

MULTI-STEP PREDICTION IN LINEARIZED LATENT STATE SPACES FOR REPRESENTATION LEARNING

A. TYTARENKO

Abstract. In this paper, we derive a novel method as a generalization over LCEs such as E2C. The method develops the idea of learning a locally linear state space by adding a multi-step prediction, thus allowing for more explicit control over the curvature. We show that the method outperforms E2C without drastic model changes which come with other works, such as PCC and P3C. We discuss the relation between E2C and the presented method and derive update equations. We provide empirical evidence, which suggests that by considering the multi-step prediction, our method – ms-E2C – allows learning much better latent state spaces in terms of curvature and next state predictability. Finally, we also discuss certain stability challenges we encounter with multi-step predictions and how to mitigate them.

Keywords: representation learning, learning controllable embedding, reinforcement learning, latent state space.

INTRODUCTION

One of the most challenging problems which the field reinforcement learning faces is learning autonomous agents capable of control in Markov Decision Processes (MDP) with complex state and action spaces. For instance, complications may arise from large action spaces [1], limited ability to interact with an environment [2], partial observability (POMDP) [3, 4], etc. Optimizing a decent policy takes a lot of samples, usually requires online interactive learning and neural networks capable of processing higher dimensional observations with large number of trainable parameters [5, 6].

There are various algorithms which try to deal with the problem of sample inefficiency, or limited amount of data. Model-based reinforcement learning algorithms [7–9] try to achieve sample efficiency by approximating transition dynamics of an MDP in online or offline mode. Offline reinforcement learning methods [2, 10] strive to extract as much useful information from limited offline data as possible, in order to learn a policy applicable to online regimes as well.

Another algorithmic framework – Learning Controllable Embedding (LCE) – approaches this problem by learning a lower dimensional latent state space and using simpler control algorithms, like iLQR [11], to perform control in this latent space. The challenge here is to make sure that the learned latent space has simpler structure (i.e. next states are easier to predict).

Some particular instances of this framework are described in [9, 12–14]. The idea of E2C [12] is to learn a locally-linear latent space, so that algorithms like LQG could be used for goal-reaching tasks. PCC [13] tries to fix some of the issues encountered in E2C by deriving losses which allow for explicit minimization of latent space's curvature. P3C [14] improves upon PCC mainly by replacing reconstruction loss, needed to make sure the learned state space carries enough

information to generate (i.e. decode) observations from latent states. P3C uses predictive coding instead.

In this paper, we seek an alternative approach to enforce lower latent space's curvature and predictability. We generalize E2C by considering multiple transitions at a time, making sure the local linearity is not just preserved between neighbouring states. We inherit the idea of minimization of a joint log likelihood of a transition, generalize it to multiple transitions, and derive a variational bound for further minimization. We then compare the results with LCE approaches and demonstrate a visual representation of learned latent state spaces for a benchmark common among LCE papers.

PRELIMINARIES

We denote a **Markov Decision Process** (MDP) M as a tuple (S, A, r, T) , where S — state space; A — action space; $r: S \times A \rightarrow \mathbb{R}$ — reward function; $T = P(s_{t+1} | s_t, a_t)$ — probability of state s_{t+1} given current state s_t and action taken a_t .

A state of an MDP is a sufficient statistics for a transition kernel, possessing a Markov property.

A task of **Reinforcement Learning** algorithm is for a given MDP M find a policy π , such that it maximizes the expected return. We are interested in a discounted return objective:

$$\pi^* = \arg \max_{\pi} E_{\tau \sim p(\cdot)} \sum_{t=0}^{\infty} \gamma^t r(s_t, a_t),$$

where τ denotes a trajectory $(x_0, a_0, x_1, a_1, \dots)$ obtained by sampling actions using a stochastic policy π .

In particular, we consider a specific class of Reinforcement Learning algorithms – **Model Based Reinforcement Learning** [7, 8]. Algorithms of this kind usually possess higher sample efficiency, but they involve some sort of an approximation of a transition kernel. LCE algorithms involve parametric models (i.e. neural nets) to learn a good model with desirable properties, like linearity and predictability. This allows to use even the simplest Model-based control (and RL) algorithms like iLQR [11].

THE MULTI-STEP EMBED TO CONTROL MODEL

Consider an internal transition dynamics of an MDP $M = (S, A, r, T)$:

$$s_{t+1} = f(s_t, a_t) + \omega, \quad \omega \sim P_{\omega}.$$

As we discussed previously, a function f may be highly nonlinear, thus being tricky to optimize with model-based RL or control algorithms. LCE approaches therefore try to learn a mapping from a state space S to some latent space Z such that its latent dynamics

$$z_{t+1} = \hat{f}(z_t, a_t) + \hat{\omega}$$

has some desired properties like local linearity, low curvature, predictability, etc.

In order to learn the mapping $Q_\phi : S \rightarrow Z$, Variational Inference framework is employed to derive a tractable algorithm of maximization a likelihood of known data points under the mapping we want to learn.

Optimization problem

As follows from the Fig. 1, we consider a dataset

$$D = \{(s_t, a_t, s_{t+1}, a_{t+1}, s_{t+2}, \dots, a_{t+K-1}, s_{t+K}) \mid i = 1, \dots, N\}$$

containing samples from real trajectories gathered before training. To follow and generalize [12], we define




Fig. 1. Graphical models for E2C and ms-E2C(K): dashed lines – state reconstruction process

$$Q_\phi = P_\phi(z_t | s_t)$$

as a generative model which for a given state s_t specifies a distribution over the latent space Z . Basically, it plays a role of the mapping from S to Z parametrized with a parameter vector ϕ . And

$$Q_\psi = P_\psi(\hat{z}_{t+1} | \hat{z}_t, a_t)$$

as a generative model which for a given latent state z_t and an action a_t predicts the distribution for the next latest state z_{t+1} . The model is also parametrized with a parameter vector ψ . Also, we denote

$$Q_\psi^j = P_\psi(\hat{z}_{t+j} | \hat{z}_{t+j-1}, a_t).$$

In order to find ϕ and ψ , we maximize the likelihood of a dataset of trajectory samples of length K with respect to the aforementioned parameter vectors:

$$\phi^*, \psi^* = \arg \max_{\phi, \psi} \prod_{i=1}^N P(s_t^i, a_t^i, s_{t+1}^i, a_{t+1}^i, s_{t+2}^i, \dots, a_{t+K-1}^i, s_{t+K}^i).$$

For the sake of readability, we denote s_t, \dots, s_{t+K} as $s_{t:t+K}$ and a_t, \dots, a_{t+K} as $a_{t:t+K}$. Thus our objective is:

$$\phi^*, \psi^* = \arg \max_{\phi, \psi} \prod_{i=1}^N P(s_{t:t+K}^i, a_{t:t+K-1}^i).$$

A corresponding graphical model is depicted on Fig. 1.

Optimization objective

The objective we defined in a previous subsection is known to be intractable and difficult to optimize. Therefore, LCE approaches employ Variational Inference to find a lower bound to the log-likelihood objective. In this section, we derive this bound for the proposed probabilistic model. Variational lower bound:

$$\begin{aligned} & -\log P(s_{t:t+K}, a_{t:t+K-1}) \leq \\ & \leq E_{\substack{z_t \sim Q_\phi \\ \hat{z}_{t+j} \sim Q_\psi^j \\ j=1, \dots, K}} \left[-\sum_{j=1, \dots, K} \log P(s_{t+j} | \hat{z}_{t+j}) - \log P(s_t | z_t) \right] + D_{KL}[Q_\phi \| P(Z)] \end{aligned}$$

Here $D_{KL}[P \| Q]$ denotes Kullback–Leibler divergence functional:

$$D_{KL}[P \| Q] = E_{x \sim P} \log \frac{P(x)}{Q(x)}.$$

Proof.

$$\begin{aligned} -\log P(s_{t:t+K}, a_{t:t+K-1}) &= -\log \int_{z_t, \hat{z}_{t+1:t+K}} P(s_{t:t+K}, a_{t:t+K-1}, z_t, \hat{z}_{t+1:t+K}) dz_t d\hat{z}_{t+1:t+K} = \\ &= -\log \int_{z_t, \hat{z}_{t+1:t+K}} P(s_{t:t+K} | a_{t:t+K-1}, z_t, \hat{z}_{t+1:t+K}) P(z_t, \hat{z}_{t+1:t+K} | a_{t:t+K-1}) \times \\ & \quad \times P(a_{t:t+K-1}) dz_t d\hat{z}_{t+1:t+K} \\ &= -\log \int_{z_t, \hat{z}_{t+1:t+K}} P(s_{t:t+K} | a_{t:t+K-1}, z_t, \hat{z}_{t+1:t+K}) P(z_t, \hat{z}_{t+1:t+K} | a_{t:t+K-1}) \times \\ & \quad \times P(a_{t:t+K-1}) \frac{Q_\phi}{Q_\phi} dz_t d\hat{z}_{t+1:t+K} \\ &= -\log \int_{z_t, \hat{z}_{t+1:t+K}} P(s_{t:t+K} | a_{t:t+K-1}, z_t, \hat{z}_{t+1:t+K}) \prod_{j=1}^K (Q_\psi^j) P(z_t) \times \\ & \quad \times P(a_{t:t+K-1}) \frac{Q_\phi}{Q_\phi} dz_t d\hat{z}_{t+1:t+K} \\ &= -\log \int_{z_t, \hat{z}_{t+1:t+K}} P(s_t | z_t) \prod_{j=1}^K (Q_\psi^j P(s_{t+j} | \hat{z}_{t+j})) P(z_t) P(a_{t:t+K-1}) \frac{Q_\phi}{Q_\phi} dz_t d\hat{z}_{t+1:t+K} = \\ &= -\log E_{\substack{z_t \sim Q_\phi \\ \hat{z}_{t+j} \sim Q_\psi^j \\ j=1, \dots, K}} \left[P(s_t | z_t) \prod_{j=1}^K (P(s_{t+j} | \hat{z}_{t+j})) P(a_{t:t+K-1}) \frac{P(z_t)}{Q_\phi} \right] \leq \end{aligned}$$

$$\begin{aligned}
 &\leq E_{\substack{z_t \sim Q_\phi \\ \hat{z}_{t+j} \sim Q_\psi^j \\ j=1, \dots, K}} \left[- \sum_{j=1, \dots, K} \log P(s_{t+j} | \hat{z}_{t+j}) - \log P(s_t | z_t) - \log \frac{P(z_t)}{Q_\phi} \right] = \\
 &= E_{\substack{z_t \sim Q_\phi \\ \hat{z}_{t+j} \sim Q_\psi^j \\ j=1, \dots, K}} \left[- \sum_{j=1, \dots, K} \log P(s_{t+j} | \hat{z}_{t+j}) - \log P(s_t | z_t) \right] + D_{KL}[Q_\phi \| P(Z)].
 \end{aligned}$$

Multi-step embed-to-control model (ms-E2C)

In this section we instantiate a model for learning a latent locally-linear state space. We use a previously derived upper bound for negative loglikelihood over the multi-step trajectory samples. A graphical models for both E2C and ms-E2C(K) are shown on the Fig. 1.

First, we instantiate **parametric models for encoding and dynamics function** as:

$$Q_\phi = P(z_t | s_t) = N(\mu_\phi(s_t), \Sigma_\phi(s_t)) \text{ --- encoder;}$$

$$\mu_\phi(s_t), \Sigma_\phi(s_t) = \text{NeuralNet}(s_t; \phi);$$

$$A_t, B_t, o_t = \text{NeuralNet}(z_t; \psi) \text{ --- latent dynamics;}$$

$$Q_\psi^j = P(\hat{z}_{t+j} | \hat{z}_{t+j-1}, a_{t+j-1}) = N(\mu_\psi^j, \Sigma_\psi^j) \text{ --- dynamics;}$$

$$\mu_\psi^j = A_t \mu_\psi^{j-1} + B_t a_{t+j-1} + o_t \text{ for } j = 2, \dots, K;$$

$$\mu_\psi^j = A_t \mu_\psi + B_t a_t + o_t \text{ for } j = 1;$$

$$\Sigma_\psi^j = A_t \Sigma_\psi^{j-1} A_t^T + \Sigma_W \text{ for } j = 2, \dots, K;$$

$$\Sigma_\psi^j = A_t \Sigma_\psi A_t^T + \Sigma_W \text{ for } j = 1.$$

Thus, given an optimal model would imply a locally linear latent space, in which curvature (i.e. linearity) is explicitly controlled by changing the number of steps per sample. Choosing a large K would recover a globally linear model and setting $K=1$ recovers an E2C model.

As it follows from the figure, ms-E2C is a generalization of E2C, which one recovers by setting $K=1$.

We also have to specify a **parametrized decoding model**, which is needed to compute the upper bound, and to enforce a "reconstruction" constraint, introduced in [12] and generalized for our multi-step model:

$$P_\theta = P_\theta(s_t | z_t) = \text{Bernoulli}(p(z_t));$$

$$p(z_t) = \text{NeuralNet}(z_t; \theta);$$

$$P_\theta^j = P_\theta(s_t | \hat{z}_{t+j}) = \text{Bernoulli}(p(\hat{z}_{t+j}));$$

$$p(\hat{z}_{t+j}) = \text{NeuralNet}(\hat{z}_{t+j}; \theta).$$

Bernoulli distribution is chosen for comparability with E2C model on shared benchmark MDPs, i.e. Planar, where the original state space consists of black-and-white images of a grid world with white obstacles and a white circle denoting the position of the agent.

Loss function. In order to complete the model's specification, we have to provide a loss function optimizable via stochastic gradient descent. In msE2C it consists of three terms: an estimation of the derived upper bound, consistency term, and stability term.

$$L_{upper}(D_i; \phi, \psi, \theta) = E_{\substack{z_t \sim Q_\phi \\ \hat{z}_{t+j} \sim Q_\psi^j \\ j=1, \dots, K}} \left[- \sum_{j=1, \dots, K} \log P_\theta(s_{t+j} | \hat{z}_{t+j}) - \log P_\theta(s_t | z_t) \right] + D_{KL}[Q_\phi \| P(Z)]$$

The expectation is estimated using a one-sample estimate and a reparametrization trick widely used in variational auto-encoders:

$$L_{consistency}(D_i; \phi, \psi) = \sum_{j=1}^K D_{KL}[Q_\psi^j \| Q_\phi];$$

$$L_{stability}(D_i; \phi, \psi) = \text{Gersh}(A_t(z_t)) + \text{Gersh}(B_t(z_t)).$$

Here $\text{Gersh}(X)$ denotes **Gershgorin loss** [15, 16]:

$$\text{Gersh}(X) = \sum_{i=1}^n \max(0, X_{i,i} + \sum_{j \neq i} |X_{i,j}| + \varepsilon),$$

where $|X_{i,j}|$ denotes a minor of a matrix X , and $\varepsilon > 0$ is a small constant.

According to the Theorem 1 from [15], if the loss value is non-positive, all eigenvalues of a matrix X are guaranteed to have a negative real part, thus ensuring dynamical system stability. The usage of Gershgorin loss in composite loss function is mandatory, as ms-E2C(K) diverges for larger K .

Algorithm

Now, we summarize an algorithm for fitting the instance of ms-E2C model we described earlier.

1. Sample a dataset of sub-trajectories using a pretrained or random policy:

$$D = \{(s_t, a_t, s_{t+1}, a_{t+1}, s_{t+2}, \dots, a_{t+K-1}, s_{t+K})_i \mid i = 1, \dots, N\}.$$

2. Initialize the weights of neural nets ϕ, ψ, θ .
3. Repeat for c epochs:
 - a) Retrieve a sample D_i from the dataset D

b) Compute updated weights using a stochastic gradient descent step:

$$\phi' = \phi - \gamma \nabla_{\phi} (L_{upper}(D_i) + \lambda_1 L_{consistency}(D_i) + \lambda_2 L_{stability}(D_i));$$

$$\psi' = \psi - \gamma \nabla_{\psi} (L_{upper}(D_i) + \lambda_1 L_{consistency}(D_i) + \lambda_2 L_{stability}(D_i));$$

$$\theta' = \theta - \gamma \nabla_{\theta} L_{upper}(D_i);$$

c) Update neural networks' parameters:

$$\phi = \phi', \quad \psi = \psi', \quad \theta = \theta'.$$

Here λ_1, λ_2 are tunable hyperparameters.

One might notice that unlike [15] we do not introduce an inner optimization loop to ensure stability of the internal latent space dynamics. Instead, we add the stability loss to the composed loss function. We found that although the difference is apparent during a few first epochs, it becomes negligible after a while. Stability condition does not get violated and the general results are almost the same.

EXPERIMENTAL VALIDATION

Planar system

Following [12–14], we use a Planar benchmark to compare the performance of the algorithms. In it, a state space is represented as a black-and-white image of a grid world with obstacles. In order to collect a dataset, we sample a random initial state and perform a series of random actions to obtain a trajectory of length K .

As in [12, 13], we use a deconvolutional network architecture [17] for image reconstruction from the latent state. For the sake of comparability, we chose the same architecture as in other papers on the topic.

The visualizations of the obtained latent state spaces are provided on a Fig. 2. The numerical results are summarized in a Table.

Comparison of reconstruction and prediction losses

Method	State Loss $\log P_{\theta}(s_t z_t)$	Next State Loss $\log P(s_{t+1} s_t, a_t)$
Non-linear E2C	9.2 ± 4.5	11.7 ± 8.8
Global E2C	7.6 ± 5.7	10.6 ± 5.2
E2C	7.6 ± 2.3	10.1 ± 2.7
ms-E2C(3)	<u>7.3 ± 1.7</u>	8.7 ± 1.9
ms-E2C(5)	7.6 ± 2.1	7.5 ± 1.6
ms-E2C(7)	7.7 ± 2.0	<u>6.3 ± 0.9</u>

State loss is a regular reconstruction loss. As we observe, ms-E2C(K) give only slight average improvements on it, which is entirely expected. The introduced method does not change the architecture of a decoding network nor does it add any improvements to the algorithm regarding this matter. An important thing to notice though, is that our generalization does not make the reconstruction per-

formance much worse, which might be expected as representation is influenced by addition prediction constraints. Next state is computed by encoding the state

$s_t \xrightarrow{Q_\phi} z_t$, predicting the next latent state $z_t \xrightarrow{Q_\psi^1} z_{t+1}$, and decoding the predicted regular state $z_{t+1} \xrightarrow{P_\theta} s_{t+1}$. Now, the results for previous methods were reproduced with slight perturbations, as we used our own codebase for it.

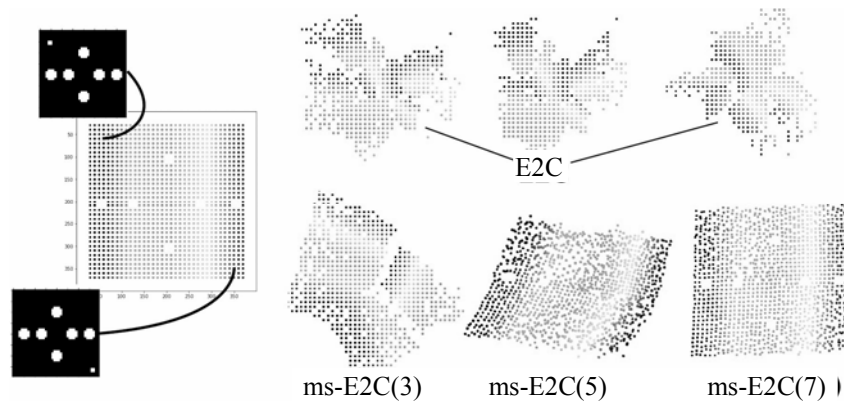


Fig 2. A comparison of latent state spaces learned by E2C and ms-E2C methods.

It's worth noting that E2C results coincide with other papers which involved reproduction of E2C [13, 14], while the original paper provides better visuals. A visualization is obtained by transforming all possible environment states with the network Q_ϕ . See the scheme on the left for details.

CONCLUSION

In this paper, a novel method had been derived as a generalization over the previous works on LCEs. We demonstrate, how the method improves upon E2C without drastic model changes which come with other works, such as PCC and P3C. We empirically show, that by considering the multistep prediction ms-E2C allows to learn a much better latent state spaces in terms of curvature and predictability, by adding a simple yet efficient way to explicitly control the desired curvature of a resulting space. At implementation is available at [18].

Moreover, our work introduces a new dimension to the LCE family of algorithms. Our future work will focus on using the approaches from the state of the art LCE methods, like predictive coding to make LCEs applicable to the higher dimensional real-world MDPs with limited amount of data to learn dynamics embedding from. We will also explore an intriguing possibility to not only encode the state, but also the action space, which sometimes has the complex structure. Lastly, we would like to study various extensions of the method to imitation learning and model-based reinforcement learning.

REFERENCES

1. G. Dulac-Arnold et al., "Deep reinforcement learning in large discrete action spaces," *arXiv preprint arXiv:1512.07679*, 2015. doi: 10.48550/arXiv.1512.07679.

2. S. Levine, A. Kumar, G. Tucker, and J. Fu, “Offline reinforcement learning: Tutorial, review, and perspectives on open problems,” *arXiv preprint arXiv:2005.01643*, 2020. doi: 10.48550/arXiv.2005.01643.
3. E.A. Feinberg, P.O. Kasyanov, and M.Z. Zgurovsky, “Partially observable total-cost markov decision processes with weakly continuous transition probabilities,” *Mathematics of Operations Research*, vol. 41, no. 2, pp. 656–681, 2016. doi: 10.1287/moor.2015.0746.
4. E.A. Feinberg, P.O. Kasyanov, and M.Z. Zgurovsky, “Convergence of probability measures and markov decision models with incomplete information,” *Proceedings of the Steklov Institute of Mathematics*, vol. 287, no. 1, pp. 96–117, 2014. doi: 10.1134/S0081543814080069.
5. O. Vinyals et al., “Grandmaster level in starcraft ii using multi-agent reinforcement learning,” *Nature*, vol. 575, no. 7782, pp. 350–354, 2019. doi: 10.1038/s41586-019-1724-z.
6. S. Reed et al., “A generalist agent,” *arXiv preprint arXiv:2205.06175*, 2022.
7. D. Ha and J. Schmidhuber, “World models,” *arXiv preprint arXiv:1803.10122*, 2018. doi: 10.48550/arXiv.1803.10122.
8. T.M. Moerland, J. Broekens, and C.M. Jonker, “Model-based reinforcement learning: A survey,” *arXiv preprint arXiv:2006.16712*, 2020. doi: 10.48550/arXiv.2006.16712.
9. D. Hafner et al., “Learning latent dynamics for planning from pixels,” in *International conference on machine learning*, PMLR, 2019, pp. 2555–2565. doi: 10.48550/arXiv.1811.04551.
10. R.F. Prudencio, M.R. Maximo, and E.L. Colombini, “A survey on offline reinforcement learning: Taxonomy, review, and open problems,” *arXiv preprint arXiv:2203.01387*, 2022. doi: 10.48550/arXiv.2203.01387.
11. W. Li and E. Todorov, “Iterative linear quadratic regulator design for nonlinear biological movement systems,” in *ICINCO (1)*, Citeseer, 2004, pp. 222–229. doi:10.5220/0001143902220229.
12. M. Watter, J. Springenberg, J. Boedecker, and M. Riedmiller, “Embed to control: A locally linear latent dynamics model for control from raw images,” *Advances in neural information processing systems*, vol. 28, 2015.
13. N. Levine, Y. Chow, R. Shu, A. Li, M. Ghavamzadeh, and H. Bui, “Prediction, consistency, curvature: Representation learning for locally linear control,” *arXiv preprint arXiv:1909.01506*, 2019.
14. R. Shu et al., “Predictive coding for locally-linear control,” in *International Conference on Machine Learning*, PMLR, 2020, pp. 8862–8871. doi: 10.5555/3524938.3525760.
15. M. Lechner, R. Hasani, D. Rus, and R. Grosu, “Gershgorin loss stabilizes the recurrent neural network compartment of an end-to-end robot learning scheme,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2020, pp. 5446–5452. doi: 10.1109/ICRA40945.2020.9196608.
16. R.A. Horn and C.R. Johnson, *Matrix analysis*. Cambridge university press, 2012. doi: 10.5555/2422911.
17. M.D. Zeiler, D. Krishnan, G.W. Taylor, and R. Fergus, “Deconvolutional networks,” in *2010 IEEE Computer Society Conference on computer vision and pattern recognition*, pp. 2528–2535. doi: 10.1109/CVPR.2010.5539957
18. A. Tytarenko, *RL-research*. Available: <https://github.com/titardrew/rl-research>, 2022.

Received 31.08.2022

INFORMATION ON THE ARTICLE

Andrii M. Tytarenko, ORCID: 0000-0002-8265-642X, Educational and Research Institute for Applied System Analysis of the National Technical University of Ukraine “Igor Sikorsky Kyiv Polytechnic Institute”, Ukraine, e-mail: titarenkoan@gmail.com

БАГАТОКРОКОВЕ ПРОГНОЗУВАННЯ В ЛІНЕАРИЗОВАНИХ ЛАТЕНТНИХ ПРОСТОРАХ ДЛЯ НАВЧАННЯ РЕПРЕЗЕНТАЦІЙ / А.М. Титаренко

Анотація. Запропоновано новий метод, що узагальнює підходи LSE, такі як E2C. Метод розвиває ідею вивчення локально-лінійного простору станів шляхом розглядання багатокрокового прогнозування, що дає змогу чіткіше контролювати кривизну шуканого простору. Продемонстровано, що метод перевершує E2C без суттєвих змін загальної моделі, на відміну від інших робіт, таких як РСС і РЗС. Розглянуто зв'язок між E2C і запропонованим методом та між їх відповідними рівняннями оновлень. Подано емпіричні докази, які свідчать, що ms-E2C дозволяє набагато краще вивчати простори прихованих станів з точки зору кривизни та прогнозованості наступних станів. Крім того, висвітлено певні проблеми стабільності, пов'язані з багатокроковими прогнозами, та способи їх вирішення.

Ключові слова: навчання репрезентацій, навчання керованих просторів, навчання з підкріпленням, латентний простір станів.