

PODNet: A Neural Network for Discovery of Plannable Options

Ritwik Bera,¹ Vinicius G. Goecks,¹ Gregory M. Gremillion,²

John Valasek,¹ and Nicholas R. Waytowich^{2,3}

¹Texas A&M University, ²Army Research Laboratory, ³Columbia University
{ritwik, vinicius.goecks, valasek}@tamu.edu,
{gregory.m.gremillion, nicholas.r.waytowich}.civ@mail.mil

Abstract

Learning from demonstration has been widely studied in machine learning but becomes challenging when the demonstrated trajectories are unstructured and follow different objectives. This short-paper proposes PODNet, Plannable Option Discovery Network, addressing how to segment an unstructured set of demonstrated trajectories for option discovery. This enables learning from demonstration to perform multiple tasks and plan high-level trajectories based on the discovered option labels. PODNet combines a custom categorical variational autoencoder, a recurrent option inference network, option-conditioned policy network, and option dynamics model in an end-to-end learning architecture. Due to the concurrently trained option-conditioned policy network and option dynamics model, the proposed architecture has implications in multi-task and hierarchical learning, explainable and interpretable artificial intelligence, and applications where the agent is required to learn only from observations.

Introduction

Learning from demonstrations to perform a single task has been widely studied in the machine learning literature (Argall et al. 2009; Ross, Gordon, and Bagnell 2011; Ross et al. 2013; Bojarski et al. 2016; Goecks et al. 2018). In these approaches, demonstrations are carefully curated in order to exemplify a specific task to be carried out by the learning agent. The challenge arises when the demonstrator is performing more than one task, or multiple hierarchical sub-tasks of a complex objective, also called *options*, where the same set of observations can be mapped to a different set of actions depending on the option being performed (Sutton, Precup, and Singh 1999; Stolle and Precup 2002). This is a challenge for traditional behavior cloning techniques that focus on learning a single mapping between observation and actions in a single option scenario.

This paper presents Plannable Option Discovery Network (PODNet), attempting to enable agents to learn the semantic structure behind those complex demonstrated tasks by using a meta-controller operating in the option-space instead of directly operating in the action-space. The main hypothesis is that a meta-controller operating in the option-space

can achieve much faster convergence on imitation learning and reinforcement learning benchmarks than an action-space policy network due to the significantly smaller size of the option-space. Our contribution, PODNet, is a custom categorical variational autoencoder (Jang, Gu, and Poole 2016) that is composed of several constituent networks that not only segment demonstrated trajectories into options, but concurrently trains an option dynamics model that can be used for downstream planning tasks and training on simulated rollouts to minimize interaction with the environment while the policy is maturing. Unlike previous imitation-learning based approaches to option discovery, our approach does not require the agent to interact with the environment in its option discovery process as it trains offline on just behavior cloning data. Moreover, being able to infer the option label for the current behavior executed by the learning agent, essentially, allowing the agent to broadcast the option it is currently pursuing, has implications in explainable and interpretable artificial intelligence.

Related Work

This work addresses how to segment an unstructured set of demonstrated trajectories for option discovery. The one-shot imitation architecture developed by Wang et al. (Wang et al. 2017) using conditional GAIL (cGAIL) maps trajectories into a set of latent codes that capture the semantics and context of the trajectories. This is analogous to word2vec (Mikolov et al. 2013) in natural language processing (NLP) where words are embedded into a vector space that preserves linguistic relationships.

In InfoGAN (Chen et al. 2016), a generative adversarial network (GAN) maximizes the mutual information between the latent variables and the observation, learning a discriminator that confidently predict the observation labels. InfoRL (Hayat, Singh, and Namboodiri 2019) and InfoGAIL (Li, Song, and Ermon 2017) utilized the concept of mutual information maximization to map latent variables to solution trajectories (generated by RL) and expert demonstrations respectively. Directed-InfoGAIL (Sharma et al. 2018) introduced the concept of directed information. It maximized the mutual information between the trajectory observed so far and the consequent option label. This modification to the In-

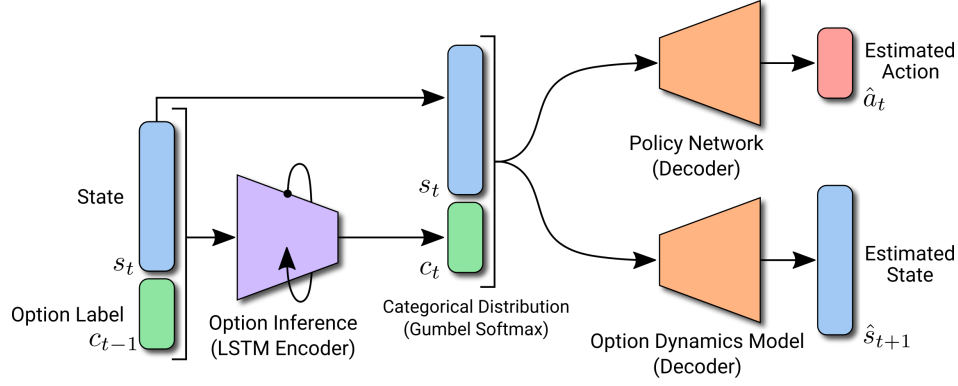


Figure 1: Proposed encoder-decoder architecture. Note that the Policy Network decoder could also be a recurrent neural network (RNN) if we wish to make the behavior label dependent on all preceding states and labels instead of just the previous state and corresponding behavior label.

foGAIL architecture allowed it to segment demonstrations and reproduce option. However, it assumed a prior knowledge of the number of options to be discovered. *Diversity Is All You Need (DIAYN)* (Eysenbach et al. 2018a) recovers distinctive sub-behaviors (from random exploration) by generating random trajectories and maximising mutual information between the states and the behavior label.

Variational Autoencoding Learning of Options by Reinforcement (VALOR) (Achiam et al. 2018) used β -VAEs (Higgins et al. 2017) to encode labels into trajectories, thus also implicitly maximising mutual information between behavior labels and corresponding trajectories. *DIAYN*’s, mutual information maximisation objective function is also implicitly solved in a β -VAE setting. Both VAEs and InfoGANs maximize mutual information between latent states and the input data. The difference is that VAE’s have access to the true data distribution while InfoGANs also have to learn to model the true data distribution. More recently, *COMPILE* (Kipf et al. 2019) employed a VAE based approach to infer not only option labels at every trajectory step but also infer option start and termination points in the given trajectory. However, once inferred to be completed, options are masked out. Thus while inferring options in the future, the agent loses track of critical options that might have happened in the past.

Most of the related works mentioned so far do not learn a dynamics model, and as a result, the discovered options cannot be used for downstream planning via model-based RL techniques. In our work, we utilize the fact that the demonstration data has state-transition information embedded within the demonstration trajectories and thus can be used to learn a dynamics model while simultaneously learning options. We also present a technique to identify the number of distinguishable options to be discovered from the demonstration data.

Plannable Option Discovery Network

Our proposed approach, Plannable Option Discovery Network (PODNet), is a custom categorical variational autoencoder (Jang, Gu, and Poole 2016) which consists of several

constituent networks: a recurrent option inference network, an option-conditioned policy network, and an option dynamics model, as seen in Figure 1. The categorical VAE allows the network to map each trajectory segment into a latent code and intrinsically perform soft k-means clustering on the inferred option labels. The following subsections explain the constituent components of PODNet.

Constituent Neural Networks

Recurrent option inference network In a complex task, the choice of an option at any time depends on both the current state, as well a history of the current and previous options that have been executed. For example, in a door opening task, an agent would decide to open a door only if it had already fetched the key earlier. We utilize a recurrent encoder using short long-term memory (LSTM) (Hochreiter and Schmidhuber 1997) to ensure the current option’s dependence on both the current state and the preceding options is captured. This helps overcome the problem where different options that contain similar or overlapping states are mapped to the same option label, as was observed in *DIAYN* (Eysenbach et al. 2018a). Our option inference network P is an LSTM that takes as input the current state s_t as well as the previous option label c_{t-1} and predicts the next option label for time step t .

Option-conditioned policy network Approaches such as *InfoGAIL* (Li, Song, and Ermon 2017), achieve the disentanglement into latent variables by imitating the demo trajectories while having access only to the inferred latent variable and not the demonstrator actions. We achieve this goal by concurrently training a option label conditioned policy network π that takes in the current predicted option c_t as well as the current state s_t and predicts the action a_t that minimizes the behavior cloning loss \mathcal{L}_{BC} of the demonstration trajectories.

Option dynamics model The main novelty of PODNet is the inclusion of an option dynamics model. The options dynamics model Q takes in as input the current state s_t and option label c_t and predicts the next state s_{t+1} . In other words,

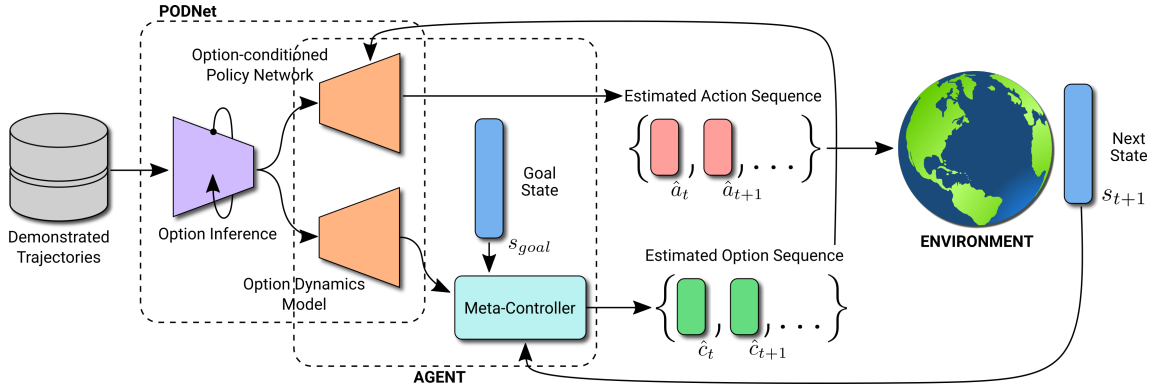


Figure 2: Complete PODNet diagram illustrating how the option dynamics model is integrated to meta-controllers to plan trajectories. Given a goal state s_{goal} , the meta-controller simulate trajectories using the option dynamics model and output the best estimated sequence of options to achieve the goal state. This sequence is then passed to the option-conditioned policy network, which outputs the sequence of estimated actions required to follow the planned option sequence.

the option dynamics model is an option-conditioned state-transition function, or dynamics model, that is dependent on the current option being executed instead of using the current action as traditional state-transition models would. The option dynamics model is trained simultaneously with the other policy and option inference networks by adding the *option dynamics consistency loss* to the overall training objective. The benefit of training an option dynamics model in this way is twofold: first, it ensures the the system dynamics can be completely defined by the option label, potentially allowing for easier recovery of option labels; second, it ensures that the recovered option labels c_t allow for modeling the environment dynamics in terms of the options themselves. This not only provides the ability to incorporate planning, but it allows planning to be performed at the options level instead of the action level, which will allow for more efficient planning on longer time-scales.

Training

The training process occurs offline and starts by collecting a dataset D consisting of unstructured demonstrated trajectories, which can be generated from any source as, for example, human experts, optimal controllers, or pre-trained reinforcement learning agents. The overall training loss-function is given as,

$$\mathcal{L}(\theta, \phi, \psi) = \mathbb{E}_{\pi_E} [\mathbb{E}_{c_t \sim P_\psi(\cdot | s_t, c_{t-1})} [(s_{t+1} - Q_\phi(s_t, c_t))^2 + (a_t - \pi_\theta(s_t, c_t))^2] - \beta D_{KL}(P_\psi(c_t | s_t, c_{t-1}) || p(c))]$$

Hence,

$$\mathcal{L}(\theta, \phi, \psi) = \text{Option Dynamics Consistency Loss}(\mathcal{L}_{ODC}) + \text{Behavior Cloning Loss}(\mathcal{L}_{BC}) + \text{Entropy Regularization}$$

Ensuring smooth backpropagation To ensure that the gradients flow through differentiable functions only during backpropagation, c_t is represented by a Gumbel-Softmax distribution, as illustrated in the literature on Categorical VAEs (Jang, Gu, and Poole 2016). Using argmax to select

the option with highest conditional probability would lead to having a discrete operation in the neural network and prohibit backpropagation in PODNet. Solo softmax is only used during the backward pass to allow backpropagation. For the forward pass, the softmax output is further subject to the argmax operator to obtain a one-hot encoded label vector.

Entropy Regularization The categorical distribution arising from the encoder network is forced to have minimal KL divergence with a uniform categorical distribution. This is done to ensure that all inputs are not encoded into the same sub-behavior cluster and are meaningfully separated into separate clusters. Entropy-driven regularization encourages exploration of the label space. This exploration can be modulated by tuning the hyperparameter β .

Downsampling of demonstration data For accurate prediction of option labels that concur with human intuition, it is important to downsample the state sequences since high-level dynamic changes have a low-frequency. Downsampling also decreases training time due to fewer samples being processed.

Prediction horizon To ensure that the option dynamics model does not simply learn an identity projection, the dynamics model is made to predict more than one time step ahead. This prediction horizon hyperparameter could be manually tuned depending on the situation.

Discovery of number of options The number of options can be obtained by having a held-out part of the demonstration, on which the behavior cloning loss \mathcal{L}_{BC} is evaluated, similar to how validation loss is. We start with an initial number of options, K , to be discovered and increment/decrement it to move towards decreasing \mathcal{L}_{BC} .

Planning Option Sequences

Although the main motivation for PODNet is to segment unstructured trajectories, the learned option dynamics model combined with the option-conditioned policy network can be used for planning option sequences. As shown in Figure

2, the option dynamics model learned with PODNets can be integrated to meta-controllers to plan trajectories. Given a goal state s_{goal} , the meta-controller simulate trajectories using the option dynamics model and output the best estimated sequence of options to achieve the goal state. This sequence is then passed to the option-conditioned policy network, which outputs the sequence of estimated actions required to follow the planned option sequence.

Conclusion

In this paper we presented PODNet, a neural network architecture for discovery of plannable options. Our approach combines a custom categorical variational autoencoder, a recurrent option inference network, option-conditioned policy network, and option dynamics model for end-to-end training and segmentation of an unstructured set of demonstrated trajectories for option discovery. PODNet’s architecture implicitly utilizes prior knowledge about options being dynamically consistent (plannable and representable by a skill dynamics model), being temporally extended and definitive of the agent’s actions at a particular state (as enforced by a option-conditioned policy network). This leads to discovery of plannable options that enable predictable behavior in AI agents when they adapt to newer tasks in a transfer learning setting. The proposed architecture has implications in multi-task and hierarchical learning, explainable and interpretable artificial intelligence.

References

- [Achiam et al. 2018] Achiam, J.; Edwards, H.; Amodei, D.; and Abbeel, P. 2018. Variational option discovery algorithms. *arXiv preprint arXiv:1807.10299*.
- [Argall et al. 2009] Argall, B. D.; Chernova, S.; Veloso, M.; and Browning, B. 2009. A survey of robot learning from demonstration. *Robotics and autonomous systems* 57(5):469–483.
- [Bojarski et al. 2016] Bojarski, M.; Del Testa, D.; Dworakowski, D.; Firner, B.; Flepp, B.; Goyal, P.; Jackel, L. D.; Monfort, M.; Muller, U.; Zhang, J.; et al. 2016. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*.
- [Chen et al. 2016] Chen, X.; Duan, Y.; Houthoofd, R.; Schulman, J.; Sutskever, I.; and Abbeel, P. 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*, 2172–2180.
- [Eysenbach et al. 2018a] Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2018a. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*.
- [Eysenbach et al. 2018b] Eysenbach, B.; Gupta, A.; Ibarz, J.; and Levine, S. 2018b. Diversity is All You Need: Learning Skills without a Reward Function.
- [Goecks et al. 2018] Goecks, V. G.; Gremillion, G. M.; Lawhern, V. J.; Valasek, J.; and Waytowich, N. R. 2018. Efficiently combining human demonstrations and interventions for safe training of autonomous systems in real-time. *CoRR* abs/1810.11545.
- [Hayat, Singh, and Namboodiri 2019] Hayat, A.; Singh, U.; and Namboodiri, V. P. 2019. InfoRL: Interpretable Reinforcement Learning using Information Maximization.
- [Higgins et al. 2017] Higgins, I.; Matthey, L.; Pal, A.; Burgess, C.; Glorot, X.; Botvinick, M.; Mohamed, S.; and Lerchner, A. 2017. beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR* 2(5):6.
- [Hochreiter and Schmidhuber 1997] Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- [Jang, Gu, and Poole 2016] Jang, E.; Gu, S.; and Poole, B. 2016. Categorical Reparameterization with Gumbel-Softmax.
- [Kipf et al. 2019] Kipf, T.; Li, Y.; Dai, H.; Zambaldi, V.; Sanchez-Gonzalez, A.; Grefenstette, E.; Kohli, P.; and Battaglia, P. 2019. Compile: Compositional imitation learning and execution. In *International Conference on Machine Learning*, 3418–3428.
- [Li, Song, and Ermon 2017] Li, Y.; Song, J.; and Ermon, S. 2017. InfoGAIL: Interpretable Imitation Learning from Visual Demonstrations.
- [Mikolov et al. 2013] Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- [Ross et al. 2013] Ross, S.; Melik-Barkhudarov, N.; Shankar, K. S.; Wendel, A.; Dey, D.; Bagnell, J. A.; and Hebert, M. 2013. Learning monocular reactive uav control in cluttered natural environments. In *2013 IEEE international conference on robotics and automation*, 1765–1772. IEEE.
- [Ross, Gordon, and Bagnell 2011] Ross, S.; Gordon, G.; and Bagnell, D. 2011. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 627–635.
- [Sharma et al. 2018] Sharma, A.; Sharma, M.; Rhinehart, N.; and Kitani, K. M. 2018. Directed-Info GAIL: Learning Hierarchical Policies from Unsegmented Demonstrations using Directed Information.
- [Stolle and Precup 2002] Stolle, M., and Precup, D. 2002. Learning options in reinforcement learning. In *International Symposium on abstraction, reformulation, and approximation*, 212–223. Springer.
- [Sutton, Precup, and Singh 1999] Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence* 112(1-2):181–211.
- [Wang et al. 2017] Wang, Z.; Merel, J.; Reed, S.; Wayne, G.; de Freitas, N.; and Heess, N. 2017. Robust Imitation of Diverse Behaviors.

Appendices

VAEs and mutual information

Disregarding the regularization terms, it can be shown that the objective function solved by training variation autoencoders (VAEs) is equivalent to maximising the mutual information term, $\mathcal{I}(x; z)$ where $x \sim p_{data}(x)$ and z is the bottleneck latent state.

The encoder network $p_\theta(z | x)$ is and the decoder network is $q_\phi(x | z)$, which also represents the approximate posterior.

$$\begin{aligned}\mathcal{I}(z; x) &= \mathcal{H}(z) - \mathcal{H}(z | x) \\ &= \mathcal{H}(z) + \mathbb{E}_{x \sim p_{data}(x)} \mathbb{E}_{z \sim p(z|x)} \log p(z | x) \\ &= \mathcal{H}(z) + \mathbb{E}_{x \sim p_{data}(x)} \mathbb{E}_{z \sim p(z|x)} \log \frac{p(x | z)p(z)}{p(x)} \\ &\geq \mathcal{H}(z) + \mathbb{E}_{x \sim p_{data}(x)} \mathbb{E}_{z \sim p_\theta(z|x)} \log q_\phi(x | z) \\ &\quad - \mathcal{H}(z) - \log p(x) \\ &\geq \mathbb{E}_{x \sim p_{data}(x)} \mathbb{E}_{z \sim p_\theta(z|x)} \log q_\phi(x | z)\end{aligned}$$

The final term is the evidence lower bound (ELBO) (minus the KL divergence regularization term) that is maximised in VAEs as a part of the objective function. The KL divergence regularization term ensures exploration and packing in the latent variable space.

Equivalence between VAE and *Diversity Is All You Need*'s training objective

Let us consider the objective function defined in *Diversity Is All You Need* (Eysenbach et al. 2018b),

$$\begin{aligned}\mathcal{R}(\pi) &= \mathcal{I}(c; s) - \mathcal{I}(a; s, c) \\ &= \mathcal{H}(c) - \mathcal{H}(c | s) + \mathcal{H}(a) - \mathcal{H}(a | s, c) \\ &\geq -\mathcal{H}(c | s) + \mathcal{H}(a | s, c) \\ &\geq \mathbb{E}_{c, \{s, a\}} (\log q_\phi(c | s) - \log p(c) - \log \pi(a | s, c)) \\ &\geq \mathbb{E}_{c, \{s, a\}} (\log q_\phi(c | s)) + \mathcal{H}(c) + \mathcal{H}(\pi(c)) \\ &\geq \mathbb{E}_{c, \{s, a\}} (\log q_\phi(c | s)) + \mathcal{H}(\pi(c)) + \text{constant}\end{aligned}$$

Note that,

$$p(\tau | c) = p(s_o) \sum_{t=1}^T \pi(a | s, c) p(s_{t+1} | s_t, a_t)$$

Since system dynamics and distribution of initial states is assumed fixed, it can be said that

$$\begin{aligned}p(\tau | c) &\propto \sum_{t=1}^T \pi(a | s, c) \\ \implies \mathcal{H}(\tau | c) &\propto \mathcal{H}(\pi(c))\end{aligned}$$

Now let us analyze the Kullback-Leibler divergence term,

$$\begin{aligned}\mathbb{E}_{c \sim p(c)} D_{KL}(p(\tau | c) || p(\tau_{\pi_o})) &= \mathbb{E}_{c \sim p(c)} \mathbb{E}_{\tau \sim p(\cdot | \pi(c))} \log \frac{p(\tau | c)}{p(\tau_{\pi_o})} \\ &= -\mathcal{H}(\tau | c) + \mathcal{H}(\tau_{\pi_o}) \\ &= -\mathcal{H}(\tau | c) + \text{constant term}\end{aligned}$$

Hence, we can write the objective function as,

$$\mathcal{R}(\pi) \geq \mathbb{E}_c [\mathbb{E}_{s \sim p(\cdot | c)} (\log q_\phi(c | s)) - D_{KL}(p(\tau | c) || p(\tau_{\pi_o}))] \quad (1)$$

Recomputing the expectation over trajectories as expectations over states

$$\begin{aligned}\mathbb{E}_{\tau \sim \pi(c)} \log Q(c | \tau) &= \sum_{t=0}^T \mathbb{E}_{s_t \sim \pi(c)} \log q(c | s_t) \\ &= \mathbb{E}_{s \sim \pi(c)} \log q(c | s)\end{aligned}$$

Hence, we can re-write the objective function (1) as,

$$\mathcal{R}(\pi) \geq \mathbb{E}_c [\mathbb{E}_{\tau \sim \pi(c)} (\log Q(c | \tau)) - D_{KL}(p(\tau | c) || p(\tau_{\pi_o}))]$$

Hence, we can see the equivalence between the ELBO objective in variational autoencoders and the objective defined in DIAYN. Note that the regularization term ensures that policy network explores the action space and generates a diverse distribution of possible output trajectories.