
Imagination-Augmented Agents for Deep Reinforcement Learning

Théophane Weber* **Sébastien Racanière*** **David P. Reichert*** **Lars Buesing**
Arthur Guez **Danilo Rezende** **Adria Puigdomènech Badia** **Oriol Vinyals**
Nicolas Heess **Yujia Li** **Razvan Pascanu** **Peter Battaglia**
David Silver **Daan Wierstra**
DeepMind

Abstract

We introduce Imagination-Augmented Agents (I2As), a novel architecture for deep reinforcement learning combining model-free and model-based aspects. In contrast to most existing model-based reinforcement learning and planning methods, which prescribe how a model should be used to arrive at a policy, I2As learn to interpret predictions from a learned environment model to construct implicit plans in arbitrary ways, by using the predictions as additional context in deep policy networks. I2As show improved data efficiency, performance, and robustness to model misspecification compared to several baselines.

1 Introduction

A hallmark of an intelligent agent is its ability to rapidly adapt to new circumstances and "achieve goals in a wide range of environments" [1]. Progress has been made in developing capable agents for numerous domains using deep neural networks in conjunction with *model-free* reinforcement learning (RL) [2–4], where raw observations directly map to values or actions. However, this approach usually requires large amounts of training data and the resulting policies do not readily generalize to novel tasks in the same environment, as it lacks the behavioral flexibility constitutive of general intelligence.

Model-based RL aims to address these shortcomings by endowing agents with a model of the world, synthesized from past experience. By using an internal model to reason about the future, here also referred to as *imagining*, the agent can seek positive outcomes while avoiding the adverse consequences of trial-and-error in the real environment – including making irreversible, poor decisions. Even if the model needs to be learned first, it can enable better generalization across states, remain valid across tasks in the same environment, and exploit additional unsupervised learning signals, thus ultimately leading to greater data efficiency. Another appeal of model-based methods is their ability to scale performance with more computation by increasing the amount of internal simulation.

Model learning and mental simulation has been hypothesized and demonstrated in animal and human learning [5–8]. Its successful deployment in artificial model-based agents however has hitherto been limited to settings where an exact transition model is available [9] or in domains where models are easy to learn – e.g. symbolic environments or low-dimensional systems [10–13]. In complex domains for which a simulator is not available to the agent, recent successes are dominated by model-free methods [2, 14]. In such domains, the performance of model-based agents employing standard planning methods usually suffers from model errors resulting from function approximation [15, 16]. These errors compound during planning, causing over-optimism and poor agent performance. There are currently no planning or model-based methods that are robust against model imperfections which are inevitable in complex domains, thereby preventing them from matching the success of their model-free counterparts.

*Equal contribution, corresponding authors: {theophane, sracaniere, reichert}@google.com.

We seek to address this shortcoming by proposing Imagination-Augmented Agents, which use approximate environment models by "learning to interpret" their imperfect predictions. Our algorithm can be trained directly on low-level observations with little domain knowledge, similarly to recent model-free successes. Without making any assumptions about the structure of the environment model and its possible imperfections, our approach learns in an end-to-end way to extract useful knowledge gathered from model simulations – in particular not relying exclusively on simulated returns. This allows the agent to benefit from model-based imagination without the pitfalls of conventional model-based planning. We demonstrate that our approach performs better than model-free baselines in various domains including Sokoban. It achieves better performance with less data, even with imperfect models, a significant step towards delivering the promises of model-based RL.

2 The I2A architecture

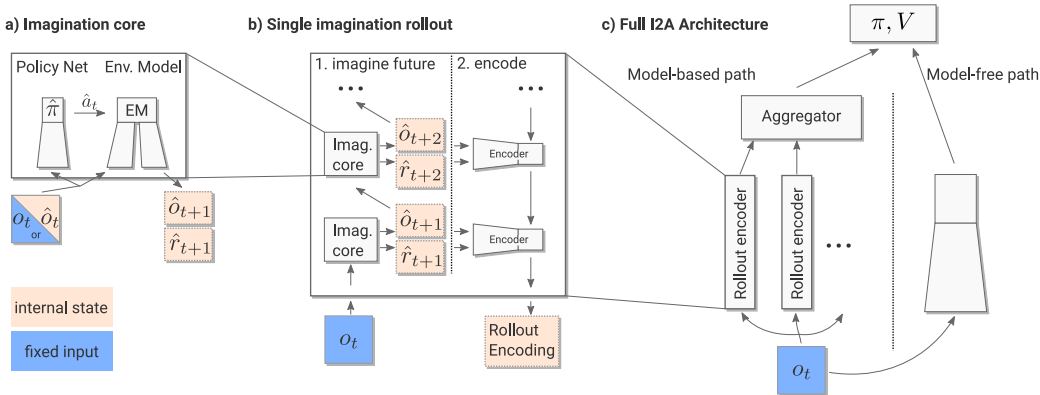


Figure 1: *I2A architecture*. $\hat{\cdot}$ notation indicates imagined quantities. *a)*: the imagination core (IC) predicts the next time step conditioned on an action sampled from the rollout policy $\hat{\pi}$. *b)*: the IC imagines trajectories of features $\hat{f} = (\hat{o}, \hat{r})$, encoded by the rollout encoder. *c)*: in the full I2A, aggregated rollout encodings and input from a model-free path determine the output policy π .

In order to augment model-free agents with imagination, we rely on environment models – models that, given information from the present, can be queried to make predictions about the future. We use these environment models to simulate *imagined trajectories*, which are interpreted by a neural network and provided as additional context to a policy network.

In general, an environment model is any recurrent architecture which can be trained in an unsupervised fashion from agent trajectories: given a past state and current action, the environment model predicts the next state and any number of signals from the environment. In this work, we will consider in particular environment models that build on recent successes of action-conditional next-step predictors [17–19], which receive as input the current observation (or history of observations) and current action, and predict the next observation, and potentially the next reward. We roll out the environment model over multiple time steps into the future, by initializing the imagined trajectory with the present time real observation, and subsequently feeding simulated observations into the model.

The actions chosen in each rollout result from a rollout policy $\hat{\pi}$ (explained in Section 3.1). The environment model together with $\hat{\pi}$ constitute the imagination core module, which predicts next time steps (Fig 1a). The imagination core is used to produce n trajectories $\hat{\mathcal{T}}_1, \dots, \hat{\mathcal{T}}_n$. Each imagined trajectory $\hat{\mathcal{T}}$ is a sequence of features $(\hat{f}_{t+1}, \dots, \hat{f}_{t+\tau})$, where t is the current time, τ the length of the rollout, and \hat{f}_{t+i} the output of the environment model (i.e. the predicted observation and/or reward).

Despite recent progress in training better environment models, a key issue addressed by I2As is that a learned model cannot be assumed to be perfect; it might sometimes make erroneous or nonsensical predictions. We therefore do not want to rely solely on predicted rewards (or values predicted from predicted states), as is often done in classical planning. Additionally, trajectories may contain information beyond the reward sequence (a trajectory could contain an informative subsequence – for

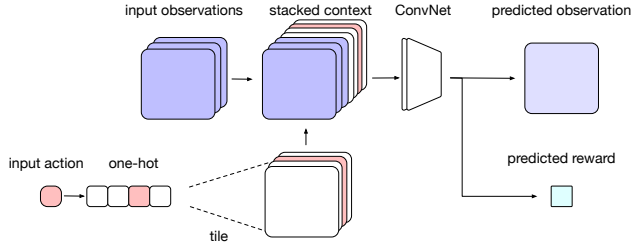


Figure 2: *Environment model*. The input action is broadcast and concatenated to the observation. A convolutional network transforms this into a pixel-wise probability distribution for the output image, and a distribution for the reward.

instance solving a subproblem – which did not result in higher reward). For these reasons, we use a rollout encoder \mathcal{E} that processes the imagined rollout as a whole and *learns to interpret it*, i.e. by extracting any information useful for the agent’s decision, or even ignoring it when necessary (Fig 1b). Each trajectory is encoded separately as a rollout embedding $e_i = \mathcal{E}(\hat{T}_i)$. Finally, an aggregator \mathcal{A} converts the different rollout embeddings into a single imagination code $c_{ia} = \mathcal{A}(e_1, \dots, e_n)$.

The final component of the I2A is the policy module, which is a network that takes the information c_{ia} from model-based predictions, as well as the output c_{mf} of a model-free path (a network which only takes the real observation as input; see Fig 1c, right), and outputs the imagination-augmented policy vector π and estimated value V . The I2A therefore learns to combine information from its model-free and imagination-augmented paths; note that without the model-based path, I2As reduce to a standard model-free network [3]. I2As can thus be thought of as augmenting model-free agents by providing additional information from model-based planning, and as having strictly more expressive power than the underlying model-free agent.

3 Architectural choices and experimental setup

3.1 Rollout strategy

For our experiments, we perform one rollout for each possible action in the environment. The first action in the i^{th} rollout is the i^{th} action of the action set \mathcal{A} , and subsequent actions for all rollouts are produced by a shared rollout policy $\hat{\pi}$. We investigated several types of rollout policies (random, pre-trained) and found that a particularly efficient strategy was to distill the imagination-augmented policy into a model-free policy. This distillation strategy consists in creating a small model-free network $\hat{\pi}(o_t)$, and adding to the total loss a cross entropy auxiliary loss between the imagination-augmented policy $\pi(o_t)$ as computed on the current observation, and the policy $\hat{\pi}(o_t)$ as computed on the same observation. By imitating the imagination-augmented policy, the internal rollouts will be similar to the trajectories of the agent in the real environment; this also ensures that the rollout corresponds to trajectories with high reward. At the same time, the imperfect approximation results in a rollout policy with higher entropy, potentially striking a balance between exploration and exploitation.

3.2 I2A components and environment models

In our experiments, the encoder is an LSTM with convolutional encoder which sequentially processes a trajectory \mathcal{T} . The features \hat{f}_t are fed to the LSTM in reverse order, from $\hat{f}_{t+\tau}$ to \hat{f}_{t+1} , to mimic Bellman type backup operations². The aggregator simply concatenates the summaries. For the model-free path of the I2A, we chose a standard network of convolutional layers plus one fully connected one [e.g. 3]. We also use this architecture on its own as a baseline agent.

Our environment model (Fig. 2) defines a distribution which is optimized by using a negative log-likelihood loss l_{model} . We can either pretrain the environment model before embedding it (with frozen weights) within the I2A architecture, or jointly train it with the agent by adding l_{model} to the total loss as an auxiliary loss. In practice we found that pre-training the environment model led to faster runtime of the I2A architecture, so we adopted this strategy.

For all environments, training data for our environment model was generated from trajectories of a partially trained standard model-free agent (defined below). We use partially pre-trained agents

²The choice of forward, backward or bi-directional processing seems to have relatively little impact on the performance of the I2A, however, and should not preclude investigating different strategies.

because random agents see few rewards in some of our domains. However, this means we have to account for the budget (in terms of real environment steps) required to pretrain the data-generating agent, as well as to then generate the data. In the experiments, we address this concern in two ways: by explicitly accounting for the number of steps used in pretraining (for Sokoban), or by demonstrating how the same pretrained model can be reused for many tasks (for MiniPacman).

3.3 Agent training and baseline agents

Using a fixed pretrained environment model, we trained the remaining I2A parameters with asynchronous advantage actor-critic (A3C) [3]. We added an entropy regularizer on the policy π to encourage exploration and the auxiliary loss to distill π into the rollout policy $\hat{\pi}$ as explained above. We distributed asynchronous training over 32 to 64 workers; we used the RMSprop optimizer [20]. We report results after an initial round of hyperparameter exploration (details in Appendix A). Learning curves are averaged over the top three agents unless noted otherwise.

A separate hyperparameter search was carried out for each agent architecture in order to ensure optimal performance. In addition to the I2A, we ran the following baseline agents (see Appendix B for architecture details for all agents).

Standard model-free agent. For our main baseline agent, we chose a model-free standard architecture similar to [3], consisting of convolutional layers (2 for MiniPacman, and 3 for Sokoban) followed by a fully connected layer. The final layer, again fully connected, outputs the policy logits and the value function.

Copy-model agent. The I2A uses a larger architecture than the standard agent and more parameters. To demonstrate that this alone is not enough to explain its performance, we also tested the following model-free baseline. We replaced the environment model in the I2A with a ‘copy’ model that simply returns the input observation. Lacking a model, this agent does not use imagination, but uses the same architecture and the same number of learnable parameters (the environment model is kept constant in the I2A). It also benefits from the same number of computations (which in both cases increases linearly with the length of the rollouts). This model effectively corresponds to an architecture where policy logits and value are the final output of an LSTM network with skip connection.

4 Sokoban experiments

We now demonstrate the performance of I2A over baselines in a puzzle environment, Sokoban. We address the issue of dealing with imperfect models, highlighting the strengths of our approach over planning baselines. We also analyze the importance of the various components of the I2A.

Sokoban is a classic planning problem, where the agent has to push a number of boxes onto given target locations. Because boxes can only be pushed (as opposed to pulled), many moves are irreversible, and mistakes can render the puzzle unsolvable. A human player is thus forced to plan moves ahead of time. We expect that artificial agents will similarly benefit from internal simulation. Our implementation of Sokoban procedurally generates a new level each episode (see Appendix D.4 for details, Fig. 3 for examples). This means an agent cannot memorize specific puzzles.³ Together with the planning aspect, this makes for a very challenging environment for our model-free baseline agents, which solve less than 60% of the levels after a billion steps of training (details below). We provide videos of agents playing our version of Sokoban online [21].




Figure 3: *Random examples of procedurally generated Sokoban levels.* The player (green sprite) needs to push all 4 boxes onto the red target squares to solve a level, while avoiding irreversible mistakes. Our agents receive sprite graphics as observations as shown.

³Out of 40 million levels generated, less than 0.7% were repeated. Training an agent on 1 billion frames requires less than 20 million episodes.

While the underlying game logic operates in a 10×10 grid world, our agents were trained directly on RGB sprite graphics as shown in Fig. 4 (image size 80×80 pixels). There are no aspects of I2As that make them specific to grid world games.

4.1 I2A performance vs. baselines on Sokoban

Figure 4 (left) shows the learning curves of the I2A architecture and various baselines explained throughout this section. First, we compare I2A against the standard model-free agent. I2A clearly outperforms the latter, reaching a performance of 85% of levels solved vs. a maximum of under 60% for the baseline. Similarly, for Sokoban, I2A far outperforms the copy-model.

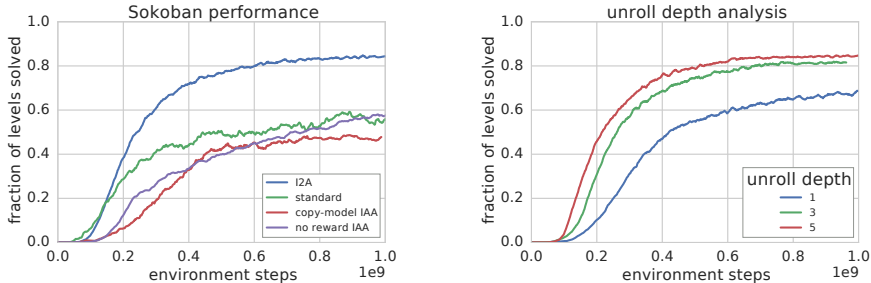


Figure 4: *Sokoban learning curves. Left:* training curves of I2A and baselines. Note that I2A use additional environment observations to pretrain the environment model, see main text for discussion. *Right:* I2A training curves for various values of imagination depth.

Since using imagined rollouts is helpful for this task, we investigate how the length of individual rollouts affects performance. The latter was one of the hyperparameters we searched over. A breakdown by number of unrolling/imagination steps in Fig. 4 (right) shows that using 3 unrolling steps improves speed of learning and top performance significantly over 1 unrolling step, and 5 is optimal. We also tested longer unrolls but found no increase in performance for this task (not shown). It is noteworthy that 5 steps is relatively small compared to the number of steps taken to solve a level, for which our best agents need about 50 steps on average. This implies that even such short rollouts can be highly informative. For example, they allow the agent to learn about moves it cannot recover from (such as pushing boxes against walls, in certain contexts).

In terms of data efficiency, it should be noted that the environment model in the I2A was pretrained (see Section 3.2). We conservatively measured the total number of frames needed for pretraining to be lower than $1e8$. Thus, even taking pretraining into account, I2A outperforms the baselines after seeing about $3e8$ frames in total (compare again Fig. 4 (left)). Of course, data efficiency is even better if the environment model can be reused to solve multiple tasks in the same environment (Section 5).

4.2 Learning with imperfect models

One of the key strengths of I2As is being able to handle learned and thus potentially imperfect environment models. However, for the Sokoban task, our learned environment models actually perform quite well when rolling out imagined trajectories. To demonstrate that I2As can deal with less reliable predictions, we ran another experiment where the I2A used an environment model that had showed much worse performance (due to a smaller number of parameters), with strong artifacts accumulating over iterated rollout predictions (Fig. 5, left). As Fig. 5 (right) shows, even with such a clearly flawed environment model, I2A performs similarly well. Unlike in the earlier experiment with the better environment model, we also found that an unroll length of 5 no longer outperforms a length of 3. This implies that I2As can learn to ignore the latter parts of the rollout as errors accumulate, but still use initial predictions when errors are less severe.

Learning a rollout encoder is what enables I2As to deal with imperfect model predictions. We can further demonstrate this point by comparing them to a setup without a rollout encoder: as in the classic Monte-Carlo search algorithm of Tesauro and Galperin [22], we now explicitly estimate the value of each action from rollouts, rather than learning an arbitrary encoding of the rollouts, as in I2A. We then select actions according to those values. Specifically, we estimate returns by summing the discounted predicted rewards and the estimated value of the final imagined frame. This value

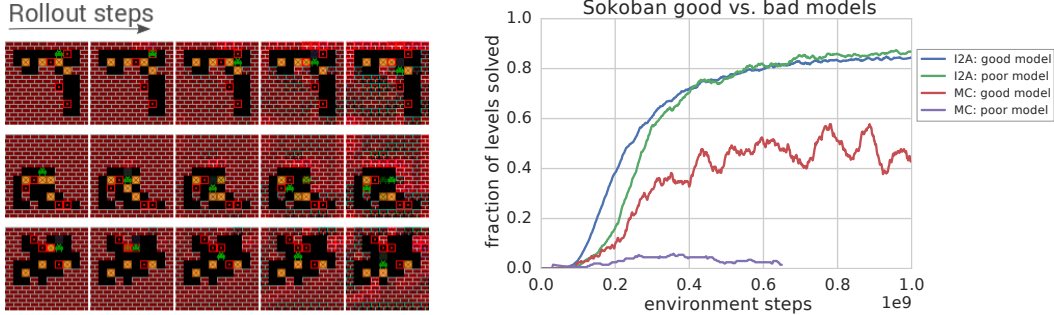


Figure 5: *Experiments with a noisy environment model.* Left: each row shows an example 5-step rollout after conditioning on an environment observation. Errors accumulate and lead to various artefacts, including missing or duplicate sprites. Right: comparison of Monte-Carlo (MC) search and I2A when using either the accurate or the noisy model for rollouts.

function is computed by a network that is trained separately to predict values from true observations. Actions are selected according to a softmax over the resulting estimated action values.⁴

We ran this rollout encoder-free agent on Sokoban with both the accurate and the noisy environment model. For the results, we chose the length of the rollout to be optimal for each environment model (from the same range as for I2A, i.e. from 1 to 5). As can be seen in Fig. 5 (right),⁵ when using the high accuracy environment model, the performance of the encoder-free agent is similar to that of the baseline standard agent. However, unlike I2A, its performance degrades catastrophically when using the poor model, showcasing the susceptibility to model misspecification.

4.3 Further insights into the workings of the I2A architecture

So far, we have studied the role of the rollout encoder. To show the importance of various other components of the I2A, we performed additional control experiments. Results are plotted in Fig. 4 (left) for comparison. First, I2A with the copy model (Section 3.3) performs far worse, demonstrating that the environment model is indeed crucial. Second, we trained an I2A where the environment model was predicting no rewards, only observations. This also performed worse. However, after much longer training ($3e9$ steps), these agents did recover the performance of the original I2A, which was never the case for the baseline agent even with that many steps. Hence, reward prediction is helpful but not absolutely necessary in this task, and imagined observations alone are informative enough to obtain high performance on Sokoban. Note this is in contrast to many classical planning and model-based reinforcement learning methods, which often rely on reward prediction.

4.4 Imagination efficiency and comparison with perfect-model planning methods

In previous sections, we illustrated that I2As can be used to efficiently solve planning problems and can be robust in the face of model misspecification. Here, we ask a different question – if we *do* assume a nearly perfect model, how does I2A compare to competitive planning methods? Beyond raw performance we focus particularly on the efficiency of planning, i.e. the number of imagination steps required to solve a fixed ratio of levels. We compare our regular I2A agent to a variant of Monte Carlo Tree Search (MCTS), which is a modern guided tree search algorithm [9, 23]. For our MCTS implementation, we aimed to have a strong baseline by using recent ideas: we include transposition tables [24], and evaluate the returns of leaf nodes by using a value network (in this case, a deep residual value network; see Appendix D.3 for further details).

Running MCTS on Sokoban, we find that it can achieve high performance, but at a cost of a much higher number of necessary environment model simulation steps: MCTS reaches the I2A performance of 87% of levels solved when using 25k model simulation steps on average to solve a level, compared

⁴The rollout policy is still obtained by distillation from the output policy, and we remove the model-free path to evaluate solely the imagination path for this experiment. We use a learned temperature in the action softmax.

⁵Note: the MC curves in Fig. 5 only used a single agent rather than averages.

I2A@87	~ 1400
I2A MC search @95	~ 4000
MCTS@87	~ 25000
MCTS@95	~ 100000
Random search	~ millions

Table 1

Boxes	1	2	3	4	5	6	7
I2A (%)	99.5	97	92	88	77	66	53
Standard (%)	97	87	72	60	47	32	23

Table 2

to 1.4k environment model calls for I2A. Using even more simulation steps, MCTS performance increases further, e.g. reaching 95% with 100k steps.

If we assume access to a high-accuracy environment model (including the reward prediction), we can also push I2A performance further, by performing basic Monte-Carlo search with a trained I2A for the rollout policy: we let the agent play whole episodes in simulation (where I2A itself uses the environment model for short-term rollouts, hence corresponding to using a model-within-a-model), and execute a successful action sequence if found, up to a maximum number of retries; this is reminiscent of nested rollouts [25]. With a fixed maximum of 10 retries, we obtain a score of 95% (up from 87% for the I2A itself). The total average number of model simulation steps needed to solve a level, including running the model in the outer loop, is now 4k, again much lower than the corresponding MCTS run with 100k steps. See Table 1 for a summary of the imagination efficiency for the different methods.

4.5 Generalization experiments

Lastly, we probe the generalization capabilities of I2As, beyond handling random level layouts in Sokoban. Our agents were trained on levels with 4 boxes. Table 2 shows the performance of I2A when such an agent was tested on levels with different numbers of boxes, and that of the standard model-free agent for comparison. We found that I2As generalizes well; at 7 boxes, the I2A agent is still able to solve more than half of the levels, nearly as many as the standard agent on 4 boxes.

5 Learning one model for many tasks in MiniPacman

In our final set of experiments, we demonstrate how a single model, which provides the I2A with a general understanding of the dynamics governing an environment, can be used to solve a collection of different tasks. We designed a simple, light-weight domain called MiniPacman, which allows us to easily define multiple tasks in an environment with shared state transitions and which enables us to do rapid experimentation.

In MiniPacman (Fig. 6, left), the player explores a maze that contains food while being chased by ghosts. The maze also contains power pills; when eaten, for a fixed number of steps, the player moves faster, and the ghosts run away and can be eaten. These dynamics are common to all tasks. Each task is defined by a vector $w_{\text{rew}} \in \mathbb{R}^5$, associating a reward to each of the following five events: moving, eating food, eating a power pill, eating a ghost, and being eaten by a ghost. We consider five different reward vectors inducing five different tasks. Empirically we found that the reward schemes were sufficiently different to lead to very different high-performing policies⁶ (for more details on the game and tasks, see appendix C, gameplay videos can be found at [21]).

To illustrate the benefits of model-based methods in this multi-task setting, we train a single environment model to predict both observations (frames) and events (as defined above, e.g. "eating a ghost"). Note that the environment model is effectively shared across all tasks, so that the marginal cost of learning the model is nil. During training and testing, the I2As have access to the frame and reward predictions generated by the model; the latter was computed from model event predictions and the task reward vector w_{rew} . As such, the reward vector w_{rew} can be interpreted as an ‘instruction’ about which task to solve in the same environment [cf. the Frostbite challenge of 8]. For a fair comparison, we also provide all baseline agents with the event variable as input.⁷

⁶For example, in the ‘avoid’ game, any event is negatively rewarded, and the optimal strategy is for the agent to clear a small space from food and use it to continuously escape the ghosts.

⁷It is not necessary to provide the reward vector w_{rew} to the baseline agents, as it is equivalent a constant bias.

We trained baseline agents and I2As separately on each task. Results in Fig. 6 (right) indicate the benefit of the I2A architecture, outperforming the standard agent in all tasks, and the copy-model baseline in all but one task. Moreover, we found that the performance gap between I2As and baselines is particularly high for tasks 4 & 5, where rewards are particularly sparse, and where the anticipation of ghost dynamics is especially important. We posit that the I2A agent can leverage its environment and reward model to explore the environment much more effectively.

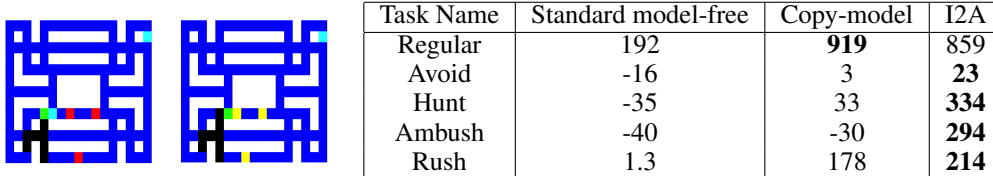


Figure 6: *Minipacman environment*. *Left*: Two frames from a minipacman game. Frames are 15×19 RGB images. The player is green, dangerous ghosts red, food dark blue, empty corridors black, power pills in cyan. After eating a power pill (right frame), the player can eat the 4 weak ghosts (yellow). *Right*: Performance after 300 million environment steps for different agents and all tasks. Note I2A clearly outperforms the other two agents on all tasks with sparse rewards.

6 Related work

Some recent work has focused on applying deep learning to model-based RL. A common approach is to learn a neural model of the environment, including from raw observations, and use it in classical planning algorithms such as trajectory optimization [26–28]. These studies however do not address a possible mismatch between the learned model and the true environment. The problem of making use of partially misspecified models has attracted attention in robotics, when transferring policies from simulation to real environments [29–31]. There, the environment model is given, not learned, and used for pretraining, not planning at test time. Liu et al. [32] also learn to extract information from trajectories, but in the context of imitation learning.

A principled way to deal with imperfect models is to capture model uncertainty, e.g. by using Gaussian Process models of the environment Deisenroth and Rasmussen [12]. The disadvantage of this method is its high computational cost; it also assumes that the model uncertainty is well calibrated and lacks a mechanism that can learn to compensate for possible miscalibration of uncertainty. Cutler et al. [33] consider RL with a hierarchy of models of increasing (known) fidelity. A recent multi-task GP extension of this study can further help to mitigate the impact of model misspecification, but again suffers from high computational burden in large domains [34].

Tamar et al. [35], Silver et al. [36], and Oh et al. [37] all present neural networks whose architectures mimic classical iterative planning algorithms, and which are trained by reinforcement learning or to predict user-defined, high-level features; in these, there is no explicit environment model. In our case, we use explicit environment models that are trained to predict low-level observations, which allows us to exploit additional unsupervised learning signals for training. This procedure is expected to be beneficial in environments with sparse rewards, where unsupervised modelling losses can complement return maximization as learning target as recently explored in Jaderberg et al. [38] and Mirowski et al. [39].

Henaff et al. [40] also learn models of discrete actions environments, and exploit the effective differentiability of the model with respect to the actions by applying continuous control planning algorithms to derive a plan. Kansky et al. [41] learn symbolic networks models of the environment and use them for planning, but are given the relevant abstractions from a hand-crafted vision system.

Close to our work is a study by Hamrick et al. [42]: they present a neural architecture that queries learned expert models, but focus on meta-control for continuous contextual bandit problems. Pascanu et al. [43] extends this work by focusing on explicit planning in sequential environments, and learn how to construct a plan iteratively.

The general idea of learning to leverage an internal model in arbitrary ways was also discussed by Schmidhuber [44].

7 Discussion

We presented I2A, an approach combining model-free and model-based ideas to implement *imagination-augmented RL*: learning to interpret environment models to augment model-free decisions. I2A outperforms model-free baselines on MiniPacman and on the challenging, combinatorial domain of Sokoban. We demonstrated that, unlike classical model-based RL and planning methods, I2A is able to successfully use imperfect models (including models without reward predictions), hence significantly broadening the applicability of model-based RL concepts and ideas.

As all model-based RL methods, I2As trade-off environment interactions for computation by pondering before acting. This is essential in irreversible domains, where actions can have catastrophic outcomes, such as in Sokoban. In our experiments, the I2A was always less than an order of magnitude slower per interaction than the model-free baselines. The amount of computation can be varied (it grows linearly with the number and depth of rollouts); we therefore expect I2As to greatly benefit from advances on dynamic compute resource allocation (e.g. Graves [45]). Another avenue for future research is on abstract environment models: learning predictive models at the "right" level of complexity and that can be evaluated efficiently at test time will help to scale I2As to richer domains.

Remarkably, on Sokoban I2As compare favourably to a strong planning baseline (MCTS) with a perfect environment model: at comparable performance, I2As require far fewer function calls to the model than MCTS, because their model rollouts are guided towards relevant parts of the state space by a learned rollout policy. This points to further potential improvement by training rollout policies that "learn to query" imperfect models in a task-relevant way.

Acknowledgements

We thank Victor Valdes for designing and implementing the Sokoban environment, Joseph Modayil for reviewing an early version of this paper, and Ali Eslami, Hado Van Hasselt, Neil Rabinowitz, Tom Schaul, Yori Zwols for various help and feedback.

References

- [1] Shane Legg and Marcus Hutter. Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4):391–444, 2007.
- [2] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [3] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pages 1928–1937, 2016.
- [4] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 1889–1897, 2015.
- [5] Edward C Tolman. Cognitive maps in rats and men. *Psychological Review*, 55(4):189, 1948.
- [6] Anthony Dickinson and Bernard Balleine. *The Role of Learning in the Operation of Motivational Systems*. John Wiley & Sons, Inc., 2002.
- [7] Brad E Pfeiffer and David J Foster. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447):74–79, 2013.
- [8] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *arXiv preprint arXiv:1604.00289*, 2016.
- [9] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [10] Jing Peng and Ronald J Williams. Efficient learning and planning within the dyna framework. *Adaptive Behavior*, 1(4):437–454, 1993.

- [11] Pieter Abbeel and Andrew Y Ng. Exploration and apprenticeship learning in reinforcement learning. In *Proceedings of the 22nd international conference on Machine learning*, pages 1–8. ACM, 2005.
- [12] Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472, 2011.
- [13] Sergey Levine and Pieter Abbeel. Learning neural network policies with guided policy search under unknown dynamics. In *Advances in Neural Information Processing Systems*, pages 1071–1079, 2014.
- [14] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *ICLR*, 2016.
- [15] Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous deep q-learning with model-based acceleration. In *Proceedings of The 33rd International Conference on Machine Learning*, pages 2829–2838, 2016.
- [16] Erik Talvitie. Agnostic system identification for monte carlo planning. In *AAAI*, pages 2986–2992, 2015.
- [17] Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard L Lewis, and Satinder Singh. Action-conditional video prediction using deep networks in atari games. In *Advances in Neural Information Processing Systems*, pages 2863–2871, 2015.
- [18] Silvia Chiappa, Sébastien Racaniere, Daan Wierstra, and Shakir Mohamed. Recurrent environment simulators. In *5th International Conference on Learning Representations*, 2017.
- [19] Felix Leibfried, Nate Kushman, and Katja Hofmann. A deep learning approach for joint video frame and reward prediction in atari games. *CoRR*, abs/1611.07078, 2016. URL <http://arxiv.org/abs/1611.07078>.
- [20] Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5-RMSprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2), 2012.
- [21] <https://drive.google.com/open?id=0B4tKsKnCCZtQY2tT0ThucHVxUTQ>.
- [22] Gerald Tesauro and Gregory R Galperin. On-line policy improvement using monte-carlo search. In *NIPS*, volume 96, pages 1068–1074, 1996.
- [23] Rémi Coulom. Efficient selectivity and backup operators in monte-carlo tree search. In *International Conference on Computers and Games*, pages 72–83. Springer, 2006.
- [24] Benjamin E Childs, James H Brodeur, and Levente Kocsis. Transpositions and move groups in monte carlo tree search. In *Computational Intelligence and Games, 2008. CIG'08. IEEE Symposium On*, pages 389–395. IEEE, 2008.
- [25] Christopher D Rosin. Nested rollout policy adaptation for monte carlo tree search. In *Ijcai*, pages 649–654, 2011.
- [26] Manuel Watter, Jost Springenberg, Joschka Boedecker, and Martin Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in Neural Information Processing Systems*, pages 2746–2754, 2015.
- [27] Ian Lenz, Ross A Knepper, and Ashutosh Saxena. DeepMPC: Learning deep latent features for model predictive control. In *Robotics: Science and Systems*, 2015.
- [28] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [29] Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(Jul):1633–1685, 2009.
- [30] Eric Tzeng, Coline Devin, Judy Hoffman, Chelsea Finn, Xingchao Peng, Sergey Levine, Kate Saenko, and Trevor Darrell. Towards adapting deep visuomotor representations from simulated to real environments. *arXiv preprint arXiv:1511.07111*, 2015.
- [31] Paul Christiano, Zain Shah, Igor Mordatch, Jonas Schneider, Trevor Blackwell, Joshua Tobin, Pieter Abbeel, and Wojciech Zaremba. Transfer from simulation to real world through learning deep inverse dynamics model. *arXiv preprint arXiv:1610.03518*, 2016.

- [32] YuXuan Liu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Imitation from observation: Learning to imitate behaviors from raw video via context translation. *arXiv preprint arXiv:1707.03374*, 2017.
- [33] Mark Cutler, Thomas J Walsh, and Jonathan P How. Real-world reinforcement learning via multifidelity simulators. *IEEE Transactions on Robotics*, 31(3):655–671, 2015.
- [34] Alonso Marco, Felix Berkenkamp, Philipp Hennig, Angela P Schoellig, Andreas Krause, Stefan Schaal, and Sebastian Trimpe. Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with bayesian optimization. *arXiv preprint arXiv:1703.01250*, 2017.
- [35] Aviv Tamar, Yi Wu, Garrett Thomas, Sergey Levine, and Pieter Abbeel. Value iteration networks. In *Advances in Neural Information Processing Systems*, pages 2154–2162, 2016.
- [36] David Silver, Hado van Hasselt, Matteo Hessel, Tom Schaul, Arthur Guez, Tim Harley, Gabriel Dulac-Arnold, David Reichert, Neil Rabinowitz, Andre Barreto, et al. The predictron: End-to-end learning and planning. *arXiv preprint arXiv:1612.08810*, 2016.
- [37] Junhyuk Oh, Satinder Singh, and Honglak Lee. Value prediction network. *arXiv preprint arXiv:1707.03497*, 2017.
- [38] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.
- [39] Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andy Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, et al. Learning to navigate in complex environments. *arXiv preprint arXiv:1611.03673*, 2016.
- [40] Mikael Henaff, William F Whitney, and Yann LeCun. Model-based planning in discrete action spaces. *arXiv preprint arXiv:1705.07177*, 2017.
- [41] Ken Kanksy, Tom Silver, David A Mély, Mohamed Eldawy, Miguel Lázaro-Gredilla, Xinghua Lou, Nimrod Dorfman, Szymon Sidor, Scott Phoenix, and Dileep George. Schema networks: Zero-shot transfer with a generative causal model of intuitive physics. *Accepted at International Conference for Machine Learning, 2017*, 2017.
- [42] Jessica B. Hamrick, Andy J. Ballard, Razvan Pascanu, Oriol Vinyals, Nicolas Heess, and Peter W. Battaglia. Metacontrol for adaptive imagination-based optimization. In *Proceedings of the 5th International Conference on Learning Representations (ICLR 2017)*, 2017.
- [43] Razvan Pascanu, Yujia Li, Oriol Vinyals, Nicolas Heess, David Reichert, Theophane Weber, Sebastien Racaniere, Lars Buesing, Daan Wierstra, and Peter Battaglia. Learning model-based planning from scratch. *arXiv preprint*, 2017.
- [44] Jürgen Schmidhuber. On learning to think: Algorithmic information theory for novel combinations of reinforcement learning controllers and recurrent neural world models. *arXiv preprint arXiv:1511.09249*, 2015.
- [45] Alex Graves. Adaptive computation time for recurrent neural networks. *arXiv preprint arXiv:1603.08983*, 2016.
- [46] Leemon C Baird III. Advantage updating. Technical report, Wright Lab. Technical Report WL-TR-93-1146., 1993.
- [47] John Schulman, Nicolas Heess, Theophane Weber, and Pieter Abbeel. Gradient estimation using stochastic computation graphs. In *Advances in Neural Information Processing Systems*, pages 3528–3536, 2015.
- [48] Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *European conference on machine learning*, pages 282–293. Springer, 2006.
- [49] Sylvain Gelly and David Silver. Combining online and offline knowledge in uct. In *Proceedings of the 24th international conference on Machine learning*, pages 273–280. ACM, 2007.
- [50] Joshua Taylor and Ian Parberry. Procedural generation of sokoban levels. In *Proceedings of the International North American Conference on Intelligent Games and Simulation*, pages 5–12, 2011.
- [51] Yoshio Murase, Hitoshi Matsubara, and Yuzuru Hiraga. Automatic making of sokoban problems. *PRICAI'96: Topics in Artificial Intelligence*, pages 592–600, 1996.

A Training and rollout policy distillation details

Each agent used in the paper defines a stochastic policy, i.e. a categorical distribution $\pi(a_t|o_t; \theta)$ over discrete actions a . The logits of $\pi(a_t|o_t; \theta)$ are computed by a neural network with parameters θ , taking observation o_t at timestep t as input. During training, to increase the probability of rewarding actions being taken, A3C applies an update $\Delta\theta$ to the parameters θ using policy gradient $g(\theta)$:

$$g(\theta) = \nabla_{\theta} \log \pi(a_t|o_t; \theta) A(o_t, a_t)$$

where $A(o_t, a_t)$ is an estimate of the advantage function [46]. In practice, we learn a value function $V(o_t; \theta_v)$ and use it to compute the advantage as the difference of the bootstrapped k -step return and the current value estimate:

$$A(o_t, a_t) = \left(\sum_{t \leq t' \leq t+k} \gamma^{t'-t} r_{t'} \right) + \gamma^{k+1} V(o_{t+k+1}; \theta_v) - V(o_t; \theta_v).$$

The value function $V(o_t; \theta_v)$ is also computed as the output of a neural network with parameters θ_v . The input to the value function network was chosen to be the second to last layer of the policy network that computes π . The parameter θ_v are updated with $\Delta\theta_v$ towards bootstrapped k -step return:

$$g(\theta_v) = -A(o_t, a_t) \partial_{\theta_v} V(o_t; \theta_v)$$

In our numerical implementation, we express the above updates as gradients of a corresponding surrogate loss [47]. To this surrogate loss, we add an entropy regularizer of $\lambda_{\text{ent}} \sum_{a_t} \pi(a_t|o_t; \theta) \log \pi(a_t|o_t; \theta)$ to encourage exploration, with $\lambda_{\text{ent}} = 10^{-2}$ throughout all experiments. Where applicable, we add a loss for policy distillation consisting of the cross-entropy between π and $\hat{\pi}$:

$$l_{\text{dist}}(\pi, \hat{\pi})(o_t) = \lambda_{\text{dist}} \sum_a \bar{\pi}(a|o_t) \log \hat{\pi}(a|o_t),$$

with scaling parameter λ_{dist} . Here $\bar{\pi}$ denotes that we do not backpropagate gradients of l_{dist} wrt. to the parameters of the rollout policy through the behavioral policy π . Finally, even though we pre-trained our environment models, in principle we can also learn it jointly with the I2A agent by adding an appropriate log-likelihood term of observations under the model. We will investigate this in future research. We optimized hyperparameters (learning rate and momentum of the RMSprop optimizer, gradient clipping parameter, distillation loss scaling λ_{dist} where applicable) separately for each agent (I2A and baselines).

B Agent and model architecture details

We used rectified linear units (ReLUs) between all hidden layers of all our agents. For the environment models, we used leaky ReLUs with a slope of 0.01.

B.1 Agents

Standard model-free baseline agent

The standard model-free baseline agent, taken from [3], is a multi-layer convolutional neural network (CNN), taking the current observation o_t as input, followed by a fully connected (FC) hidden layer. This FC layer feeds into two heads: into a FC layer with one output per action computing the policy logits $\log \pi(a_t|o_t, \theta)$; and into another FC layer with a single output that computes the value function $V(o_t; \theta_v)$. The sizes of the layers were chosen as follows:

- for MiniPacman: the CNN has two layers, both with 3x3 kernels, 16 output channels and strides 1 and 2; the following FC layer has 256 units
- for Sokoban: the CNN has three layers with kernel sizes 8x8, 4x4, 3x3, strides of 4, 2, 1 and number of output channels 32, 64, 64; the following FC has 512 units

I2A

The model free path of the I2A consists of a CNN identical to one of the standard model-free baseline (without the FC layers). The rollout encoder processes each frame generated by the environment model with another identically sized CNN. The output of this CNN is then concatenated with the reward prediction (single scalar broadcast into frame shape). This feature is the input to an LSTM with 512 (for Sokoban) or 256 (for MiniPacman) units. The same LSTM is used to process all 5 rollouts (one per action); the last output of the LSTM for all rollouts are concatenated into a single vector c_{ia} of length 2560 for Sokoban, and 1280 on MiniPacman. This vector is concatenated with the output c_{mf} of the model-free CNN path and is fed into the FC layers computing policy logits and value function as in the baseline agent described above.

Copy-model

The copy-model agent has the exact same architecture as the I2A, with the exception of the environment model being replaced by the identity function (constantly returns the input observation).

B.2 Environment models

For the I2A, we pre-train separate auto-regressive models of order 1 for the raw pixel observations of the MiniPacman and Sokoban environments (see figures 7 and 8). In both cases, the input to the model consisted of the last observation o_t , and a broadcasted, one-hot representation of the last action a_t . Following previous studies, the outputs of the models were trained to predict the next frame o_{t+1} by stochastic gradient descent on the Bernoulli cross-entropy between network outputs and data o_{t+1} .

MiniPacman model

The input and output frames were of size $15 \times 19 \times 3$ (width x height x RGB). The model is depicted in figure 7. It consisted of a size preserving, multi-scale CNN architecture with additional FC layers for reward prediction. In order to capture long-range dependencies across pixels, we also make use of a layer we call *pool-and-inject*, which applies global max-pooling over each feature map and broadcasts the resulting values as feature maps of the same size and concatenates the result to the input. Pool-and-inject layers are therefore size-preserving layers which communicate the max-value of each layer globally to the next convolutional layer.

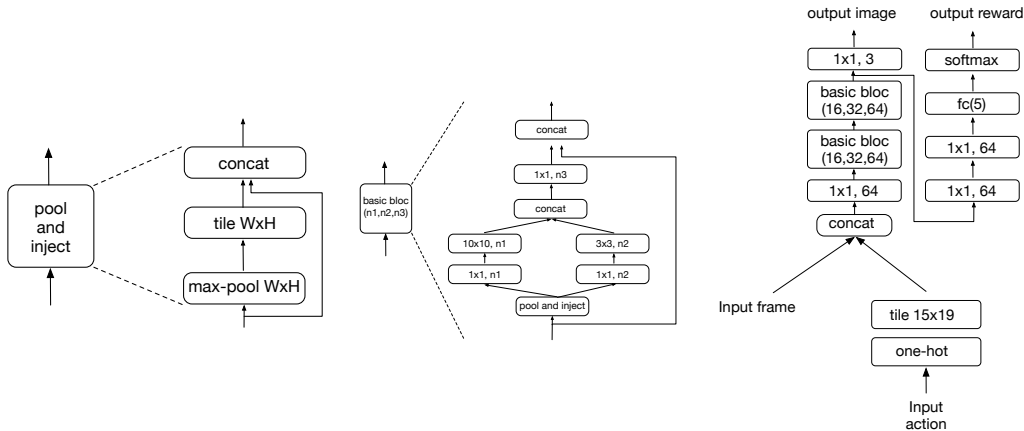


Figure 7: The MiniPacman environment model. The overview is given in the right panel with blow-ups of the basic convolutional building block (middle panel) and the pool-and-inject layer (left panel). The basic build block has three hyperparameters n_1, n_2, n_3 determining the number of channels in the convolutions; their numeric values are given in the right panel.

Sokoban model

The Sokoban model was chosen to be a residual CNN with an additional CNN / fully-connected MLP pathway for predicting rewards. The input of size $80 \times 80 \times 3$ was first processed with convolutions

with a large 8x8 kernel and stride of 8. This reduced representation was further processed with two size preserving CNN layers before outputting a predicted frame by a 8x8 convolutional layer.

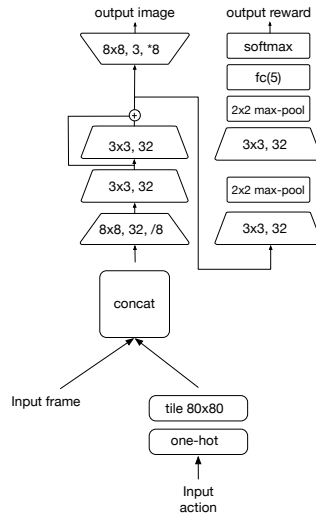


Figure 8: The sokoban environment model.

C MiniPacman additional details

MiniPacman is played in a 15×19 grid-world. Characters, the ghosts and Pacman, move through a maze. Walls positions are fixed. At the start of each level 2 power pills, a number of ghosts, and Pacman are placed at random in the world. Food is found on every square of the maze. The number of ghosts on level k is $1 + \frac{\text{level}-1}{2}$ rounded down, where level = 1 on the first level.

Game dynamics

Ghosts always move by one square at each time step. Pacman usually moves by one square, except when it has eaten a power pill, which makes it move by two squares at a time. When moving by 2 squares, if Pacman's new position ends up inside a wall, then it is moved back by one square to get back to a corridor.

We say that Pacman and a ghost meet when they either end up at the same location, or when their path crosses (even if they do not end up at the same location). When Pacman moves to a square with food or a power pill, it eats it. Eating a power pill gives Pacman super powers, such as moving at double speed and being able to eat ghosts. The effects of eating a power pill last for 19 time steps. When Pacman meets a ghost, either Pacman dies eaten by the ghost, or, if Pacman has recently eaten a power pill, the ghost dies eaten by Pacman.

If Pacman has eaten a power pill, ghosts try to flee from Pacman. They otherwise try to chase Pacman. A more precise algorithm for the movement of a ghost is given below in pseudo code, see Algorithm 1.

Task collection

We used 5 different tasks available in MiniPacman. They all share the same environment dynamics (layout of maze, movement of ghosts, ...), but vary in their reward structure and level termination. The rewards associated with various events for each task are given in the table below.

Algorithm 1 move ghost

```
1: function MOVEGHOST
2:   Inputs: Ghost object ▷ Contains position and some helper methods
3:   PossibleDirections  $\leftarrow$  [DOWN, LEFT, RIGHT, UP]
4:   CurrentDirection  $\leftarrow$  Ghost.current_direction
5:   AllowedDirections  $\leftarrow$  []
6:   for dir in PossibleDirections do
7:     if Ghost.can_move(dir) then
8:       AllowedDirections += [dir]
9:     if len(AllowedDirections) == 2 then ▷ We are in a straight corridor, or at a bend
10:      if Ghost.current_direction in AllowedDirections then
11:        return Ghost.current_direction
12:      if opposite(Ghost.current_direction) == AllowedDirections[0] then
13:        return AllowedDirections[1]
14:      return AllowedDirections[0]
15:     else ▷ We are at an intersection
16:       if opposite(Ghost.current_direction) in AllowedDirections then
17:         AllowedDirections.remove(opposite(Ghost.current_direction)) ▷ Ghosts do
not turn around
18:       X = normalise(Pacman.position - Ghost.position)
19:       DotProducts = []
20:       for dir in AllowedDirections do
21:         DotProducts += [dot_product(X, dir)]
22:       if Pacman.ate_super_pill then
23:         return AllowedDirections[argmin(DotProducts)] ▷ Away from Pacman
24:       else
25:         return AllowedDirections[argmax(DotProducts)] ▷ Towards Pacman
```

Task	At each step	Eating food	Eating power pill	Eating ghost	Killed by ghost
Regular	0	1	2	5	0
Avoid	0.1	-0.1	-5	-10	-20
Hunt	0	0	1	10	-20
Ambush	0	-0.1	0	10	-20
Rush	0	-0.1	10	0	0

When a level is cleared, a new level starts. Tasks also differ in the way a level was *cleared*.

- Regular: level is cleared when all the food is eaten;
- Avoid: level is cleared after 128 steps;
- Hunt: level is cleared when all ghosts are eaten or after 80 steps.
- Ambush: level is cleared when all ghosts are eaten or after 80 steps.
- Rush: level is cleared when all power pills are eaten.

There are no lives, and episode ends when Pacman is eaten by a ghost.

The time left before the effect of the power pill wears off is shown using a pink shrinking bar at the bottom of the screen as in Fig. 9.

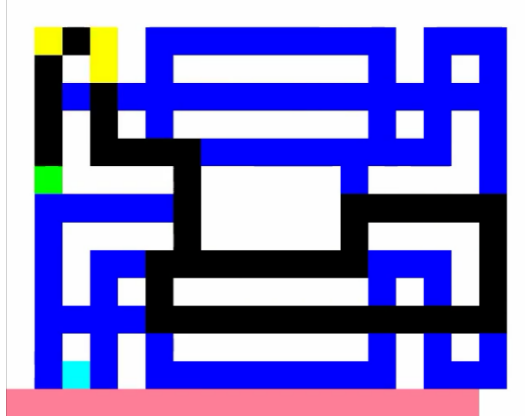


Figure 9: The pink bar appears when Pacman eats a power pill, and it decreases in size over the duration of the effect of the pill.

Training curves

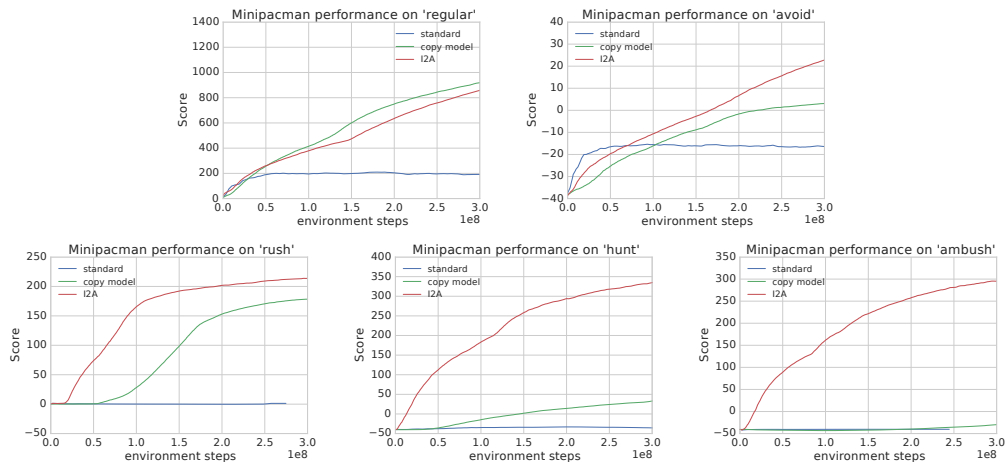


Figure 10: Learning curves for different agents and various tasks

D Sokoban additional details

D.1 Sokoban environment

In the game of Sokoban, random actions on the levels would solve levels with vanishingly small probability, leading to extreme exploration issues for solving the problem with reinforcement learning. To alleviate this issue, we use a shaping reward scheme for our version of Sokoban:

- Every time step, a penalty of -0.1 is applied to the agent.
- Whenever the agent pushes a box on target, it receives a reward of +1.
- Whenever the agent pushes a box off target, it receives a penalty of -1.
- Finishing the level gives the agent a reward of +10 and the level terminates.

The first reward is to encourage agents to finish levels faster, the second to encourage agents to push boxes onto targets, the third to avoid artificial reward loop that would be induced by repeatedly pushing a box off and on target, the fourth to strongly reward solving a level. Levels are interrupted after 120 steps (i.e. agent may bootstrap from a value estimate of the last frame, but the level resets to a new one). Identical levels are very rarely encountered during training or testing (out of 40 million

levels generated, less than 0.7% were repeated). Note that with this reward scheme, it is always optimal to solve the level (thus our shaping scheme is valid). An alternative strategy would have been to have the agent play through a curriculum of increasingly difficult tasks; we expect both strategies to work similarly.

D.2 Details for: Imagination efficiency and comparison with perfect-model planning methods

We applied the mental retry method using our I2A architecture, and our near perfect environment model of Sokoban. We let the agent try to solve the levels up to 16 times within its internal model. The base I2A architecture was solving around 87% of levels; nested rollouts boosted its performance to around 95% of levels solved. Although the agent was allowed up to 16 mental retries, in practice all the performance increase was obtained within the first 10 mental retries. Exact percentage gain by each mental retry is shown in Fig. 11. Note in Fig. 11, only 83% of the levels are solved on the first mental attempt, even though the I2A architecture could solve around 87% of levels. The gap is explained by the use of an environment model: although it looks nearly perfect to the naked eye, the model is not actually equivalent to the environment.

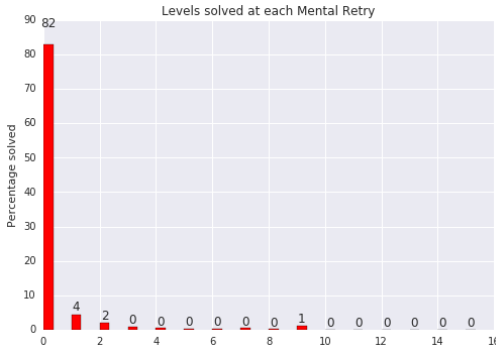


Figure 11: Gain in percentage by each additional mental retry using a near perfect environment model.

D.3 Planning with the perfect model and Monte-Carlo Tree Search in Sokoban

We first trained a value network that estimates the value function of a trained model-free policy using Monte-Carlo returns from one million different episodes. This policy solves approximately 50% of the levels in Sokoban. The value network architecture is a residual network which stacks one convolution layer and 3 convolution blocks with a final fully-connected layer of 128 hidden units. The first convolution is 1×1 convolution with 128 feature maps. Each of the three residual convolution block is composed of two convolutional layers; the first is a 1×1 convolution with 32 feature maps, the second a 3×3 convolution with 32 feature maps, and the last a 1×1 layer with 128 feature maps. To help the value networks, we trained them not on the pixel representation, but on a $10 \times 10 \times 4$ symbolic representation.

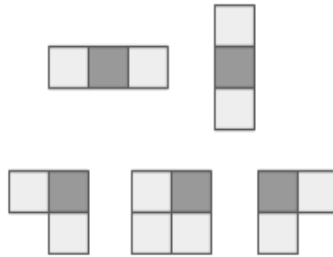
The trained value network is then employed during search to evaluate leaf-nodes — similar to [9], replacing the role of traditional random rollouts in MCTS. The tree policy uses [48, 49] with a fine-tuned exploration constant of 1. Depth-wise transposition tables for the tree nodes are used to deal with the symmetries in the Sokoban environment. External actions are selected by taking the max Q value at the root node. The tree is reused between steps but selecting the appropriate subtree as the root node for the next step.

Reported results are obtained by averaging the results over 250 episodes.

D.4 Level Generation for Sokoban

We detail here our procedural generation for Sokoban levels - we follow closely methods described in [50, 51].

The generation of a Sokoban level involves three steps: room topology generation, position configuration and room reverse-playing. Topology generation: Given an initial width*height room entirely constituted by wall blocks, the topology generation consists in creating the ‘empty’ spaces (i.e. corridors) where boxes, targets and the player can be placed. For this simple random walk algorithm with a configurable number of steps is applied: a random initial position and direction are chosen. Afterwards, for every step, the position is updated and, with a probability $p = 0.35$, a new random direction is selected. Every ‘visited’ position is emptied together with a number of surrounding wall blocks, selected by randomly choosing one of the following patterns indicating the adjacent room blocks to be removed (the darker square represents the reference position, that is, the position being visited). Note that the room ‘exterior’ walls are never emptied, so from a width×height room only a $(width-2) \times (height-2)$ space can actually be converted into corridors. The random walk approach guarantees that all the positions in the room are, in principle, reachable by the player. A relatively small probability of changing the walk direction favours the generation of longer corridors, while the application of a random pattern favours slightly more convoluted spaces. Position configuration:



Once a room topology is generated, the target locations for the desired N boxes and the player initial position are randomly selected. There is the obvious prerequisite of having enough empty spaces in the room to place the targets and the player but no other constraints are imposed in this step.

Reverse playing: Once the topology and targets/player positions are generated the room is reverse-played. In this case, on each step, the player has eight possible actions to choose from: simply moving or moving+pulling from a box in each possible direction (assuming for the latter, that there is a box adjacent to the player position).

Initially the room is configured with the boxes placed over their corresponding targets. From that position a depth-first search (with a configurable maximum depth) is carried out over the space of possible moves, by ‘expanding’ each reached player/boxes position by iteratively applying all the possible actions (which are randomly permuted on each step). An entire tree is not explored as there are different combinations of actions leading to repeated boxes/player configurations which are skipped.

Statistics are collected for each boxes/player configuration, which is, in turn, scored with a simple heuristic:

$$\text{RoomScore} = \text{BoxSwaps} \times \sum_i \text{BoxDisplacement}_i$$

where BoxSwaps represents the number of occasions in which the player stopped pulling from a given box and started pulling from a different one, while BoxDisplacement represents the Manhattan distance between the initial and final position of a given box. Also whenever a box or the player are placed on top of one of the targets the RoomScore value is set to 0. While this scoring heuristic doesn’t guarantee the complexity of the generated rooms it’s aimed to a) favour room configurations where overall the boxes are further away from their original positions and b) increase the probability of a room requiring a more convoluted combination of box moves to get to a solution (by aiming for solutions with higher boxSwaps values). This scoring mechanism has empirically proved to generate levels with a balanced combination of difficulties.

The reverse playing ends when there are no more available positions to explore or when a predefined maximum number of possible room configurations is reached. The room with the higher RoomScore is then returned.

Default parameters:

- A maximum of 10 room topologies and for each of those 10 boxes/player positioning are retried in case a given combination doesn't produce rooms with a score > 0 .
- The room configuration tree is by default limited to a maximum depth of 300 applied actions.
- The total number of visited positions is by default limited to 1000000.
- Default random-walk steps: $1.5 \times (\text{room width} + \text{room height})$.