

Stephen H. Muggleton
s.muggleton@imperial.ac.uk

¹ University of Oxford, Oxford, UK

² KU Leuven, Leuven, Belgium

³ Imperial College London, London, UK

To illustrate ILP¹ suppose you want to learn a string transformation program from the following examples:

Input	Output
inductive	e
logic	c
programming	g

Most forms of ML would use a table to represent these examples. Each row would be an example. Each column would be a feature, such as a one-hot-encoding representation of the string. By contrast, in ILP, we would represent these examples as logical atoms, such as $f([i, n, d, u, c, t, i, v, e], e)$, where f is the target predicate that we want to learn (the relation to generalise). We would also provide auxiliary information as background knowledge (BK), also represented as a logic program. For instance, we could provide BK that contains logical definitions for string operations, such as `empty(A)`, which holds when the list A is empty; `head(A, B)`, which holds when B is the head of the list A ; and `tail(A, B)`, which holds when B is the tail of the list A . Given the aforementioned examples and BK, an ILP system could induce the hypothesis (a logic program):

```
f(A,B):- tail(A,C),empty(C),head(A,B).
f(A,B):- tail(A,C),f(C,B).
```

Each line of the program is a rule. The first rule says that the relation $f(A, B)$ holds when the three literals `tail(A, C)`, `empty(C)`, and `head(A, B)` hold. In other words, the first rule says that B is the last element of A when the tail of A is empty and B is the head of A . The second rule is recursive and says that the relation $f(A, B)$ holds when the two literals `tail(A, C)` and `f(C, B)` hold. In other words, the second rule says that $f(A, B)$ holds when the same relation holds for the tail of A .

1.1 Why ILP?

Compared to most ML approaches, ILP has several attractive features (Cropper et al., 2020a; Cropper & Dumancic, 2020a).

Data efficiency Many forms of ML are notorious for their inability to generalise from small numbers of training examples, notably deep learning (Marcus, 2018; Chollet, 2019). As Evans and Grefenstette (2018) point out, if we train a neural system to add numbers with 10 digits, it might generalise to numbers with 20 digits, but when tested on numbers with 100 digits, the predictive accuracy drastically decreases (Reed & de Freitas, 2016; Kaiser & Sutskever, 2016). By contrast, ILP can induce hypotheses from small numbers of examples, often from a single example (Lin et al., 2014; Muggleton et al., 2018a).

¹ We do not introduce ILP in detail and refer the reader to the introductory paper of Cropper and Dumancic (2020a) or the textbooks of Nienhuys-Cheng and Wolf (1997) and De Raedt (2008).

Table 1 A simplified comparison of old and new ILP systems

	Old ILP	New ILP
Search method	Top-down and bottom-up	Meta-level
Recursion	Limited	Yes
Predicate invention	No	Limited
Hypotheses	First-order	ASP, higher-order, probabilistic
Optimality	No	Yes
Technology	Prolog	Prolog, ASP, NNs

Background knowledge ILP learns using BK represented as a logic program. Using logic programs to represent data allows ILP to learn with complex relational information, such as constraints about causal networks (Inoue et al., 2013), the axioms of the event calculus when learning to recognise events (Katzouris et al., 2015, 2016), and using a theory of light to understand images (Muggleton et al., 2018a). Moreover, because hypotheses are symbolic, hypotheses can be added to BK, and thus ILP systems naturally support lifelong and transfer learning (Lin et al., 2014; Cropper, 2019, 2020).

Expressivity Because of the expressivity of logic programs, ILP can learn complex relational theories, such as cellular automata (Inoue et al., 2014; Evans et al., 2021), event calculus theories (Katzouris et al., 2015, 2016), Petri nets (Bain & Srinivasan, 2018), answer set programs (ASP) (Law et al., 2014), and general algorithms (Cropper & Morel, 2021a). Because of the symbolic nature of logic programs, ILP can reason about hypotheses, which allows it to learn *optimal* programs, such as minimal time-complexity programs (Cropper & Muggleton, 2019) and secure access control policies (Law et al., 2020a).

Expainability Because of logic’s similarity to natural language, logic programs can be easily read by humans, which is crucial for explainable AI. For instance, Muggleton et al., (2018b) provide the first demonstration of *ultra-strong ML* (Michie, 1988), where a learned hypothesis is expected to not only be accurate but to also demonstrably improve the performance of a human when provided with the learned hypothesis.

1.2 Recent advances

Some of the aforementioned advantages come from developments in the last decade of ILP research, which we survey in this paper.² To aid the reader, we coarsely compare old and new ILP systems, where new represents systems from the past decade. We use FOIL (Quinlan, 1990), Progol (Muggleton, 1995), Aleph (Srinivasan, 2001), TILDE (Blokceel & De Raedt, 1998), and HYPER (Bratko, 1999) as representative old systems and ILASP (Law et al., 2014), Metagol (Cropper & Muggleton, 2016), ∂ ILP (Evans & Grefenstette, 2018), and Popper (Cropper & Morel, 2021a) as representative new systems. This comparison, shown in Table 1, is, of course, vastly oversimplified, and there are many exceptions. In the rest of this paper, we survey these developments (each row in the table) in turn. After

² This paper extends the paper of Cropper et al. (2020a).

discussing these new ideas, we discuss recent application areas (Sect. 8) before concluding by proposing directions for future research.

2 Search methods

The fundamental ILP problem is to efficiently search a large hypothesis space. Most older ILP approaches search in either a *top-down* or *bottom-up* fashion. These methods rely on notions of generality (typically using theta-subsumption (Plotkin, 1971)), where one program is more *general* or more *specific* than another. A third new search approach has recently emerged called *meta-level* ILP (Inoue et al., 2013; Muggleton et al., 2015; Inoue, 2016; Law et al., 2020b; Cropper & Morel, 2021a). We discuss these approaches in turn.

2.1 Top-down and bottom-up

Top-down approaches (Quinlan, 1990; Blockeel & De Raedt, 1998; Bratko, 1999) start with a general hypothesis and then specialise it. HYPER, for instance, searches a tree in which the nodes correspond to hypotheses and each child of a hypothesis in the tree is more specific than or equal to its predecessor in terms of theta-subsumption. An advantage of top-down approaches is that they can often learn recursive programs (although not all do). A disadvantage is that they can be prohibitively inefficient because they can generate many hypotheses that do not cover the examples.

Bottom-up approaches, by contrast, start with the examples and generalise them (Muggleton, 1987; Muggleton & Buntine, 1988; Muggleton & Feng, 1990; Inoue et al., 2014). For instance, Muggleton and Feng (1990) generalises pairs of examples based on relative least-general generalisation (Nienhuys-Cheng & Wolf, 1997). Bottom-up approaches can be seen as being *data-* or *example-driven*. An advantage of these approaches is that they are typically fast. As Bratko (1999) points out, disadvantages include (i) they typically use unnecessarily long hypotheses with many clauses, (ii) it is difficult for them to learn recursive hypotheses and multiple predicates simultaneously, and (iii) they do not easily support predicate invention.

Muggleton (1995), which inspired many other ILP approaches (Srinivasan, 2001; Ray, 2009; Ahlgren & Yuen, 2013; Schüller & Benz, 2018), combines both top-down and bottom-up approaches. Starting with an empty program, Progol picks an uncovered positive example to generalise. To generalise an example, Progol uses mode declarations to build the *bottom clause* (Muggleton, 1995), the logically most-specific clause that explains the example. The bottom clause bounds the search from below (the bottom clause) and above (the empty set). Progol then uses an A* algorithm to generalise the bottom clause in a top-down (general-to-specific) manner and uses the other examples to guide the search.

2.1.1 Meta-level

Top-down and bottom-up approaches refine and revise a single hypothesis. A third approach has recently emerged called *meta-level* ILP (Inoue et al., 2013; Muggleton et al.,

2015; Inoue, 2016; Law et al., 2020b; Cropper & Morel, 2021a; Patsantzis & Muggleton, 2021). There is no standard definition for *meta-level ILP*. Most approaches encode the ILP problem as a meta-level logic program, i.e. a program that reasons about programs. Meta-level approaches then often delegate the search for a hypothesis to an off-the-shelf solver (Corapi et al., 2011; Cropper & Muggleton, 2016; Law et al., 2014; Kaminski et al., 2018; Schüller & Benz, 2018; Evans et al., 2021; Cropper & Morel, 2021a) after which the meta-level solution is translated back to a standard solution for the ILP task. In other words, instead of writing a procedure to search in a top-down or bottom-up manner, most meta-level approaches formulate the learning problem as a declarative search problem. For instance, ASPAL (Corapi et al., 2011) translates an ILP task into a meta-level ASP program that describes every example and every possible rule in the hypothesis space. ASPAL then delegates the search to an ASP system to find a subset of the rules that covers all the positive but none of the negative examples.

The main advantage of meta-level approaches is that they can more easily learn recursive programs and optimal programs (Corapi et al., 2011; Law et al., 2014; Cropper & Muggleton, 2016; Kaminski et al., 2018; Evans et al., 2021; Cropper & Morel, 2021a), which we discuss in Sects. 3 and 6 respectively. Moreover, whereas classical ILP systems were almost entirely based on Prolog, meta-level approaches use diverse techniques and technologies, such as ASP solvers (Corapi et al., 2011; Law et al., 2014; Kaminski et al., 2018; Cropper & Morel, 2021a; Evans et al., 2021), which we expand on in Sect. 7.

The development of meta-level ILP approaches has, therefore, diversified ILP from the standard clause refinement approach of earlier ILP systems.

Most meta-level approaches encode the ILP learning task as a single static meta-level program (Corapi et al., 2011; Law et al., 2014; Kaminski et al., 2018; Evans et al., 2021).

A major issue with this approach is that the meta-level program can be very large so these approaches can struggle to scale to problems with non-trivial domains and programs with large clauses.

Two related approaches try to overcome this limitation by continually revising the meta-level program.

ILASP3 (Law, 2018) employs a counter-example-driven select-and-constrain loop. ILASP3 first pre-computes every clause in the hypothesis space defined by a set of given mode declarations (Muggleton, 1995). ILASP3 then starts its select-and-constrain loop. With each iteration, ILASP3 uses an ASP solver to find the best hypothesis (a subset of the rules) it can. If the hypothesis does not cover one of the examples, ILASP3 finds a reason why and then generates constraints (boolean formulas over the rules) which it adds to the meta-level program to guide subsequent search. Another way of viewing ILASP3 is that it uses a counter-example-guided approach and translates an uncovered example e into a constraint that is satisfied if and only if e is covered.

Popper (Cropper & Morel, 2021a) adopts a similar approach but differs in that it (i) does not precompute every possible rule in the hypothesis space, and (ii) translates a *hypothesis*, rather than an uncovered example, into a set of constraints. Popper works in three repeating stages: *generate*, *test*, and *constrain*. Popper first constructs a meta-level logic program where its models correspond to hypotheses. In the generate stage, Popper asks an ASP solver to find a model (a hypothesis). In the test stage, Popper tests the hypothesis against the examples. A hypothesis *fails* when it is incomplete (does not entail all the positive examples) or inconsistent (entails a negative example). If a hypothesis fails, Popper learns constraints from the failure, which it then uses to restrict subsequent generate stages. For instance, if a hypothesis is inconsistent, then Popper generates a generalisation constraint to prune all generalisations of the hypothesis and adds the constraint to the

meta-level program, which eliminates models and thus prunes the hypothesis space. This process repeats until Popper finds a complete and consistent program.

For more information about meta-level learning, we suggest the work of Inoue (2016) and Law et al. (2020b).

3 Recursion

Learning recursive programs has long been considered a difficult problem for ILP (Muggleton et al., 2012; Cropper & Dumancic, 2020a). The power of recursion is that an infinite number of computations can be described by a finite recursive program (Wirth, 1985). To illustrate the importance of recursion, reconsider the string transformation problem from the introduction. Without recursion, an ILP system would need to learn a separate clause to find the last element for each list of length n , such as this program for when $n = 3$:

```
f(A,B):- tail(A,C),empty(C),head(A,B).
f(A,B):- tail(A,C),tail(C,D),empty(D),head(C,B).
f(A,B):- tail(A,C),tail(C,D),tail(D,E),empty(E),head(D,B).
```

This program does not generalise to lists of arbitrary lengths. Moreover, most ILP systems would need examples of lists of each length to learn such a program. By contrast, an ILP system that supports recursion can learn the compact program:

```
f(A,B):- tail(A,C),empty(C),head(A,B).
f(A,B):- tail(A,C),f(C,B).
```

Because of the symbolic representation and the recursive nature, this program generalises to lists of arbitrary length and which contain arbitrary elements (e.g. integers and characters). In general, without recursion, it can be difficult for an ILP system to generalise from small numbers of examples (Cropper et al., 2015).

Older ILP systems struggle to learn recursive programs, especially from small numbers of training examples. A common limitation with existing approaches is that they rely on *bottom clause* construction (Muggleton, 1995). In this approach, for each example, an ILP system creates the most specific clause that entails the example and then tries to generalise the clause to entail other examples. However, this sequential covering approach requires examples of both the base and inductive cases. The classical ILP system FOIL (Quinlan, 1990) also struggles to learn recursive programs because it induces programs one clause at a time.

Interest in recursion has resurged with the introduction of meta-interpretive learning (MIL) (Muggleton et al., 2014, 2015; Cropper et al., 2020c) and the MIL system Metagol (Cropper & Muggleton, 2016). The key idea of MIL is to use *metarules* (Cropper &

Touret, 2020), or program templates, to restrict the form of inducible programs, and thus the hypothesis space.³ A metarule is a higher-order clause. For instance, the *chain* metarule is $P(A, B) \leftarrow Q(A, C), R(C, B)$, where the letters P , Q , and R denote higher-order variables and A , B and C denote first-order variables. The goal of a MIL system, such as Metagol, is to find substitutions for the higher-order variables. For instance, the *chain* metarule allows Metagol to induce programs such as $f(A, B) :- tail(A, C), head(C, B)$.⁴ Metagol induces recursive programs using recursive metarules, such as the *tailrec* metarule $P(A, B) \leftarrow Q(A, C), P(C, B)$.

Following MIL, many meta-level ILP systems can learn recursive programs (Law et al., 2014; Evans & Grefenstette, 2018; Kaminski et al., 2018; Cropper & Morel, 2021a). With recursion, ILP systems can now generalise from small numbers of examples, often a single example (Lin et al., 2014). Moreover, the ability to learn recursive programs has opened up ILP to new application areas, including learning string transformations programs (Lin et al., 2014), answer set grammars (Law et al., 2019), and general algorithms (Cropper & Morel, 2021a).

4 Predicate invention

A key characteristic of ILP is the use of BK, which contains facts and rules (extensional and intensional definitions) in the form of a logic program. For instance, when learning string transformation programs, we may provide helper background relations, such as *head/2* and *tail/2*. For other domains, we may supply more complex BK, such as a theory of light to understand images (Muggleton et al., 2018a) or higher-order operations, such as *map/3*, *filter/3*, and *fold/4*, to solve programming puzzles (Cropper et al., 2020c).

Choosing appropriate BK is crucial for good learning performance. ILP has traditionally relied on hand-crafted BK, often designed by domain experts. This approach is limited because obtaining suitable BK can be difficult and expensive. Indeed, the over-reliance on hand-crafted BK is a common criticism of ILP (Evans & Grefenstette, 2018).

Rather than expecting a user to provide all the necessary BK, the goal of *predicate invention* (PI) (Muggleton & Buntine, 1988; Stahl, 1995) is for an ILP system to automatically invent new auxiliary predicate symbols. This idea is similar to when humans create new functions when manually writing programs, to reduce code duplication or to improve readability. Whilst PI has attracted interest since the beginnings of ILP (Muggleton & Buntine, 1988), and has subsequently been repeatedly stated as a major challenge (Kok & Domingos, 2007; Muggleton et al., 2012; Kramer, 2020), most ILP systems do not support it, including classical systems, such as Progol (Muggleton, 1995), TILDE (Blockeel & De Raedt, 1998), and Aleph (Srinivasan, 2001), and modern systems, such as ATOM (Ahlgren & Yuen, 2013) and LFIT (Inoue et al., 2014).

³ The idea of using metarules to restrict the hypothesis space has been widely adopted by many approaches (Wang et al., 2014; Albarghouthi et al., 2017; Rocktäschel & Riedel, 2017; Evans & Grefenstette, 2018; Bain & Srinivasan, 2018; Kaminski et al., 2018). However, despite their now widespread use, there is little work determining which metarules to use for a given learning task (Cropper & Touret, 2020 is an exception), which future work must address.

⁴ Metagol can induce longer clauses though predicate invention, which is described in Sect. 4.

A key challenge faced by ILP systems is deciding when and how to invent a new symbol. As Kramer (1995) points out, PI is difficult because it is unclear how many arguments an invented predicate should have, how the arguments should be ordered, etc. Several PI approaches try to address this challenge, which we discuss in turn.

4.1 Placeholders

A classical approach to PI is to predefine invented symbols through mode declarations, which (Leban et al., 2008) call *placeholders*. However, this placeholder approach is limited because it requires that a user manually specify the arity and argument types of a symbol (Law et al., 2014), which rather defeats the point, or requires generating all possible invented predicates (Evans & Grefenstette, 2018; Evans et al., 2021), which is computationally expensive.

4.2 Metarules

Interest in *automatic* PI (where a user does not need to predefine an invented symbol) has resurged with the introduction of MIL. MIL avoids the issues of older ILP systems by using metarules to define the hypothesis space and in turn reduce the complexity of inventing a new predicate symbol. For instance, the *chain* metarule ($P(A, B) \leftarrow Q(A, C), R(C, B)$) allows Metagol to induce programs such as $f(A, B) :- tail(A, C), tail(C, D)$, which would drop the first two elements from a list. To induce longer clauses, such as to drop first three elements from a list, Metagol uses the same metarule but invents a predicate symbol to chain their application, such as to induce the program:

```
f(A,B):- tail(A,C),inv1(C,B).
inv1(A,B):- tail(A,C),tail(C,B).
```

To learn this program, Metagol invents the predicate symbol `inv1` and induces a definition for it using the *chain* metarule. Metagol uses this new predicate symbol in the definition for the target predicate `f`.

A side-effect of this metarule-driven approach is that problems are forced to be decomposed into reusable solutions. For instance, to learn a program that drops the first four elements of a list, Metagol learns the following program, where the invented predicate symbol `inv1` is used twice:

```
f(A,B):- inv1(A,C),inv1(C,B).
inv1(A,B):- tail(A,C),tail(C,B).
```

PI has been shown to help reduce the size of target programs, which in turn reduces sample complexity and improves predictive accuracy (Cropper, 2019). Several new ILP systems support PI using a metarule-guided approach (Evans & Grefenstette, 2018; Kaminski et al., 2018; Hocquette & Muggleton, 2020).

4.3 Pre/post-processing

Metarule-driven PI approaches perform PI during the learning task. A recent trend is to perform PI as a pre/post-processing step to improve knowledge representation (Dumančić & Blockeel, 2017; Dumančić et al., 2019; Cropper, 2019; Hocquette & Muggleton, 2020).

CUR²LED (Dumančić & Blockeel, 2017) performs PI by clustering constants and relations in the provided BK, turning each identified cluster into a new BK predicate. The key insight of CUR²LED is not to use a single similarity measure, but rather a set of various similarities. This choice is motivated by the fact that different similarities are useful for different tasks, but in the unsupervised setting the task itself is not known in advance. CUR²LED performs PI by producing different clusterings according to the features of the objects, community structure, and so on.

ALPs (Dumančić et al., 2019) perform PI using an auto-encoding principle: they learn an *encoding* logic program that maps the provided data to a new, compressive latent representation (defined in terms of the invented predicates), and a *decoding* logic program that can reconstruct the provided data from its latent representation. This approach shows improved performance on supervised tasks, even though the PI step is task-agnostic.

Knorf (Dumancic et al., 2020) pushes the idea of ALPs even further. Knorf compresses a program by removing redundancies in it. If the learnt program contains invented predicates, Knorf revises them and introduces new ones that would lead to a smaller program. The refactored program is smaller in size and contains less redundancy in clauses, both of which lead to improved performance. The authors experimentally demonstrate that refactoring improves learning performance in lifelong learning and that Knorf substantially reduces the size of the BK program, reducing the number of literals in a program by 50% or more.

4.4 Lifelong learning

An approach to acquiring BK is to learn it in a lifelong learning setting. The general idea is to reuse knowledge gained from solving one problem to help solve a different problem.

Metagol_{DF} is an ILP system (Lin et al., 2014) which given a set of tasks, uses Metagol to try to learn a solution for each task using at most one clause. If Metagol finds a solution for a task, it adds the solution to the BK and removes the task from the set. Metagol_{DF} then asks Metagol to find solutions for the rest of the tasks but can now (i) use an additional clause, and (ii) reuse solutions from previously solved tasks. This process repeats until Metagol_{DF} solves all the tasks or reaches a maximum program size. In this approach, Metagol_{DF} automatically identifies easier problems, learn programs for them, and then reuses the solutions to help learn programs for more difficult problems. The authors experimentally show that their multi-task approach performs substantially better than a single-task approach because learned programs are frequently reused and leads to a hierarchy of induced programs.

Metagol_{DF} saves all learned programs (including invented predicates) to the BK, which can be problematic because too much irrelevant BK is detrimental to learning performance (Srinivasan et al., 2003). To address this problem, Forgetgol (Cropper, 2020) introduces the idea of *forgetting*. In this approach, Forgetgol continually grows and shrinks its hypothesis space by adding and removing learned programs to and from its BK. The authors show

that forgetting can reduce both the size of the hypothesis space and the sample complexity of an ILP learner when learning from many tasks.

4.5 Limitations

The aforementioned techniques have improved the ability of ILP to invent high-level concepts. However, PI is still difficult and there are many challenges to overcome, notably that (i) many systems struggle to perform PI at all, and (ii) those that do support PI mostly need much user-guidance, such as metarules to restrict the space of invented symbols or that a user specifies the arity and argument types of invented symbols. There are notable exceptions. Ferilli (2016) describe an PI approach based on the ideal of specialising a theory to account for negative examples, similar to early work in non-monotonic ILP (Bain & Mugleton, 1992). POPPI (Cropper & Morel, 2021b) is an ILP system that supports automatic predicate invention, i.e. does not require metarules nor requires a user to predefine invented symbols.

By developing better approaches for PI, we can make progress on existing challenging problems. For instance, in *inductive general game playing* (IGGP) (Cropper et al., 2020b), the task is to learn the symbolic rules of games from observations of gameplay, such as learning the rules of *connect four*. The target solutions, which come from the general game playing competition (Genesereth & Björnsson, 2013), often contain auxiliary predicates. For instance, the rules for *connect four* are defined in terms of definitions for lines which are themselves defined in terms of columns, rows, and diagonals. Although these auxiliary predicates are not strictly necessary to learn the target solution, inventing such predicates significantly reduces the size of the solution, which in turn makes them easier to learn. Although new methods for PI can invent high-level concepts, they are not yet sufficiently powerful enough to perform well on the IGGP dataset. Making progress in this area would constitute a major advancement in ILP.

5 Hypotheses

ILP systems have traditionally induced definite and normal logic programs, typically represented as Prolog programs. A recent development has been to use different hypothesis representations.

5.1 Datalog

Datalog is a syntactical subset of Prolog which disallows complex terms as arguments of predicates and imposes restrictions on the use of negation.

Datalog is a truly declarative language, whereas in Prolog reordering clauses can change the program. Moreover, Datalog query is guaranteed to terminate, though this guarantee is at the expense of not being a Turing-complete language, which Prolog is. Several works (Albarghouthi et al., 2017; Evans & Grefenstette, 2018; Kaminski et al., 2018) induce Datalog programs. The general motivation for reducing the expressivity of the representation language from Prolog to Datalog is to allow the problem to be encoded as a satisfiability problem, particularly to leverage recent developments in SAT and SMT. We discuss the advantages of this approach more in Sect. 7.1.

5.2 Answer set programming

ASP (Gebser et al., 2012a) is a logic programming paradigm based on the stable model semantics of normal logic programs that can be implemented using the latest advances in SAT solving technology. Law et al. (2018) discuss some of the advantages of learning ASP programs, rather than Prolog programs, which we reiterate. When learning Prolog programs, the procedural aspect of SLD-resolution must be taken into account. For instance, when learning Prolog programs with negation, programs must be stratified; otherwise, they may loop under certain conditions. By contrast, as ASP is a truly declarative language, no such consideration needs to be taken into account when learning ASP programs. Compared to Datalog and Prolog, ASP supports additional language constructs, such as disjunction in the head of a clause, choice rules, and hard and weak constraints. A key difference between ASP and Prolog is semantics. A definite logic program has only one model (the least Herbrand model). By contrast, an ASP program can have one, many, or even no stable models (answer sets). Due to its non-monotonicity, ASP is particularly useful for expressing common-sense reasoning (Law, 2018).

To illustrate the benefits of learning ASP programs, we reuse an example from Law et al. (2020b). Given a sufficient examples of Hamiltonian graphs, ILASP (Law et al., 2014) can learn a program to definite them:

```

0 { in(V0, V1) } 1 :- edge(V0, V1).
reach(V0) :- in(1, V0).
reach(V1) :- reach(V0), in(V0, V1).
:- not reach(V0), node(V0).
:- V1 != V2, in(V0, V2), in(V0, V1).

```

This program illustrates useful language features of ASP. The first rule is a *choice* rule, which means that an atom *can* be true. In this example, the rule indicates that there can be an in edge from the vertex $V1$ to $V0$. The last two rules are *hard constraints*, which essentially enforce integrity constraints. The first hard constraint states that it is impossible to have a node that is not reachable. The second hard constraint states that it is impossible to have a vertex with two in edges from distinct nodes. For more information about ASP we recommend the book by Gebser et al. (2012a).

Approaches to learning ASP programs can mostly be divided into two categories: *brave learners*, which aim to learn a program such that at least one answer set covers the examples, and *cautious learners*, which aim to find a program which covers the examples in all answer sets. ILASP is notable because it supports both brave and cautious learning, which are both needed to learn some ASP programs (Law et al., 2018). Moreover, ILASP differs from most Prolog-based ILP systems because it learns ASP programs, including programs with normal rules, choice rules, and both hard and weak constraints, which classical ILP systems cannot. Learning ASP programs allows for ILP to be used for new problems, such as inducing answer set grammars (Law et al., 2019).

5.3 Higher-order programs

Imagine learning a *droplasts* program, which removes the last element of each sublist in a list, e.g. $[alice, bob, carol] \mapsto [alic, bo, caro]$. Given suitable input data, Metagol can learn this first-order recursive program:




Fig. 1 ILP systems struggle with structured examples that exhibit observational noise. All three examples clearly spell the word “ILP”, with some alterations: 3 noisy pixels, shifted and elongated letters. If we would be to learn a program that simply draws “ILP” in the middle of the picture, without noisy pixels and elongated letters, that would be a correct program

```
f(A,B):- empty(A),empty(B).
f(A,B):- head(A,C),tail(A,D),head(B,E),tail(B,F),f1(C,E),f(D,F).
f1(A,B):- reverse(A,C),tail(C,D),reverse(D,B).
```

Although semantically correct, the program is verbose. To learn smaller programs, Metagol_{ho} (Cropper et al., 2020c) extends Metagol to support learning higher-order programs, where predicate symbols can be used as terms. For instance, for the same *droplasts* problem, Metagol_{ho} learns the higher-order program:

```
f(A,B):- map(A,B,f1).
f1(A,B):- reverse(A,C),tail(C,D),reverse(D,B).
```

To learn this program, Metagol_{ho} invents the predicate symbol $f1$, which is used twice in the program: as term in the $\text{map}(A,B,f1)$ literal and as a predicate symbol in the $f1(A,B)$ literal. Compared to the first-order program, this higher-order program is smaller because it uses $\text{map}/3$ (predefined in the BK) to abstract away the manipulation of the list and to avoid the need to learn an explicitly recursive program (recursion is implicit in $\text{map}/3$). Metagol_{ho} has been shown to reduce sample complexity and learning times and improve predictive accuracies (Cropper et al., 2020c).

5.4 Probabilistic logic programs

A major limitation of logical representations, such as Prolog and its derivatives, is the implicit assumption that the BK is perfect. That is, most ILP systems assume that atoms are true or false, leaving no room for uncertainty. This assumption is problematic if data is noisy, which is often the case.

ILP systems have limited capabilities for dealing with noise. If a perfect program is not in the hypothesis space, the most common strategy is to find a program that covers as many positive and as few negative examples as possible. Though this approach helps us handle mislabelled examples, is not a good way for dealing with observational noise generally. The limitations of the approach become obvious when examples are structured (e.g., complex relations, images, or lists) rather than simple labels. Consider the example in Fig. 1, in which we want to learn a program drawing simple images spelling “ILP”. Each of the examples displays a clear concept but with some alterations: every image contains

three ‘noisy’ pixels: the second example has the first letter shifted to the left, while the arc in the letter “P” is elongated in the third example. All of these alterations are noise and not something we want our program to explicitly represent—the ground truth program is the one that draws “ILP” in the middle of the figure, without any noise. However, ILP systems based on entailment would consider a solution to be correct only if it models all of the noisy aspects. An ILP system capable of handling such noise is Brute (Cropper & Dumančić, 2020b), which uses a distance to the target solution (e.g., pixel distance in case of images) as an optimisation criterion instead of entailment.

The most principle way of handling noise is to integrating probabilistic aspects into logical representations so that uncertainties in data can be directly modelled. This integration is the focus of statistical relational artificial intelligence (StarAI) (De Raedt & Kersting, 2008; De Raedt et al., 2016). In essence, StarAI hypothesis representations extend BK with probabilities or weights indicating the degree of confidence in the correctness of parts of BK. StarAI is a big and prolific field; for that reason, we will not cover it in entirety but rather briefly introduce the main ideas that overcome limitations of logic programming-based ILP systems.

Generally, StarAI techniques are based on two ideas: *distribution semantics* and *maximum entropy*. Distribution semantics approaches (Sato, 1995), including Problog (De Raedt et al., 2007) and PRISM (Sato & Kameya, 2001), explicitly annotate uncertainties in BK. To allow such annotation, they extend Prolog with two primitives for stochastic execution: probabilistic facts and annotated disjunctions. Probabilistic facts are the most basic stochastic primitive and they take the form of logical facts labelled with a probability p . Each probabilistic fact represents a Boolean random variable that is `true` with probability p and `false` with probability $1 - p$. For instance, the following probabilistic fact states that there is 1% chance of an earthquake in Naples.

```
0.01::earthquake(naples).
```

An alternative interpretation of this statement is that 1% of executions of the probabilistic program would observe an earthquake. The second type of stochastic primitive is an annotated disjunction. Whereas probabilistic facts introduce non-deterministic behaviour on the level of facts, annotated disjunctions introduce non-determinism on the level of clauses. Annotated disjunctions allow for multiple literals in the head, where only one of the head literals can be `true` at a time. For instance, the following annotated disjunction states that a ball can be either green, red, or blue, but not a combination of colours:

```
 $\frac{1}{3}$ ::colour(B,green);  $\frac{1}{3}$ ::colour(B,red);  $\frac{1}{3}$ ::colour(B,blue) :- ball(B).
```

By contrast, maximum entropy approaches annotate uncertainties only at the level of a logical theory. That is, they assume that the predicates in the BK are labelled as either true or false, but the label may be incorrect. These approaches are not based on logic programming, but rather on first-order logic. Consequently, the underlying semantics are different: rather than consider proofs, these approaches consider models or groundings of a theory. This difference primarily changes what uncertainties represent. For instance, Markov Logic Networks (MLN) (Richardson & Domingos, 2006) represent programs as a set of weighted clauses. The weights in MLN do not correspond to probabilities of a formula being true

but, intuitively, to a log odds between a possible world (an interpretation) where the clause is true and a world where the clause is false. For instance, a clause that is true in 80% of the worlds would have a weight of 1.386 ($\log \frac{0.8}{0.2}$).

The techniques from learning such probabilistic programs are typically direct extensions of ILP techniques. For instance, ProbFOIL (De Raedt et al., 2015) extends FOIL (Quinlan, 1990) with probabilistic clauses. Similarly, SLIPCOVER (Bellodi & Riguzzi, 2015) is a bottom-up approach, similar to Aleph (Srinivasan, 2001) and Progol (Muggleton, 1995). Huynh and Mooney (2008) use Aleph to find interesting clauses and then learn the corresponding weights. Kok and Domingos (2009) use relational pathfinding over BK to identify useful clauses. That is, they interpret the BK as a hypergraph in which constants form vertices and atoms form hyper-edges and perform random walks. Frequently occurring walks, or their subparts, are then turned into clauses. Such random walks could be seen as an approximate way to construct bottom clauses.

It is worth noting that StarAI also considers an alternative learning problem—that of learning the probabilistic parameters of a given program. We do not survey these approaches here as the problem is different in nature from the ILP problem: whereas ILP searches from a program solving the tasks, parameter learning methods assume that such program is given.

6 Optimality

There are often multiple (sometimes infinitely many) hypotheses that explain the data. Deciding which hypothesis to choose has long been a difficult problem. Many systems aim for maximum classification accuracy. For instance, Aleph, by default, aims to maximum *coverage* of each clause it adds to a hypothesis, where coverage is measured as $P - N$, where P and N are the number of positive and negative examples covered by the clause respectively. Note that Aleph supports various evaluation metrics, such as *compression*, measured as $P - N - L + 1$, where P and N are as before and L is the number of literals in the clause. However, older ILP systems are typically not guaranteed to induce optimal programs/theories. A key reason for this limitation was that most search techniques learned a single clause at a time, leading to the construction of sub-programs that are sub-optimal in terms of program size and coverage. For instance, Aleph offers no guarantee of optimality with respect to the program size and coverage.

Newer ILP systems try to address this limitation. As with the ability to learn recursive programs, the main development is to take a global view of the induction task by using meta-level search techniques. In other words, rather than induce a single clause at a time from a single example, the idea is to induce multiple clauses from multiple examples. For instance, ILASP uses ASP's optimisation abilities to provably learn the program with the fewest literals. ILASP3 (Law, 2018) adopts a similar approach to account for noise.

The ability to learn optimal programs opens up ILP to new problems. For instance, learning efficient logic programs has long been considered a difficult problem in ILP (Muggleton & De Raedt, 1994; Muggleton et al., 2012), mainly because there is no declarative difference between an efficient program, such as mergesort, and an inefficient program, such as bubble sort. To address this issue, Metaopt (Cropper & Muggleton, 2019) extends Metagol to support learning efficient programs. Metaopt maintains a cost during the hypothesis search and uses this cost to prune the hypothesis space. To learn minimal time complexity logic programs, Metaopt minimises the number of resolution steps. For

instance, imagine trying to learn a *find duplicate* program, which finds any duplicate element in a list e.g. $[p,r,o,g,r,a,m] \mapsto r$, and $[i,n,d,u,c,t,i,o,n] \mapsto i$. Given suitable input data, Metagol can induce the program:

```
f(A,B):- head(A,B),tail(A,C),element(C,B).
f(A,B):- tail(A,C),f(C,B).
```

This program goes through the elements of the list checking whether the same element exists in the rest of the list. Given the same input, Metaopt induces the program:

```
f(A,B):- mergesort(A,C),f1(C,B).
f1(A,B):- head(A,B),tail(A,C),head(C,B).
f1(A,B):- tail(A,C),f1(C,B).
```

This program first sorts the input list and then goes through the list to check whether for duplicate adjacent elements. Although larger, both in terms of clauses and literals, the program learned by Metaopt is more efficient $O(n \log n)$ than the program learned by Metagol $O(n^2)$. Metaopt has been shown to learn efficient robot strategies, efficient time complexity logic programs, and even efficient string transformation programs.

FastLAS (Law et al., 2020a) is an ASP-based ILP system that takes as input a custom scoring function and computes an optimal solution with respect to the given scoring function when learning non-recursive programs without PI. The authors show that this approach allows a user to optimise domain-specific performance metrics on real-world datasets, such as access control policies.

7 Technologies

Older ILP systems mostly use Prolog for reasoning. Recent work considers using different technologies.

7.1 Constraint satisfaction and satisfiability

There have been tremendous recent advances in SAT (Heule et al., 2016). To leverage these advances, much recent work in ILP uses related techniques, notably ASP (Corapi et al., 2011; Muggleton et al., 2014; Law et al., 2014; Katzouris et al., 2015, 2016; Schüller & Benz, 2018; Kaminski et al., 2018; Evans et al., 2021; Cropper & Morel, 2021a). The main motivations for using ASP are to leverage (i) the language benefits of ASP (Sect. 5.2), and (ii) the efficiency and optimisation techniques of modern ASP solvers, such as CLASP (Gebser et al., 2012b), which supports conflict propagation and learning. With similar motivations, other approaches encode the ILP problem as SAT (Ahlgren & Yuen, 2013) or SMT (Albarghouthi et al., 2017) problems. These approaches have been shown able to reduce learning times compared to standard Prolog-based approaches. However, some unresolved issues remain. A key issue is that most approaches encode an ILP problem as a single (often very large) satisfiability

problem. These approaches therefore often struggle to scale to very large problems (Cropper et al., 2020c), although preliminary work attempts to tackle this issue (Cropper & Morel, 2021a).

7.2 Neural networks

With the rise of deep learning, several approaches have explored using gradient-based methods to learn logic programs. These approaches all replace discrete logical reasoning with a relaxed version that yields continuous values reflecting the confidence of the conclusion.

The various neural approaches can be characterised along four orthogonal dimensions. The first dimension is whether the neural network implements forward or backward inference. While some (Rocktäschel & Riedel, 2017) use backward (goal-directed) chaining with a neural implementation of unification, most approaches (Evans & Grefenstette, 2018; Yang et al., 2017; Dong et al., 2019a) use forward chaining. The second dimension is whether the network is designed for big data problems (Yang et al., 2017; Rocktäschel & Riedel, 2017) or for data-efficient learning from a handful of data items (Evans & Grefenstette, 2018). Few neural systems to date are capable of handling both big data and small data, with the notable exception of Dong et al. (2019a). The third dimension is whether the neural system jointly learns embeddings (mapping symbolic constants to continuous vectors) along with the logical rules (Rocktäschel & Riedel, 2017). The advantage of jointly learning embeddings is that it enables fuzzy unification between constants that are similar but not identical. The challenge for these approaches that jointly learn embeddings is how to generalize appropriately to constants that have not been seen at training time. The fourth dimension is whether or not the neural system is designed to allow explicit human-readable logical rules to be extracted from the weights of the network. While most neural ILP systems (Yang et al., 2017; Rocktäschel & Riedel, 2017; Evans & Grefenstette, 2018) do produce explicit logic programs, some Dong et al. (2019a) do not. It is perhaps moot whether implicit systems that do not produce explicit programs count as ILP systems at all—but note that even in the implicit neural systems, the weight sharing of the neural net is designed to achieve strong generalisation by performing the same computation on all tuples of objects.

Currently, most neural approaches to ILP require the use of metarules or templates to make the search space tractable and fail to support predicate invention, recursion and abduction. This severely limits the applicability of these approaches, as the user cannot always be expected to provide suitable and complete background knowledge and metarules for a new problem. The only approach that avoids the use of metarules or templates is Neural Logic Machines (Dong et al., 2019a), and the only one to fully integrate neural net learning with predicate invention, recursion, and abduction is Abductive Meta-Interpretive Learning (Dai & Muggleton, 2021).

8 Applications

We now survey recent application areas for ILP.

Scientific discovery Perhaps the most prominent application of ILP is in scientific discovery. ILP has, for instance, been used to identify and predict ligands (substructures responsible for medical activity) (Kaalia et al., 2016) and infer missing pathways in protein signalling networks (Inoue et al., 2013). There has been much recent work on applying ILP in ecology (Bohan et al., 2011, 2017; Tamaddoni-Nezhad et al., 2014). For instance, Bohan et al. (2011) use ILP to generate plausible and testable hypotheses for trophic relations ('who eats whom') from ecological data.

Program analysis Due to the expressivity of logic programs as a representation language, ILP systems have found successful applications in software design. ILP systems have proven effective in learning SQL queries (Albarghouthi et al., 2017; Sivaraman et al., 2019), programming language semantics (Bartha & Cheney, 2019), and code search (Sivaraman et al., 2019).

Robotics Robotics applications often require incorporating domain knowledge or imposing certain requirements on the learnt programs. For instance, The Robot Engineer (Sammut et al., 2015) uses ILP to design tools for robots and even complete robots, which are tests in simulations and real-world environments. *Metagol_o* (Cropper & Muggleton, 2015) learns robot strategies considering their resource efficiency and Antanas et al. (2015) recognise graspable points on objects through relational representations of objects.

Vision Background knowledge is also valuable in Computer Vision. Recent work Muggleton et al. (2018a) demonstrated that *Logical Vision*, which employs MIL, can outperform state-of-the-art statistical machine learning in particular image recognition tasks, given general Newtonian physics background knowledge concerning reflection of light.

Games Inducing game rules has a long history in ILP, where chess has often been the focus (Muggleton et al., 2009). Legras et al., (2018) show that Aleph and TILDE can outperform an SVM learner in the game of Bridge. Law et al. (2014) use ILASP to induce the rules for Sudoku and show that this more expressive formalism allows for game rules to be expressed more compactly. Cropper et al. (2020b) introduce the ILP problem of *inductive general game playing*: the problem of inducing game rules from observations, such as *Checkers*, *Sokoban*, and *Connect Four*. Muggleton and Hocquette (2019) show the MIL system MIGO consistently outperforms deep reinforcement learning for both Noughts-and-Crosses and Hexapawn.

Data curation and transformation Another successful application of ILP is in data curation and transformation, which is again largely because ILP can learn executable programs. The most prominent example of such tasks is string transformations, such as the example given in the introduction. There is much interest in this topic, largely due to success in synthesising programs for end-user problems, such as string transformations in Microsoft Excel (Gulwani, 2011). String transformations have become a standard benchmark for recent ILP papers (Lin et al., 2014; Cropper et al., 2020c; Cropper, 2019; Cropper & Dumančić, 2020b). Other transformation tasks include extracting values from semi-structured data (e.g. XML files or medical records), extracting relations from ecological papers, and spreadsheet manipulation (Cropper et al., 2015).

Learning from trajectories Learning from interpretation transitions (LFIT) (Inoue et al., 2014) automatically constructs a model of the dynamics of a system from the observation of its state transitions. Given time-series data of discrete gene expression, it can learn gene interactions, thus allowing to explain and predict states changes over time (Ribeiro et al., 2020). LFIT has been applied to learn biological models, like Boolean Networks, under several semantics: memory-less deterministic systems (Inoue et al., 2014; Ribeiro & Inoue, 2014), and their multi-valued extensions (Ribeiro et al., 2015; Martínez et al., 2016). Martínez et al. (2016) combine LFIT with a reinforcement learning algorithm to learn probabilistic models with exogenous effects (effects not related to any action) from scratch. The learner was notably integrated with a robot to perform the task of clearing the tableware on a table. In this task external agents interacted, people brought new tableware continuously and the manipulator robot had to cooperate with mobile robots to take the tableware to the kitchen. The learner was able to learn a usable model in just five episodes of 30 action executions. Evans et al. (2021) apply the *Apperception Engine* to explain sequential data, such as cellular automata traces, rhythms and simple nursery tunes, image occlusion tasks, game dynamics, and sequence induction intelligence tests. Surprisingly, they show that their system can achieve human-level performance on the sequence induction intelligence tests in the zero-shot setting (without having been trained on lots of other examples of such tests, and without hand-engineered knowledge of the particular setting). At a high level, these systems take the unique selling point of ILP systems (the ability to strongly generalise from a handful of data), and apply it to the self-supervised setting, producing an explicit human-readable theory that explains the observed state transitions.

9 Summary and future work

In a survey paper from a decade ago, Muggleton et al. (2012) proposed directions for future research. In the decade since, there have been major advances on many of the topics, notably in predicate invention (Sect. 4), using higher-order logic as a representation language (Sect. 4.2) and to represent hypotheses (Sect. 5.3), and applications in learning actions and strategies (Sect. 8). Despite the advances, there are still many limitations in ILP that future work should address.

9.1 Limitations and future research

Better systems Muggleton et al. (2012) argue that a problem with ILP is the lack of well-engineered tools. They state that whilst over 100 ILP systems have been built, less than a handful of systems can be meaningfully used by ILP researchers. In the decade since the authors highlighted this problem, little progress has been made: most ILP systems are not easy to use. In other words, ILP systems are still notoriously difficult to use and you often need a PhD in ILP to use any of the tools. Even then, it is still often only the developers of a system that know how to properly use it. By contrast, driven by industry, other forms of ML now have reliable and well-maintained implementations, such as PyTorch and TensorFlow, which has helped drive research. A frustrating issue with ILP systems is that they use many different language biases or even different syntax for the same biases. For instance, the way of specifying a learning task in Progol, Aleph, TILDE, and ILASP varies considerably despite them all using mode declarations. If it is difficult for ILP researchers

to use ILP tools, then what hope do non-ILP researchers have? For ILP to be more widely adopted both inside and outside of academia, we must develop more standardised, user-friendly, and better-engineered tools.

Language biases As Cropper et al. (2020a) state, one major issue with ILP is choosing an appropriate language bias. For instance, Metagol uses metarules (Sect. 4.2) to restrict the syntax of hypotheses and thus the hypothesis space. If a user can provide suitable metarules, then Metagol is extremely efficient. However, if a user cannot provide suitable metarules (which is often the case), then Metagol is almost useless. This same brittleness applies to ILP systems that employ mode declarations (Muggleton, 1995). In theory, a user can provide very general mode declarations, such as only using a single type and allowing unlimited recall. In practice, however, weak mode declarations often lead to very poor performance. For good performance, users of mode-based systems often need to manually analyse a given learning task to tweak the mode declarations, often through a process of trial and error. Moreover, if a user makes a small mistake with a mode declaration, such as giving the wrong argument type, then the ILP system is unlikely to find a good solution. Even for ILP experts, determining a suitable language bias is often a frustrating and time-consuming process. We think the need for an almost perfect language bias is severely holding back ILP from being widely adopted. By contrast, there are some neural net architectures (e.g. the transformer; Vaswani et al., 2017) that can be applied successfully to a large range of diverse problems without requiring any domain-specific tuning. We think that an important direction for future work in ILP is to develop techniques for automatically identifying suitable language biases. Although there is some work on mode learning (McCreath & Sharma, 1995; Ferilli et al., 2004; Picado et al., 2017, 2021) and work on identifying suitable metarules (Cropper & Tourret, 2020), this area of research is largely under-researched.

Better datasets Interesting problems, alongside usable systems, drive research and attract interest in a research field. This relationship is most evident in the deep learning community which has, over a decade, grown into the largest AI community. This community growth has been supported by the constant introduction of new problems, datasets, and well-engineered tools. Challenging problems that push the state-of-the-art to its limits are essential to sustain progress in the field; otherwise, the field risks stagnation through only small incremental progress. ILP has, unfortunately, failed to deliver on this front: most research is still evaluated on 20-year old datasets. Most new datasets that have been introduced often come from toy domains and are designed to test specific properties of the introduced technique. To an outsider, this sends a message that ILP is not applicable to real-world problems. We think that the ILP community should learn from the experiences of other AI communities and put significant efforts into developing datasets that identify limitations of existing methods as well as showcase potential applications of ILP. After all, it is no coincidence that SAT solving performance increased dramatically after the introduction of the SAT solving competitions (Järvisalo et al., 2012).

Relevance New methods for predicate invention (Sect. 4) have improved the abilities of ILP systems to invent high-level concepts. These techniques raise the potential for ILP to be used in lifelong learning settings. However, inventing and acquiring new BK could lead to a problem of too much BK, which can overwhelm an ILP system (Srinivasan et al., 2003; Cropper, 2020). On this issue, a key under-explored topic is that of *relevancy*. Given

a new induction problem with large amounts of BK, how does an ILP system decide which BK is relevant? One emerging technique is to train a neural network to score how relevant programs are in the BK and to then only use BK with the highest score to learn programs (Balog et al., 2017; Ellis et al., 2018). However, the empirical efficacy of this approach has yet to be demonstrated. Moreover, these approaches have only been demonstrated on small amounts of BK and it is unclear how they scale to BK with thousands of relations. Without efficient methods of relevance identification, it is unclear how efficient lifelong learning can be achieved.

Handling mislabelled and ambiguous data A major open question in ILP is how best to handle noisy and ambiguous data. Neural ILP systems (Rocktäschel & Riedel, 2017; Evans & Grefenstette, 2018) are designed from the start to robustly handle mislabelled data. Although there has been work in recent years on designing ILP systems that can handle noisy mislabelled data, there is much less work on the even harder and more fundamental problem of designing ILP systems that can handle *raw ambiguous data*. ILP systems typically assume that the input has already been preprocessed into symbolic declarative form (typically, a set of ground atoms representing positive and negative examples). But real-world input does not arrive in symbolic form. Consider e.g. a robot with a video camera, where the raw input is a sequence of pixel images. Converting each pixel image into a set of ground atoms is a challenging non-trivial achievement that should not be taken for granted. For ILP systems to be widely applicable in the real world, they need to be redesigned so they can handle raw ambiguous input from the outset (Evans & Grefenstette, 2018; Dong et al., 2019b).

Probabilistic ILP Real-world data is often noisy and uncertain. Extending ILP to deal with such uncertainty substantially broadens its applicability. While StarAI is receiving growing attention, learning probabilistic programs from data is still largely under-investigated due to the complexity of joint probabilistic and logical inference. When working with probabilistic programs, we are interested in the probability that a program covers an example, not only whether the program covers the example. Consequently, probabilistic programs need to compute all possible derivations of an example, not just a single one. Despite added complexity, probabilistic ILP opens many new challenges. Most of the existing work on probabilistic ILP considers the minimal extension of ILP to the probabilistic setting, by assuming that either (i) BK facts are uncertain, or (ii) that learned clauses need to model uncertainty. These assumptions make it possible to separate structure from uncertainty and simply reuse existing ILP techniques. Following this minimal extension, the existing work focuses on discriminative learning in which the goal is to learn a program for a single target relation. However, a grand challenge in probabilistic programming is generative learning. That is, learning a program describing a generative process behind the data, not a single target relation. Learning generative programs is a significantly more challenging problem, which has received very little attention in probabilistic ILP.

Explainability Explainability is one of the claimed advantages of a symbolic representation. Recent work Muggleton et al. (2018b) and Ai et al. (2020) evaluates the comprehensibility of ILP hypotheses using Michie's Michie (1988) framework of *ultra-strong machine learning*, where a learned hypothesis is expected to not only be accurate but to also demonstrably improve the performance of a human being provided with the learned hypothesis. Muggleton et al. (2018b) empirically demonstrate improved human understanding directly

through learned hypotheses. However, given the demonstration of both beneficial and harmful effects of explainability (Ai et al., 2020) more work is required to better understand the conditions under which this can be achieved, especially given the rise of PI.

Unifying ILP with neural methods It has often been noted (Evans & Grefenstette, 2018) that the strengths and weaknesses of neural networks and ILP are *complementary*: neural networks (1) scale to huge datasets, (2) are robust to mislabelled data, (3) are robust to ambiguous (raw, undiscretised) data, but (4) are very data hungry, (5) often struggle to generalise outside the training distribution, and (6) are uninterpretable. ILP systems, by contrast (1) often fail to scale to large datasets, (2) sometimes fail to handle mislabelled data, (3) almost always fail to handle raw undiscretised data, but (4) are very data efficient, (5) often generalise well outside the training distribution, and (6) produce human-readable programs. Given that the strengths and weaknesses of the two approaches are complementary, many people have advocated some sort of unification of the two (De Raedt et al., 2016; Garcez & Lamb, 2020; Evans & Grefenstette, 2018). There is much activity in this area, and much work still to do to produce a truly convincing unification of these very different paradigms.

9.2 Summary

As ILP approaches 30, we think that the advances made in the last decade, surveyed in this paper, have opened up new areas of research for ILP to explore. Moreover, we hope that the next decade sees developments on the numerous limitations we have discussed so that ILP can have a significant impact on AI.

Author contributions AC wrote Sects. 1–7, and 9. SD wrote Sects. 4.3, 5.4, and 8. RE wrote Sects. 7.2, 8, and 9. SM wrote Sect. 9.

Funding Not applicable.

Availability of data and materials Not applicable.

Code availability Not applicable.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not

permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ahlgren, J., & Yuen, S. Y. (2013). Efficient program synthesis using constraint satisfaction in inductive logic programming. *The Journal of Machine Learning Research*, 14(1), 3649–3682.
- Ai, L., Muggleton, S., Hocquette, C., Gromowski, M., & Schmid, U. (2020). Beneficial and harmful explanatory machine learning. *Machine Learning*, 110, 695–721.
- Albarghouthi, A., Koutris, P., Naik, M., & Smith, C. (2017). Constraint-based synthesis of datalog programs. In *23rd international conference on principles and practice of constraint programming, CP 2017. Lecture notes in computer science* (Vol. 10416, pp. 689–706). Springer.
- Antanas, L., Moreno, P., & De Raedt, L. (2015). Relational kernel-based grasping with numerical features. In *25th international conference on inductive logic programming, ILP 2015. Lecture notes in computer science* (Vol. 9575, pp. 1–14). Springer.
- Bain, M., & Muggleton, S. (1992). Non-monotonic learning. In *Inductive logic programming* (pp. 145–161). Academic Press.
- Bain, M., & Srinivasan, A. (2018). Identification of biological transition systems using meta-interpreted logic programs. *Machine Learning*, 107(7), 1171–1206.
- Balog, M., Gaunt, A. L., Brockschmidt, M., Nowozin, S., & Tarlow, D. (2017). Deepcoder: Learning to write programs. In *5th international conference on learning representations, ICLR 2017*. OpenReview.net.
- Bartha, S., & Cheney, J. (2019). Towards meta-interpretive learning of programming language semantics. In *29th international conference on inductive logic programming, ILP 2019. Lecture notes in computer science* (Vol. 11770, pp. 16–25). Springer.
- Bellodi, E., & Riguzzi, F. (2015). Structure learning of probabilistic logic programs by searching the clause space. *Theory and Practice of Logic Programming*, 15(2), 169–212.
- Blockeel, H., & De Raedt, L. (1998). Top-down induction of first-order logical decision trees. *Artificial Intelligence*, 101(1–2), 285–297.
- Bohan, D. A., Caron-Lormier, G., Muggleton, S., Raybould, A., & Tamaddoni-Nezhad, A. (2011). Automated discovery of food webs from ecological data using logic-based machine learning. *PLoS ONE*, 6(12), e29028.
- Bohan, D. A., Vacher, C., Tamaddoni-Nezhad, A., Raybould, A., Dumbrell, A. J., & Woodward, G. (2017). Next-generation global biomonitoring: Large-scale, automated reconstruction of ecological networks. *Trends in Ecology & Evolution*, 32(7), 477–487.
- Bratko, I. (1999). Refining complete hypotheses in ILP. In *9th international workshop on inductive logic programming, ILP-99. Lecture notes in computer science* (Vol. 1634, pp. 44–55). Springer.
- Chollet, F. (2019). *On the measure of intelligence*. CoRR, [arXiv:1911.01547](https://arxiv.org/abs/1911.01547)
- Corapi, D., Russo, A., & Lupu, E. (2011). Inductive logic programming in answer set programming. In *21st international conference on inductive logic programming, ILP 2011. Lecture notes in computer science* (Vol. 7207, pp. 91–97). Springer.
- Cropper, A. (2019). Playgol: Learning programs through play. In *Proceedings of the twenty-eighth international joint conference on artificial intelligence, IJCAI 2019* (pp. 6074–6080). ijcai.org.
- Cropper, A. (2020). Forgetting to learn logic programs. In *The thirty-fourth AAAI conference on artificial intelligence* (pp. 3676–3683). AAAI Press.
- Cropper, A., & Dumančić, S. (2020a). *Inductive logic programming at 30: A new introduction*. CoRR, arxiv.org/abs/2008.07912
- Cropper, A., & Dumančić, S. (2020b). Learning large logic programs by going beyond entailment. In *Proceedings of the twenty-ninth international joint conference on artificial intelligence, IJCAI 2020* (pp. 2073–2079). ijcai.org.
- Cropper, A., & Morel, R. (2021a). Learning programs by learning from failures. *Machine Learning*, 110(4), 801–856. <https://doi.org/10.1007/s10994-020-05934-z>.
- Cropper, A., & Morel, R. (2021b). *Predicate invention by learning from failures*. CoRR, arxiv.org/abs/2104.14426
- Cropper, A., & Muggleton, S. H. (2015). Learning efficient logical robot strategies involving composable objects. In *Proceedings of the twenty-fourth international joint conference on artificial intelligence, IJCAI 2015* (pp. 3423–3429). AAAI Press.
- Cropper, A., & Muggleton, S. H. (2016). *Metagol system*. <https://github.com/metagol/metagol>

- Cropper, A., & Muggleton, S. H. (2019). Learning efficient logic programs. *Machine Learning*, 108(7), 1063–1083.
- Cropper, A., & Tourret, S. (2020). Logical reduction of metarules. *Machine Learning*, 109(7), 1323–1369.
- Cropper, A., Tamaddoni-Nezhad, A., & Muggleton, S. H. (2015). Meta-interpretive learning of data transformation programs. In *25th international conference on inductive logic programming, ILP 2015. Lecture notes in computer science* (Vol. 9575, pp 46–59). Springer.
- Cropper, A., Dumančić, S., & Muggleton, S. H. (2020a). Turning 30: New ideas in inductive logic programming. In *Proceedings of the twenty-ninth international joint conference on artificial intelligence, IJCAI 2020* (pp. 4833–4839). ijcai.org.
- Cropper, A., Evans, R., & Law, M. (2020b). Inductive general game playing. *Machine Learning*, 109(7), 1393–1434.
- Cropper, A., Morel, R., & Muggleton, S. (2020c). Learning higher-order logic programs. *Machine Learning*, 109(7), 1289–1322.
- Dai, W. Z., & Muggleton, S. H. (2021). Abductive knowledge induction from raw data. In *Proceedings of the 35th conference on artificial intelligence (IJCAI 2021), IJCAI (in Press)*.
- De Raedt, L. (2008). *Logical and relational learning. Cognitive technologies*. Springer.
- De Raedt, L., Dries, A., Thon, I., den Broeck, G. V., & Verbeke, M. (2015). Inducing probabilistic relational rules from probabilistic examples. In *Proceedings of the twenty-fourth international joint conference on artificial intelligence, IJCAI 2015* (pp. 1835–1843). AAAI Press.
- De Raedt, L., & Kersting, K. (2008). *Probabilistic inductive logic programming* (pp. 1–27). Springer.
- De Raedt, L., Kersting, K., Natarajan, S., & Poole, D. (2016). *Statistical relational artificial intelligence: Logic, probability, and computation. Synthesis lectures on artificial intelligence and machine learning*. Morgan & Claypool Publishers.
- De Raedt, L., Kimmig, A., & Toivonen, H. (2007). Problog: A probabilistic prolog and its application in link discovery. In *IJCAI 2007, Proceedings of the 20th international joint conference on artificial intelligence, Hyderabad, India, January 6–12, 2007*, pp. 2462–2467.
- Dong, H., Mao, J., Lin, T., Wang, C., Li, L., & Zhou, D. (2019a). Neural logic machines. In *ICLR*.
- Dong, H., Mao, J., Lin, T., Wang, C., Li, L., & Zhou, D. (2019b). Neural logic machines. In *7th international conference on learning representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*, OpenReview.net. <https://openreview.net/forum?id=BlxY-hRctX>
- Dumančić, S., & Blockeel, H. (2017). Clustering-based relational unsupervised representation learning with an explicit distributed representation. In *Proceedings of the twenty-sixth international joint conference on artificial intelligence, IJCAI 2017* (pp. 1631–1637). ijcai.org.
- Dumancic, S., Guns, T., & Cropper, A. (2020). Knowledge refactoring for inductive program synthesis. In *AAAI*.
- Dumančić, S., Guns, T., Meert, W., & Blockeel, H. (2019). Learning relational representations with auto-encoding logic programs. In *Proceedings of the twenty-eighth international joint conference on artificial intelligence, IJCAI 2019* (pp. 6081–6087). ijcai.org.
- Ellis, K., Morales, L., Sablé-Meyer, M., Solar-Lezama, A., & Tenenbaum, J. (2018). Learning libraries of subroutines for neurally-guided Bayesian program induction. *NeurIPS*, 2018, 7816–7826.
- Evans, R., & Grefenstette, E. (2018). Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61, 1–64.
- Evans, R., Hernández-Orallo, J., Welbl, J., Kohli, P., & Sergot, M. (2021). Making sense of sensory input. *Artificial Intelligence*, 293, 103438.
- Ferilli, S. (2016). Predicate invention-based specialization in inductive logic programming. *Journal of Intelligent Information Systems*. <https://doi.org/10.1007/s10844-016-0412-9>.
- Ferilli, S., Esposito, F., Basile, T. M. A., & Mauro, N. D. (2004). Automatic induction of first-order logic descriptors type domains from observations. In *14th international conference on inductive logic programming, ILP 2004. Lecture notes in computer science* (Vol. 3194, pp. 116–131). Springer.
- Garcez, Ad., & Lamb, L. C. (2020). *Neurosymbolic ai: The 3rd wave*. arXiv preprint [arXiv:2012.05876](https://arxiv.org/abs/2012.05876)
- Gebser, M., Kaminski, R., Kaufmann, B., & Schaub, T. (2012a). *Answer set solving in practice. Synthesis lectures on artificial intelligence and machine learning*. Morgan & Claypool Publishers.
- Gebser, M., Kaufmann, B., & Schaub, T. (2012b). Conflict-driven answer set solving: From theory to practice. *Artificial Intelligence*, 187, 52–89.
- Genesereth, M. R., & Björnsson, Y. (2013). The international general game playing competition. *AI Magazine*, 34(2), 107–111.
- Gulwani, S. (2011). Automating string processing in spreadsheets using input-output examples. In *Proceedings of the 38th ACM SIGPLAN-SIGACT symposium on principles of programming languages, POPL 2011* (pp. 317–330). ACM.

- Heule, M. J. H., Kullmann, O., & Marek, V. W. (2016). Solving and verifying the boolean pythagorean triples problem via cube-and-conquer. In N. Creignou & D. L. Berre (Eds.), *19th international conference on theory and applications of satisfiability testing—SAT 2016, Bordeaux, France, July 5–8, 2016, Proceedings. Lecture notes in computer science* (Vol. 9710, pp. 228–245). Springer. https://doi.org/10.1007/978-3-319-40970-2_15
- Hocquette, C., & Muggleton, S. H. (2020). Complete bottom-up predicate invention in meta-interpretive learning. In *Proceedings of the twenty-ninth international joint conference on artificial intelligence, IJCAI 2020* (pp. 2312–2318). ijcai.org.
- Huynh, T. N., & Mooney, R. J. (2008). Discriminative structure and parameter learning for Markov logic networks. In *Proceedings of the 25th international conference on machine learning* (pp. 416–423). Association for Computing Machinery. <https://doi.org/10.1145/1390156.1390209>
- Inoue, K. (2016). Meta-level abduction. *FLAP*, 3(1), 7–36.
- Inoue, K., Doncescu, A., & Nabeshima, H. (2013). Completing causal networks by meta-level abduction. *Machine Learning*, 91(2), 239–277.
- Inoue, K., Ribeiro, T., & Sakama, C. (2014). Learning from interpretation transition. *Machine Learning*, 94(1), 51–79.
- Järvisalo, M., Le Berre, D., Roussel, O., & Simon, L. (2012). The international sat solver competitions. *Ai Magazine*, 33(1), 89–92.
- Kaalia, R., Srinivasan, A., Kumar, A., & Ghosh, I. (2016). ILP-assisted de novo drug design. *Machine Learning*, 103(3), 309–341.
- Kaiser, L., & Sutskever, I. (2016). Neural gpu learn algorithms. In *4th international conference on learning representations, ICLR 2016*.
- Kaminski, T., Eiter, T., & Inoue, K. (2018). Exploiting answer set programming with external sources for meta-interpretive learning. *Theory and Practice of Logic Programming*, 18(3–4), 571–588.
- Katzouris, N., Artikis, A., & Paliouras, G. (2015). Incremental learning of event definitions with inductive logic programming. *Machine Learning*, 100(2–3), 555–585.
- Katzouris, N., Artikis, A., & Paliouras, G. (2016). Online learning of event definitions. *Theory and Practice of Logic Programming*, 16(5–6), 817–833.
- Kok, S., & Domingos, P. (2009). Learning Markov logic network structure via hypergraph lifting. In *Proceedings of the 26th international conference on machine learning* (pp. 505–512). Association for Computing Machinery. <https://doi.org/10.1145/1553374.1553440>
- Kok, S., & Domingos, P. M. (2007). Statistical predicate invention. In *Machine Learning, Proceedings of the twenty-fourth international conference (ICML 2007), ACM international conference proceeding series* (Vol. 227, pp. 433–440). ACM.
- Kramer, S. (1995). *Predicate invention: A comprehensive view*. Rapport technique OFAI-TR-95-32, Austrian Research Institute for Artificial Intelligence.
- Kramer, S. (2020). A brief history of learning symbolic higher-level representations from data (and a curious look forward). In *Proceedings of the twenty-ninth international joint conference on artificial intelligence, IJCAI 2020* (pp. 4868–4876). ijcai.org.
- Law, M. (2018). *Inductive learning of answer set programs*. PhD thesis, Imperial College London, UK.
- Law, M., Russo, A., & Broda, K. (2014). Inductive learning of answer set programs. In *14th European conference on logics in artificial intelligence, JELIA 2014. Lecture notes in computer science* (Vol. 8761, pp. 311–325). Springer.
- Law, M., Russo, A., & Broda, K. (2018). The complexity and generality of learning answer set programs. *Artificial Intelligence*, 259, 110–146.
- Law, M., Russo, A., Bertino, E., Broda, K., & Lobo, J. (2019). Representing and learning grammars in answer set programming. In *The thirty-third AAAI conference on artificial intelligence, AAAI 2019* (pp. 2919–2928). AAAI Press.
- Law, M., Russo, A., Bertino, E., Broda, K., & Lobo, J. (2020a). Fastlas: Scalable inductive logic programming incorporating domain-specific optimisation criteria. In *The thirty-fourth AAAI conference on artificial intelligence, AAAI 2020* (pp. 2877–2885). AAAI Press.
- Law, M., Russo, A., & Broda, K. (2020b). *The ilasp system for inductive learning of answer set programs*. The Association for Logic Programming Newsletter.
- Leban, G., Zabkar, J., & Bratko, I. (2008). An experiment in robot discovery with ILP. In *18th international conference inductive logic programming, ILP 2008. Lecture notes in computer science* (Vol. 5194, pp. 77–90). Springer.
- Legras, S., Rouveïrol, C., & Ventos, V. (2018). The game of bridge: A challenge for ILP. In *28th international conference inductive logic programming, ILP 2018. Lecture notes in computer science* (Vol. 11105, pp. 72–87). Springer.

- Lin, D., Dechter, E., Ellis, K., Tenenbaum, J. B., & Muggleton, S. (2014). Bias reformulation for one-shot function induction. In *ECAI 2014—21st European Conference on Artificial Intelligence, 18–22 August 2014, frontiers in artificial intelligence and applications* (Vol. 263, pp. 525–530). IOS Press.
- Marcus, G. (2018). *Deep learning: A critical appraisal*. CoRR, arXiv:1801.00631
- Martínez, D., Alenyà, G., Torras, C., Ribeiro, T., & Inoue, K. (2016). Learning relational dynamics of stochastic domains for planning. In *Proceedings of the twenty-sixth international conference on automated planning and scheduling, ICAPS 2016* (pp. 235–243). AAAI Press.
- McCreath, E., & Sharma, A. (1995). Extraction of meta-knowledge to restrict the hypothesis space for ilp systems. In *Eighth Australian joint conference on artificial intelligence*, pp. 75–82.
- Michie, D. (1988). Machine learning in the next five years. In D. H. Sleeman (Ed.), *Proceedings of the third European Working Session on Learning, EWSL 1988* (pp. 107–122). Turing Institute, Pitman Publishing.
- Muggleton, S. (1987). Duce, an oracle-based approach to constructive induction. In *Proceedings of the 10th International joint conference on artificial intelligence* (pp. 287–292). Morgan Kaufmann.
- Muggleton, S. (1991). Inductive logic programming. *New Generation Computing*, 8(4), 295–318.
- Muggleton, S. (1995). Inverse entailment and Progol. *New Generation Computing*, 13(3 & 4), 245–286.
- Muggleton, S., & Buntine, W. L. (1988). Machine invention of first order predicates by inverting resolution. In *Machine Learning, Proceedings of the fifth international conference on machine learning* (pp. 339–352). Morgan Kaufmann.
- Muggleton, S., & De Raedt, L. (1994). Inductive logic programming: Theory and methods. *The Journal of Logic Programming*, 19(20), 629–679.
- Muggleton, S., De Raedt, L., Poole, D., Bratko, I., Flach, P. A., Inoue, K., & Srinivasan, A. (2012). ILP turns 20—Biography and future challenges. *Machine Learning*, 86(1), 3–23.
- Muggleton, S., Dai, W., Sammut, C., Tamaddoni-Nezhad, A., Wen, J., & Zhou, Z. (2018a). Meta-interpretive learning from noisy images. *Machine Learning*, 107(7), 1097–1118.
- Muggleton, S., & Feng, C. (1990). Efficient induction of logic programs. In *First international workshop on algorithmic learning theory, ALT '90*, pp. 368–381.
- Muggleton, S., & Hocquette, C. (2019). Machine discovery of comprehensible strategies for simple games using meta-interpretive learning. *New Generation Computing*, 37, 203–217.
- Muggleton, S. H., Lin, D., Pahlavi, N., & Tamaddoni-Nezhad, A. (2014). Meta-interpretive learning: Application to grammatical inference. *Machine Learning*, 94(1), 25–49.
- Muggleton, S. H., Lin, D., & Tamaddoni-Nezhad, A. (2015). Meta-interpretive learning of higher-order dyadic Datalog: Predicate invention revisited. *Machine Learning*, 100(1), 49–73.
- Muggleton, S., Paes, A., Costa, V. S., & Zaverucha, G. (2009). Chess revision: Acquiring the rules of chess variants through FOL theory revision from examples. In *19th international conference inductive logic programming, ILP 2009. Lecture notes in computer science* (Vol. 5989, pp. 123–130). Springer.
- Muggleton, S. H., Schmid, U., Zeller, C., Tamaddoni-Nezhad, A., & Besold, T. R. (2018b). Ultra-strong machine learning: Comprehensibility of programs learned with ILP. *Machine Learning*, 107(7), 1119–1140.
- Nienhuys-Cheng, S. H., & Wolf, R. (1997). *Foundations of inductive logic programming*. Springer.
- Patsantzis, S., & Muggleton, S. (2021). Top program construction and reduction for polynomial time meta-interpretive learning. *Machine Learning*, 110, 755–778.
- Picado, J., Termehchy, A., Fern, A., & Pathak, S. (2017). Towards automatically setting language bias in relational learning. In *Proceedings of the 1st workshop on Data Management for End-to-End Machine Learning, DEEM@SIGMOD 2017* (pp. 3:1–3:4). ACM.
- Picado, J., Termehchy, A., Fern, A., Pathak, S., Ilango, P., & Davis, J. (2021). Scalable and usable relational learning with automatic language bias. In G. Li, Z. Li, S. Idreos, & D. Srivastava (Eds.), *SIGMOD '21: International Conference on Management of Data, Virtual Event, China, June 20–25, 2021* (pp. 1440–1451). ACM. <https://doi.org/10.1145/3448016.3457275>
- Plotkin, G. (1971). *Automatic methods of inductive inference*. PhD thesis, Edinburgh University.
- Quinlan, J. R. (1990). Learning logical definitions from relations. *Machine Learning*, 5, 239–266.
- Ray, O. (2009). Nonmonotonic abductive inductive learning. *Journal of Applied Logic*, 7(3), 329–340.
- Reed, S. E., & de Freitas, N. (2016). Neural programmer-interpreters. In *4th international conference on learning representations, ICLR 2016*.
- Ribeiro, T., & Inoue, K. (2014). Learning prime implicant conditions from interpretation transition. In *24th international conference on inductive logic programming, ILP 2014. Lecture notes in computer science* (Vol. 9046, pp. 108–125). Springer.
- Ribeiro, T., Folschette, M., Magnin, M., & Inoue, K. (2020). *Learning any semantics for dynamical systems represented by logic programs*. Working paper or preprint.

- Ribeiro, T., Magnin, M., Inoue, K., & Sakama, C. (2015). Learning multi-valued biological models with delayed influence from time-series observations. In *14th IEEE international conference on machine learning and applications, ICMLA 2015* (pp. 25–31). IEEE.
- Richardson, M., & Domingos, P. M. (2006). Markov logic networks. *Machine Learning*, *62*(1–2), 107–136. <https://doi.org/10.1007/s10994-006-5833-1>.
- Rocktäschel, T., & Riedel, S. (2017). End-to-end differentiable proving. In *Advances in neural information processing systems 30: Annual conference on neural information processing systems 2017*, 4–9 December 2017, pp. 3788–3800.
- Sammur, C., Sheh, R., Haber, A., & Wicaksono, H. (2015). The robot engineer. In *Late breaking papers of the 25th international conference on inductive logic programming, CEUR Workshop Proceedings* (Vol. 1636, pp. 101–106). CEUR-WS.org.
- Sato, T. (1995). A statistical learning method for logic programs with distribution semantics. In L. Sterling (Ed.), *Logic programming, Proceedings of the twelfth international conference on logic programming*, Tokyo, Japan, June 13–16, 1995 (pp. 715–729). MIT Press.
- Sato, T., & Kameya, Y. (2001). Parameter learning of logic programs for symbolic-statistical modeling. *Journal of Artificial Intelligence Research*, *15*, 391–454. <https://doi.org/10.1613/jair.912>.
- Schüller, P., & Benz, M. (2018). Best-effort inductive logic programming via fine-grained cost-based hypothesis generation—The inspire system at the inductive logic programming competition. *Machine Learning*, *107*(7), 1141–1169.
- Sivaraman, A., Zhang, T., den Broeck, G. V., & Kim, M. (2019). Active inductive logic programming for code search. In *Proceedings of the 41st international conference on software engineering, ICSE 2019* (pp. 292–303). IEEE/ACM.
- Srinivasan, A. (2001). *The ALEPH manual*. Machine Learning at the Computing Laboratory, Oxford University.
- Srinivasan, A., King, R. D., & Bain, M. (2003). An empirical study of the use of relevance information in inductive logic programming. *The Journal of Machine Learning Research*, *4*, 369–383.
- Stahl, I. (1995). The appropriateness of predicate invention as bias shift operation in ILP. *Machine Learning*, *20*(1–2), 95–117.
- Tamaddoni-Nezhad, A., Bohan, D., Raybould, A., & Muggleton, S. (2014). Towards machine learning of predictive models from ecological data. In *24th international conference on inductive logic programming, ILP 2014. Lecture notes in computer science* (Vol. 9046, pp. 154–167). Springer.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008.
- Wang, W. Y., Mazaitis, K., & Cohen, W. W. (2014). Structure learning via parameter learning. In *Proceedings of the 23rd ACM international conference on conference on information and knowledge management, CIKM 2014* (pp. 1199–1208). ACM.
- Wirth, N. (1985). *Algorithms and data structures*. Prentice Hall.
- Yang, F., Yang, Z., & Cohen, W. W. (2017). Differentiable learning of logical rules for knowledge base reasoning. In *NIPS 2017*.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.