



Modular design patterns for hybrid learning and reasoning systems

a taxonomy, patterns and use cases

Michael van Bekkum¹ · Maaike de Boer¹ · Frank van Harmelen² · André Meyer-Vitali¹ · Annette ten Teije²

Accepted: 27 March 2021 / Published online: 18 June 2021
© The Author(s) 2021

Abstract

The unification of statistical (data-driven) and symbolic (knowledge-driven) methods is widely recognized as one of the key challenges of modern AI. Recent years have seen a large number of publications on such hybrid neuro-symbolic AI systems. That rapidly growing literature is highly diverse, mostly empirical, and is lacking a unifying view of the large variety of these hybrid systems. In this paper, we analyze a large body of recent literature and we propose a set of *modular design patterns* for such hybrid, neuro-symbolic systems. We are able to describe the architecture of a very large number of hybrid systems by composing only a small set of elementary patterns as building blocks. The main contributions of this paper are: 1) a taxonomically organised vocabulary to describe both processes and data structures used in hybrid systems; 2) a set of 15+ design patterns for hybrid AI systems organized in a set of elementary patterns and a set of compositional patterns; 3) an application of these design patterns in two realistic use-cases for hybrid AI systems. Our patterns reveal similarities between systems that were not recognized until now. Finally, our design patterns extend and refine Kautz's earlier attempt at categorizing neuro-symbolic architectures.

Keywords Neuro-symbolic systems · Design patterns

1 Introduction

It is widely acknowledged in recent AI literature that the data-driven and knowledge-driven approaches to AI have complementary strengths and weaknesses [16]. This has led to an explosion of publications that propose different architectures to combine both symbolic and statistical techniques. Surveys exist on narrow families of such systems [3, 61, 62], but to date no conceptual framework is available in which such hybrid symbolic-statistical systems can be discussed, compared, configured and combined. In this paper, we propose a set of modular design patterns

for Hybrid AI systems that combine learning and reasoning (using combinations of data-driven and knowledge-driven AI components). With this set of design patterns, we aim to achieve the following goals. First, we provide high-level descriptions of the architectures of hybrid AI systems. Such abstract descriptions should enable us to better understand the commonalities and differences between different systems, while abstracting from their specific technical details. We will show for a number of our design patterns that they describe systems which have not been recognised in the literature as essentially doing the same task. Secondly, our set of design patterns is intended to bridge the gap between the different communities that are currently studying hybrid approaches to AI systems. AI communities such as machine learning and knowledge-based systems as well as other communities (such as cognitive science) often use very different terminologies which hamper the communication about the systems under study. Finally, and perhaps most importantly, our design patterns are modular, and are intended as a tool for engineering hybrid AI systems out of reusable components. In this respect, our design patterns for hybrid AI systems have the same goals as the design patterns which are well known from Software Engineering [25].

This article belongs to the Topical Collection: *30th Anniversary Special Issue*

✉ Frank van Harmelen
Frank.van.Harmelen@vu.nl

Annette ten Teije
Annette.ten.Teije@vu.nl

¹ TNO, The Hague, The Netherlands

² Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

Darwiche [16] draws attention to the distinction between two types of components in AI systems, for which he uses the terms function-based and model-based. Similarly, Pearl [47] uses the term “model-free” for the representations typically used in many learning systems. Other names used for inferences at this layer are: “model-blind,” “black-box,” or “data-centric” [47], to emphasize that the main task performed by machine learning systems is function-fitting: fitting data by a complex function defined by a neural network architecture. Such “function-based” or “model-free” representations are in contrast to the “model-based” representations typically used in reasoning systems. We will use a similar distinction in our design patterns. Both Darwiche and Pearl argue for combining components of these two types: “the question is not whether it is functions or models but how to profoundly integrate and fuse functions with models” [16] and “Our general conclusion is that human-level AI cannot emerge solely from model-blind learning machines; it requires the symbiotic collaboration of data and models” [47]. However, neither of the cited works discuss how such combinations must be made. This is exactly what we set out to do in this paper by proposing a set of modular design patterns for such combinations.

Lamb & Garcez [17] ask for “*The trained network and the logic [to] become communicating modules of a hybrid system [...]. This distinction between having neural and symbolic modules that communicate in various ways and having translations from one representation to the other in a more integrative approach to reasoning and learning should be at the centre of the debate in the next decade.*” Our proposal is a concrete step in precisely the direction that they call for, providing a set of composition patterns that can be used as modular building blocks in hybrid systems.

The main contributions of this paper are as follows: (i) a taxonomically organised vocabulary to describe both the processes that constitute hybrid AI systems as well as the data structures that such processes produce and consume; (ii) a set of modular design patterns for hybrid AI systems, organised in a set of elementary patterns, plus more sophisticated patterns that can be constructed by composing such elementary patterns; we will show how a number of systems from the recent literature can be described as such compositional patterns; these patterns are a further elaboration of our first proposal for such design patterns in [60]; (iii) two realistic use-cases for hybrid AI systems (one for skill matching, one for robot action selection), showing how the architecture for each of these use-cases can be described in terms of our compositional design patterns.

The paper is structured as follows, first we describe our taxonomical vocabulary in Section 2, in Section 3 we describe a set of elementary patterns, followed by a set of compositional patterns in Section 4. Section 5 describes two realistic use-cases in terms of the patterns from Sections 3

and 4. Section 7 concludes and discusses directions for future work.

2 A taxonomical vocabulary

In order to describe design patterns, a terminology is required that defines a taxonomy of both processes and their inputs and outputs on various levels of abstraction. On the highest level of abstraction, we define instances, models, processes and actors. In the pattern diagrams, instances are represented as rectangular boxes, models as hexagonal boxes, processes as ovals and actors as triangles. More specific concepts will be used, when necessary and useful, using a colon-separated notation. For example, *model:stat:NN* refers to a neural network model. The level of abstraction depends on the use of the pattern and the stage of design and implementation: the closer to implementation, the more specific the concepts will be. The design patterns should abstract from implementation details, but be specific enough to document applicable design choices. As abbreviated notation, we will not name the highest abstraction level, because that is implied by the type of box. In the models, we always indicate whether a specific model is statistical (*stat*) or semantic (*sem*). The full taxonomy with definitions can be found in the [Appendix](#).

2.1 A taxonomy of instances

Instances are the basic building blocks of “things”, examples or single occurrences of something. The two main classes of instances are data and symbols. The precise distinction between symbols and “non-symbols” remains a contentious issue in modern philosophy. We are following [7] by imposing the following requirements before any token can be deemed a symbol: (1) a symbol must designate an object, a class or a relation in the world, and such a designated object, class or relation is then called the interpretation of the symbol; (2) symbols can be either atomic or complex, in which case they are composed of other symbols according to a formal set of compositional rules; and (3) there must be system of operations that, when applied to a symbol, generates new symbols, that again must have a designation. Thus, the tokens p_1 and p_2 may designate particular persons, with the symbol r designating some relation between these persons; then $r(p_1, p_2)$ is a complex symbol made out of these atomic symbols, and the operation $r(p_1, p_2) \models_T r(p_2, p_1)$ defines an operation constructing one complex symbol out of another. In logic, such “operations” correspond to logical inference \models , and this logical view is most relevant in this paper, but in another context such operations may be transitions between the symbols that denote the states of a finite state machine. All

this makes the symbol p_1 different from a data item, say a picture of a person, where the collection of pixels may be an accurate image of a person, but does not designate the person in a way that allows the construction of more complex designations and operations that transform these into other designations. Simply put: a symbol p designates a person, whereas a picture is just itself. Such tokens which are not symbols are what we will call “data” in this paper.

The types of data that appear in Hybrid AI systems include

- numbers: numerical measurements;
- texts: sequences of words, sentences;
- tensors: multi-dimensional spaces, including bitmaps;
- streams: (real-time) sequences of data, including video and other sensor inputs.

The types of symbols, on the other hand, include

- labels: short descriptions;
- relations: connections between data items, such as triples and other n -ary relations;
- traces: (historical) records of data and events, such as proof traces for explanations.

2.2 A taxonomy of models

Models are descriptions of entities and their relationships. They are useful for inferring data and knowledge. Models in Hybrid AI systems can be either (1) statistical or (2) semantic models. Statistical models represent dependencies between statistical variables. Examples are (Deep) Neural Networks, Bayesian Networks and Markov Models. Semantic models represent the implicit meaning of symbols by specifying their concepts, attributes and relationships. Examples are Taxonomies, Ontologies, Knowledge Graphs, Rulebases and Differential Equations. We summarise these semantic models under the umbrella term “knowledge base” (KB).

2.3 A taxonomy of processes

In order to perform operations on instances and models, processes define the steps that lead from inputs to results. Three main types of processes are: (i) the generation of instances and models, (ii) their transformation and (iii) inferencing thereupon. Generation of models is performed either via training or, “manually” by knowledge engineering with experts. Many forms of transformations exist, such as transforming a knowledge graph to a vector space. Inferences are made using induction or deduction. Induction is constructing a generalisation out of specific instances. Such a generalisation (a “model”) can take many different forms, ranging from the trained weights in a neural network

to clauses in a learned Logic Program, but in every case, such models are created by inductions on instances.

Using such models, we can apply deductive inferencing in order to reach conclusions about specific instances of data. Commonly, deduction is associated with logical inference, but the standard definition, namely as inference in which the conclusion is of no greater generality than the premises, equally applies to the forward pass of a neural network, where a general model (the trained network) is applied to an instance in order to arrive at a conclusion. Both inductive and deductive processes can be either symbolic or statistical, but in every case induction reasons from the specific (the instances) to the general (a model), and deduction does the converse. Thus, the distinction between induction and deduction corresponds precisely to the distinction between learning and reasoning. Both classification (“the systematic arrangement in groups or categories according to established criteria”, see our [Appendix](#)) and prediction (“to calculate some future event or condition as a result of analysis of available data”, again see our [Appendix](#)) are deductive processes, since both use a generic model (obtained through an inductive process of learning or engineering) to derive information about specific instances (the objects to be classified or the events to be predicted).

2.4 A taxonomy of actors

Processes in an AI system are initiated by autonomous actors, based on their intentions and goals. They interact with each other using many protocols and behaviours, such as collaboration, negotiation or competition. These interactions lead to collective intelligence and emergent social behaviour, such as in swarms, multi-agent systems or human-agent teams. Autonomy is a gradual property, ranging from remotely controlled to selfish behaviour and all forms of cooperation in between. Actors can be humans, (software) agents or robots (physically embedded agents). Examples are proactive software components, apps, services, mobile agents, drones or (parts of) autonomous vehicles. Actors are not yet used explicitly in the current collection of patterns, but will be included in the future, when we will also dive into distributed AI and human-agent interaction.

3 Elementary patterns

The taxonomy of instances, models and processes from the previous section gives rise to a number of elementary building blocks for hybrid systems that combine learning and reasoning. The “train” process consumes either data or symbols to produce a model (Fig. 1a and b, all these terms taken from the taxonomy)

Additionally, an actor (e.g., a domain expert or knowledge engineer) can create a model, such as an ontology or rule-base (Fig. 1c).

In the processing, often a transformation step is needed to create the right type of data, either from symbol or data (Fig. 1d).

Of course models are only trained in order to be subsequently used in a downstream task (predicting labels, deriving conclusions, predicting links, etc). This process is captured in the patterns (Fig. 2a–c), depending on the symbolic or statistical nature of the data. Following our taxonomy, an infer step uses such models in combination with either data or symbols to infer conclusions:

Finally, sometimes an operation on a semantic model is neither a logical induction or deduction, but a transformation into another datastructure. This is captured by the final elementary pattern (Fig. 2d):

As is encoded in these diagrams, the types of models involved (symbolic or statistical) and the type of results derived is constrained by the type of the inputs of these elementary processes.

These elementary patterns allow us to give a more precise definition of the concept of “hybrid systems”, which is often used rather nebulously in the literature:

Definition *Machine Learning systems* are systems that combine pattern (Fig. 1a) with (Fig. 2a), yielding pattern (Fig. 3a), see below; *Knowledge Representation systems* are

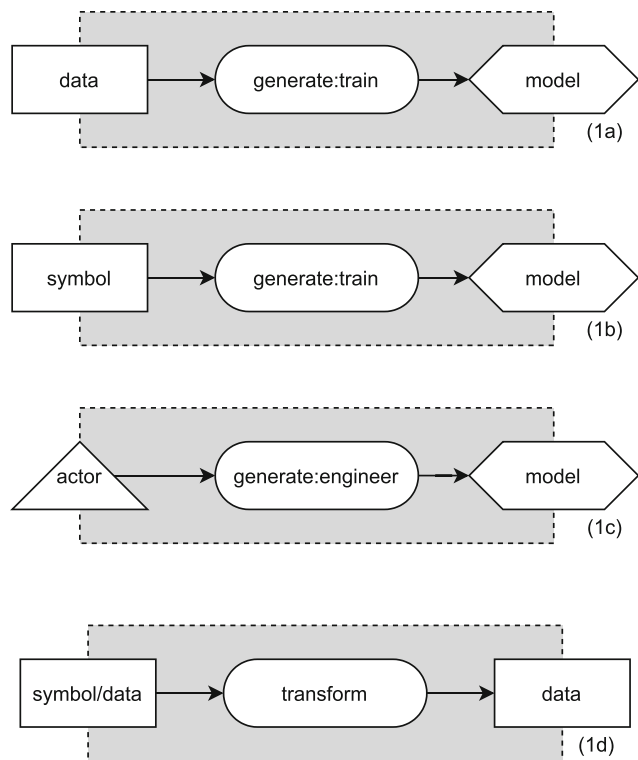


Fig. 1 Elementary patterns to generate a model

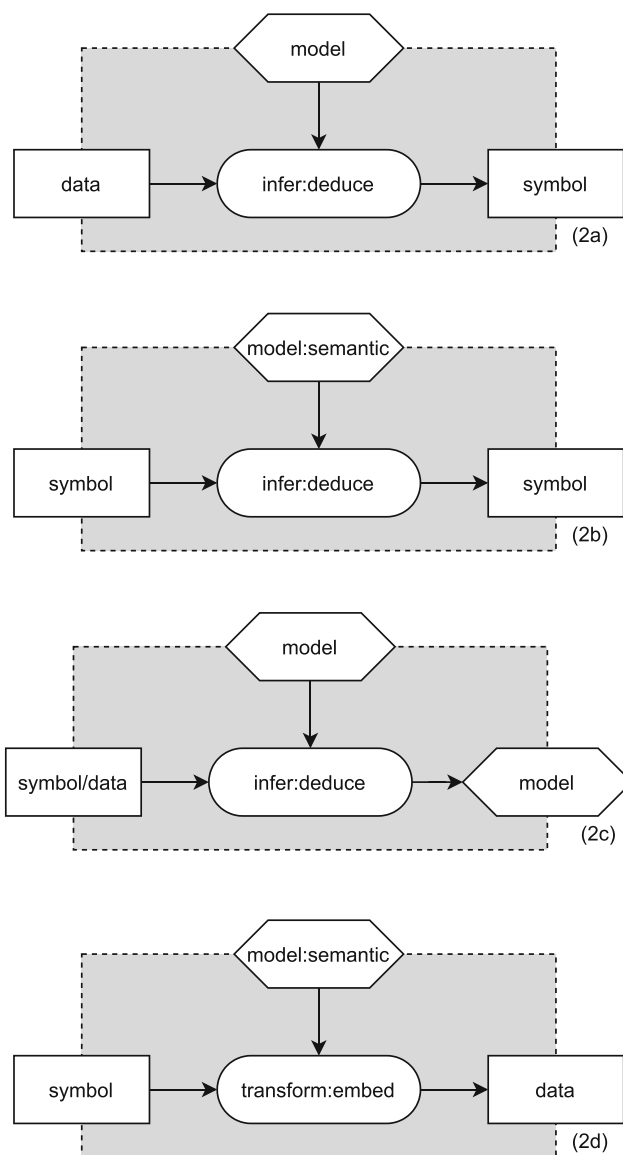


Fig. 2 Elementary patterns to use a model

systems that follow pattern (Fig. 2b); *Hybrid systems* are systems that form any other combination of the elementary patterns (Figs. 1a–2d).

Already these elementary patterns (Figs. 1a–d, 2a–d), even in their almost trivial simplicity, can be used to group together a large number of very different approaches from the literature: even though the algorithms and representations of Inductive Logic Programming (ILP) [34], Markov Logic Networks [50], and Probabilistic Soft Logic [4, 29] are completely different, the architecture pattern Fig. 1b applies to all of them, showing that they are all aimed at the same goal: learning over symbolic structures.

Similarly, learning a symbolic rule-set that captures rules for knowledge graph completion [43] is captured by this

pattern. Constructing knowledge graph embeddings into a high-dimensional vector space [44, 46, 62] is also captured by Fig. 1b. So ILP and KG embedding would each be captured more specifically by adding type annotations to the constructed model: *model:sem* for ILP, and *model:stat* for KG embedding. Many “classical” learning algorithms such as decision tree learning and rule mining, as well as deep learning systems, are covered by architectural pattern Fig. 1a. The learning patterns (Fig. 1a–b) must be combined with the prediction patterns (Fig. 2a–c) to give a model for the full learning and prediction task (Fig. 3a):

This model is precisely the composition of the elementary processes for train and infer given above. Even learning with a regular neural network is captured by this diagram (although it is not typically recognized that the feed-forward phase, when a trained neural network is applied to new data, is actually a deductive task, namely reasoning from given premises (input data plus the learned network weights) to a conclusion.

Analogously to learning from *data*, it is also possible to learn from *symbols*, using elementary patterns (Figs. 1b and 2b) instead of (Fig. 1a) and (Fig. 2a):

As mentioned above, this pattern then describes the learning (and subsequent inference) in Inductive Logic Programming, Knowledge Graph embeddings, Probabilistic Soft Logic and Markov Logic Networks.

Using our hierarchical taxonomy, the two patterns (Fig. 3a and b) can be abstracted into a single pattern, replacing all the boxes labelled with “data” or “symbol” by the generic term “instance”. The specific diagrams (Fig. 3a and b) can then be recovered by adding type annotations “instance:sym” or “instance:data”. Such type-specialisations maintain the insight that many of these very different approaches (ILP, MLN, PSL, Knowledge Graph embeddings) actually follow the same schema.

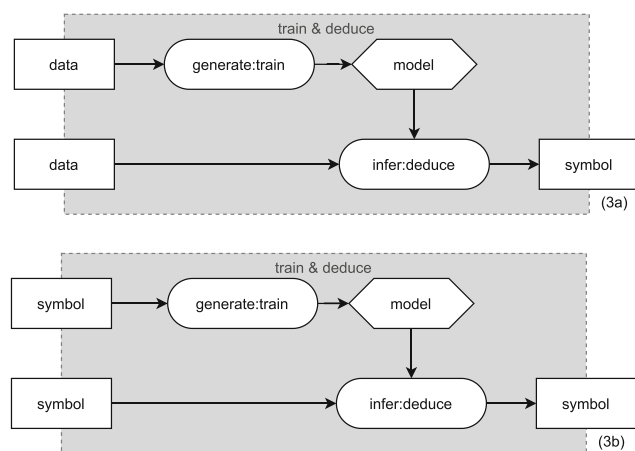


Fig. 3 Pattern for basic machine learning systems

4 A collection of compositional patterns

In this section, we describe compositional patterns based on the elementary pattern described in the previous section. We combined papers from several fields into one pattern. From the elementary patterns, we create compositions in two ways: (1) we can create a more complex pattern by connecting or ‘stitching’ elementary patterns; (2) we can go more specific or more abstract (only showing the boxes of Fig. 1a – b for example); in a specific pattern we can specify the type of, for example, a symbol block in terms of *symbol:relations*.

4.1 Learning from data with symbolic output

In ontology learning, a symbolic ontology is learned from data in the form of text [3, 10, 11, 22, 35, 65]. The text is first translated into (subject, verb, object) relations using a statistical model such as the Stanford Parser [12, 18]. These relations are an intermediate representation. A semantic model, for example rules for Hearst patterns, can then infer the relations that form a full ontology including relation hierarchies and axioms. This pattern (Fig. 4) combines patterns (Fig. 2a and b). Whereas in other cases an ontology can play the role of a model on the basis of which properties of instances are deduced, in this case, we represent the ontology as a set of relations, because it is the output of a process, and not a model which is input to a process.

A related but different instantiation of this pattern is the use of text-mining not to learn full-blown ontologies, but to learn just the class/instance distinction (which is always problematic in ontology modelling), as done in [45]. As concerns the architectural patterns, this work only differs in the actual content of the symbolic output: a full-blown ontology, or only a class/instance label.

In contrast, other ontology-learning systems [10, 35] start from a given set of relations (the “A-box” of description logic) and then infer an ontological hierarchy. These systems only apply the second half of the above pipeline, pattern (Fig. 2b).

An entirely different application domain is found in [2], where symbolic first-order logic representations are generated to describe the content of images.

4.2 Explainable learning systems through rational reconstruction

Hybrid symbolic-statistical systems are seen as one possible way to remedy the “black-box” problem of many modern machine learning systems [63]. Pattern (Fig. 5a) shows one of the hybrid architectures that have been proposed in the literature for this purpose. A standard machine learning system is trained (generate:train) to construct a model




Fig. 4 Pattern for learning with symbolic output

which is then applied to input data in order to produce (infer:deduce) a prediction (for example, a label for a given input image). The result of this process (in the form of the pairs of image + label) is then passed on to a symbolic reasoning system which then uses background knowledge (model:semantic) to produce a “rational reconstruction” of a reason to justify the input/output-pair of the learning system. An example of this is the work by [59] who uses large knowledge graphs to reconstruct the justification of temporal patterns learned from Google Trends data. It is important to emphasize that the justification found by the symbolic system is unrelated to the inner workings of the black box of the machine learning system. The symbolic system produces a post-hoc justification that is not necessarily reflecting the statistical computation. This architecture is also used in [53], where a description logic reasoner is used to come up with a logical justification of classifications produced by a deep learning system. Notice that pattern (Fig. 5a) is a straightforward combination of elementary patterns (Figs. 1 a, 2a and b).

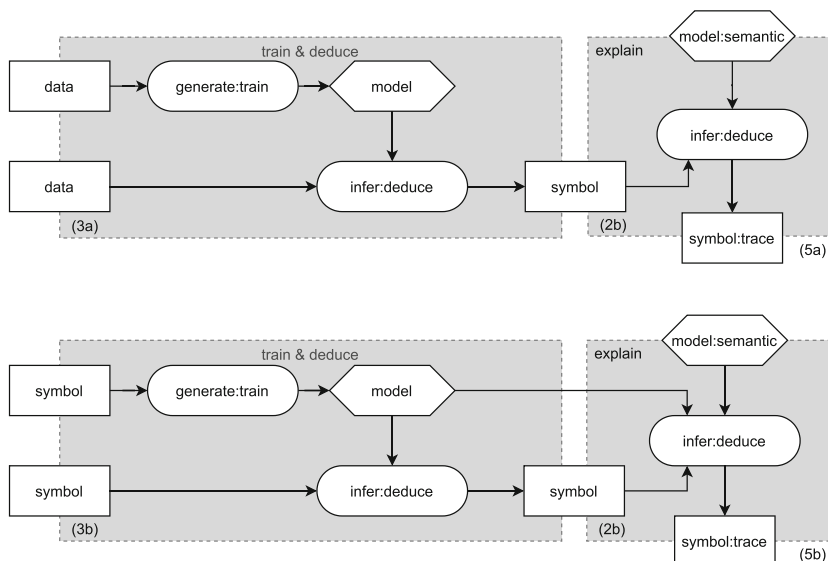
Pattern (5a) captures so-called “instance-level explanations”, where a separate explanation is generated for every specific result of applying the learned model. In contrast,

it is also possible to generate “model-level explanations”, where a generic explanation is constructed that captures the structure of the entire learned model (Fig. 5b). An example of this is [13], which trains a second neural network that uses the input/output behaviour of a classifier network to general first-order logic formulas that can then be used to explain the behaviour of the classifier. This results in a modification of the above pattern in which the subsystem labelled with “2b” is replaced by a learning system that takes the learned model from 1a as input and produces a symbolic explanation of that model. In other words: the explanation is generated based on the trained model, and not just on the derived individual result of applying that model to a given piece of data.

4.3 Learning an intermediate abstraction

Intermediate abstraction for learning A well known machine learning benchmark is to recognise sets of handwritten digits [37]. This digit-recognition task could be extended to perform addition on such handwritten digits through end-to-end training on the bitmap representations of the handwritten digits. This would correspond to our basic

Fig. 5 Patterns for explainable learning systems through rational reconstruction



pattern (Fig. 3a). However, many authors have observed that it is more efficient to learn an intermediate symbolic representation (mapping the pixel representation of the handwritten digit into a symbolic representation), and then train a model to solve the symbolically represented addition task. This pattern is represented in Fig. 6a, where two standard train+deduce patterns (pattern Fig. 3a, b) are chained together through a symbolic intermediate representation which serves as output for the first pattern and as input for the second. This pattern is exploited in Deep Problog [40] where a system is trained to add handwritten digits by first recognising the digit and then doing the addition. This then turns out to be a much more robust approach than simple end-to-end training going from the digit bitmaps to the summed result in one step. This same pattern also captures the DeepMind experiment [27] where a reinforcement learning agent is trained to not just navigate on a bitmap representation of its world, but to first learn an intermediate symbolic representation of the world and then use that for navigation.

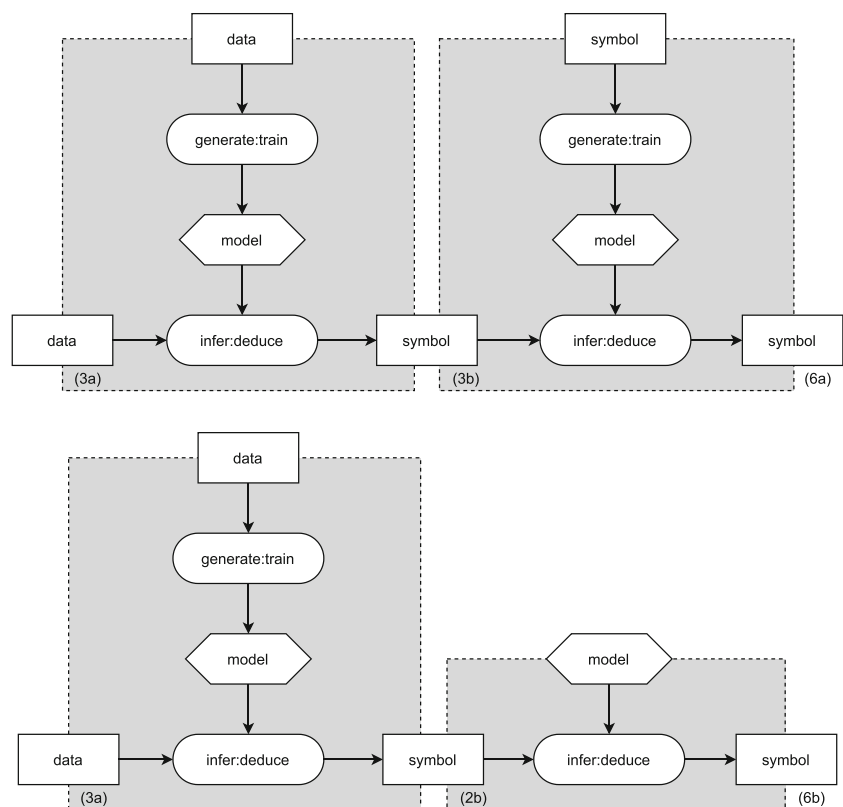
Besides learning a spatial abstraction (as in [27]), the work in [32] uses the same architecture pattern for deriving a temporal abstraction of sequence of subtasks, which are then input to reinforcement learning agents. One of the advantages of such an intermediate representation is the much higher rates of transfer learning that can be obtained

after making trivial changes to the input distribution, be they handwritten digits or bitmaps of floor spaces.

Intermediate abstraction for reasoning Whereas pattern (Fig. 6a) consists of a composition of two patterns for learning, pattern (Fig. 3) (first deriving an intermediate abstraction on the basis of a trained model and then using this intermediate abstraction as the basis for a further derivation on the basis of a second trained model), it is also possible to use the derived abstraction as the input for a deductive reasoning task, composing pattern (Fig. 3) with pattern (Fig. 2b), creating pattern (Fig. 6b). A classic example of this pattern is the AlphaGo system [56], where machine learning is used to train an evaluation function that gives rise to a symbolic search tree which is then traversed using deductive techniques (Monte Carlo tree search) in order to derive a (symbolically represented) next move on the Go board. Notice that pattern Fig. 5a is a specialisation of this pattern (Fig. 6b).

Pattern (Fig. 4) above (ontology learning from text) can now be seen to also be a variation on the general theme of “learning an intermediate abstraction”, where the set of relations extracted by linguistic analysis is the intermediate abstract that is input for the set of rules that constructs the final ontology out of these relations. In pattern (Fig. 4) the models (model:semantic) are assumed to be given (e.g. by

Fig. 6 Patterns for learning an intermediate abstraction



using the pretrained Stanford parser) and hence the training phases are omitted.

4.4 Informed learning with prior knowledge

In [61] a large collection of over 100 different systems is discussed which are all captured by pattern (Fig. 7). In this pattern, the training phase of a learning system is guided by information that is obtained from a symbolic inference system (pattern Fig. 2b). For this purpose, the training step from elementary pattern (Fig. 1a) is extended with a further input to allow for this guidance by inferred symbolic information. A particular example is where domain knowledge (such as captured in modern large knowledge graphs) is used as a symbolic prior to constrain the search space of the training phase [8]. In general, this patterns also captures all systems with a so-called semantic loss function [66], where (part of) the loss-function is formulated in terms of the degree to which the symbolic background knowledge is violated. Such a semantic loss-function is also used in [38], where the semantic loss is calculated by weighted model counting. In [20] and [41] the semantic loss-function is realised through approximate constraint satisfaction. Another example is [19] where logical rules are used as background knowledge for a gradient descent learning task in a high-dimensional real-valued vector space. In the same spirit, [52] exploits a type-hierarchy to inform an embedding in hyperbolic space.

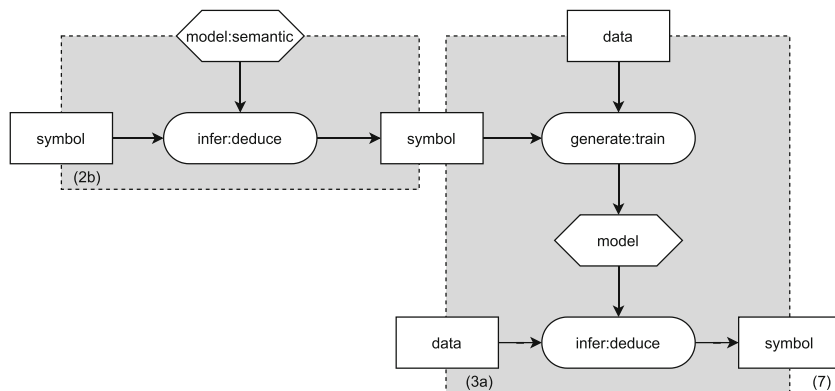
Logic Tensor Networks [21] also fall in this category, since they jointly minimise both the loss function of a neural network and maximise the degree to which a first-order logic theory is satisfied. The fact those LTN’s are captured by the same design pattern as semantic loss functions suggests an analogy between the two (namely that the maximisation of first-order satisfiability in LTN’s can be regarded as a semantic loss-function). This analogy between these two systems was not mentioned in their original papers, but only comes to light through our analysis in terms of high-level design patterns.

An entirely different category of systems that is captured by the same pattern are constrained reinforcement learners (e.g. [28]), where the exploration behaviour of a reinforcement learning agent is constrained through symbolic constraints that enforce safety conditions. Similarly, [33] uses high-level symbolic plans to guide a reinforcement learner towards efficiently learning a policy. Silvestri et al. [57] shows how adding domain knowledge in the form of symbolic constraints greatly improves the sampling-frequency of a neural network trained to solve a combinatorial problem. The LYRICS system [42] proposes a generic interface layer that allows to define arbitrary first order logic background knowledge, allowing a learning system to learn its weights under the constraints imposed by the prior knowledge.

The full design pattern (Fig. 7) requires that the symbolic prior is derived by a symbolic reasoning system, but it is of course also possible that this symbolic prior (or “inductive bias”, using the terminology from [6]) is simply in the form of an explicit knowledge-base for which no further derivation is possible. This would lead to a simplified version of pattern (Fig. 7) where the “Infer” step would be omitted. An example of this is [36], where input data is first abstracted with the help of a symbolic ontology, and is then fed into a classifier, which performs better on the abstracted symbolic data than on the original raw data. A similar example is given in [5], where knowledge graphs are successfully used as priors in a scene description task.

An interesting variation is presented in [67]. This work exploits the fact that pattern (Fig. 7) uses prior knowledge in symbolic as input for learning, while pattern (Fig. 4) produces symbolic results as the output of learning. The system described in [67] then iterates between producing symbolic knowledge, then using this symbolic knowledge as input for informed machine learning, followed again by using the learned model to produce better symbolic knowledge, hence iterating between patterns (Fig. 7) and (Fig. 4). The IterefinE system [1] is another example of this pattern.

Fig. 7 Pattern for informed learning with prior knowledge



4.5 From symbols to data and back again

Link prediction (or: graph completion) in knowledge graphs [44, 46, 62], has been a very active area on the boundary of symbolic and statistical representations, and is an example of what is captured in pattern (Fig. 8). Almost all graph completion algorithms perform this task by first translating the knowledge graph to a representation in a high-dimensional vector space (a process called “embedding”, this is captured in pattern (Fig. 1d)), and then use this representation to predict additional edges which are deemed to be true based on geometric regularities in the vector space, even though they are missing from the original graph. This can be expressed in a variant of pattern Fig. 2a.

4.6 Learning logical representations for statistical inferencing

Integrating knowledge representations into a machine learning system is a long standing challenge in Hybrid AI, since it allows logic calculus to be carried out by a neural network in an efficient and robust manner. Encoding prior knowledge also allows for better training results on fewer data. This pattern (Fig. 9) describes the integration of logic into a machine learning model through tensorization of the logic involved [26] by applying prior (semantic) knowledge representations as constraints for machine learning. The pattern transforms a semantic model into vector/tensor representations and uses these to train a neural network in order to learn. The machine learner can make inferences based on its embedded logic, for example Logic Tensor Networks [54, 68] where relational data is embedded in a (convolutional) neural network. Graph Neural Networks (GNNs, [9]) embed a semantic graph model by transforming a neighbourhood structure of its nodes and edges into vectors and using these to train a neural network. In [14] a reified knowledge base is provided as input to train neural

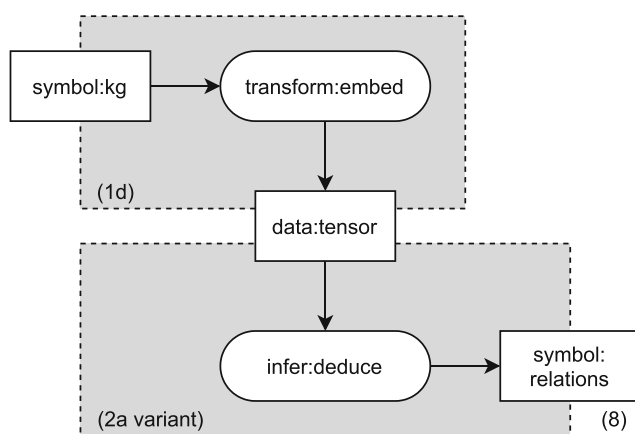


Fig. 8 From symbols to data and back again

modules to perform multi-hop inferencing tasks. The pattern in itself is an extension of pattern Fig. 3a, where the training input data is a (transformed) representation of relational data.

4.7 Learning to reason

Whereas pattern (Fig. 9) provides representation learning by embedding vector/tensor representations of logical structures in a neural network, there are also attempts to learn the reasoning process itself in neural networks. This is motivated by the ability of neural networks to provide higher scalability and better robustness when dealing with noise in the input data (incomplete, contradictory, and erroneous). The focus of pattern (Fig. 10) is on reasoning with first-order logic on knowledge graphs. This pattern learns specific reasoning tasks based on symbolic input tuples and the inferencing results from the symbolic reasoner. Pattern (Fig. 10) is a combination of our basic patterns for symbolic reasoning (Fig. 2b) and training to produce a statistical model (Fig. 1a).

This pattern for training a neural network to do logical reasoning captures a wide variety of approaches such as reasoning over RDF knowledge bases [22], Description Logic Reasoning [30], and logic programming [51]. Relational Tensor Networks (RTNs) [30] use a recurrent neural network to learn two kinds of predictions, namely the membership of individuals to classes and the existence of relations. In a somewhat different application, [22] takes a set of normalized triples and a normalized query to learn classification of the entailment of the query from the statements in the current knowledge graph. Whereas current efforts focus on deductive reasoning in knowledge bases of FOL and DL (or fragments thereof), the pattern can in theory be applied to other inference tasks and mechanisms.

4.8 Meta reasoning for control

There is a long-standing tradition in both AI [15] and in the field of cognitive architectures (e.g. [48]) to investigate so-called meta-reasoning systems, where one system reasons about (or:learns from) the behaviour of another system.

It is widely recognised that the configuration of machine learning systems is highly non-trivial, ranging from choosing an appropriate neural network architecture to setting a multitude of hyper-parameters for that architecture. The field of AutoML [31] aims to automate that configuration task. This is often done through applying machine learning techniques to this task (ie the system is learning the right hyper-parameter settings for the target learning system), but the configuration of the target system can also be done by capturing the knowledge of machine learning engineers. This is done in a system such as

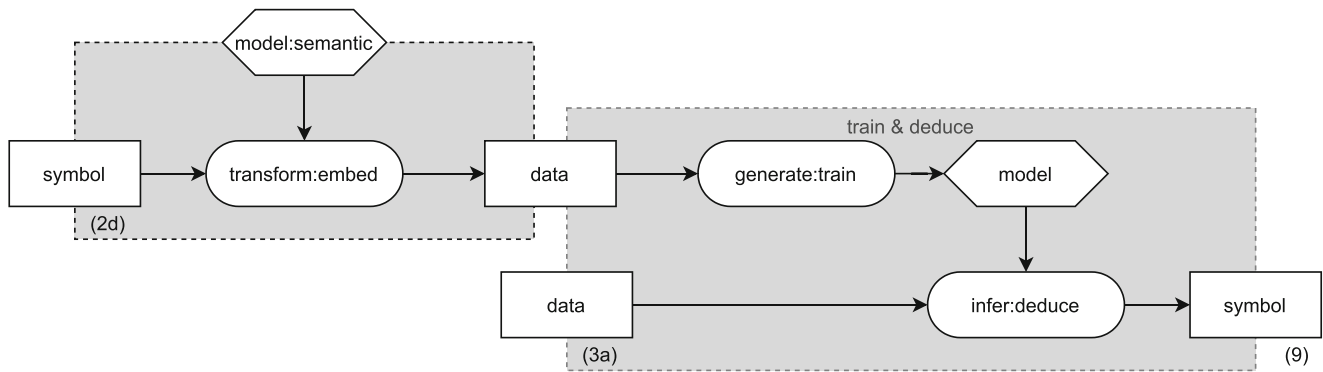


Fig. 9 Pattern for learning logical representations for statistical inferencing

Alpine Meadow [55] and is captured in pattern (Fig. 11): a knowledge based of ML configuration knowledge is used to deduce appropriate hyper-parameter settings for a learning system (sub-pattern (Fig. 2b)), these parameters are then used to train a model (requiring a slightly modified version of sub-pattern (Fig. 1a)), and the resulting performance of this model is inspected by the knowledge base which may

give rise to adaptations of the hyper-parameters in the next iteration.

Another class of systems which at first sight may seem very different, but which are an instantiation of the same pattern are so called “curriculum-guided learning” systems. Curriculum learning problem can be defined as the problem of finding the most efficient sequence of learning situations

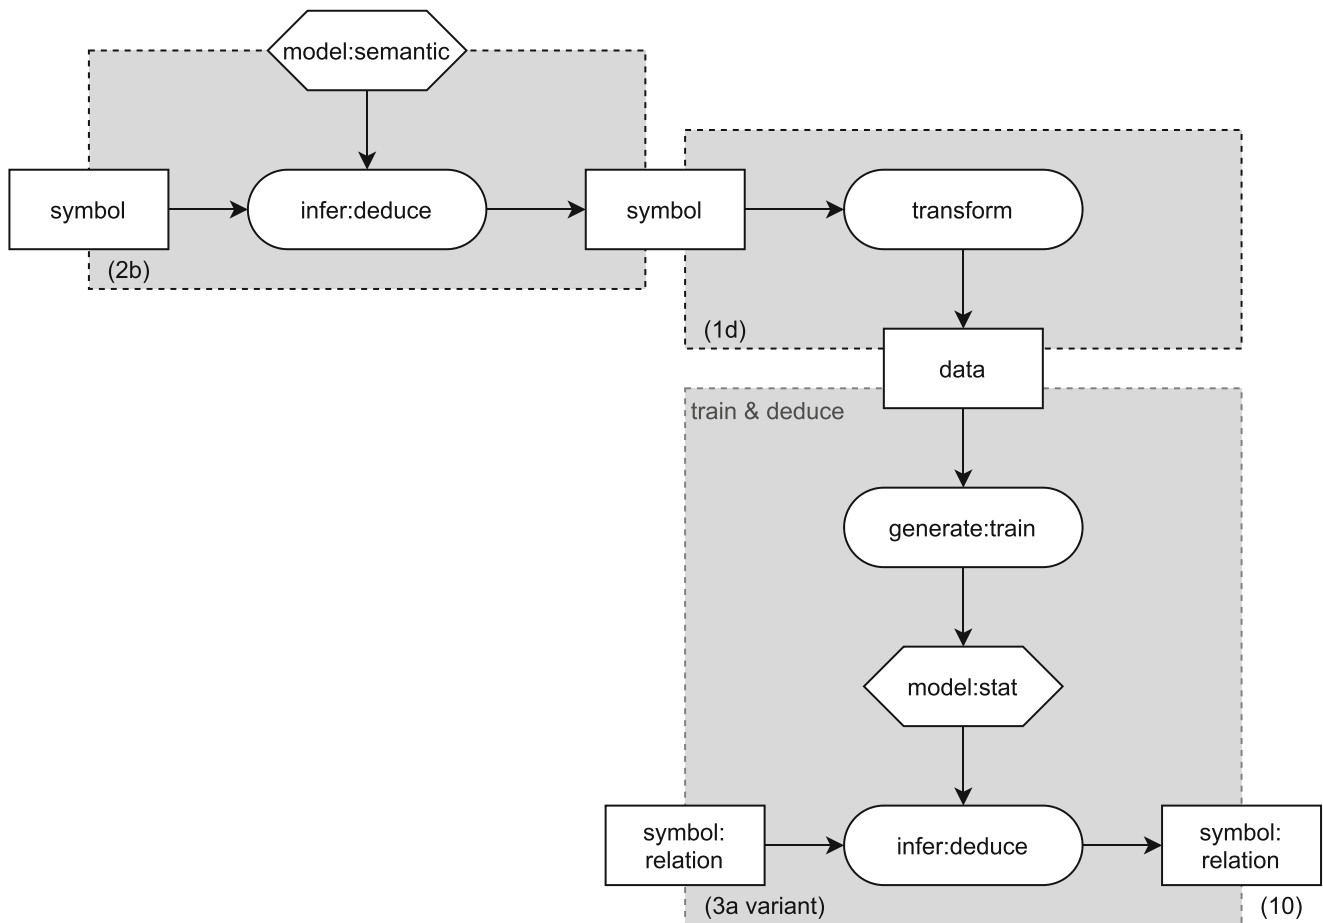


Fig. 10 Learning to reason




Fig. 11 Meta-reasoning for control

in these various tasks so as to maximize its learning speed [23, 24]. In terms of pattern (Fig. 11), the task of the (Fig. 2b) subsystem is to feed the learner in subsystem (Fig. 1a) its training instances in the optimal sequence.

Notice that this pattern closely resembles pattern (Fig. 7) (informed learning with prior knowledge). In that pattern, the symbolic component deduces prior knowledge as input for the training component once, but the resulting training model is not subsequently inspected to possibly adjust this input. Whereas in pattern (Fig. 11) a symbolic system is used to guide the learning behaviour of a subsymbolic system, the converse is also possible. In [39], a system is presented where a subsymbolic system learns the search strategies needed to guide a symbolic theorem prover. This line of work has a long history, dating back to the 1990's [58].

5 Two use-cases

In this section, we describe the use of our boxology patterns in two real world use cases.

We have identified the need for a pattern-based system of systems approach towards design and evaluation of hybrid AI systems. Using our approach can increase the level of trustworthiness of such systems. Trust in AI systems emerges from a number of factors, such as transparency, reproducibility, predictability and explainability. A Hybrid AI system is not to be seen as a monolithic component, but communicating modules of such a hybrid system [17]. Insight into the individual modules and components and their relationships and dependencies is essential, in particular in a decentralised system. The specification and verification of each component and their interactions enable a system-wide validation of expected behaviour. The definition and use of best practices and design patterns supports the generation of trustworthy AI systems - either when building new systems or when understanding existing systems by dissection and reverse-engineering.

Our method allows for step-wise refinement of a system design by starting with a high level of abstraction and drilling down towards implementation details and reusable

components by specifying more and more concrete choices, such as which models to use. Starting from generic patterns, an implementation can be derived and deployed, based on the experience and best practices in Hybrid AI.

5.1 Skills matching

In the first use case, the goal is to create a piece of software that is able to match open vacancies or job descriptions with CVs of job seekers. In this specific use case skills, defined as the ability to perform a task, are used to do the matchmaking. In the first part of the project, a large architecture picture was created with a lot of boxes, arrows, and terms. The distinction between processes and data was not very clear and the type of model was also not explicitly defined.

With use of the boxology, the architecture picture is better readable (see Fig. 12). This figure made it possible to on the one hand talk about the bigger boxes (elementary patterns) and on the other hand go into more detail about the specific implementation of for example a model box and the specific input and output type needed. The figure makes it also possible to think about the future of the project in terms of patterns that have to be added, substituted or removed.

To go into more detail for this specific use case: in the training phase, vacancies are transformed to the data type tensor using an (word2vec) embedding. A Skills ontology, which is a semantic model, engineered by a human (pattern Fig. 1c), is also transformed to a tensor using an embedding (variant of pattern Fig. 1d). Both the vacancy and the skill ontology are used to train a neural network (pattern Fig. 1a), a statistical model. This model learns which sentences or parts of sentences of a vacancy text contain a skill and match to which specific skill in the ontology. When this model is applied to a new vacancy, transformed using an embedding to a tensor (pattern Fig. 1d), it predicts the most probable skill through deduction (pattern Fig. 2a). An additional step is to use the ontology to deduce (pattern Fig. 2b) the label skill to a more understandable skill, for example a skill with a description.

5.2 Robot in action

In the second use case, the goal is to get a robot to choose the best action to perform. The robot has access to a camera. In a training phase, an object detector and a tracker are trained using statistical models, whereas an ontology is handcrafted by an expert actor. The object detector is a neural network trained using images of the environment (pattern Fig. 1a). The tracker uses the video stream of the camera and uses a different type of neural network (pattern Fig. 1a). A semantic model in the form of an ontology about the world is engineered by a human (pattern Fig. 1c). At

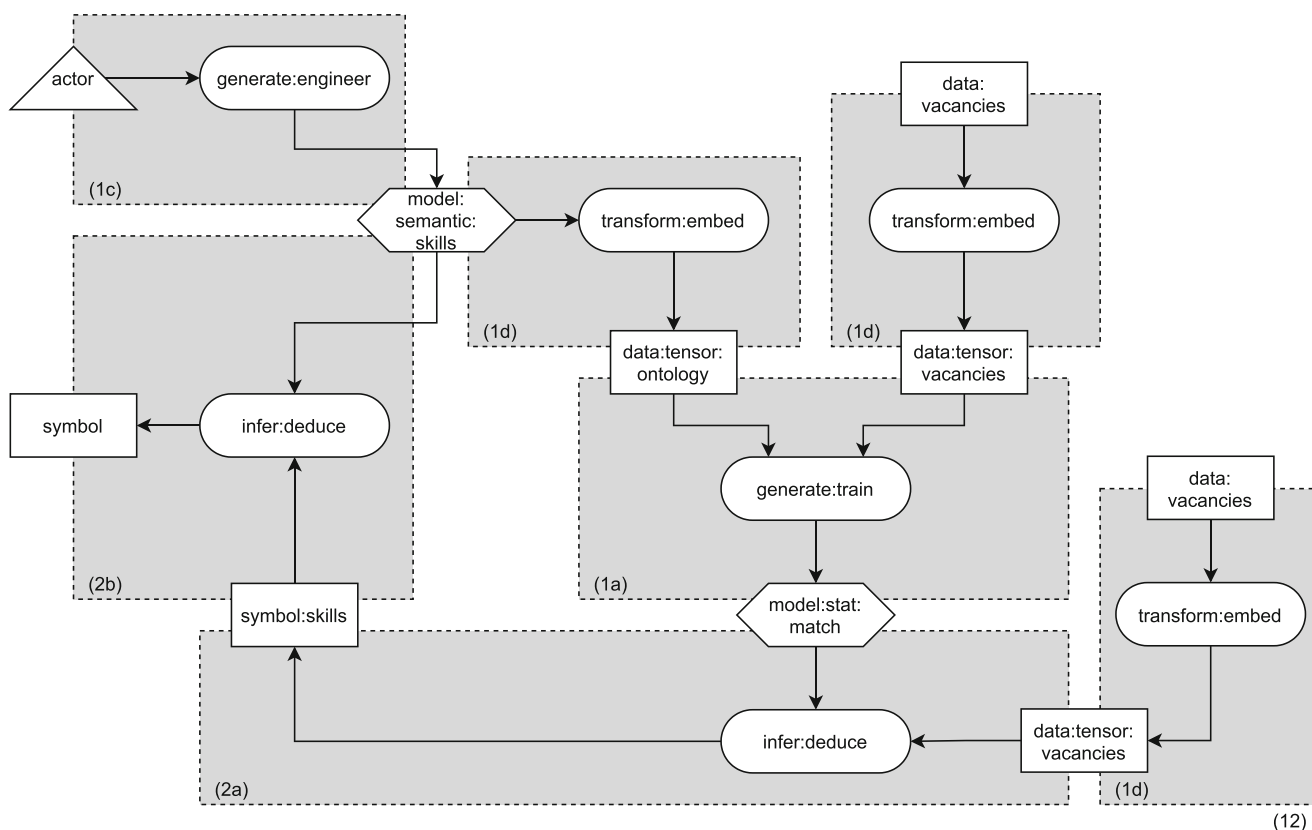


Fig. 12 Use-case for skills matching

each timestep, the robot obtains new information from the camera. The image is used to predict which objects are visible in the environment (pattern Fig. 2a). The tracker is used to track objects through time (pattern Fig. 2a). The tracks and the world model are used to induce semantic rules (pattern Fig. 2c), such as reasoning what will happen next. Then these rules and the detected object(s) are used to predict what the best action is (pattern Fig. 2b).

In this (Fig. 13), the boxology helped to see the similarity between the object detector and the tracker and to determine how the detailed process flow should be.

6 Related work

In [61] a large collection of over 100 different systems for “informed machine learning” is presented. The survey paper provides a broad overview of how many different learning algorithms can be enriched with prior knowledge. Their Fig. 2 provides a taxonomy of “informed machine learning” across three dimensions: which source of knowledge is integrated (e.g. expert knowledge or common sense world knowledge), how that knowledge is represented (e.g. as knowledge graphs, logic rules, algebraic equations and 5 other types of representation) and where this knowledge is

integrated in the machine learning pipeline (in the training data, the hypothesis set, the learning algorithm or the final hypothesis). The second and third of these dimensions are used to categorise well over 100 different published systems into an 8x4 grid. All of these systems are captured in one of our design patterns (pattern Fig. 7), so while our proposal covers a much broader set of hybrid systems, the result of [61] is a very detailed analysis of one of our patterns.

In his invited address to the AAAI 2020 conference, Henry Kautz introduced a taxonomy for neural-symbolic systems¹ The proposed taxonomy consists of a flat set of 6 types. We will briefly summarise (our understanding of) these informally characterised types (partly based on the explanation in [17]), and we will show how Kautz’s types relate to the design patterns that we will propose in this paper. Type 1 systems (informally written by Kautz as “symbolic Neuro symbolic”) are learning systems that have symbols as input and output. This directly corresponds to our elementary pattern Fig. 3b. Type 2 systems (informal notation: “Symbolic[Neuro]”) are symbolic reasoning systems that are guided by a

¹<https://www.cs.rochester.edu/u/kautz/talks/index.html>

learned search strategy. These directly correspond to a variation of our pattern (Fig. 11). Type 3 systems (informal notation “Neuro;Symbolic”) consist of a sequence of a neural learning system that performs abstraction from data to symbols, followed by a symbolic system. This corresponds to our patterns (Fig. 6a and b), showing that in this case we make a more fine-grained distinction. Type 4 systems (informal notation “Neuro:Symbolic → Neuro”) use symbolic I/O training pairs to teach a neural system. These correspond partly to our elementary pattern

(Fig. 3b) (for example: inductive logic programming), partly to our pattern (Fig. 8) (eg link prediction) and partly to pattern (Fig. 10) (eg learning to reason), again showing that we propose a much more fine-grained distinction. Type 5 systems (informal notation “Neuro”) use symbolic rules that inform neural learning. These correspond to our pattern Fig. 7. Finally, Type 6 systems (informally “Neuro[Symbolic]”) remains somewhat unclear (as also acknowledged in [17]), and we refrain from interpreting this type.

Kautz types:	Type 1	Type 2	Type 3	Type 4	Type 5	Type 6
Our Patterns:	3b	11	6a,6b	3b, 8, 10	7	-

The above table shows that there are substantial differences between our proposed design patterns and the system types from Kautz. Kautz’s taxonomy has similar goals to ours, namely to identify different interaction patterns between neural and symbolic components in a modular hybrid architecture, but our proposal goes beyond Kautz’s proposal because (a) Kautz proposes a taxonomy of systems without describing the internal architectures of the types of systems in his taxonomy, and (b) we make more fine-grained distinctions than Kautz, refining his 6 categories into distinctive subtypes, each with their own internal modular architecture (= design pattern).

In [49] the authors survey hybrid (“neural-symbolic”) systems along eight different dimensions. We briefly describe each of these, and discuss their relationship to the distinction made in our own work.

Directed vs undirected graphical models. This is a finer grained distinction about representations than we make.

In our patterns, these are both captured by the same “semantic model” component

Model-based vs. proof-based inference (better known as model-theoretic vs. proof-theoretic inference). Note this use of “model-based” is using “model” as the term is used by logicians, which is unlike how Darwiche uses the term model-based, using the term model as used by machine learning, showing the highly ambiguous use of the term “model” in Computer Science in general and in AI in particular. Again, this is a finer grained distinction than we make, and again both of these forms of inference are captured in our single KR component.

Logic vs. Neural. This corresponds to our distinction between ML and KR components

Boolean vs. probabilistic semantics. Similar to above, both of these would be captured by the KR component without making the distinction

Structure vs. parameter learning. This is captured in our notation by ML components that have either a statistical or a semantic model as their result

Symbols vs. Sub-symbols. This corresponds to our distinction between symbols and data. Unfortunately, and similar to our work, De Raedt et al. do not give a precise distinction between the two categories.

Type of Logic. This is another finer distinction than we make, these different types of logic are all captured in terms of our KR component

Summarising, on the one hand, De Raedt et al. make a number of finer distinctions than our boxology, mostly in terms of the details inside our components (different variations of KR components, different variations of models), while on the other hand de Raedt et al. do not discuss how these components should be configured into larger systems in order to achieve a particular functionality, which is the goal of our boxology. Whereas our boxology is a refinement of the 6 types proposed by Kautz (both aiming to describe modular architectures of interacting components), the work by de Raedt et al. is a refinement of some of our components, and could be combined with our work in a future version. The same is true for [61].

7 Conclusion and future work

In our paper, we have presented a visual language (boxology) to represent learning and reasoning systems. The taxonomical vocabulary and a collection of patterns expressed in this language aim to foster a better understanding of Hybrid AI systems and support communication between AI communities. A practical application in two use cases demonstrates that we can use the boxology to create a communicable blueprint of rather complex Hybrid AI systems that integrate various AI technologies.

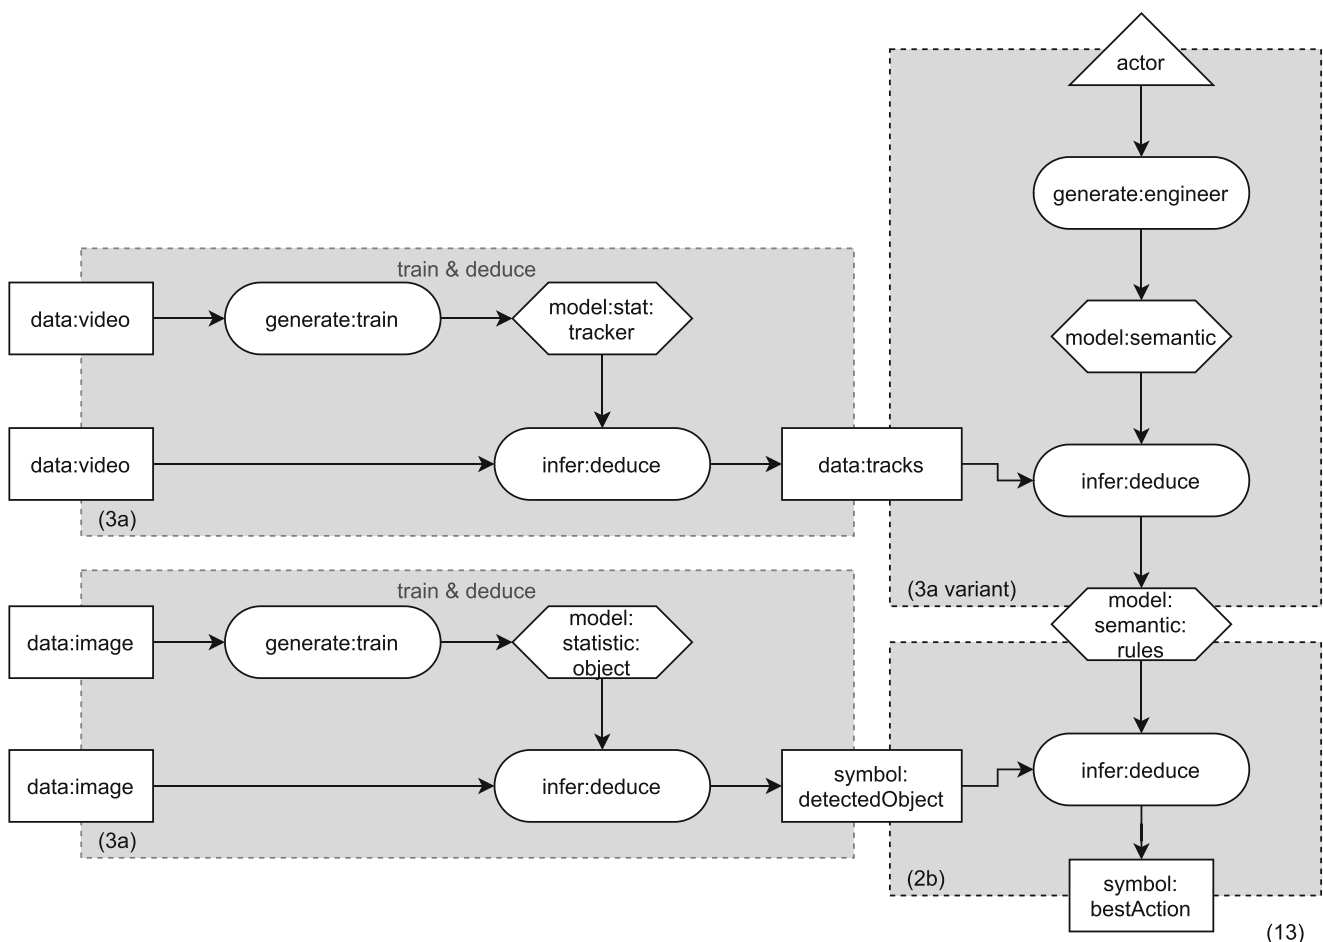


Fig. 13 Use-case for robot in action

The work presented here provides ample opportunities for additional features and uses. We expect to apply the taxonomy and visual language in many more use cases and is likely to evolve further as a result. New examples of AI systems will contribute to extending and improving the taxonomy, which in turn allows us to cover more use cases. Using this approach, an increasingly more mature visual language will evolve.

As a first extension to the current boxology, the concept of actors can be defined, along with the corresponding interaction processes and models. Actors are necessary for modelling interactions among autonomous entities, such as software agents or robots, whether they are physically or logically distributed. They also allow for specifying systems with humans in the loop and human-machine interaction in general. Use cases for actors include federated learning and reasoning, multi-robot coordination or hybrid human-agent teams. In [64] the authors propose an extension of our boxology [60] with two abstract patterns for humans-in-the-loop systems, namely where the human agent either

performs the role of a feedback-provider or a feedback-consumer.

Future work also includes developing the boxology from a means of representing system functionality towards an architectural tool-set of reusable components for design, implementation and deployment of hybrid AI systems. A more coherent methodology for complex AI systems based on the boxology allows these systems to be easier to understand in terms of functionality. This in turn provides a basis for more explainable and trustworthy AI systems design. An interesting topic to pursue in this respect is the creation and development of a generative grammar and logic calculus for composing and verifying patterns. This would facilitate the above-mentioned goals of and allow for formal verification at the component, pattern and system levels.

When using a coherent methodology for complex Hybrid AI system design, it is expected that such a design becomes easier to understand and maintain. In addition, hybrid AI systems will become better explainable, responsible, reliable, and predictable. It is our aim to develop such

systems as being trustworthy by design. This could provide a framework for system quality control by evaluation and certification.

Finally, the methodology needs to be further documented using guidelines for specifying increasingly concrete implementations of the concepts.

Appendix: Taxonomy

Table 1 Hybrid AI taxonomy

Concept	Definition
Instance	An example or single occurrence of something. (Collins)
Data	Factual information (such as measurements or statistics) used as a basis for reasoning, discussion, or calculation. (MW)
Number	A numerical quantity that is assigned or is determined by calculation or measurement. (MW)
Text	Words and form of a written or printed work. Something (such as a story or movie) considered as an object to be examined, explicated, or deconstructed. (MW)
Tensor	A set of components, functions of the coordinates of any point in space, that transform linearly between coordinate systems. For three-dimensional space there are 3 r components, where r is the rank. A tensor of zero rank is a scalar, of rank one, a vector. (Collins) A picture (bitmap) can be represented as a tensor of rank 2.
Stream	A stream of things is a large number of them occurring one after another. (Collins) Digital data (such as audio or video material) that is continuously delivered one packet at a time and is usually intended for immediate processing or playback. (MW)
Symbol	Something that stands for or suggests something else by reason of relationship, association, convention, or accidental resemblance. An arbitrary or conventional sign used in writing or printing relating to a particular field to represent operations, quantities, elements, relations, or qualities (MW)
Label	A descriptive or identifying word or phrase (MW)
Trace	An aspect or quality (such as resemblance) that connects two or more things or parts as being or belonging or working together or as being of the same kind. (MW)
Model	A sign or evidence of some past thing. (MW)
Statistical Model	A system of postulates, data, and inferences presented as a mathematical description of an entity or state of affairs. (MW)
Semantic Model	A statistical model is usually specified as a mathematical relationship between one or more random variables and other non-random variables. (MW)
Processing	A conceptual model represents 'concepts' (entities) and relationships between them. (MW) Semantic technologies formally represent the meaning involved in information. For example, ontology can describe concepts, relationships between things, and categories of things. (MW)
Generation	A series of actions or operations conducting to an end. (MW)
Training	A process of coming or bringing into being. (MW) Defining or originating (something, such as a mathematical or linguistic set or structure) by the application of one or more rules or operations. (MW)
Engineering	The process of learning the skills that you need for a particular job or activity. (Collins) In particular: model building, the process of preparing a machine learning model to be useful by feeding it data from which it can learn, ie. detect statistical patterns.
Transformation	The application of science and mathematics by which the properties of matter and the sources of energy in nature are made useful to people. The design and manufacture of complex products. (MW) In particular, software engineering: the methodical design, implementation, and maintenance of models and components of software architectures.
Inference	The operation of changing (as by rotation or mapping) one configuration or expression into another in accordance with a mathematical rule. (MW)
Induction	A conclusion or opinion that is formed because of known facts or evidence. (A) The act of passing from one proposition, statement, or judgment considered as true to another whose truth is believed to follow from that of the former. (B) The act of passing from statistical sample data to generalisations, usually with calculated degrees of certainty. (MW)
Deduction	Inference of a generalised conclusion from particular instances. (MW)
Classification	The deriving of a conclusion by reasoning. Inference in which the conclusion about particulars follows necessarily from general or universal premises. (MW)
	systematic arrangement in groups or categories according to established criteria. (MW)

Table 1 (continued)

Concept	Definition
Prediction	To calculate (some future event or condition) usually as a result of study and analysis of available pertinent data. (MW)
Actor	Autonomous entity acting proactively to initiate processes, based on intentions and goals. Interaction among actors leads to emergent system behaviour.
Human	Human beings interacting with AI system components and other actors.
Agent	Active software components that are proactive, rather than merely reactive. Agents can be based on many internal reasoning mechanisms.
Robot	Physically embedded agents that can sense their environment and act in it.

Acknowledgements This research was partially funded by the Hybrid Intelligence Center, a 10-year programme funded by the Dutch Ministry of Education, Culture and Science through the Netherlands Organisation for Scientific Research NWO, <https://hybrid-intelligence-centre.nl>.

Author Contributions Authors have contributed equally, authors are listed alphabetically.

Funding The work by Van Harmelen and Ten Teije is part of the research programme Hybrid Intelligence with project number 024.004.022, which is (partly) financed by the Dutch Ministry of Education, Culture and Science (OCW).

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

- Arora S, Bedathur S, Ramanath M, Sharma D (2020) Iterefine: Iterative KG refinement embeddings using symbolic knowledge. CoRR, arXiv:2006.04509
- Asai M (2019) Unsupervised grounding of plannable first-order logic representation from images. In: Benton J, Lipovetzky N, Onaindia E, Smith DE, Srivastava S (eds) Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling, ICAPS 2018. AAAI Press, Berkeley, pp 583–591
- Asim MN, Wasim M, Khan MUG, Mahmood W, Abbasi HM (2018) A survey of ontology learning techniques and applications. Database 2018
- Bach SH, Broecheler M, Huang B, Getoor L (2017) Hinge-loss markov random fields and probabilistic soft logic. J Mach Learn Res 18:109:1–109:67
- Baier S, Ma Y, Tresp V (2017) Improving visual relationship detection using semantic modeling of scene descriptions. In: d'Amato C, Fernández M, Tamma V. A. M., Lécué F, Cudré-Mauroux P, Sequeda J. F., Lange C, Heflin J (eds) The semantic web - ISWC 2017 - 16th International semantic web conference, proceedings, Part I, volume 10587 of Lecture notes in computer science. Springer, Vienna, pp 53–68
- Battaglia PW, Hamrick JB, Bapst V, Sanchez-Gonzalez A, Zambaldi VF, Malinowski M, Tacchetti A, Raposo D, Santoro A, Faulkner R, Gülçehre Ç, Song F, Ballard AJ, Gilmer J, Dahl GE, Vaswani A, Allen K, Nash C, Langston V, Dyer C, Heess N, Wierstra D, Kohli P, Botvinick M, Vinyals O, Li Y, Pascanu R (2018) Relational inductive biases, deep learning, and graph networks. CoRR, arXiv:1806.01261
- Berkeley I (2008) What the $< 0.70, 1.17, 0.99, 1.07 >$ is a symbol? Minds Mach 18:93–105
- Bian J, Gao B, Liu T (2014) Knowledge-powered deep learning for word embedding. In: Calders T, Esposito F, Hüllermeier E, Meo R (eds) Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2014, Nancy, France, September 15–19, 2014. Proceedings, Part I, volume 8724 of Lecture Notes in Computer Science, pp 132–148
- Borgwardt K, Ghisu E, Llinares-López F, O'Bray L, Rieck B (2020) Graph kernels: State-of-the-art and future challenges
- Bouraoui Z, Jameel S, Schockaert S (2017) Inductive reasoning about ontologies using conceptual spaces. In: Singh SP, Markovitch S (eds) Proceedings of the Thirty-First AAAI conference on artificial intelligence. AAAI Press, San Francisco, pp 4364–4370
- Brewster CA (2008) Mind the gap : bridging from text to ontological knowledge. PhD thesis, University of Sheffield, UK
- Cimiano P, Mädche A, Staab S, Völker J (2009) Ontology learning. In: Handbook on ontologies. Springer, 245–267
- Ciravegna G, Giannini F, Gori M, Maggini M, Melacci S (2020) Human-driven FOL explanations of deep learning. In: Bessiere C (ed) Proceedings of the twenty-ninth international joint conference on artificial intelligence, IJCAI 2020, pp 2234–2240. ijcai.org
- Cohen WW, Sun H, Hofer RA, Siegler M (2020) Scalable neural methods for reasoning with a symbolic knowledge base. arXiv:2002.06115
- Costantini S (2002) Meta-reasoning: A survey. In: Kakas AC, Sadri F (eds) Computational Logic: Logic Programming and Beyond, Essays in Honour of Robert A. Kowalski, Part II, volume 2408 of Lecture Notes in Computer Science. Springer, pp 253–288
- Darwiche A (2018) Human-level intelligence or animal-like abilities? Commun ACM 61(10):56–67
- d'Avila Garcez A, Lamb LC (2020) Neurosymbolic AI: The 3rd wave
- de Boer MH, Verhoosel JP (2019) Creating and evaluating data-driven ontologies. Int J Adv Softw 12(3 and 4):300–309
- Demeester T, Rocktäschel T, Riedel S (2016) Lifted rule injection for relation embeddings. In: Su J, Carreras X, Duh K (eds) Proceedings of the 2016 conference on empirical methods in natural language processing, EMNLP. The Association for Computational Linguistics, Austin, pp 1389–1399

20. Detassis F, Lombardi M, Milano M (2020) Teaching the old dog new tricks Supervised learning with constraints. CoRR, arXiv:2002.10766
21. Donadello I, Serafini L, d'Avila Garcez AS (2017) Logic tensor networks for semantic image interpretation. In: Sierra C (ed) Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017. ijcai.org, Melbourne, pp 1596–1602
22. Ebrahimi M, Sarker MK, Bianchi F, Xie N, Doran D, Hitzler P (2018) Reasoning over RDF knowledge bases using deep learning. CoRR, arXiv:1811.04132
23. Fang M, Zhou T, Du Y, Han L, Zhang Z (2019) Curriculum-guided hindsight experience replay. In: Wallach HM, Larochelle H, Beygelzimer A, d'Alché-Buc F, Fox E. B., Garnett R. (eds) Advances in neural information processing systems 32: Annual conference on neural information processing systems 2019, NeurIPS 2019, Vancouver, pp 12602–12613
24. Fournier P, Sigaud O, Chetouani M, Oudeyer P (2018) Accuracy-based curriculum learning in deep reinforcement learning. CoRR, arXiv:1806.09614
25. Gamma E, Helm R, Johnson R, Vlissides JM (1994) Design patterns: elements of reusable object-oriented software, 1st edn. Addison-Wesley Professional, Boston
26. Garcez A. d., Gori M, Lamb LC, Serafini L, Spranger M, Tran SN (2019) Neural-symbolic computing: An effective methodology for principled integration of machine learning and reasoning. arXiv:1905.06088
27. Garnelo M, Arulkumaran K, Shanahan M (2016) Towards deep symbolic reinforcement learning. CoRR, arXiv:1609.05518
28. Geibel P (2006) Reinforcement learning for mdps with constraints. In: Fürnkranz J, Scheffer T, Spiliopoulou M (eds) Machine Learning: ECML 2006, 17th European conference on machine learning, 2006, Proceedings, volume 4212 of Lecture Notes in Computer Science. Springer, Berlin, pp 646–653
29. Getoor L (2013) Probabilistic soft logic: A scalable approach for markov random fields over continuous-valued variables - (abstract of keynote talk). In: Morgenstern L, Stefanescu PS, Lévy F, Wyner AZ, Paschke A (eds) Theory, practice, and applications of rules on the Web - 7th International Symposium, RuleML 2013, Proceedings, volume 8035 of Lecture Notes in Computer Science. Springer, Seattle, p 1
30. Hohenecker P, Lukasiewicz T (2017) Deep learning for ontology reasoning. CoRR, arXiv:1705.10342
31. Hutter F., Kotthoff L., Vanschoren J. (eds) (2019) Automatic machine learning: methods, systems, challenges. Springer, Berlin
32. Icarte RT, Klassen TQ, Valenzano RA, McIlraith SA (2018) Teaching multiple tasks to an RL agent using LTL. In: André E, Koenig S, Dastani M, Sukthankar G (eds) Proceedings of the 17th international conference on autonomous agents and multiagent systems, AAMAS 2018, International foundation for autonomous agents and multiagent systems. ACM, Stockholm, pp 452–461
33. Illanes L, Yan X, Icarte RT, McIlraith SA (2020) Symbolic plans as high-level instructions for reinforcement learning. In: Beck JC, Buffet O, Hoffmann J, Karpas E, Sohrabi S (eds) Proceedings of the thirtieth international conference on automated planning and scheduling. AAAI Press, Nancy, pp 540–550
34. Inoue K, Ohwada H, Yamamoto A (2017) Special issue on inductive logic programming
35. Konstantopoulos S, Charalambidis A (2010) Formulating description logic learning as an inductive logic programming task. In: FUZZ-IEEE 2010, IEEE International Conference on Fuzzy Systems, 2010, Proceedings. IEEE, Barcelona, pp 1–7
36. Kop R et al (2016) And Predictive modeling of colorectal cancer using a dedicated pre-processing pipeline on routine electronic medical records. Comp Bio Med 76:30–38
37. Lecun Y, Bottou L, Bengio Y, Haffner P (1998) Gradient-based learning applied to document recognition. Proc IEEE 86(11):2278–2324
38. Liello LD, Ardino P, Gobbi J, Morettin P, Teso S, Passerini A (2020) Efficient generation of structured objects with constrained adversarial networks. CoRR, arXiv:2007.13197
39. Loos SM, Irving G, Szegedy C, Kaliszcyk C (2017) Deep network guided proof search. In: Eiter T, Sands D (eds) LPAR-21, 21st International conference on logic for programming, artificial intelligence and reasoning, volume 46 of EPiC Series in Computing. EasyChair, Maun, pp 85–105
40. Manhaeve R, Dumancic S, Kimmig A, Demeester T, Raedt LD (2018) Deepprolog: Neural probabilistic logic programming. In: Bengio S, Wallach HM, Larochelle H, Grauman K, Cesa-Bianchi N, Garnett R (eds) Advances in neural information processing systems 31: Annual conference on neural information processing systems 2018, NeurIPS 2018, Montréal, pp 3753–3763
41. Marra G, Diligenti M, Giannini F, Gori M, Maggini M (2020) Relational neural machines. In: Giacomo G. D., Catalá A, Dilkina B, Milano M, Barro S, Bugarín A, Lang J (eds) ECAI 2020 - 24th European conference on artificial intelligence, 2020 - Including 10th conference on prestigious applications of artificial intelligence (PAIS 2020), volume 325 of Frontiers in artificial intelligence and Applications. IOS Press, Santiago de Compostela, pp 1340–1347
42. Marra G, Giannini F, Diligenti M, Gori M (2019) LYRICS: A general interface layer to integrate logic inference and deep learning. In: Brefeld U, Fromont É, Hotho A, Knobbe AJ, Maathuis MH, Robardet C (eds) Machine learning and knowledge discovery in databases - European conference, ECML PKDD 2019, Proceedings, Part II, volume 11907 of Lecture notes in computer science. Springer, Würzburg, pp 283–298
43. Meilicke C, Chekol MW, Ruffinelli D, Stuckenschmidt H (2019) Anytime bottom-up rule learning for knowledge graph completion. In: Proceedings of the twenty-eighth international joint conference on artificial intelligence, IJCAI-19. International Joint Conferences on Artificial Intelligence Organization, pp 3137–3143
44. Nickel M, Murphy K, Tresp V, Gabrilovich E (2016) A review of relational machine learning for knowledge graphs. Proc IEEE 104(1):11–33
45. Padiá A, Martin D, Patel-Schneider PF (2018) Automating class/instance representational choices in knowledge bases. In: Faron-Zucker C, Ghidini C, Napoli A, Toussaint Y (eds) Knowledge engineering and knowledge management - 21st International conference, EKAW 2018, 2018 Proceedings, volume 11313 of Lecture Notes in Computer Science. Springer, pp 273–288
46. Paulheim H (2017) Knowledge graph refinement: A survey of approaches and evaluation methods. Semant Web 8(3):489–508
47. Pearl J (2018) Theoretical impediments to machine learning with seven sparks from the causal revolution. In: Chang Y, Zhai C, Liu Y, Maarek Y (eds) Proceedings of the Eleventh ACM International conference on web search and data mining, WSDM 2018. ACM, Marina Del Rey, p 3
48. Perlis D, Brody J, Kraus S, Miller MJ (2017) The internal reasoning of robots. In: Gordon AS, Miller R, Turán G (eds) Proceedings of the Thirteenth International Symposium on Commonsense Reasoning, COMMONSENSE, volume 2052 of CEUR Workshop Proceedings. CEUR-WS.org, London, p 2017
49. Raedt LD, Dumancic S., Manhaeve R, Marra G (2020) From statistical relational to neuro-symbolic artificial intelligence. In: Bessiere C (ed) Proceedings of the twenty-ninth international joint conference on artificial intelligence, IJCAI 2020. ijcai.org, pp 4943–4950
50. Richardson M, Domingos P (2006) Markov logic networks. Mach Learn 62(1-2):107–136

51. Rocktäschel T, Riedel S, Vishwanathan SVN (2017) End-to-end differentiable proving. In: Guyon I, von Luxburg U, Bengio S, Wallach HM, Fergus R, Garnett R (eds) *Advances in neural information processing systems 30: Annual conference on neural information processing systems 2017*, 4–9 December 2017, Long Beach, CA, USA, pp 3791–3803
52. Sala F, Sa CD, Gu A, Ré C (2018) Representation tradeoffs for hyperbolic embeddings. In: Dy JG, Krause A (eds) *Proceedings of the 35th international conference on machine learning, ICML 2018*, volume 80 of *Proceedings of machine learning research*. PMLR, Stockholm, pp 4457–4466
53. Sarker MK, Xie N, Doran D, Raymer M, Hitzler P (2017) Explaining trained neural networks with semantic web technologies: First steps. In: Besold TR, d’Avila Garcez AS, Noble I (eds) *Proceedings of the twelfth international workshop on neural-symbolic learning and reasoning, NeSy 2017*, volume 2003 of *CEUR workshop proceedings*. CEUR-WS.org, London
54. Serafini L, Garcez A. d. (2016) Logic tensor networks: Deep learning and logical reasoning from data and knowledge. arXiv:1606.04422
55. Shang Z, Zraggen E, Kraska T (2019) Alpine meadow: A system for interactive automl. In: *NEURIPS*
56. Silver D, Schrittwieser J, Simonyan K, Antonoglou I, Huang A, Guez A, Hubert T, Baker L, Lai M, Bolton A, Chen Y, Lillicrap TP, Hui F, Sifre L, van den Driessche G, Graepel T, Hassabis D (2017) Mastering the game of go without human knowledge. *Nat* 550(7676):354–359
57. Silvestri M, Lombardi M, Milano M (2020) Injecting domain knowledge in neural networks: a controlled experiment on a constrained problem. *CoRR*, arXiv:2002.10742
58. Suttner CB, Ertel W (1990) Automatic acquisition of search guiding heuristics. In: Stickel ME (ed) *10th international conference on automated deduction, Kaiserslautern, FRG, 1990 Proceedings*, volume 449 of *Lecture Notes in Computer Science*. Springer, pp 470–484
59. Tididi I, d’Aquin M, Motta E (2015) Data patterns explained with linked data. In: Bifet A, May M, Zadrozny B, Gavaldà R, Pedreschi D, Bonchi F, Cardoso JS, Spiliopoulou M (eds) *Machine learning and knowledge discovery in databases - European conference, ECML PKDD 2015, 2015 Proceedings, Part III*, volume 9286 of *Lecture Notes in Computer Science*. Springer, Porto, pp 271–275
60. van Harmelen F, ten Teije A (2019) A boxology of design patterns for hybrid learning and reasoning systems. *J Web Eng* 18(1-3):97–124
61. Von Rueden L, Mayer S, Garcke J, Bauckhage C, Schuecker J (2019) Informed machine learning—towards a taxonomy of explicit integration of knowledge into machine learning. *Learning* 18:19–20
62. Wang Q, Mao Z, Wang B, Guo L (2017) Knowledge graph embedding: A survey of approaches and applications. *IEEE Trans Knowl Data Eng* 29(12):2724–2743, 12
63. Weld DS, Bansal G (2018) Intelligible artificial intelligence. *CoRR*, arXiv:1803.04263
64. Witschel HF, Pande C, Martin A, Laurenzi E, Hinkelmann K (2021) *Visualization of patterns for hybrid learning and reasoning with human involvement*. Springer International Publishing, Cham, pp 193–204
65. Wong W, Liu W, Bennamoun M (2012) Ontology learning from text: A look back and into the future. *ACM Comput Surv (CSUR)* 44(4):1–36
66. Xu J, Zhang Z, Friedman T, Liang Y, den Broeck GV (2018) A semantic loss function for deep learning with symbolic knowledge. In: Dy JG, Krause A (eds) *Proceedings of the 35th international conference on machine learning, ICML 2018*, volume 80 of *JMLR Workshop and conference proceedings*. JMLR.org, Stockholm, pp 5498–5507
67. Zhang W, Paudel B, Wang L, Chen J, Zhu H, Zhang W, Bernstein A, Chen H (2019) Iteratively learning embeddings and rules for knowledge graph reasoning. In: Liu L, White RW, Mantrach A, Silvestri F, McAuley JJ, Baeza-Yates R, Zia L (eds) *The World Wide Web Conference, WWW 2019*. ACM, San Francisco, pp 2366–2377
68. Zhou J, Cui G, Zhang Z, Yang C, Liu Z, Wang L, Li C, Sun M (2018) *Graph neural networks: A review of methods and applications*. arXiv:1812.08434

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



Dr. Michael van Bekkum (MSc) is a scientist at TNO, The Netherlands, with a master’s degree in electrotechnical engineering from Twente University. His interest is in engineering human-inspired AI, i.e. bringing human mechanisms of knowledge representation, learning and reasoning to AI systems.



Dr. Maaike de Boer is a scientist at TNO, The Netherlands, and obtained her PhD at the Radboud University (Nijmegen). Her interests are text mining as well as combining data-driven methods and knowledge-driven methods (such as ontologies) with the goal to get the best of both worlds.



Prof. Frank van Harmelen is professor of Knowledge Representation and Reasoning at the Vrije Universiteit Amsterdam, with a PhD from the University of Edinburgh. He is also adjunct professor at Wuhan University and Wuhan University of Science and Technology. He has a specific interest in combining symbolic and statistical methods in AI and is principle investigator the Hybrid Intelligence project (<http://hybrid-intelligence-center.nl/>).



Dr. André Meyer-Vitali is a senior scientist at TNO, The Netherlands, and obtained his PhD at the University of Zürich, Switzerland. His interests are in software engineering, distributed systems/AI, social simulation, multi-agent systems, knowledge representation and reasoning and combinations with learning and self-organisation.



Dr. Annette ten Teije is an associate professor at the Vrije Universiteit Amsterdam, and obtained her PhD at the University of Amsterdam. Her interests are in knowledge modelling, representation and reasoning in particular in the medical domain.