



Attentive Multi-Task Deep Reinforcement Learning

Timo Bräm, Gino Brunner ^[0000-0002-4341-2940],
Oliver Richter ^[0000-0001-7886-5176], and Roger Wattenhofer*

Department of Information Technology and Electrical Engineering
ETH Zurich
Switzerland

{brunnegi,richtero,wattenhofer}@ethz.ch
braem.timo@gmail.com

Abstract. Sharing knowledge between tasks is vital for efficient learning in a multi-task setting. However, most research so far has focused on the easier case where knowledge transfer is not harmful, i.e., where knowledge from one task cannot negatively impact the performance on another task. In contrast, we present an approach to multi-task deep reinforcement learning based on attention that does not require any a-priori assumptions about the relationships between tasks. Our attention network automatically groups task knowledge into sub-networks on a state level granularity. It thereby achieves positive knowledge transfer if possible, and avoids negative transfer in cases where tasks interfere. We test our algorithm against two state-of-the-art multi-task/transfer learning approaches and show comparable or superior performance while requiring fewer network parameters.

1 Introduction

Humans are often excellent role models for machines. Unlike machines, humans have been interacting with their environment since time immemorial, and this extensive experience should not be ignored. So how are we humans learning, and what can machines learn from us?

First, humans learn with a limited amount of training data, as we cannot afford to first train for an unreasonably long time before becoming active. Also, we usually do not require labeled training data, but instead rely on experience gained from interactions with our world. This situation is well represented by the reinforcement learning paradigm: We observe the environment and take actions to hopefully maximize our cumulative reward. Second, humans learn many tasks concurrently, not only because there is no time to learn all possible tasks sequentially, but also because tasks are often similar in nature, and useful strategies can be transferred between comparable tasks. This is a fundamental aspect of intelligence known as multi-task learning. Third, our brain would be overwhelmed if

* Authors listed in alphabetical order.

it had to focus on all skills acquired over the span of our lives at every point in time. Therefore, we focus our attention to a set of skills useful at the moment. If we were not able to relate similar tasks and attend to skills based on extrinsic or intrinsic cues, our brain would not be able to learn much. Recent advances in neuroscience [17] also suggest that the attention mechanisms of humans are themselves learned through reinforcement learning.

In this paper, we investigate the combination of these three paradigms, in other words, we study *attentive multi-task deep reinforcement learning*. More specifically, we employ the insight of human attention by developing a simple yet effective architecture for model-free multi-task reinforcement learning. We use a neural network based attention mechanism to focus on sub-networks depending on the current state of the environment and the task to be solved. Most recent work [29, 12, 15, 2] in the multi-task/transfer deep reinforcement learning setting capitalize on some shared property between tasks. In contrast, our approach makes no assumptions about the similarity between tasks. Instead, possible relations are automatically inferred during training.

An additional advantage of using an attention based architecture is that unrelated tasks can effectively be separated and learned in different sub-parts of the architecture. We thereby automatically embrace the *negative transfer* problem (the effect that training one task might actually harm performance on another task) which most related approaches omit in their evaluation. We show that our approach scales economically with an increasing number of tasks as the attention mechanism automatically learns to group related skills in the same part of the architecture. We back our claims by comparing against two state of the art algorithms [29, 26] on a large set of grid world tasks with different amounts of transferable knowledge. We show that our method scales better in the number of parameters per task, while achieving comparable or superior performance in terms of steps to convergence. Especially, when the action spaces of the tasks are not aligned we outperform [29, 26].¹

2 Related Work

Transfer learning in classical reinforcement learning [28] is a well established research area. Even though Lin [18] already used neural networks in combination with reinforcement learning, a renewed interest in this combination came with the recent success on Atari (DQN, [21]), followed by an increased interest in developing transfer learning techniques specific to deep learning. Parisotto et al. [22] train a neural network to predict the features and outputs of several expert DQNs and use multi-task network weights as initialization for a target task DQN. Rusu et al. [25] use a single network to match expert DQN policies from different games by policy distillation. Yin et al. [30] improve policy distillation by making the convolutional layers task specific and by using hierarchical experience replay. Schmitt et al. [27] also build on the idea of policy distillation

¹ To stimulate future research in this area, our source code is available at: <https://github.com/braemt/attentive-multi-task-deep-reinforcement-learning>.

but additionally propose to anneal the teacher signal such that the student can surpass the teacher’s performance. Further, in [11, 7, 23] knowledge is transferred from human expert demonstrations, while the algorithm of Aytar et al. [1] learns from YouTube video demonstrations. Gupta et al. [9] transfer knowledge from source to target agent by training matched feature spaces. Closely related to our approach is the work of Rajendran et al. [24] who also incorporate several sub-networks and an attention mechanism to transfer knowledge from an expert network. In contrast to the architecture described in [24] and all related work mentioned so far, our algorithm learns multiple tasks simultaneously from scratch, without guidance from any demonstrations or experts. This makes our approach self-sustained and as such more general than mentioned related work.

Glatt et al. [8] train a DQN on a source task and investigate how the learned weights, which are used as initialization for a target task, alter the performance. In a similar manner, [4, 6, 10] show that some transfer is possible by simply training one network on multiple tasks. However, since these algorithms do not incorporate any task-specific weights, the best that can be done is to interpolate between conflicting tasks. In contrast, our method allows conflicting tasks to be learned in separate networks.

One interesting line of research [15, 31, 3, 2, 16] capitalizes on transferring knowledge based on successor features, i.e., shared environment dynamics. In contrast, our method does not rely on shared environment dynamics nor action alignment across tasks.

Czarnecki et al. [5] use multiple networks similar to our approach. However, their focus is on automated curriculum learning. Therefore they adjust the policy mixing weights through population based training [13] while we learn attention weights conditioned on the task state.

Rusu et al. [26] introduce *Progressive Neural Networks* (PNN), an effective approach for learning in a sequential multi-task setting. In PNN, a new network and lateral connections for each additional task are added in order to enable knowledge transfer, which speeds up the training of subsequent tasks. The additional network parts let the architecture grow super-linearly, while our network scales economically with an increasing number of tasks. Another strong approach is introduced by Teh et al. [29]. Their algorithm, *Distral*, learns multiple tasks at once by sharing knowledge through a distillation process of an additional shared policy network. In contrast to our approach, this requires an aligned action space and a separate network for each task. We compare against Distral and PNN in our experiments.

3 Background

In reinforcement learning, an agent learns through interactions with an environment. The agent repeatedly chooses an action $a_t \in \mathcal{A}$ at step t and observes a reward $r_t \in \mathbb{R}$ and the next state $s_{t+1} \in \mathcal{S}$, where \mathcal{A} and \mathcal{S} denote the sets of possible actions and states, respectively. The agent chooses the actions according to a policy $\pi(a_t|s_t) : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ which indicates the probability of




Fig. 1: Our architecture consists of an attention network (blue), which decides the weighting of the sub-network outputs (red) to generate the policy and value function on a task and state basis. The first convolutional layers are shared between the attention network and the sub-networks. The weighted sub-network policies are transformed by a task-specific layer (in green) to account for different numbers of possible actions in different tasks.

choosing action a_t in state s_t . The objective is to find a policy that maximizes the expected discounted return, i.e., to find

$$\pi^* = \max_{\pi} \left(\mathbb{E}_{\pi} \left[\sum_{t'=0}^{\infty} \gamma^{t'} r_{t'} \right] \right)$$

where $\gamma \in [0, 1]$ is the discount factor for future rewards.

In this work, we train on this objective using asynchronous advantage actor-critic training (A3C, [20]), a well established policy gradient method that uses multiple asynchronous actors for experience collection. However, our approach is general and can be readily applied to most on- and off-policy deep reinforcement learning algorithms.

In multi-task reinforcement learning, the goal is to solve a set of tasks \mathcal{T} simultaneously by training a policy $\pi(a_t|s_t, \tau)$ and value function $V(s_t, \tau)$, also referred to as critic, for each task $\tau \in \mathcal{T}$. While the objective to maximize the discounted rewards in each of the tasks remains unchanged, an additional goal is to share knowledge between tasks to accelerate training.

4 Architecture

Our network architecture, as shown in Figure 1, consists of a number of independent sub-networks and an attention module that weights the output of all sub-networks to generate a weighted policy and value function per task. The policies are then used to choose the next action in each of the environments. The attention and sub-networks all operate on top of a shared CNN that extracts high-level features of the environments. The attention network determines whether sub-networks become specialized on certain tasks, or whether they learn

features that are shared across a group of tasks. However, we do not explicitly enforce this. Thus, we do not require any a-priori knowledge about the nature of the tasks or about their similarity. In other words, we do not make any assumptions about whether potential for positive or negative transfer exists.

4.1 Shared Feature Extractor

The first stage of our architecture consists of a CNN that outputs a state-embedding $\phi(s)$. The embedding $\phi(s)$ is shared among all following sub-networks as well as the attention network. Thus, $\phi(s)$ will learn general high-level features that are relevant for all subsequent parts of the architecture. Since we do not decrease the dimensionality of the input in these layers, the architecture can in the (worst) case, where no information can be shared, learn an approximate identity mapping from s to $\phi(s)$ and leave the specialization to the sub-networks.

4.2 Attention Network

One could think of several ways how to combine the different sub-network outputs into a policy per task. One way would be to choose in each time step one of the sub-networks directly as policy. However, this sort of hard attention leads to noisy gradients (since a stochastic sampling operation would be added to the computation graph) and no complex interactions of several sub-networks could be learned. Therefore we employ a soft attention mechanism, where the final output is a linear combination of the sub-networks' outputs. Intuitively, this allows all sub-networks that are helpful to contribute to the policy and value function. This can also be seen as an ensemble, where different sub-networks with possibly different specializations vote on the next action, but where the final decision is governed by an attention network.

More concretely, the attention network consists of a CNN that operates on the shared embedding $\phi(s)$. The output of the CNN is fed into a fully connected network (FCN) that projects the output into a latent vector. This vector is then concatenated with a one-hot encoding of the task ID τ from which the input s originates, and processed further in the fully connected network. Finally, a linear layer with softmax activation produces the attention weights $w_i(s, \tau)$, which decide the contribution of the policy and value functions of each sub-network i in state s of task τ .

4.3 Sub-Networks

We use N sub-networks that contribute to the final weighted policy and value function. The number of sub-networks can be chosen based on resource requirements and/or availability. In a practical application of our method, one would choose the maximum number of networks for which the entire model still fits into memory. Unused sub-networks can be automatically ignored by the attention network (see Section 6.4), and could potentially be pruned to reduce the

overall number of parameters. In our experiments we choose a small number of networks to show that we can achieve comparable or superior performance to state of the art methods while requiring substantially fewer parameters. Specifically, we chose the number of sub-networks N depending on the number of tasks. That is, we roughly add one sub-network for four tasks. More precisely we let $N = \lfloor (|\mathcal{T}| + 2)/4 \rfloor + 1$, as we found this scaling to work well in our experiments. The sub-networks can act independently, as in an ensemble, or specialize on certain types of (sub-)tasks. The exact mode of operation depends on the nature of the tasks and is governed by the attention network. In other words, if the attention network decides that specialization is most beneficial, then the sub-networks will be encouraged to specialize, and vice versa.

The sub-networks all have the same architecture and get the embedding $\phi(s)$ as input. First, a CNN learns to extract sub-network specific features from $\phi(s)$ that are then passed to a FCN. From the last hidden representation of the FCN, a linear layer directly outputs the value function estimate $V_i(s)$ for the i -th sub-network. A softmax layer maps the last hidden representation of the FCN to a $|\mathcal{A}_{max}|$ -dimensional vector $p_i(a|s)$, where $|\mathcal{A}_{max}|$ is the largest action space size across all tasks.

4.4 Attentive Multi-Task Network

The attention weighted $\pi_i(a|s)$ is in the end fed to a task-specific linear layer that maps it to the action dimension of each task, and a final softmax normalization is applied to generate a valid probability distribution over actions, i.e., a policy. More formally, the sub-network outputs $\pi_i(a|s)$ are combined into the final policy as

$$\pi(a|s, \tau) = \text{softmax} \left(W_\tau \cdot \left(\sum_i^N \pi_i(a|s) w_i(s, \tau) \right) + b_\tau \right)$$


where $W_\tau \in \mathbb{R}^{|\mathcal{A}_\tau| \times |\mathcal{A}_{max}|}$ is a task-specific weight matrix and $b_\tau \in \mathbb{R}^{|\mathcal{A}_\tau|}$ is a task-specific bias. Note that W_τ and b_τ are shared across the sub-networks and only depend on the task.

Putting everything together, we use the attention weights $w_i(s, \tau)$ to also compute the final value function $V(s, \tau)$ from the outputs of the sub-networks as

$$V(s, \tau) = \sum_{i=1}^N w_i(s, \tau) V_i(s)$$

5 Task Environments

To evaluate our approach we create a set of environments which are designed to have the potential for positive as well as negative knowledge transfer. Since



(a) # are walls and X is the target. $+$ ($-$) denote objects giving positive (negative) rewards. If only one bonus (penalty) object is present, it is located at $+_1$ ($-_1$).

(b) A state space of a grid world task. The player is the white square. Walls, the target, three bonus objects and one penalty object are present.

(c) Example state of the connect four task. The agent controls the dark tokens, the random opponent controls the gray tokens.

Fig. 2: (a) Grid world environment template. (b) Example of a concrete grid world instantiation. (c) Example of a concrete connect four state.

we aim at evaluating our approach on a large set of tasks, we opt for simple, easy to generate environments, even though our initial results on the Arcade Learning Environment [19] (not reported here) were promising as well. We leave the adaption of our methodology to more complex environments to future work as we aim to show the evolution of transfer depending on the number of tasks in this report, which was not feasible on more complex tasks within our resource constraints due to the large amount of networks trained (600 for Figure 4 alone) and experiments conducted.

5.1 Grid Worlds

The first set of environments contains 20 grid world tasks. The environments of this set consist of 8×8 gray-scale images representing the state of the environment. The agent is a single pixel in the grid and the possible actions are moving up, down, left or right. For all tasks, the goal is to reach a target pixel where a positive reward is received and the episode terminates. The environments can also contain additional objects that represent positive/negative rewards, as well as impassable walls. All objects in the environments are at fixed locations, and only the starting location of the player is random. Figure 2a shows the template for all tasks and Figure 2b shows an example of such an environment as seen by the agent.

In the following we give a detailed description of every variation, each defining a task. In the first task, the goal is to find the target as fast as possible. No walls or additional rewards are put into the environment, just the agent and the target. To encourage speed, the agent is penalized with a small negative reward at every

step. In the other tasks there is no such penalty, but if the player leaves the board, a negative reward of -0.5 is observed and the episode is terminated. This is also the only difference between task one and two: The goal of the second task is to reach the target without leaving the board. In the third task, a bonus object is added at location marked $+_1$ in Figure 2a that yields a positive reward when collected. In the fourth task, additionally to the bonus object of task 3, another bonus object at location $+_2$ and a penalty object at the location marked with $-_1$ are added. The penalty object yields a negative reward when collected. The fifth and sixth task both contain three bonus objects at the locations marked with a $+$ in Figure 2a, where the sixth task additionally contains another penalty object at location $-_2$. Tasks 7 and 8 are visually indistinguishable from tasks 4 and 5, but we invert the rewards of the bonus and penalty objects in order to test negative transfer. Similar to these two tasks, tasks 9 and 10 consist of three objects looking like penalty objects (at locations marked with $-$) but yielding positive reward. Task 10 additionally contains an object that looks like a bonus object at location $+_1$ (see Figure 2a) yielding a negative reward. Tasks 11 to 20 are the same as tasks 1 to 10 but additionally contain impassable walls. The maximum achievable reward is set to 1.0 for all tasks, distributed equally among bonus objects and target. For example, if there are three bonus objects, the target and bonus objects yield rewards of 0.25 each. The penalty objects give a negative reward that is equal in magnitude to the bonus objects’ positive reward. In addition, walking into a wall yields a reward of -0.5 . Furthermore, if the agent does not reach the target after 200 steps, the task terminates without any additional reward.

5.2 Connect Four

To test the behavior of our model on unrelated tasks with little to no potential for knowledge transfer, we generate environments from a completely different domain. We implement a two-player game based on connect four. Each location or token is represented by a single pixel. The agent drops in a token from the top, which then appears on top of the top most token in the column, or in the bottom row if the column was empty. The goal of this task is to have four tokens in a horizontal, vertical or diagonal line. Our connect four tasks consist of 8 rows and 8 columns, and thus looks visually similar to the grid world tasks, but has otherwise no relation to them. The agent has 8 different actions to choose from, indicating in which column the token is to be dropped. An example of this is shown in Figure 2c. If the agent plays an invalid action, i.e., if the chosen column is already full, the agent loses the game immediately. When the agent wins the game it receives a reward of 1, and -1 if it loses. In case of a tie the reward is 0. The opponent chooses a *valid* action uniformly at random. We additionally implement three variations of this basic connect four task. The goal of the first variation is to connect five tokens instead of four. The second and third variation rotate the state of the connect four and connect five tasks by 90 degrees, such that the players now choose rows and not columns.

Table 1: Architecture details for the policy networks (value output is omitted for readability). The base network is the basic network building block for Distral and PNN, each having one such base network per task. Additionally PNN has lateral connections and Distral has an additional base network for the shared policy. The columns *Shared CNN*, *Sub-networks* and *Attention network* describe our architecture (see Section 4 and Figure 1). The + in the attention network Layer 4 indicates concatenation of the task embedding.

	Base network	Shared CNN	Sub-networks	Attention network
Layer 1	3x3x16, stride 2	3x3x32, stride 2	-	-
Layer 2	3x3x16, stride 1	3x3x32, stride 1	-	-
Layer 3	3x3x16, stride 1	-	3x3x16, stride 1	3x3x16, stride 1
Layer 4	FC 256	-	FC 256	FC $N \cdot \mathcal{T} + \mathcal{T} $
Layer 5	Softmax $ \mathcal{A}_\tau $	-	Softmax $ \mathcal{A}_{max} $	FC 256
Layer 6	-	-	Softmax $ \mathcal{A}_\tau $	Softmax N

6 Experiments and Results

We evaluate the performance of our architecture on the set of grid worlds described before and compare the results to two state of the art architectures: Progressive Neural Networks (PNN) [26] and Distral [29]. PNN learns tasks sequentially by freezing already trained networks and adding an additional network for each new task. The new networks are connected to previous ones to allow knowledge transfer. The order in which the tasks are trained with our PNN implementation is sampled randomly for all experiments. In contrast to PNN, but similar to our approach, Distral learns all tasks simultaneously. Here, a distilled policy $\hat{\pi}_0$ is used for sharing and transferring knowledge, while each task also has its own network to learn task-specific policies $\hat{\pi}_\tau$. We implement the *KL+ent 2col* approach (see [29]). The distilled policy network and the task-specific networks have the same network architecture as the base PNN model which is listed as *Base network* in Table 1. Note that even though our architecture starts with more filters in the shared CNN when compared to the base architecture, this does not give us a parameter advantage, since those filters are shared across all tasks while Distral and PNN get additional CNN parameters for each additional task. For all approaches and all experiments, we use the same hyper parameters which are summarized in Table 2. We chose these hyper parameters based on the performance of all three approaches on multiple grid world tasks such that no approach has an unfair advantage. We use the smallest multiple of $|\mathcal{T}|$ (the number of tasks) which is equal or larger than 24 for the number of parallel workers in the A3C training and distribute tasks equally over the workers. The loss function is minimized with the Adam optimizer [14]. For PNN we had to reduce the number of workers to 16, as the memory consumption for a large number of tasks was too high. For Distral, we set $\alpha = 0.5$ and $\beta = 10^4$ and compute the policy as

$$\hat{\pi}_i(a|s) = \text{softmax}(\alpha h(a|s) + f(a|s))$$

Table 2: Hyper parameters used for the experiments.

Stacked input frames:	1	Discount factor γ :	0.99
Adam learning rate:	1e-4	Rollout length:	5
Adam β_1 :	0.9	Entropy regularization:	0.02
Adam β_2 :	0.999	Distral α :	0.5
Adam ϵ :	1e-08	Distral β :	10^4




Fig. 3: Number of parameters of the different architecture choices with an increasing number of tasks. *Linear* represents training each task in a separate network of the size of PNNs base model.

where h is the output of the distilled network and f is the β -scaled output of the task-specific network (see Appendix B.2 of [29]).

6.1 Model Size

First, we compare the model sizes of our Attentive Multi-Task (AMT) architecture, PNN, Distral and *Linear*. *Linear* simply represents training a separate network (same size as the base network) on each task which leads to a linear increase in parameters with each additional task. The results are shown in Figure 3. In our experiments we add a new sub-network to AMT for every fourth task, thus the number of network parameters grows more slowly with the number of tasks than in the other approaches. Depending on memory requirements we can easily increase or decrease the total number of parameters since we do not assign sub-networks to tasks a-priori; more difficult tasks can automatically be assigned more effective network capacity by the attention network.

Distral uses slightly more parameters than having a separate network for each task due to the additional distilled policy network. The only way to reduce the number of total parameters would be to decrease the size of the task networks. However, unlike our approach, doing so could more strongly affect difficult tasks that require more network capacity to be solved, or tasks that cannot profit from the distilled policy due to a lack of transfer potential. One could tune each




Fig. 4: The number of steps required when trained on a set of tasks to reach an average score of at least 0.9 and a score of at least 0.8 on each task separately over 10^5 steps. The median over 10 runs and for task set sizes between 1 and 20 is shown, where the tasks are sampled randomly from all grid worlds tasks. The shaded area represents 30% to 70% performance of the runs. The number of steps for *Linear* is calculated by extrapolating the median of all runs from the other three approaches that are trained on one task only.

task network individually and, e.g., use larger networks for more difficult tasks, but this would require a substantial tuning effort. In contrast, our method assigns effective network capacity automatically, and can thus utilize the available network parameters more efficiently.

PNN also adds a new sub-network for each task and additionally connects all existing sub-networks to the newly added one. Thus, the number of total parameters grows super-linearly in the number of tasks. This parameter explosion causes high memory consumption and high computational costs, which can quickly become a problem when training on an increasing number of tasks with limited hardware.

6.2 Sample Efficiency vs. Number of Tasks

In this section, we compare the performance of AMT to PNN and Distral when trained on an increasing number of tasks. We perform 10 runs for each approach and every number of tasks (from 1 to 20). For each of the 10 runs, the tasks are chosen uniformly at random without replacement from all 20 grid world tasks. The tasks are considered solved if the average score over 10^5 steps is at least 0.9 and each individual task has a score of at least 0.8. The results are shown in Figure 4. The number of steps required to solve a given number of tasks scales sub-linearly for all three approaches, i.e., training on multiple tasks requires fewer interactions with the environment than training every task separately. This means that knowledge is shared between different tasks in all approaches as expected. For a larger amount of tasks, our approach is faster than PNN and only slightly worse than Distral in terms of steps required to reach the given





Fig. 5: Median scores over 10 runs. The shaded area represents 30% to 70% performance of the runs. Training is performed on 6 tasks, where we switch the action dimensions for 3 of them. This means that the agent goes to the left instead of the right, to the top instead of the bottom, and vice versa. For AMT and Distral, the solid line represents the average score of all tasks and the dashed line represents the task where the lowest score is observed. For PNN (dotted), a threshold score of 0.95 has to be reached over 10^5 steps before switching to the next task. For each approach, the vertical line is drawn at the first point where the score exceeds 0.9 over 10^5 steps.

performance threshold. Note however that our approach has substantially fewer parameters than the other approaches in this large number of tasks setup.


6.3 Unaligned Action Spaces

To see whether the approaches can handle transfer between domains where the action spaces are not aligned, we take the second, third and fourth grid world task and switch their action dimensions, meaning that the agent goes to the left instead of the right, to the top instead of the bottom, and vice versa. We combine these new tasks with the original grid world tasks 2, 3 and 4 and train the three different approaches to solve these six tasks simultaneously. Figure 5 compares the number of steps required to reach a score of 0.9 on all tasks separately and on average. Our approach clearly outperforms PNN and Distral in the number of steps required to reach the target performance on all tasks. We see two explanations for this: either two of our sub-networks specialize to the two sets of tasks and allow fast transfer as such, or the task specific linear layer (W_τ, b_τ) in our architecture effectively learns to invert the action space of some tasks such that two tasks from the two different sets look similar to a sub-network in our architecture. Most likely, the improvement is due to an entangled combination of both explanations. The results of PNN are comparable to the previous experiment as PNN is also able to deal well with unaligned action spaces. In contrast to multi-task approaches however, PNN is bound to define a threshold



(a) First sub-network. (b) Second sub-network. (c) Third sub-network.

Fig. 6: A smoothed average of the attention weights for 8 different grid world tasks. Set 1 (blue) consists of tasks 3 to 6 and set 2 (orange) consists of tasks 7 to 10, for which the reward associated with the collectable objects are inverted.



(a) First sub-network. (b) Second sub-network. (c) Third sub-network.

Fig. 7: Smoothed average of the attention weights for 8 different tasks of two domains. Grid World (blue) consists of grid world tasks 3 to 6 and Connect Four (orange) consists of the four connect-four/five tasks. The connect-four/five tasks are almost exclusively learned into the second sub-network, while the grid world tasks use mostly the third sub-network. We can see that there is a clear separation between the weights of the tasks of the two domains.

for when to freeze the current task’s network weights and move on to the next one. Further, a curriculum needs to be specified and tasks learned earlier cannot profit from knowledge discovered during learning later tasks. Therefore PNN ultimately learns slower in this setup than our approach. Distral, an approach that aligns the action space between tasks, requires more steps for these six tasks than for randomly selected six tasks like in the previous experiment, as the distilled policy cannot deal with the three environments and their counterparts at the same time. This underlines our claim that, while other approaches are effective for multi-task learning in a controlled setup, our approach is able to deal with multiple tasks even if the action spaces are not aligned.

6.4 Analyzing the Learned Attention Weights

To give an insight in how tasks are separated into sub-networks we take a subset of the grid world tasks where we expect negative transfer when knowledge is

shared. More specifically, we take tasks 3 – 6 and 7 – 10. Note that these two sets of tasks are visibly indistinguishable and equivalent apart from the fact that bonus objects yield negative rewards and penalty objects positive rewards in the second set. In Figure 6 we plot a smoothed average of the attention weights $w_i(s, \tau)$ of each task τ for all sub-networks $i \in \{1, 2, 3\}$. As can be seen in the figure, our architecture discovers that three sub-networks are not needed for these six tasks and learns to discard one of them. Further, one can see a tendency that one set of tasks is learned into one of the sub-networks while the other set of tasks is learned into the other remaining sub-network. Note however, that this distinction is not sharp since there is still a lot of transfer possible between the two sets of tasks, i.e., the agent has to stay on the board and find the target in both sets. This brings us to the interesting question how the distribution of the weights would look like if one uses two sets of tasks from completely unrelated domains. To answer this question, we train our model on connect-four/five and on grid world tasks. Figure 7 shows the weighting of the sub-networks when trained on those tasks. Clearly, the second sub-network learns to specialize on the connect-four/five task. Further, even though the connect-four/five and grid world tasks are unrelated to each other, parts of the “connect-four-knowledge” is used for the grid worlds while the non-overlapping state-action correlations are safely learned in a separate sub-network. Again, one of the sub-networks is left almost unused by all tasks, i.e., the model automatically learned that there are more sub-networks than needed for the two task domains.

7 Conclusion

We present a multi-task deep reinforcement learning algorithm based on the intuition of human attention. We show that knowledge transfer can be achieved by a simple attention architecture that does not require any a-priori knowledge of the relationship between the tasks. We show that our approach achieves transfer comparable to state of the art approaches as the number of tasks increases while using substantially fewer network parameters. Further, our approach clearly outperforms Distral and PNN when the action space between tasks is not aligned, since the task-specific weights and specialized sub-networks can account for this discrepancy. In future work, we plan to apply our approach to more complex tasks by incorporating recent, more resource efficient algorithms like [6, 10].

References

1. Aytar, Y., Pfaff, T., Budden, D., Paine, T.L., Wang, Z., de Freitas, N.: Playing hard exploration games by watching youtube. CoRR **abs/1805.11592** (2018), <http://arxiv.org/abs/1805.11592>
2. Barreto, A., Borsa, D., Quan, J., Schaul, T., Silver, D., Hessel, M., Mankowitz, D.J., Zidek, A., Munos, R.: Transfer in deep reinforcement learning using successor features and generalised policy improvement. In: Proceedings of the 35th International Conference on Machine Learning, ICML 2018 (2018), <http://proceedings.mlr.press/v80/barreto18a.html>

3. Barreto, A., Dabney, W., Munos, R., Hunt, J.J., Schaul, T., Silver, D., van Hasselt, H.P.: Successor features for transfer in reinforcement learning. In: Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017 (2017), <http://papers.nips.cc/paper/6994-successor-features-for-transfer-in-reinforcement-learning>
4. Birck, M., Corrêa, U., Ballester, P., Andersson, V., Araujo, R.: Multi-task reinforcement learning: An hybrid a3c domain approach (01 2017)
5. Czarnecki, W.M., Jayakumar, S.M., Jaderberg, M., Hasenclever, L., Teh, Y.W., Osindero, S., Heess, N., Pascanu, R.: Mix&match-agent curricula for reinforcement learning. arXiv preprint arXiv:1806.01780 (2018)
6. Espeholt, L., Soyer, H., Munos, R., Simonyan, K., Mnih, V., Ward, T., Doron, Y., Firoiu, V., Harley, T., Dunning, I., Legg, S., Kavukcuoglu, K.: IMPALA: scalable distributed deep-rl with importance weighted actor-learner architectures. In: ICML 2018 (2018), <http://proceedings.mlr.press/v80/espeholt18a.html>
7. Gao, Y., Xu, H., Lin, J., Yu, F., Levine, S., Darrell, T.: Reinforcement learning from imperfect demonstrations. CoRR **abs/1802.05313** (2018), <http://arxiv.org/abs/1802.05313>
8. Glatt, R., da Silva, F.L., Costa, A.H.R.: Towards knowledge transfer in deep reinforcement learning. In: BRACIS 2016 (2016). <https://doi.org/10.1109/BRACIS.2016.027>, <https://doi.org/10.1109/BRACIS.2016.027>
9. Gupta, A., Devin, C., Liu, Y., Abbeel, P., Levine, S.: Learning invariant feature spaces to transfer skills with reinforcement learning. CoRR **abs/1703.02949** (2017), <http://arxiv.org/abs/1703.02949>
10. Hessel, M., Soyer, H., Espeholt, L., Czarnecki, W., Schmitt, S., van Hasselt, H.: Multi-task deep reinforcement learning with popart. CoRR **abs/1809.04474** (2018), <http://arxiv.org/abs/1809.04474>
11. Hester, T., Vecerik, M., Pietquin, O., Lanctot, M., Schaul, T., Piot, B., Horgan, D., Quan, J., Sendonaris, A., Osband, I., Dulac-Arnold, G., Agapiou, J., Leibo, J.Z., Gruslys, A.: Deep q-learning from demonstrations. In: AAAI 2018 (2018), <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16976>
12. Higgins, I., Pal, A., Rusu, A.A., Matthey, L., Burgess, C., Pritzel, A., Botvinick, M., Blundell, C., Lerchner, A.: DARLA: improving zero-shot transfer in reinforcement learning. In: Proceedings of the 34th International Conference on Machine Learning, ICML 2017 (2017), <http://proceedings.mlr.press/v70/higgins17a.html>
13. Jaderberg, M., Dalibard, V., Osindero, S., Czarnecki, W.M., Donahue, J., Razavi, A., Vinyals, O., Green, T., Dunning, I., Simonyan, K., Fernando, C., Kavukcuoglu, K.: Population based training of neural networks. CoRR **abs/1711.09846** (2017), <http://arxiv.org/abs/1711.09846>
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. CoRR **abs/1412.6980** (2014), <http://arxiv.org/abs/1412.6980>
15. Larocche, R., Barlier, M.: Transfer reinforcement learning with shared dynamics. In: AAAI 2017 (2017), <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14315>
16. Lehnert, L., Littman, M.L.: Successor features support model-based and model-free reinforcement learning. CoRR **abs/1901.11437** (2019)
17. Leong, Y.C., Radulescu, A., Daniel, R., DeWoskin, V., Niv, Y.: Dynamic interaction between reinforcement learning and attention in multidimensional environments. *Neuron* **93**(2), 451–463 (2017)
18. Lin, L.J.: Reinforcement learning for robots using neural networks. Tech. rep., Carnegie-Mellon Univ Pittsburgh PA School of Computer Science (1993)

19. Machado, M.C., Bellemare, M.G., Talvitie, E., Veness, J., Hausknecht, M.J., Bowling, M.: Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents. CoRR **abs/1709.06009** (2017)
20. Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T.P., Harley, T., Silver, D., Kavukcuoglu, K.: Asynchronous methods for deep reinforcement learning. In: ICML 2016 (2016), <http://jmlr.org/proceedings/papers/v48/mniha16.html>
21. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M.A., Fidjeland, A., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D.: Human-level control through deep reinforcement learning. *Nature* (2015). <https://doi.org/10.1038/nature14236>, <https://doi.org/10.1038/nature14236>
22. Parisotto, E., Ba, L.J., Salakhutdinov, R.: Actor-mimic: Deep multitask and transfer reinforcement learning. CoRR **abs/1511.06342** (2015), <http://arxiv.org/abs/1511.06342>
23. Pohlen, T., Piot, B., Hester, T., Azar, M.G., Horgan, D., Budden, D., Barth-Maron, G., van Hasselt, H., Quan, J., Vecerik, M., Hessel, M., Munos, R., Pietquin, O.: Observe and look further: Achieving consistent performance on atari. CoRR **abs/1805.11593** (2018), <http://arxiv.org/abs/1805.11593>
24. Rajendran, J., Lakshminarayanan, A.S., Khapra, M.M., Prasanna, P., Ravindran, B.: Attend, adapt and transfer: Attentive deep architecture for adaptive transfer from multiple sources in the same domain. arXiv preprint arXiv:1510.02879 (2015)
25. Rusu, A.A., Colmenarejo, S.G., Gülçehre, Ç., Desjardins, G., Kirkpatrick, J., Pascanu, R., Mnih, V., Kavukcuoglu, K., Hadsell, R.: Policy distillation. CoRR **abs/1511.06295** (2015), <http://arxiv.org/abs/1511.06295>
26. Rusu, A.A., Rabinowitz, N.C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., Hadsell, R.: Progressive neural networks. CoRR **abs/1606.04671** (2016), <http://arxiv.org/abs/1606.04671>
27. Schmitt, S., Hudson, J.J., Zidek, A., Osindero, S., Doersch, C., Czarnecki, W.M., Leibo, J.Z., Küttler, H., Zisserman, A., Simonyan, K., Eslami, S.M.A.: Kickstarting deep reinforcement learning. CoRR **abs/1803.03835** (2018), <http://arxiv.org/abs/1803.03835>
28. Taylor, M.E., Stone, P.: Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research* **10** (2009). <https://doi.org/10.1145/1577069.1755839>, <http://doi.acm.org/10.1145/1577069.1755839>
29. Teh, Y.W., Bapst, V., Czarnecki, W.M., Quan, J., Kirkpatrick, J., Hadsell, R., Heess, N., Pascanu, R.: Distral: Robust multitask reinforcement learning. In: Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017 (2017), <http://papers.nips.cc/paper/7036-distral-robust-multitask-reinforcement-learning>
30. Yin, H., Pan, S.J.: Knowledge transfer for deep reinforcement learning with hierarchical experience replay. In: AAAI 2017 (2017), <http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14478>
31. Zhang, J., Springenberg, J.T., Boedecker, J., Burgard, W.: Deep reinforcement learning with successor features for navigation across similar environments. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2017 (2017). <https://doi.org/10.1109/IROS.2017.8206049>, <https://doi.org/10.1109/IROS.2017.8206049>