

Combined Reinforcement Learning via Abstract Representations

Vincent François-Lavet

McGill University, Mila
vincent.francois-lavet@mcgill.ca

Doina Precup

McGill University, Mila, DeepMind
dprecup@cs.mcgill.ca

Yoshua Bengio

Université de Montreal, Mila
yoshua.bengio@mila.quebec

Joelle Pineau

McGill University, Mila, Facebook AI Research
jpineau@cs.mcgill.ca

Abstract

In the quest for efficient and robust reinforcement learning methods, both model-free and model-based approaches offer advantages. In this paper we propose a new way of explicitly bridging both approaches via a shared low-dimensional learned encoding of the environment, meant to capture summarizing abstractions. We show that the modularity brought by this approach leads to good generalization while being computationally efficient, with planning happening in a smaller latent state space. In addition, this approach recovers a sufficient low-dimensional representation of the environment, which opens up new strategies for interpretable AI, exploration and transfer learning.

1 Introduction

In reinforcement learning (RL), there are two main approaches to learn how to perform sequential decision-making tasks from experience. The first approach is the model-based approach where the agent learns a model of the environment (the dynamics and the rewards) and then makes use of a planning algorithm to choose the action at each time step. The second approach, so-called model-free, builds directly a policy or an action-value function (from which an action choice is straightforward). For some tasks, the structure of the policy (or action-value function) offers more regularity and thus a model-free approach would be more efficient, whereas in other tasks it may be easier to learn the dynamics directly due to some structure of the environment in which case a model-based approach would be preferable. In practice, it is possible to develop a combined approach that incorporates both strategies.

We present a novel deep RL architecture, which we call CRAR (Combined Reinforcement via Abstract Representations). The CRAR agent combines model-based and model-free components, with the additional specificity that the proposed model forces both components to jointly infer a sufficient abstract representation of the environment. This is achieved by explicitly training both the model-based and the model-free components end-to-end, including the joint abstract representation. To ensure the expressiveness of the abstract state, we also introduce an approximate entropy maximization penalty in the objective function, at the output

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

of the encoder. As compared to previous works that build implicitly an abstract representation through model-free objectives (see Section 5 for details), the CRAR agent creates a low-dimensional representation that captures meaningful dynamics, even in the absence of any reward (thus without the model-free part). In addition, our approach is modular thanks to the explicit learning of the model-based and model-free components. The main elements of the CRAR architecture are illustrated in Figure 1.

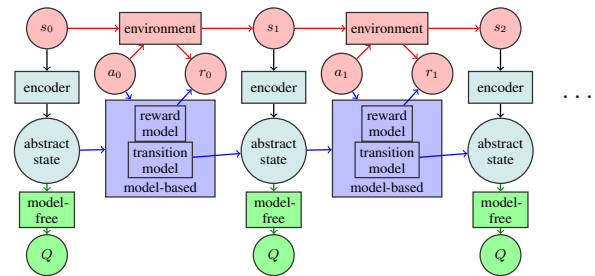


Figure 1: Illustration of the integration of model-based and model-free RL in the CRAR architecture, with a low-dimensional abstract state over which transitions and rewards are modeled. The elements related to the actual environment dynamics are in red (the state s_t , the action a_t and the reward r_t). The model-free elements are depicted in green (value function $Q(s, a)$) while the model-based elements (transition model and reward model) are in blue. The encoder and the abstract state are shared for both the model-based and model-free approaches and are depicted in light cyan. Note that the CRAR agent can learn from any off-policy data (red circles).

Learning everything through the abstract representation has the following advantages:

- it ensures that the features inferred in the abstract state provide good generalization, since they must be effective for both the model-free and the model-based predictions;
- it enables computationally efficient planning within the model-based module since planning is done over the abstract state space;
- it facilitates interpretation of the decisions taken by the agent by expressing dynamics and rewards over the abstract state;

- it allows developing new exploration strategies based on this low-dimensional representation of the environment.

In the experimental section, we show for two contrasting domains that the CRAR agent is able to build an interpretable low-dimensional representation of the task and that it can use it for efficient planning. We also show that the CRAR agent leads to effective multi-task generalization and that it can efficiently be used for transfer learning.

2 Formal setting

We consider an agent interacting with its environment over discrete time steps. The environment is modeled as an MDP (Bellman 1957), defined by (i) a state space \mathcal{S} that is discrete or continuous; (ii) a discrete action space $\mathcal{A} = \{1, \dots, N_{\mathcal{A}}\}$; (iii) the environment’s transition function $T : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, which we assume to be deterministic in this paper (although it can be extended to the stochastic case as discussed in Section 6); (iv) the environment’s reward function $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}$ where \mathcal{R} is a continuous set of possible rewards in a range $R_{\max} \in \mathbb{R}^+$ (e.g., $[0, R_{\max}]$); and (v) a general discount factor $G : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$, similarly to White (2016)¹. This setting encompasses the partially observable case if we consider that the state is a history of actions, rewards and observations.

The environment starts in a distribution of initial states $b(s_0)$. At time step t , the agent chooses an action based on the state of the system $s_t \in \mathcal{S}$ according to a policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$. After taking action $a_t \sim \pi(s_t, \cdot)$, the agent then observes a new state $s_{t+1} \in \mathcal{S}$ as well as a reward signal $r_t \in \mathcal{R}$ and a discount $\gamma_t \in \mathcal{G}$. The objective is to optimize an expected return $V^\pi(s) : \mathcal{S} \rightarrow \mathbb{R}$ such that

$$V^\pi(s) = \mathbb{E} \left[r_t + \sum_{k=1}^{\infty} \left(\prod_{i=0}^{k-1} \gamma_{t+i} \right) r_{t+k} \mid s_t = s, \pi \right], \quad (1)$$

where $r_t = \mathbb{E}_{a \sim \pi(s_t, \cdot)} R(s_t, a)$, $\gamma_t = \mathbb{E}_{a \sim \pi(s_t, \cdot)} G(s_t, a, s_{t+1})$, and $s_{t+1} = T(s_t, a_t)$.

3 The CRAR agent

We now describe in more detail the proposed CRAR approach illustrated in Figure 1.

3.1 CRAR components and notations

We define an abstract state as $x \in \mathcal{X}$ where $\mathcal{X} = \mathbb{R}^{n_{\mathcal{X}}}$ and $n_{\mathcal{X}} \in \mathbb{N}$ is the dimension of the continuous abstract state space. We define an encoder $e : \mathcal{S} \rightarrow \mathcal{X}$ as a function parametrized by θ_e , which maps the raw state s to the abstract state x . We also define the *internal* (or model) transition dynamics $\tau : \mathcal{X} \times \mathcal{A} \rightarrow \mathcal{X}$, parametrized by θ_τ :

¹The dependence of the discount factor on the transition is used for terminal states, where $G = 0$. This is necessary so that the agent captures properly the implication of the end of an episode when planning (the cumulative future rewards equal to 0 in a terminal state). Note that a biased discount factor $\Gamma : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is used during the training phase with $\Gamma \leq G, \forall (s, a, s') \in (\mathcal{S} \times \mathcal{A} \times \mathcal{S})$.

$x' = x + \tau(x, a; \theta_\tau)$. In addition, we define the internal (or model) reward function $\rho : \mathcal{X} \times \mathcal{A} \rightarrow \mathcal{R}$, parametrized by θ_ρ . For planning, we also need to fit the expected discount factor thanks to $g : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$, parametrized by θ_g .

In this paper, we investigate a model-free architecture with a Q-network $Q : \mathcal{X} \times \mathcal{A} \rightarrow \mathbb{R}$, parametrized by θ_Q : $Q(x, a; \theta_Q)$, which estimates the expected value of discounted future returns.

3.2 Learning the model

Ideally a model-free learner uses an off-policy algorithm that can use past experience (with a replay memory) that is not necessarily obtained under the current policy. We use a variant of the DQN algorithm (Mnih et al. 2015), called the double DQN algorithm (van Hasselt, Guez, and Silver 2016). The current Q-value $Q(x, a; \theta_k)$ (for the abstract state x relative to state s , when action a is performed) is updated from a set of tuples (s, a, r, γ, s') (with r and s' the observed reward and next-state), at every iteration, towards a target value:

$$Y_k = r + \gamma Q \left(e(s'; \theta_e^-), \operatorname{argmax}_{a \in \mathcal{A}} Q(e(s'; \theta_e), a; \theta_Q); \theta_Q^- \right), \quad (2)$$

where, at any step k , θ_e^- and θ_Q^- are the parameters of earlier (buffered) encoder and Q-network, which together are called the target network. The training is done by minimizing the loss

$$\mathcal{L}_{\text{mf}}(\theta_e, \theta_Q) = (Q(e(s; \theta_e), a; \theta_Q) - Y_k)^2. \quad (3)$$

These losses are back-propagated into the weights of both the encoder and the Q-network. The model-free component of the CRAR agent could benefit in a straightforward way from using any other existing variant of DQN (Hessel et al. 2017) or actor-critic architectures (Mnih et al. 2016), where the latter would be able to deal with continuous action space or stochastic policies.

The model-based part is trained using data from the sequence of tuples (s, a, r, γ, s') . We have one loss for learning the reward, one for the discount factor², and one for learning the transition³:

$$\mathcal{L}_\rho(\theta_e, \theta_\rho) = |r - \rho(e(s; \theta_e), a; \theta_\rho)|^2, \quad (4)$$

$$\mathcal{L}_g(\theta_e, \theta_g) = |\gamma - g(e(s; \theta_e), a; \theta_g)|^2, \quad (5)$$

$$\mathcal{L}_\tau(\theta_e, \theta_\tau) = |(e(s; \theta_e) + \tau(e(s; \theta_e), a; \theta_\tau) - e(s'; \theta_e))|^2. \quad (6)$$

These losses train the weights of both the encoder and the model-based components. These different components force the abstract state to represent the important low-dimensional features of the environment. The model-based and the model-free approaches are complementary and both contribute to the abstract state representation.

In practice, the problem that may appear is that a local minimum is found where too much information is lost in

²This could be extended to the current option in an option-critic architecture.

³This loss is not applied when $\gamma = 0$.

the representation $x = e(s; \theta_e)$. Keep in mind that if only the transition loss was considered, the optimal representation function would be a constant function (leading to 0 error in predicting the next abstract representation and a collapse of the representation). In practice the other loss terms prevent this but there is still a pressure to decrease the amount of information being represented (this will be clearly shown for the experiment described in Section 4.1 and in the ablation study in Appendix B.1). This loss of information mainly happens for states that are far (temporally) from any reward as the loss $\mathcal{L}_\tau(\theta_e, \theta_\tau)$ then tends to keep the transitions trivial in the abstract state space. In order to prevent that contraction, a loss that encourages some form of entropy maximization in the state representation can be added. In our model, we use:

$$\mathcal{L}_{d1}(\theta_e) = \exp(-C_d \|e(s_1; \theta_e) - e(s_2; \theta_e)\|_2), \quad (7)$$

where s_1 and s_2 are random states stored in the replay memory and C_d is a constant. As the successive states are less easily distinguished, we also introduce the same loss but with a particular sampling such that s_1 and s_2 are the successive states and we call it \mathcal{L}'_{d1} . Both losses are minimized.

The risk of obtaining very large values for the features of the state representation is avoided by the following loss that penalizes abstract states that are out of an L_∞ ball of radius 1 (other choices are possible):

$$\mathcal{L}_{d2}(\theta_e) = \max(\|e(s_1; \theta_e)\|_\infty^2 - 1, 0). \quad (8)$$

The loss $\mathcal{L}_d = \mathcal{L}_{d1} + \beta \mathcal{L}'_{d1} + \mathcal{L}_{d2}$ is called the representation loss and β is a scalar hyper-parameter that defines the proportion of resampling of successive states for the loss \mathcal{L}_{d1} .

At each iteration, a sum of the aforementioned losses are minimized using gradient descent⁴:

$$\mathcal{L} = \alpha (\mathcal{L}_{mf}(\theta_e, \theta_Q) + \mathcal{L}_\rho(\theta_e, \theta_\rho) + \mathcal{L}_g(\theta_e, \theta_g) + \mathcal{L}_\tau(\theta_e, \theta_\tau) + \mathcal{L}_d(\theta_e)), \quad (9)$$

where α is the learning rate. Details for the architecture and hyper-parameters used in the experiments are given in the appendix A⁵. In the experiments, we will show the effect of the different terms, in particular how it is possible to learn an abstract state representation only from $\mathcal{L}_\tau(\theta_e, \theta_\tau)$ and \mathcal{L}_d in a case where there is no reward, we will discuss the importance of the representation loss \mathcal{L}_d as well as the effect of α and β .

3.3 Interpretable AI

In several domains it may be useful to recover an interpretable solution, which has sufficient structure to be meaningful to a human. Interpretability in this context could mean that some (few) features of the state representation are distinctly affected by some actions. To achieve this, we add the following optional loss (which will be used in some of the experiments) to make the predicted abstract state change

⁴In practice, each term is minimized by mini-batch gradient descent (RMSprop in our case).

⁵The source code for all experiments is available at <https://github.com/VinF/deer/>

aligned with the chosen embedding vector $v(a)$:

$$\mathcal{L}_{interpret}(\theta_e, \theta_\tau) = -\cos\left(\tau(e(s; \theta_e), a; \theta_\tau)_{0:n}, v(a)\right), \quad (10)$$

where \cos stands for the cosine similarity⁶ and where the n -dimensional vector $v(a)$ ($n \in \mathbb{N} \leq n_{\mathcal{X}}$) provides the direction that is softly encouraged for the n first features of the transition in the abstract domain (when taking action a). The learning rate associated with that loss is denoted $\alpha_{interpret}$. The CRAR framework is sufficiently modular to incorporate other notions of interpretability; one could for instance think about maximizing the mutual information between the action a and the direction of the transitions $\tau(e(s; \theta_e), a; \theta_\tau)$, with techniques such as MINE (Belghazi et al. 2018).

3.4 Planning

The agent uses both the model-based and the model-free approaches to estimate the optimal action at each time step. The planning is divided into an expansion step and a backup step, similarly to Oh, Singh, and Lee (2017). One starts from the estimated abstract state \hat{x}_t and consider a number $b_d \leq N_{\mathcal{A}}$ of best actions based on $Q(\hat{x}_t, a; \theta_Q)$ (b_d is a hyper-parameter that simply depends on the planning depth d in our setting). By simulating these b_d actions with the model-based components, the agent reaches b_d new different \hat{x}_{t+1} . For each of these \hat{x}_{t+1} , the expansion continues for a number b_{d-1} of best actions, based on $Q(\hat{x}_{t+1}, a; \theta_Q)$. This expansion continues overall for a depth of d expansion steps. During the backup step, the agent then compares the simulated trajectories to select the next action.

We now formalize this process. The dynamics for some sequence of actions is estimated recursively as follows for any t' :

$$\hat{x}_{t'} = \begin{cases} e(s_t; \theta_e), & \text{if } t' = t \\ \hat{x}_{t'-1} + \tau(\hat{x}_{t'-1}, a_{t'-1}; \theta_\tau), & \text{if } t' > t \end{cases} \quad (11)$$

We define recursively the depth- d estimated expected return as

$$\hat{Q}^d(\hat{x}_t, a) = \begin{cases} \rho(\hat{x}_t, a) + g(\hat{x}_t, a) \max_{a' \in \mathcal{A}^*} \hat{Q}^{d-1}(\hat{x}_{t+1}, a'), & \text{if } d > 0 \\ Q(\hat{x}_t, a; \theta_k), & \text{if } d = 0 \end{cases} \quad (12)$$

where \mathcal{A}^* is the set of b_d best actions based on $Q(\hat{x}_t, a; \theta_Q)$ ($\mathcal{A}^* \subseteq \mathcal{A}$). To obtain the action selected at time t , we use a hyper-parameter $D \in \mathbb{N}$ which quantify the depth of planning. We then use a simple sum of the Q-values obtained with planning up to a depth D :

$$Q_{plan}^D(\hat{x}_t, a) = \sum_{d=0}^D \hat{Q}^d(\hat{x}_t, a). \quad (13)$$

The optimal action is given by $\operatorname{argmax}_{a \in \mathcal{A}} Q_{plan}^D(\hat{x}_t, a)$. Note that, using only b_d -best options at each expansion step is important for computational reasons. Indeed, planning has

⁶Given two vectors a and b , the cosine similarity is computed by $\frac{a \cdot b}{(\|a\| \|b\|) + \epsilon}$ where ϵ is a small real number used to avoid division by 0 when $\|a\| = 0$ or $\|b\| = 0$.




Figure 2: Representation of one state for a labyrinth task (without any reward).

a computational complexity that grows with the number of potential trajectories tested. In addition, it is also important to avoid overfitting to the model-based approach. Indeed, with a long planning horizon, the errors on the abstract states will usually grow due to the model approximation. When the internal model is accurate, a longer planning horizon and less pruning is beneficial, while if the model is inaccurate, one should rely less on planning.

4 Experiments

4.1 Labyrinth task

First, we consider a labyrinth MDP with four actions illustrated in Figure 2. The agent moves in the four cardinal directions (by 6 pixels) thanks to the four possible actions, except when the agent reaches a wall (block of 6×6 black pixels). This simple labyrinth MDP has no reward — $r = 0, \forall (s, a) \in (\mathcal{S}, \mathcal{A})$ — and no terminal state — $\gamma = 1, \forall (s, a) \in (\mathcal{S}, \mathcal{A})$. As a consequence, the reward loss \mathcal{L}_ρ , the discount loss \mathcal{L}_g and the model-free loss \mathcal{L}_{mf} are trivially learned and can be removed without any noticeable change.





Figure 3: Two-dimensional representation of the simple labyrinth environment using t-SNE (blue represents states where the agent is on the left part, green on the right part and orange in the junction). This plot is obtained by running the t-SNE algorithm from a dataset containing all possible states of the labyrinth task. The perplexity used is 20.

As can be seen in Fig 3, using techniques such as t-SNE (Maaten and Hinton 2008) are inefficient to represent a meaningful low-dimensional representation of this task⁷. This is because methods such as t-SNE or auto-encoders do not

⁷The implementation used in Figure 3 can be found at the address <https://lvdmaaten.github.io/tsne/>

make use of the dynamics and only provide a representation based on the similarity between the visual inputs. As opposed to this type of methods, we show in Figure 4a that the CRAR agent is able to build a disentangled 2D abstract representation of the states. The dataset used is made up of 5000 transitions obtained with a purely random policy. Details, hyper-parameters along with an ablation study are provided in Appendix B. This ablation study shows the importance of the representation loss \mathcal{L}_d and it also shows that replacing the representation loss \mathcal{L}_d by a reconstruction loss (via an auto-encoder) is not suitable to ensure a sufficient diversity in the low-dimensional abstract space.



(a) Without using the inter-pretability loss $\mathcal{L}_{interpr}$. (b) With enforcing $\mathcal{L}_{interpr}$ and $v(a_0) = [1, 0]$, the action 0 is forced to correspond to an increasing feature X_1 .

Figure 4: The CRAR agent is able to reconstruct a sensible representation of its environment in 2 dimensions.

In addition, when adding $\mathcal{L}_{interpr}$, it is shown in Figure 4b how forcing some features can be used for interpretable AI.

4.2 Catcher

The state representation is a two-dimensional array of 36×36 pixels $\in [-1, 1]$. This is illustrated in Figure 5 and details are provided in Appendix C. This environment has only a few low-dimensional underlying important features for the state representation: (i) the position of the paddle (one feature) and (ii) the position of the blocks (two features). These features are sufficient to fully define the environment at any given time. This environment illustrates that the CRAR agent is not limited to navigation tasks and the difference with the previous example is that it has an actual reward function and model-free objective.

We show in Figure 6 that all the losses behave well during training and that they can all together decrease to low values with a decreasing learning rate α . Note that all losses are learned through the abstract representation.




Figure 5: Representation of one state for the catcher environment.




Figure 6: Representation of model-based and model-free losses through training in catcher. $\alpha = 5 \times 10^{-4}$, $\beta = 0.2$ and decreasing α by 10% every 2000 training steps. All results obtained are qualitatively similar and robust to different learning rates as long as the initial learning rate α was not initialized to a too large value.

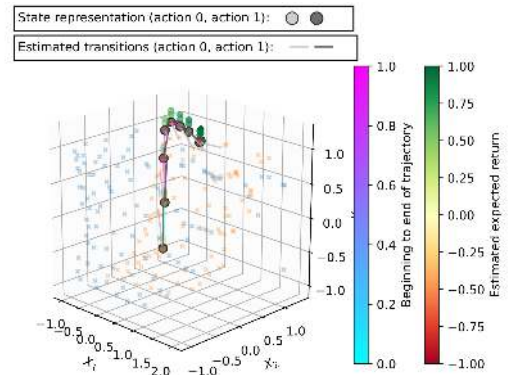
In Figure 7, it is shown that the CRAR agent is able to build a three dimensional abstract representation of its environment. Note that the CRAR agent is also able to catch the ball all the time (after 50k training steps and when following a greedy policy).

4.3 Meta-learning with limited off-policy data

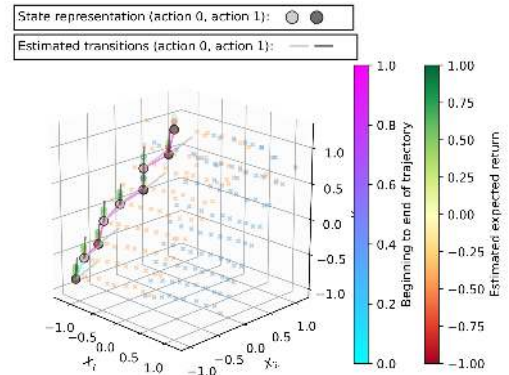
The CRAR architecture can also be used in a meta-learning setting. We consider a distribution of labyrinth tasks (over reward locations, and wall configurations), where one sample is illustrated in Figure 8. Overall, the empirical probability that two labyrinths taken randomly are the same is lower than 10^{-7} ; see details in the appendix. The reward obtained by the agent is equal to 1 when it reaches a key and it is equal to -0.1 for any other transition. We consider the batch RL setting where the agent has to build a policy offline from experience gathered following a purely random policy on a training set of 2×10^5 steps. This is equivalent to the set of transitions required in expectation to obtain the three keys by a random policy on about 500 different labyrinths (depending on the random seed). This setting makes up a more challenging task as compared to an online setting with (tens/hundreds of) millions of steps since the agent has to build a policy from a limited off-policy experience and requires strong generalization. In addition, it allows removing the exploration/exploitation influence from the experiment, thus easing the interpretation of the results.

In this context, we use a CRAR agent with an abstract state space made up of 3 channels, each of size 8×8 and the encoder is made up of CNNs only. It is shown in Figure 9 that the CRAR agent is able to achieve better data efficiency by using planning (with depth $D = 1, 3, 6$) as compared to pure model-free or pure model-based approaches.

The model-free DDQN baseline uses the same neural architecture but is trained only with the loss \mathcal{L}_{mf} . As baselines, the pure model-based approach (represented as dotted lines) performs planning similarly to the CRAR agent, but selects the branches randomly and has a constant estimate of the value function at the leafs (when $d=0$ in Equation 12). Note that, for a fair comparison, the model-based baselines



(a) Without interpretability loss



(b) We use $v(a^{(1)}) = (1, 1)$ and $v(a^{(2)}) = (-1, 1)$ such that the first feature is forced to either increase or decrease depending on the action and the second feature is forced to increase with time (for both actions).

Figure 7: Abstract representation of the domain by the CRAR agent after 50k training steps (details in the appendix). The blue and orange crosses represent respectively all possible reachable states for the ball starting respectively on the right and on the left. The trajectory is represented by the blue-purple curve (at the beginning a ball has just appeared). The colored dots represent the estimated expected return provided by $Q(x, a; \theta_Q)$. The actions taken are represented by the black/grey dots. The estimated transition are represented by straight lines (black for right, grey for left).

have similar computational cost to take a decision than the CRAR agent for a given depth d ; however, they have a worse performance due to the ablation of the model-free component.

This experiment is, to the best of our knowledge, the first that is successfully able to learn efficiently from a small set of off-policy data in a complex distribution of tasks, while using planning in an abstract state space. A discussion of other similar works are provided in Section 5.2.




Figure 8: Representation of one state for one sample labyrinth with rewards.




Figure 9: Meta-learning score on a distribution of labyrinths where the training is done with a limited number of transitions obtained off-line by a random policy. An epoch is considered to be every 2000 gradient descent steps (on all the losses). Every epoch, 200 steps on new labyrinths from the distribution are taken using different planning depths and the considered score is the running average of 10 such scores. The reported score is the mean of that running average along with the standard deviation (10 independent runs). Dotted lines represent policies without the model-free component.

4.4 Illustration of transfer learning


The CRAR architecture has the advantage of explicitly training its different components, and hence can be used for transfer learning by retraining/replacing some of its components to adjust to new tasks. In particular, one could enforce that states related to the same underlying task but with different renderings (e.g. real and simulation) are mapped into an abstract state that is close. In that case, an agent can be trained in simulation and then deployed in a realistic setting with limited retraining.

To illustrate the possibility of using the CRAR agent for transfer, we consider the setting where after 250 epochs, the high-dimensional state representations is now the negative of the previous representations⁸. The experience available to the agent is the same as previously, except that all the (high-dimensional) state representations in the replay memory are converted to the negative images. The transfer procedure consists in forcing, by supervised learning (with a MSE error


⁸We define the negative of an image as the image where all pixels have the opposite value.

and a learning rate of 5×10^{-4}), the encoder to fit the same abstract representation for 100 negative images than for the positive images (80 images are used as training set and 20 are used as validation set).

It can be seen in Figure 10a that, with the transfer procedure, no retraining is necessary in contrast to Figure 10b. Several other approaches exist to achieve this type of transfer; but this experiment demonstrates the flexibility of replacing some of the components of the CRAR agent to achieve transfer.



(a) with transfer procedure at epoch 250.



(b) without transfer procedure at epoch 250.

Figure 10: The reported score is the mean of the running average along with the standard deviation (5 independent runs). On the first 250 epochs, training and test are done on the original distribution of labyrinths while for the remaining 250 epochs, training and test are done on the same distribution of tasks but where the states are negative images as compared to the original labyrinths.

5 Related work

5.1 Building an abstract representation

The idea of building an abstract representation with a low-dimensional representation of the important features for the task at hand is key in the whole field of deep learning and also highly prevalent in reinforcement learning. One of the key advantages of using a small but rich abstract representation is to allow for improved generalization.

One approach is to first infer a factorized set of generative factors from the observations (e.g., with an encoder-decoder architecture variant (Ha and Schmidhuber 2018; Zhang, Satija, and Pineau 2018), variational auto-encoder (Higgins

et al. 2017) or using for instance t-SNE (Maaten and Hinton 2008)). Then these features can be used as input to a reinforcement learning algorithm. The learned representation can, in some contexts, greatly help for generalization as it provides a more succinct representation that is less prone to overfitting. In our setting, using an auto-encoder loss instead of the representation loss \mathcal{L}_d does not ensure a sufficient diversity in the low-dimensional abstract space. This problem is illustrated in Appendix B.1. In addition, an auto-encoder is often too strong of a constraint. On the one hand, some features may be kept in the abstract representation because they are important for the reconstruction of the observations, while they are otherwise irrelevant for the task at hand (e.g., the color of the cars in a self-driving car context). On the other hand, crucial information about the scene may also be discarded in the latent representation, particularly if that information takes up a small proportion of the observations x in pixel space (Higgins et al. 2017).

Another approach to build a set of relevant features is to share a common representation for solving a set of tasks. The reason is that learning related tasks introduce an inductive bias that causes a model to build low level features in the neural network that can be useful for the range of tasks (Jaderberg et al. 2016).

The idea of an abstract representation can be found in neuroscience where the phenomenon of access consciousness can be seen as the formation of a low-dimensional combination of a few concepts which condition planning, communication and the interpretation of upcoming observations. In this context, the abstract state could be formed using an attention mechanism able to select specific relevant variables in a context-dependent manner (Bengio 2017).

In this work, we focus on building an abstract state that provides sufficient information to simultaneously fit an internal meaningful dynamics as well as the estimation of the expected value of an optimal policy. The CRAR agent does not make use of any reconstruction loss, but instead learns both the model-free and model-based components through the state representation. By learning these components along with an approximate entropy maximization penalty, we have shown that the CRAR agent ensures that the low-dimensional representation of the task is meaningful.

5.2 Integrating model-free and model-based

Several recent works incorporate model-based and model-free RL and achieve improved sample efficiency. The closest works to CRAR include the value iteration network (VIN) (Tamar et al. 2016), the predictron (Silver et al. 2016) and the value prediction network (VPN) architecture (Oh, Singh, and Lee 2017). VIN is a fully differentiable neural network with a planning module that learns to plan from model-free objectives. As compared to CRAR, VIN has only been shown to work for navigation tasks from one initial position to one goal position. In addition, it does not work in a smaller abstract state space. The predictron is aimed at developing an algorithm that is effective in the context of planning. It works by *implicitly* learning an internal model in an abstract state space which is used for policy evaluation. The predictron is trained end-to-end to learn,

from the abstract state space, (i) the immediate reward and (ii) value functions over multiple planning depths. The initial predictron architecture was limited to policy evaluation; it was then extended to learn an optimal policy through the VPN model. Since VPN relies on n-step Q-learning, it can not directly make use of off-policy data and is limited to the online setting. As compared to these works, we show how it is possible to *explicitly* learn both the model and a value function from off-policy data while ensuring that they are based on a shared sufficient state representation. In addition, our algorithm ensures a disentanglement of the low-dimensional abstract features, which opens up many possibilities. In particular, the obtained low-dimensional representation is still effective even in the absence of any reward (thus without the model-free part).

As compared to (Kansky et al. 2017), our approach relies directly on raw features (e.g. raw images) instead of an input of entity states. As compared to I2A (Weber et al. 2017) and many model-based approaches, our approach allows to build a model in an abstract low-dimensional space. This is more computationally efficient because planning can happen in the low-dimensional abstract state space. As compared to treeQn (Farquhar et al. 2017), the learning of the model is explicit and we show how that approach allows recovering a sufficient interpretable low-dimensional representation of the environment, even in the absence of model-free objectives.

6 Discussion

In this paper, we have shown that it is possible to learn an abstract state representation thanks to both the model-free and model-based components as well as the approximate entropy maximization penalty. In addition, we have shown that the logical steps that require planning and estimating the expected return can happen in that low-dimensional abstract state space.

Our architecture could be extended to the case of stochastic environments. For the model-based component, one could use a generative model conditioned on both the abstract state and the action, e.g. by using a GAN (Goodfellow et al. 2014). The planning algorithm should then take into account the stochastic nature of the dynamics. Concerning the model-free components, it would be possible to use the distributional representation of the value function (Bellemare, Dabney, and Munos 2017).

Exploration is also one of the most important open challenges in deep reinforcement learning. The approach developed in this paper can be used as the basis for new exploration strategies as the CRAR agent provides a low-dimensional representation of the states, which can be used to more efficiently assess novelty. An illustration of this is shown in Figure 11, in the context of the labyrinth task described in Section 4.

Finally, in this paper, we have only considered a transition model for one time step. An interesting future direction of work would be to incorporate temporal abstractions such as options (see e.g., Bacon, Harb, and Precup 2017).




Figure 11: Abstract representation of the domain when the top part has not been explored (corresponding to the left part in this 2D low-dimensional representation). Thanks to the extrapolation abilities of the internal transition function, the CRAR agent can find a sequence of actions such that the expected representation of the new state is as far from any known abstract state (previously observed) for a given metric (e.g., L_2 distance). Details related to this experiment are given in Appendix B.

Appendix

The supplementary material is available online at <https://arxiv.org/abs/1809.04506>.

References

Bacon, P.-L.; Harb, J.; and Precup, D. 2017. The option-critic architecture. In *AAAI*, 1726–1734.

Belghazi, I.; Rajeswar, S.; Baratin, A.; Hjelm, R. D.; and Courville, A. 2018. MINE: Mutual information neural estimation. *arXiv preprint arXiv:1801.04062*.

Bellemare, M. G.; Dabney, W.; and Munos, R. 2017. A distributional perspective on reinforcement learning. In *International Conference on Machine Learning*, 449–458.

Bellman, R. 1957. A markovian decision process. *Journal of Mathematics and Mechanics* 679–684.

Bengio, Y. 2017. The consciousness prior. *arXiv preprint arXiv:1709.08568*.

Farquhar, G.; Rocktäschel, T.; Igl, M.; and Whiteson, S. 2017. Treeqn and atreec: Differentiable tree planning for deep reinforcement learning. *arXiv preprint arXiv:1710.11417*.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.

Ha, D., and Schmidhuber, J. 2018. World models. *arXiv preprint arXiv:1803.10122*.

Hessel, M.; Modayil, J.; van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M.; and Silver, D. 2017. Rainbow: Combining improvements in deep reinforcement learning. *arXiv preprint arXiv:1710.02298*.

Higgins, I.; Pal, A.; Rusu, A.; Matthey, L.; Burgess, C.; Pritzel, A.; Botvinick, M.; Blundell, C.; and Lerchner, A. 2017. Darla: Improving zero-shot transfer in reinforcement

learning. In *International Conference on Machine Learning*, 1480–1490.

Jaderberg, M.; Mnih, V.; Czarnecki, W. M.; Schaul, T.; Leibo, J. Z.; Silver, D.; and Kavukcuoglu, K. 2016. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*.

Kansky, K.; Silver, T.; Mély, D. A.; Eldawy, M.; Lázaro-Gredilla, M.; Lou, X.; Dorfman, N.; Sidor, S.; Phoenix, S.; and George, D. 2017. Schema networks: Zero-shot transfer with a generative causal model of intuitive physics. In *International Conference on Machine Learning*, 1809–1818.

Maaten, L. v. d., and Hinton, G. 2008. Visualizing data using t-sne. *Journal of machine learning research* 9(Nov):2579–2605.

Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *Nature* 518(7540):529–533.

Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T. P.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*.

Oh, J.; Singh, S.; and Lee, H. 2017. Value prediction network. In *Advances in Neural Information Processing Systems*, 6120–6130.

Silver, D.; van Hasselt, H.; Hessel, M.; Schaul, T.; Guez, A.; Harley, T.; Dulac-Arnold, G.; Reichert, D.; Rabinowitz, N.; Barreto, A.; et al. 2016. The predictron: End-to-end learning and planning. *arXiv preprint arXiv:1612.08810*.

Tamar, A.; Levine, S.; Abbeel, P.; WU, Y.; and Thomas, G. 2016. Value iteration networks. In *Advances in Neural Information Processing Systems*, 2146–2154.

van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep reinforcement learning with double q-learning. In *Thirtieth AAAI Conference on Artificial Intelligence*.

Weber, T.; Racanière, S.; Reichert, D. P.; Buesing, L.; Guez, A.; Rezende, D. J.; Badia, A. P.; Vinyals, O.; Heess, N.; Li, Y.; et al. 2017. Imagination-augmented agents for deep reinforcement learning. *arXiv preprint arXiv:1707.06203*.

White, M. 2016. Unifying task specification in reinforcement learning. *arXiv preprint arXiv:1609.01995*.

Zhang, A.; Satija, H.; and Pineau, J. 2018. Decoupling dynamics and reward for transfer learning. *arXiv preprint arXiv:1804.10689*.