

Abstraction for Deep Reinforcement Learning

Murray Shanahan^{1,2} and Melanie Mitchell³

¹DeepMind

²Imperial College London

³Santa Fe Institute

m.shanahan@imperial.ac.uk, mm@santafe.edu

Abstract

We characterise the problem of abstraction in the context of deep reinforcement learning. Various well established approaches to analogical reasoning and associative memory might be brought to bear on this issue, but they present difficulties because of the need for end-to-end differentiability. We review developments in AI and machine learning that could facilitate their adoption.

1 Introduction

Consider an embodied agent, either an animal (human or non-human), a physical robot, or a virtual agent in a simulated environment. Suppose the agent is exposed to a certain amount of experience from which it can learn. At every moment, the agent must address the question “How can I shape my future experience to my liking?” by answering the question “How does my present experience resemble what I have experienced in the past?”¹ How well an agent does this — that is to say how well it *generalises* from past experience — is one measure of its intelligence.

In contemporary deep reinforcement learning (DRL), answers to these two questions are incorporated in an agent’s learned policy, a function that maps observations to agent actions (or distributions over actions), maximising cumulative discounted reward. This function is realised by a deep neural network, and generalisation from past experience is achieved through the standard tools of backpropagation and stochastic gradient descent. DRL has yielded impressive results, and must surely count as one of machine learning’s most important advances of the past decade. However, in its current guise, DRL inherits the well-known shortcomings of contemporary neural network methods, which include low sample efficiency, limited transfer ability, and poor out-of-distribution generalisation [Garnelo *et al.*, 2016; Lake *et al.*, 2017].

We believe these shortcomings result from an inability to form sufficiently general abstractions. However, we see no in-principle barrier to developing DRL agents that are endowed

¹Or, to be more precise, the mechanism that underlies the agent’s choice of actions must implicitly address the first question by answering the second.

with a greater capacity for abstraction, enabling them to generalise better from past experience [Mitchell, 2021]. The purpose of this paper is to review past and present work in artificial intelligence that has the potential to advance this aim. We will discuss a number of relevant symbolic and neural-network approaches that purport to yield more abstract, more general representations. For each approach, we will make the case for its applicability, review progress, and identify the challenges it poses.

2 Abstraction: Framing the Problem

A human being’s capacity to shape the future according to their liking via abstraction exceeds anything we can achieve in artificial intelligence today. The fundamental cognitive operation behind this ability is that of *analogy-making* or *seeing similarity*. To see a similarity is to identify certain respects in which one thing (a state or a sequence) is like another, and to conclude that they are alike in certain other respects. Seeing similarities (analogies) among a *set* of things is the basis for a trio of further fundamental cognitive operations: 1) forming a concept of which those things are all instances, 2) seeing that something is an instance of a concept so formed, and 3) applying the concept to infer something new about the thing in question. We will use the umbrella term *abstraction* to denote this cluster of operations: seeing similarity (analogy-making), forming a concept, and applying a concept.

2.1 Common Sense and Transfer

Abstraction, as conceived here, is the key to transfer. Things that appear different on a superficial level can present similarities on a deeper, more abstract, level, similarities that can be exploited in the service of transfer. To achieve this, an agent must be able to form concepts that can be usefully applied in contexts different from those in which they were acquired. In other words, the domain of a concept’s application must be larger than the domain of its acquisition. The larger and more diverse the domain of a concept’s application is, compared to the domain of its acquisition, the wider the *transfer gap* it enables an agent to cross.

The concepts with the largest and most diverse application are those belonging to “foundational common sense” [Spelke, 2000; Shanahan *et al.*, 2020]. Humans and, to a lesser degree, other animals possess a repertoire of these fundamental concepts that includes, in the domain of everyday physics, such

concepts as object, path, obstacle, portal, container, and so on, and, in the social domain, such concepts as other agents, being with others, meeting, giving, taking, helping, and so on. Foundational common sense concepts, though acquired through basic sensorimotor interactions with the everyday world, also structure our thinking at a more abstract level [Lakoff and Johnson, 1980]. Life is a journey; every career has its obstacles; relationships are a matter of give and take.

Let’s consider an example. Josie is five years old. Most weekdays, she goes to the nursery with her mother, which entails leaving the house, getting in the car, a short drive, getting out of the car, and passing through the nursery entrance. At the weekend, Josie often goes to the shops with her father, which entails leaving the house, walking up the hill and round the corner, and entering the supermarket through the sliding doors. From these and other similar experiences, Josie has distilled the concept of a journey. This concept captures a recurring pattern, a patchwork of similarities, in certain of the episodes that make up Josie’s life. These episodes typically involve a point of departure, a period of motion along a path, a variety of obstacles, and a destination.

Each of the elements that features in Josie’s concept of a journey can be instantiated in many ways; the place of departure could be a house or a shop, the movement might involve a car or not, the obstacles might include a door or the corner of a building. Josie’s concept of a journey is indifferent to many other details too: the colour of the car, who else is in the street, and so on. Primarily Josie’s concept of a journey captures narrative structure: the events involved, their types, and their order. Thanks to its concern with narrative structure and its indifference to detail, this concept can usefully be applied in new and unfamiliar circumstances. Consequently, when Josie plays a new computer game that involves retrieving a magic ring from a far off castle, she knows to look for a means of transport. Perhaps the unicorn will carry her there.

2.2 The Limitations of Today’s DRL Agents

A major obstacle to achieving artificial general intelligence would be overcome if we knew how to realise abstraction in DRL agents. Currently, we don’t know how to endow them with the fundamental cognitive operations characterised in the previous section: seeing similarity (analogy-making), forming a concept, and applying a concept. Of course, our current deep learning systems do achieve a degree of generalisation. If a model is trained on a dataset drawn iid from some distribution, and the resulting model performs well on a separate test set that is drawn iid from the same distribution, then we can justifiably claim that it has generalised from the training set to the distribution as a whole. Moreover, there are numerous research papers describing systems that are claimed to achieve transfer, or systematic generalisation, or generalisation to held out combinations of features [Kirk *et al.*, 2021]. However, in all cases, the type of transfer achieved is severely limited.

Voudouris *et al.* [2022] explicitly compare the performance of children and agents submitted to the 2019 Animal-AI competition [Crosby *et al.*, 2020]. In this competition, developers submitted agents for testing on a battery of tasks inspired by the animal cognition literature (Fig. 1). Importantly,




Figure 1: An example Animal-AI task. The green sphere is the reward item. The subject (human or AI) has to push the box into the pillar to knock the green sphere off.

the test tasks were not revealed to the developers in advance. Instead, developers were supplied with a 3D environment to train their agents in, comprising an arena and an inventory of objects that included green and yellow spheres as rewarding “food” items. They were also given an indication of the general types of tasks their agents would be required to solve to obtain the food items. In the work of Voudouris *et al.*, subjects aged between six and 10 years were asked to solve a subset of these tasks of varying degrees of difficulty. They found that the children were able to solve all the tasks, significantly out-performing the AI agents submitted to the 2019 competition on all but the easiest levels.

Of course, the children and the AI agents differed enormously in the previous experience they brought to bear on these tasks. The children benefited from several years’ experience observing and interacting with the real world in all its variety and complexity, not to mention the innate endowment of human evolutionary history. By contrast, the AI agents’ only experience was of a very limited virtual world containing very few types of object (and no other agents). Nevertheless, it’s instructive to consider the considerable transfer gaps the children were able to cross to solve tasks that defeated the AI agents. In one task, for example, the food item rested on a pillar too high for the subject to reach (Fig. 1). However, it was possible to knock the food off the pillar by pushing a tall box into it. 42% of the children (and 60% of the ten-year olds) solved this task, whereas none of the AI agents did.

To pin down the transfer gap here, we need to know what relevant prior experience the children had. Most relevantly, they were shown a two-minute video that introduced the Animal-AI environment, and they were allowed to play a series of tutorial levels. However, neither the video nor the tutorial levels included objects on pillars. On the other hand, the children will certainly have encountered similar situations in the real world. They will have had the experience of wanting something that is out of reach. They will have had the experience of poking or pushing one object with another to make the first object move. They may or may not have faced a directly analogous situation, such as a ball stuck high in a tree but within reach of a stick. Either way, the children’s success will have depended on a formidable capac-

ity to map their previous experience onto a novel scenario in an unfamiliar virtual environment [Tsvividis *et al.*, 2017; Dubey *et al.*, 2018].

2.3 Exploiting Similarity for Generalisation

We would like to be able to emulate this capacity in a DRL agent. To see what this could mean in practice, let’s imagine how such an agent might solve the pillar task zero-shot. Pending the actual development of an agent that fits the description, this will necessarily be a speculative “just-so” story. But it will help us pin down the specification of the sort of architecture we are aiming for.

Suppose the agent is presented with the view we see in Fig. 1. The raw image is encoded into a representation at the level of objects and their features and relations, which triggers the retrieval or “activation” of a set of concepts that match the agent’s situation. Although the agent has never before had to get hold of a sphere perched on tall pillar, it has frequently attempted to obtain some target object that is not immediately to hand, either successfully or unsuccessfully, and the attempt has often involved overcoming (or trying to overcome) some difficulty in order to bring the object within its grasp. The agent has, accordingly, formed a concept that captures the similarities among all these experiences. This concept can be thought of as a *template* characterised by certain actors and objects, namely the agent and the target, along with landmark states and prototypical sequences of events.²

To activate a concept, on this view, is to match a template to the situation at hand by binding (a subset of) its actors, objects, states, and events to actual actors, objects, states, and events. Some of the states and events will be in the past, and some will be in the future, suggesting possible courses of action. In this case the actors and objects are, respectively, the agent itself and the green sphere. The landmark states include the present situation with the sphere on the pillar and the agent some way away from it, and a future state where the agent is in possession of the sphere. The prototypical sequence of events will include movement towards the green sphere. But the disadvantageous position of the sphere means that this template alone isn’t enough to characterise a future that makes sense, that coheres with the agent’s model of intuitive physics. When the agent arrives at the pillar, the sphere isn’t going to magically fall at its feet. So there’s a gap between the final state according to the template and the intermediate state of being next to the pillar.

To fill this gap, another concept needs to be activated. Something has to happen to knock the sphere off the pillar. But on many occasions in its past, the agent has made one object move by prodding or pushing it with another object. It has poked spheres with sticks, shoved boxes against each other, and so on. So the agent has formed a concept that distils the similarities among these experiences, the template for a “small spatial story” [Turner, 1998]. The sole actor in this

²Cf: Schank’s concept of a *script* [1977]. Note that we are not implying that templates are structures that exist explicitly in a machine’s memory. Rather, they are a way of characterising flexible, adaptive *behaviour*. In terms of underlying implementation, templates could be represented in a compressed and / or distributed form.

template is again the agent itself, and it features a tool object that the agent holds plus a target object to which that tool is applied. The prototypical sequence of events comprises the act of applying the tool object to the target object and the consequent movement of the target.

Binding the tool object to the box and the target object to the green sphere results in a match for this template, which in turn helps to fill the gaps in the overall template. Further concepts representing similar stereotypical narratives are activated to complete the story (the box must be retrieved and carried to the pillar). The resulting set of filled-in templates is blended together to form a coherent whole, an executable plan leading from the current state, with the sphere some distance away on the pillar, to a state with high value in terms of expected reward, with the sphere on the ground near the agent.

3 Approaches to Abstraction

To reiterate, the account just presented is purely speculative. We don’t yet know how to build a system capable of such abstraction, because we don’t know how to realise the underlying cognitive operations of seeing similarity, concept formation, and concept application in a machine. However, the field of artificial intelligence has developed a plethora of relevant techniques. The aim of the following sections is to present an inventory of techniques with the potential to realise the cognitive operations required for abstraction, drawing on ideas from the field’s past and on recent developments in machine learning.

3.1 Structure Mapping

Gentner and collaborators proposed that the analogical-mapping process underlying abstraction is one of *structure mapping*; human analogy-makers map systematic conceptual *structures* rather than unconnected components [Gentner, 1983; Falkenhainer *et al.*, 1989]. The structures in question can be represented as graphs whose nodes represent *objects*, *attributes*, and *actions*, and whose edges represent *relations*. A situation (or state) encountered by an agent can be mapped to a previously encountered situation by 1) generating graphs corresponding to the two situations, and 2) mapping one graph to the other by finding a set of systematic correspondences between similar structures of nodes and edges. For example, the situation depicted in (Fig. 1) could be represented by a graph including the sphere (food), pillar, and box, along with their attributes (e.g. colour) and spatial relationships, as well as the goal of the agent to obtain the food. A prior situation triggered in memory might include desired food on top of a fence, which the agent shook by pushing a rock into it, thus causing the sphere to drop. This situation could be similarly represented as a graph, in which the system of objects, attributes, actions, and relations would be mapped to the current-situation graph, allowing the agent to infer that pushing the box into the pillar might analogously yield a favourable result.

To form a concept given a set of similar states or sequences is to find an intermediate graph to which they can all be mapped in this way. Then, to apply that concept to a given

state is to 1) map the graph X representing the state to the graph Y representing the concept (matching), and 2) supplement X with nodes and edges (things and relations) present in Y but lacking in X (completion). A crucial feature of the matching operation is that it is not exact. There can be nodes in graph X (e.g., the colour of an object) for which there is no corresponding node in graph Y , and vice versa. Part of the task of cognition here is to learn what to ignore, what is and what is not relevant. Moreover, two nodes do not have to be identical to correspond, only to be similar. In this sense, seeing similarity is a recursive operation, which allows concepts to be organised into a hierarchy of abstraction by collapsing sub-graphs into nodes that feature in other graphs at a deeper level of abstraction.

The final inferential step (completion) is what makes a concept useful for guessing how one portion of a scene will look given the appearance of another, or for predicting the future from the past, or, more generally, for filling in the gaps when things are partially observable due to the constraints imposed on an agent by time and space. Given a repertoire of concepts, the task of cognition is to find a coherent subset of that repertoire that matches the sequence of states unfolding for an agent (retrieval), maximising the extent to which the agent can predict the future and, more specifically, can anticipate the effects of its actions. The greater the predictive power of its repertoire of concepts, the larger its spatial and temporal compass, the more the world can be said to “make sense” to the agent.

A significant difficulty with the structure mapping approach is the need to learn discrete representations, which are tricky to accommodate in an end-to-end differentiable processing pipeline [Crouse *et al.*, 2021]. Moreover, in the structure mapping approach, the cognitive operations of seeing similarity, forming a concept, and applying a concept take place after the world has been carved into discrete chunks with hard boundaries. Structure mapping assumes that a distinct “perceptual module” is responsible for generating the representations of the current state and prior memory. However, the representation and mapping processes are likely to be strongly interdependent in any intelligent agent; the best way to represent a new situation depends on the context of the prior situations it is mapped to.

The same consideration ought to apply to a DRL agent. As Brian Cantwell Smith asserts, “taking the world to consist of discrete intelligible mesoscale objects is an *achievement* of intelligence, not a premise on top of which intelligence runs” [Smith, 2019]. The problem would be mitigated if the task of discovering objects and relations in raw data could be integrated in an end-to-end differentiable DRL architecture (thanks to the top-down flow of influence, via gradient-based update, from reward through action to perception) [Garnelo and Shanahan, 2019]. Considerable recent progress along these lines has been reported in the deep learning literature [Burgess *et al.*, 2019; Greff *et al.*, 2019; Engelcke *et al.*, 2020; Locatello *et al.*, 2020; Creswell *et al.*, 2021], but it remains to be seen whether any of these methods are sufficiently general to support genuine abstraction.

3.2 Other Approaches to Analogy

Hofstadter *et al.*’s Active Symbol Architecture (ASA) for analogy-making [Hofstadter and Mitchell, 1994; Hofstadter, 1995] shares with structure mapping the need for discrete representations, but eschews the compartmentalisation of representation and mapping. In the ASA, the construction of each situation representation is interleaved with the construction of a mapping between situations. The ASA uses an agent-based search process, inspired by earlier “blackboard architectures” [Lesser *et al.*, 1975], and is related to ideas from global workspace theory [Baars, 2013; Shanahan, 2006; Goyal *et al.*, 2022].

All these architectures have a common character. A set of parallel, independent, specialist processes (or “agents” or “modules”) compete and co-operate to modify a shared, global memory structure (the “workspace” or “blackboard”), and the state of that memory structure is, in turn, visible to the full cohort of processes. The use of a shared workspace promotes integration among the specialist processes, which must all be capable of manipulating representations in a common format. If such an architecture is realised in a neural network, the bottleneck of a shared workspace will pressure processes to collectively learn a common representational format, which in turn leads to greater abstraction and better generalisation [Goyal *et al.*, 2022].

In addition to symbolic approaches such as structure mapping and ASA, several hybrid symbolic-connectionist approaches to analogy have also been explored, such as the ACME, LISA, and DORA systems [Holyoak and Thagard, 1989; Hummel and Holyoak, 1996; Doumas *et al.*, 2008]. Attempts have also been made to apply deep learning directly to analogical reasoning problems [Barrett *et al.*, 2018; Zhang *et al.*, 2019]. However, these deep learning models, while they can achieve high performance on specific test sets, are trained on large datasets in the domain in which the analogies are made, and the trained models have no application outside that domain. This is, in a sense, the very opposite of the kind of abstraction we are seeking, the essence of which is generic applicability.

3.3 Associative Memory

Associative memory is a long-established area of study in machine learning. Associative memories are not usually studied in the context of abstraction, yet they meet several of our requirements. A *Hopfield network*, for example, combines matching, retrieval, and pattern completion in a single operation, albeit one that requires multiple iterations to achieve convergence [Hopfield, 1982; Hopfield, 2007]. Perhaps something akin to a Hopfield network might serve as a basis for the sort of abstraction we are looking for. There are, for sure, many respects in which Hopfield networks are unsuitable for abstraction. However, there is an abundance of recent work in machine learning with the potential to bridge the gap between associative memory and the cognitive operations we have identified as central to abstraction.

A standard Hopfield network is an array of neurons with all-to-all recurrent connections. An attractor landscape of implicitly stored patterns is defined by the network’s connection weights, which can be trained using a Hebbian learning rule.

To retrieve a stored pattern, the array is initialised with an input pattern (e.g., an image), to which an update rule is repeatedly applied. This update rule minimises an energy function, ensuring that the network converges on an attractor. Under the right conditions, the attractor will correspond to the stored pattern that most closely resembles the input pattern. Because the input pattern can be partially blanked out or corrupted with noise, the retrieval process in effect performs pattern completion as well as matching.

Hopfield networks and their variants and relatives (e.g., Boltzmann machines) have been the subject of a large body of research spanning several decades. This includes a number of notable recent results. Krotov and Hopfield [2016] introduced a new parameterised energy function and corresponding update rule that dramatically increases the capacity of the Hopfield network while improving on the standard model’s ability to store and retrieve correlated patterns. Moreover, they showed that by adjusting the update rule’s parameter, the way the network functions can be smoothly varied from a distributed memory at one extreme, in which each stored pattern represents a cloud of features, to a prototype memory at the other, where each stored pattern represents a single exemplar.

Known as *dense associative memory* or the *modern Hopfield network*, this model was generalised by Demircigil *et al.* [2017], who introduced an exponential term into the energy function. This results in much faster convergence, with a high probability of convergence after a single update step. Demircigil *et al.*’s energy function and update rule, which accounted for discrete-valued patterns, is further generalised by Ramsauer *et al.* [2021] to the continuous case. Ramsauer *et al.* also show that the resulting model is mathematically equivalent to *self-attention*, as used in transformer architectures [Vaswani *et al.*, 2017]. Not only does this bring Hopfield networks, in their powerful modern guise, within the ambit of contemporary deep learning, it does so by relating them to one of the most successful innovations in the recent history of artificial intelligence.

Since they combine retrieval, matching, and pattern completion, associative memories could perhaps be used directly to see similarities between past events and present circumstances. And since modern Hopfield networks and transformer architectures are so closely related, they could potentially fulfil this role in the context of DRL. The first step would be to integrate an associative memory into a DRL architecture so that it can be used as a predictive model in the way sketched in Section 2.3.

There have been many successful attempts to incorporate a learned predictive world model into a DRL architecture [Racanière *et al.*, 2017; Ha and Schmidhuber, 2018; Gregor *et al.*, 2019; Schrittwieser *et al.*, 2020]. To date, as far as we are aware, associative memories have not been used for this purpose. However, while images are the canonical use case for associative memory, (modern) Hopfield networks can be used to store and retrieve any sort of data, including sequences of images or feature vectors. In the context of sequential data, matching and retrieval can be carried out on the beginning of a sequence, with the rest of the sequence being provided through pattern completion, effectively turning an associative memory into a predictive world model. So, given

its relationship to (differentiable) transformers, the modern Hopfield network could, in principle, play the part of a world model in a deep reinforcement learning architecture.

3.4 Latent Spaces

Recall that our aim here is to endow deep reinforcement learning agents with the ability to generalise better from their past experience by seeing similarity and making analogies. In Section 2.3, we envisaged an agent equipped with a repertoire of stereotypical templates for “small spatial stories”, each with very wide application, such as the obtain-object template. Hopfield-style associative memories look like good candidates for realising such templates, since their attractor dynamics ensures that every situation maps to one of a small, discrete set of patterns. The obtain-object template, for example, might be represented by a corresponding stored pattern, one that would match the situation facing the agent in the pillar task, even though that pattern was learned through past experiences of situations that were very different in appearance.

However, conventional Hopfield-style associative memories, as they stand, are not up to the job. One respect in which they fall short is that they exactly reproduce specific stored patterns. This is not what we are looking for when an agent matches its current situation to a template and fills it in accordingly. When the agent confronts the green sphere on top of the pillar, for example, we don’t want it to recreate, from episodic memory, a specific past experience involving other objects. Rather, we want it to imagine a future that involves the (unfamiliar) objects in the present scene – the pillar, the green sphere, and the box. In this respect, our requirements are for a matching and completion mechanism that is more “constructive”, in the manner of a generative model, such as the Kanerva machine of Wu *et al.* [2018], which combines an associative memory with a variational autoencoder.

Variational autoencoders (VAEs) [Kingma and Welling, 2014] and generative adversarial networks (GANs) [Goodfellow *et al.*, 2014] learn a relatively low-dimensional representation (the latent space) of a high-dimensional distribution (over images, say). These latent encodings are obliged to exhibit a degree of abstraction because they are compressed representations of the input observations. Moreover, as shown by Higgins *et al.* [2016; 2018], a VAE can be coaxed into learning a *disentangled* representation of its input distribution, wherein independent factors, such as shape, size, and colour, are represented by distinct variables. The level of abstraction in such latent encodings is modest compared to what we are seeking here. Nevertheless, they have been shown to facilitate zero-shot transfer in reinforcement learning [Higgins *et al.*, 2017].

The abstraction afforded by a generative model’s latent space also has the potential to improve a DRL agent’s predictive world model [Gregor *et al.*, 2019]. Conventional world models typically operate at too concrete a level to be used in the way we have envisaged in this paper. Our hypothetical obtain-object template, for example, should be indifferent to the appearance of the object being obtained, and although it has to commit to the occurrence of certain key events, it should have little or nothing to say about their duration or

about intermediate states. The world models typically used in DRL, by contrast, make predictions in observation space. They try to predict exactly how things will look to the agent in the future, on a frame by frame basis. By making predictions in latent space instead, the model described by Gregor et al [2019] exhibits a degree of both state abstraction and temporal abstraction.

The objective of minimising reconstruction error is dominant in VAE-based approaches to learning an abstract latent space. Yet an essential aspect of abstraction is the ability to ignore, or at least de-emphasise, irrelevant surface features, a requirement that is very different to reconstruction capability. Self-supervised *contrastive* learning offers a different approach to learning abstract latent spaces, one that is not based on reconstruction. Augmenting DRL architectures with an auxiliary contrastive prediction loss can improve performance in standard benchmarks [Oord et al., 2018; Guo et al., 2020]. A recent body of work has shown how to use contrastive losses to pre-train image encoders that are *general-purpose*, in the sense that they are useful for downstream tasks without further training [Grill et al., 2020; Chen et al., 2020; Mitrovic et al., 2021]. The resulting latent encodings have also been shown to improve data efficiency in a DRL setting [Banino et al., 2022].

3.5 Abstraction Through Language

Our emphasis in this paper is on abstractions, such as those of foundational common sense, that are simultaneously low-level – in the sense that they are distilled from basic sensorimotor interaction with the world – yet highly generic. Abstractions that are built upon language are higher-level, and arguably depend on a prior capacity for sensorimotor abstraction found even in animals that lack language. Humans, however, are exposed to language from infancy, and language, in a sense, pre-packages abstraction. It presents the infant with a ready-made framework for decomposing the world into objects and relations, and for combining and re-arranging those elements into representations that reflect the compositionality of the world. As such, exposure to human language in infancy can be thought of as turbocharging the process of acquiring low-level abstractions from sensorimotor interaction. So it’s natural to wonder whether the same advantage can be conferred on DRL agents, even if they are not solving language-related tasks.

A number of authors have pursued this line of thinking [Luketina et al., 2019]. For example, Andreas et al. [2018] show that the provision of linguistic task descriptions during training can improve performance on supervised classification in a few-shot multi-task setting, even when linguistic descriptions are not provided at evaluation time. Extending their ideas to multi-task reinforcement learning, the authors show that the provision of linguistic goal descriptions while pre-training on expert trajectories boosts performance during subsequent reinforcement learning on previously unseen tasks without linguistic descriptions. In both these settings, the trained system is evaluated on tasks that do not involve language. Nevertheless, exposure to linguistic descriptions during training enhances performance on those tasks. The authors hypothesise that language helps to mitigate over-fitting

in the supervised setting, and helps to structure the agent’s exploration in the reinforcement learning setting.

Further evidence that language can beneficially shape representations in reinforcement learning tasks, even when those tasks don’t inherently involve language, is provided by Lampinen et al. [2021]. In their work, an agent is required to select the odd one out from a set of objects. They study both 2D and 3D settings. The 3D setting is particularly demanding because the agent cannot see all the objects at once, obliging it to remember the features of previously encountered objects as it moves around. In Lampinen et al.’s setup, the agent has an auxiliary task, in addition to maximising reward, which is to generate a reward-relevant linguistic description of the object in front of it. During training, these are compared to ground-truth descriptions, and the resulting loss helps to structure the agent’s learned representations. But at test time no language is required. In both 2D and 3D settings, the authors report, agent performance is higher with the auxiliary language task than without it, and they speculate that this is because the linguistic descriptions “highlight abstract task structure”.

Studies like these lend support to the idea that exposure to language shortcuts the challenge of learning how to divide the world into discrete objects and relations, which is a pre-requisite for acquiring generic templates of the sort we advocate here as a basis for abstraction. Perhaps language achieves this by promoting the emergence of representational structures resembling those of classical, symbolic AI. On this view, however, in contrast to the hard-coded structures of symbolic AI, the relevant structure is present in the data, partly in the world itself, and partly in the corresponding language.

Large language models offer an intriguing glimpse of another way abstraction might be achieved through language, namely with a transformer-based predictive model trained on a large corpus of sequential data. GPT-3 is capable of few-shot generalisation on tasks it was not trained on, where the tasks in question are specified by engineering suitable prompts [Brown et al., 2020]. Indeed, there is some evidence that GPT-3 can solve certain analogical reasoning problems presented this way [Mitchell, 2020]. Perhaps a large enough corpus of sequential data with the right statistical properties (eg: written human language) can pressure a large enough transformer-based model to find a general-purpose pattern completion mechanism in its parameter space.

4 Conclusion: The Missing Pieces

As we have seen, there is plenty of work in artificial intelligence, much of it recent, that is relevant to building deep reinforcement learning agents capable of acquiring and using abstract representations. Yet none of these advances, even in combination, amounts to a blueprint for an agent with the capabilities envisioned at the start of the paper. None of the techniques described implements anything similar to the sort of template we hypothesised as a medium for seeing similarity, making analogies, or forming concepts. The representations learned by all these techniques remain to some degree tied to the domain in which they were acquired.

By contrast, the central feature of a template is its slots or variables, which can be bound to *any thing*. In formal logic terms, they are “universally quantified”. Without variables with the potential to bind to any thing, no representational structure or computational process can find open-ended application to novel situations, such as those in which the children found themselves who were so successful on the Animal-AI test suite [Marcus, 2001; Martin and Doumas, 2019; Greff *et al.*, 2020]. Variable binding is an inherent feature of any programming language, and program induction is a promising approach to learning abstractions [Lázaro-Gredilla *et al.*, 2019; Ellis *et al.*, 2020; Evans *et al.*, 2021]. However, integrating program induction with deep reinforcement learning is challenging thanks to the usual problems of reconciling discreteness and differentiability. Work on integrating binding with neural networks has a long pedigree [Smolensky, 1990], but further work on differentiable mechanisms for variable binding is critical.

However it is achieved, variable binding would also facilitate *compositionality*, another essential property of abstraction (and also a feature of both natural language and programming languages) [Higgins *et al.*, 2018]. To solve unseen problems like the pillar task in the way we have proposed requires templates to be composed. Ideally, it should be possible to bind the slots in a template to other templates, leading to a form of hierarchical compositionality.

Acknowledgments

Thanks to Richard Evans, Felix Hill, Andrew Lampinen, and the anonymous reviewers.

References

- [Andreas *et al.*, 2018] Jacob Andreas, Dan Klein, and Sergey Levine. Learning with latent language. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2166–2179, 2018.
- [Baars, 2013] Bernard J. Baars. A global workspace theory of conscious experience. In *Consciousness in philosophy and cognitive neuroscience*, pages 161–184. Psychology Press, 2013.
- [Banino *et al.*, 2022] Andrea Banino, Adria Puigdomenech Badia, Jacob C Walker, Tim Scholtes, Jovana Mitrovic, and Charles Blundell. CoBERL: Contrastive BERT for reinforcement learning. In *International Conference on Learning Representations*, 2022.
- [Barrett *et al.*, 2018] D. G. T. Barrett, F. Hill, A. Santoro, A. S. Morcos, and T. Lillicrap. Measuring abstract reasoning in neural networks. In *Proceedings of the International Conference on Machine Learning*, pages 4477–4486, 2018.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, et al. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901, 2020.
- [Burgess *et al.*, 2019] Christopher P Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. MONet: Unsupervised scene decomposition and representation. *arXiv:1901.11390*, 2019.
- [Chen *et al.*, 2020] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning (ICML)*, pages 1597–1607, 2020.
- [Creswell *et al.*, 2021] Antonia Creswell, Rishabh Kabra, Christopher P. Burgess, and Murray Shanahan. Unsupervised object-based transition models for 3D partially observable environments. In *Advances in Neural Information Processing Systems*, 2021.
- [Crosby *et al.*, 2020] Matthew Crosby, Benjamin Beyret, Murray Shanahan, José Hernández-Orallo, Lucy Cheke, and Marta Halina. The Animal-AI testbed and competition. In *Proceedings of the NeurIPS 2019 Competition and Demonstration Track*, volume 123, pages 164–176, 2020.
- [Crouse *et al.*, 2021] Maxwell Crouse, Constantine Nakos, Ibrahim Abdelaziz, and Ken Forbus. Neural analogical matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 809–817, 2021.
- [Demircigil *et al.*, 2017] Mete Demircigil, Judith Heusel, Matthias Löwe, Sven Upgang, and Franck Vermet. On a model of associative memory with huge storage capacity. *Journal of Statistical Physics*, 168(2):288–299, 2017.
- [Doumas *et al.*, 2008] Leonidas A.A. Doumas, John E. Hummel, and Catherine M. Sandhofer. A theory of the discovery and predication of relational concepts. *Psychological review*, 115(1):1, 2008.
- [Dubey *et al.*, 2018] Rachit Dubey, Pulkit Agrawal, Deepak Pathak, Tom Griffiths, and Alexei Efros. Investigating human priors for playing video games. In *Proceedings of the 35th International Conference on Machine Learning*, pages 1349–1357, 2018.
- [Ellis *et al.*, 2020] K. Ellis, C. Wong, M. Nye, M. Sablé-Meyer, L. Cary, L. Morales, L. Hewitt, A. Solar-Lezama, and J. B. Tenenbaum. DreamCoder: Growing generalizable, interpretable knowledge with wake-sleep Bayesian program learning. *arXiv:2006.08381*, 2020.
- [Engelcke *et al.*, 2020] Martin Engelcke, Adam R Kosiorek, Oivi Parker Jones, and Ingmar Posner. Genesis: Generative scene inference and sampling with object-centric latent representations. In *International Conference on Learning Representations*, 2020.
- [Evans *et al.*, 2021] R. Evans, J. Hernández-Orallo, J. Welbl, P. Kohli, and M. Sergot. Making sense of sensory input. *Artificial Intelligence*, 293:103438, 2021.
- [Falkenhainer *et al.*, 1989] Brian Falkenhainer, Kenneth D Forbus, and Dedre Gentner. The structure-mapping engine: Algorithm and examples. *Artificial intelligence*, 41(1):1–63, 1989.

- [Garnelo and Shanahan, 2019] Marta Garnelo and Murray Shanahan. Reconciling deep learning with symbolic artificial intelligence: representing objects and relations. *Current Opinion in Behavioral Sciences*, 29:17–23, 2019.
- [Garnelo *et al.*, 2016] Marta Garnelo, Kai Arulkumaran, and Murray Shanahan. Towards deep symbolic reinforcement learning. *arXiv: 1609.05518*, 2016.
- [Gentner, 1983] Dedre Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive science*, 7(2):155–170, 1983.
- [Goodfellow *et al.*, 2014] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, et al. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 28, pages 2672–2680, 2014.
- [Goyal *et al.*, 2022] Anirudh Goyal, Aniket Rajiv Didolkar, Alex Lamb, Kartikeya Badola, Nan Rosemary Ke, et al. Coordination among neural modules through a shared global workspace. In *International Conference on Learning Representations*, 2022.
- [Greff *et al.*, 2019] Klaus Greff, Raphaël Lopez Kaufman, Rishabh Kabra, Nick Watters, Christopher Burgess, Daniel Zoran, et al. Multi-object representation learning with iterative variational inference. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97, pages 2424–2433, 2019.
- [Greff *et al.*, 2020] Klaus Greff, Sjoerd van Steenkiste, and Jürgen Schmidhuber. On the binding problem in artificial neural networks. *arXiv:2012.05208*, 2020.
- [Gregor *et al.*, 2019] Karol Gregor, George Papamakarios, Frederic Besse, Lars Buesing, and Theophane Weber. Temporal difference variational auto-encoder. In *International Conference on Learning Representations*, 2019.
- [Grill *et al.*, 2020] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, et al. Bootstrap your own latent - a new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*, volume 33, pages 21271–21284, 2020.
- [Guo *et al.*, 2020] Zhaohan Daniel Guo, Bernardo Avila Pires, Bilal Piot, Jean-Bastien Grill, Florent Altché, et al. Bootstrap latent-predictive representations for multitask reinforcement learning. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pages 3875–3886, 2020.
- [Ha and Schmidhuber, 2018] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems*, volume 31, pages 2455–2467, 2018.
- [Higgins *et al.*, 2016] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, et al. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *International Conference on Learning Representations*, 2016.
- [Higgins *et al.*, 2017] Irina Higgins, Arka Pal, Andrei Rusu, Loic Matthey, Christopher Burgess, et al. DARLA: Improving zero-shot transfer in reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1480–1490, 2017.
- [Higgins *et al.*, 2018] Irina Higgins, Nicolas Sonnerat, Loic Matthey, Arka Pal, Christopher P Burgess, et al. SCAN: Learning hierarchical compositional visual concepts. In *Proceedings International Conference on Learning Representations*, 2018.
- [Hofstadter and Mitchell, 1994] D. R. Hofstadter and M. Mitchell. The Copycat project: A model of mental fluidity and analogy-making. In K. J. Holyoak and J. A. Barnden, editors, *Advances in Connectionist and Neural Computation Theory*, volume 2, pages 31–112. Ablex Publishing Corporation, 1994.
- [Hofstadter, 1995] D. R. Hofstadter. *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought*. Basic Books, New York, 1995.
- [Holyoak and Thagard, 1989] K. J. Holyoak and P. Thagard. Analogical mapping by constraint satisfaction. *Cognitive Science*, 13(3):295–355, 1989.
- [Hopfield, 1982] John J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.
- [Hopfield, 2007] John J. Hopfield. Hopfield network. *Scholarpedia*, 2(5):1977, 2007.
- [Hummel and Holyoak, 1996] J. E. Hummel and K. J. Holyoak. LISA: A computational model of analogical inference and schema induction. In *Proceedings of the Eighteenth Annual Meeting of the Cognitive Science Society*, pages 352–357, 1996.
- [Kingma and Welling, 2014] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *arXiv:1312.6114*, 2014.
- [Kirk *et al.*, 2021] Robert Kirk, Amy Zhang, Edward Grefenstette, and Tim Rocktäschel. A survey of generalisation in deep reinforcement learning. *arXiv:2111.09794*, 2021.
- [Krotov and Hopfield, 2016] Dmitry Krotov and John J. Hopfield. Dense associative memory for pattern recognition. In *Advances in Neural Information Processing Systems*, volume 29, pages 1172–1180, 2016.
- [Lake *et al.*, 2017] Brenden M. Lake, Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, 2017.
- [Lakoff and Johnson, 1980] George Lakoff and Mark Johnson. *Metaphors We Live By*. University of Chicago Press, 1980.
- [Lampinen *et al.*, 2021] Andrew K. Lampinen, Nicholas A. Roy, Ishita Dasgupta, Stephanie C. Y. Chan, Allison C.

- Tam, et al. Tell me why! - explanations support learning of relational and causal structure. *arXiv:2112.03753*, 2021.
- [Lázaro-Gredilla *et al.*, 2019] Miguel Lázaro-Gredilla, Dianhuan Lin, J. Swaroop Guntupalli, and Dileep George. Beyond imitation: Zero-shot task transfer on robots by learning concepts as cognitive programs. *Science Robotics*, 4(26), 2019.
- [Lesser *et al.*, 1975] Victor Lesser, Richard Fennell, Lee Erman, and D. Reddy. Organization of the HEARSAY II speech understanding system. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 23(1):11–24, 1975.
- [Locatello *et al.*, 2020] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, et al. Object-centric learning with slot attention. In *Advances in Neural Information Processing Systems*, volume 33, pages 11525–11538, 2020.
- [Luketina *et al.*, 2019] Jelena Luketina, Nantas Nardelli, Gregory Farquhar, Jakob N. Foerster, Jacob Andreas, et al. A survey of reinforcement learning informed by natural language. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 6309–6317, 2019.
- [Marcus, 2001] Gary F Marcus. *The Algebraic Mind*. MIT Press, 2001.
- [Martin and Doumas, 2019] Andrea E Martin and Leonidas A A Doumas. Predicate learning in neural systems: using oscillations to discover latent structure. *Current Opinion in Behavioral Sciences*, 29:77–83, 2019.
- [Mitchell, 2020] Melanie Mitchell. Can GPT-3 make analogies? <https://medium.com/@melaniemitchell.me/can-gpt-3-make-analogies-16436605c446>, 2020. Accessed: 2022-05-09.
- [Mitchell, 2021] Melanie Mitchell. Abstraction and analogy-making in artificial intelligence. *Annals of the New York Academy of Sciences*, 1505:79–101, 2021.
- [Mitrovic *et al.*, 2021] Jovana Mitrovic, Brian McWilliams, Jacob C Walker, Lars Holger Buesing, and Charles Blundell. Representation learning via invariant causal mechanisms. In *International Conference on Learning Representations*, 2021.
- [Oord *et al.*, 2018] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv:1807.03748*, 2018.
- [Racanière *et al.*, 2017] Sébastien Racanière, Theophane Weber, David Reichert, Lars Buesing, Arthur Guez, et al. Imagination-augmented agents for deep reinforcement learning. In *Advances in Neural Information Processing Systems 30*, pages 5690–5701, 2017.
- [Ramsauer *et al.*, 2021] Hubert Ramsauer, Bernhard Schäfl, Johannes Lehner, Philipp Seidl, Michael Widrich, et al. Hopfield networks is all you need. In *International Conference on Learning Representations*, 2021.
- [Schank and Abelson, 1977] Roger C. Schank and Robert P. Abelson. *Scripts, Plans, Goals, and Understanding: An Inquiry Into Human Knowledge Structures*. Psychology Press, 1977.
- [Schrittwieser *et al.*, 2020] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, et al. Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, 588:604–609, 2020.
- [Shanahan *et al.*, 2020] Murray Shanahan, Benjamin Beyret, Matthew Crosby, and Lucy Cheke. Artificial intelligence and the common sense of animals. *Trends in Cognitive Sciences*, 24(11):862–872, 2020.
- [Shanahan, 2006] Murray Shanahan. A cognitive architecture that combines internal simulation with a global workspace. *Consciousness and cognition*, 15(2):433–449, 2006.
- [Smith, 2019] Brian Cantwell Smith. *The Promise of Artificial Intelligence: Reckoning and Judgment*. MIT Press, 2019.
- [Smolensky, 1990] Paul Smolensky. Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial Intelligence*, 46(1-2):159–216, 1990.
- [Spelke, 2000] Elizabeth S. Spelke. Core knowledge. *American Psychologist*, November 2000:1233–1243, 2000.
- [Tsvividis *et al.*, 2017] Pedro A Tsvividis, Thomas Pouncy, Jacqueline L Xu, Joshua B Tenenbaum, and Samuel J Gershman. Human learning in Atari. In *AAAI Spring Symposium on Science of Intelligence: Computational Principles of Natural and Artificial Intelligence*, 2017.
- [Turner, 1998] Mark Turner. *The Literary Mind: The Origins of Thought and Language*. Oxford University Press, 1998.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, et al. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, pages 6000–6010, 2017.
- [Voudouris *et al.*, 2022] Konstantinos Voudouris, Matthew Crosby, Benjamin Beyret, José Hernández-Orallo, Murray Shanahan, et al. Direct human-AI comparison in the Animal-AI environment. *Frontiers in Psychology*, 2022.
- [Wu *et al.*, 2018] Yan Wu, Greg Wayne, Alex Graves, and Timothy Lillicrap. The kanerva machine: A generative distributed memory. In *International Conference on Learning Representations*, 2018.
- [Zhang *et al.*, 2019] C. Zhang, F. Gao, B. Jia, Y. Zhu, and S.-C. Zhu. RAVEN: A dataset for relational and analogical visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5312–5322, 2019.