

Failure-aware Policy Learning for Self-assessable Robotics Tasks

Kechun Xu, Runjian Chen, Shuqi Zhao, Zizhang Li, Hongxiang Yu, Ci Chen, Yue Wang, Rong Xiong

Abstract—Self-assessment rules play an essential role in safe and effective real-world robotic applications, which verify the feasibility of the selected action before actual execution. But how to utilize the self-assessment results to re-choose actions remains a challenge. Previous methods eliminate the selected action evaluated as failed by the self-assessment rules, and re-choose one with the next-highest affordance (*i.e.* process-of-elimination strategy [1]), which ignores the dependency between the self-assessment results and the remaining untried actions. However, this dependency is important since the previous failures might help trim the remaining over-estimated actions. In this paper, we set to investigate this dependency by learning a failure-aware policy. We propose two architectures for the failure-aware policy: by representing the self-assessment results of previous failures as the variable state, and leveraging recurrent neural networks to implicitly memorize the previous failures. Experiments conducted on three tasks demonstrate that our method can achieve better performances with higher task success rates by less trials. Moreover, when the actions are correlated, learning a failure-aware policy can achieve better performance than the process-of-elimination strategy.

I. INTRODUCTION

A crucial problem for robotic applications in real world is how to promise action safety and effectiveness, especially for the applications of robot learning policies in unseen testing scenarios. A common way to deal with this problem is to utilize pre-defined self-assessment rules, which verify the feasibility of the selected actions before actual robot execution. For example, most of works in autonomous driving [2], [3], [4], [5], [6], [7] utilize the pre-built global map to forecast potential collision with the selected action. Once the selected action fails to pass the self-assessment tests, the learned policy has to re-choose an action. However, since there is no actual action execution, the observation stays invariant. Thus, the same action will be re-chosen by the learned policy and fail again, which raises a problem: how to re-choose actions among the remaining untried ones?

Considering that the self-assessment results enable the re-decision of a sequence of actions, in this paper, we formally state this re-choosing process as a sequential decision making problem under an invariant observation. Specifically, as shown in Fig. 1(a), given an observation o , a learned policy $\pi_0(o)$ generates the initial action a_0 , followed by a self-assessment module SA to indicate success or failure. If failed, the failure-aware policy π_{FA} is activated for re-choosing the action. Conditioned on the observation o , $\pi_{FA}(s_t|o)$ takes the self-assessment results of the previous

This work was supported in part by the National Key R & D Program of China under Grant 2021ZD0114500. Kechun Xu, Shuqi Zhao, Zizhang Li, Hongxiang Yu, Ci Chen, Yue Wang, Rong Xiong are with Zhejiang University, Hangzhou, China. Runjian Chen is with The University of Hong Kong. Corresponding author, wangyue@ipc.zju.edu.cn.

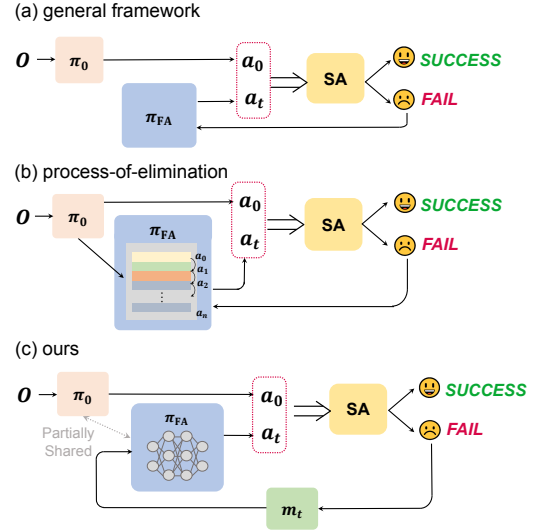


Fig. 1. (a) The general framework to utilize the self-assessment module. A learned policy π_0 takes as input the observation o , and selects the initial action a_0 . Then a self-assessment module SA evaluates the selected action. If failed, a failure-aware policy π_{FA} will re-choose the action a_t until getting successful feedback. (b) A common pipeline (process-of-elimination [1]) to design π_{FA} , which re-chooses action a_t in a sorting way using the action affordance map generated from π_0 . (c) Our pipeline, which constructs a failure memory representation m_t and uses a learning-based π_{FA} .

actions as state s_t to generate a new action a_t at step t until getting successful feedback from SA. Given the action affordance distribution predicted from π_0 , one intuitive way to design π_{FA} , which is shown in Fig. 1(b), is to choose the action with the next-highest affordance if the previously selected action is evaluated as unqualified, and so forth (*i.e.* process-of-elimination strategy [1], [2], [3], [4], [5], [6], [7]). However, there raises a further question of whether it is an optimal failure-aware policy. In other words, *under invariant observation, does the equality hold between the action with the second-highest affordance and the action with the highest affordance conditioned on previous failures?* The process-of-elimination strategy gives a positive answer to this question, which means previous action failures cause no influence on the affordance distribution of the remaining untried actions. However, we argue that the previous failure is an important prior for the action re-choosing. Therefore, the equality might not always hold.

In this paper, we set to investigate the dependency between the self-assessment results and the remaining untried actions by learning the failure-aware policy π_{FA} (Fig. 1(c)). We define self-assessable robotics tasks as those where the robot can evaluate itself by some self-assessment rules before actual action execution. Our key insight is to integrate the self-assessment results during the observation-invariant process

into the training of π_{FA} . We represent the results of the previous failure verified by the self-assessment module as m_t , which serves as the variable state s_t of π_{FA} . Also, Recurrent Neural Networks [8], [9] are helpful for the implicit representation to memorize the previous failures. Based on these points, we propose two architectures for the failure-aware policy. One tends to explicitly degenerate actions similar to failed ones, while another uses recurrent network to implicitly represent failure memory of the trial sequence. Experiments conducted on three tasks demonstrate that our method can achieve better performances with higher task success rates by less trials. Moreover, we find that when the actions are correlated, learning a failure-aware policy can achieve better performance than the process-of-elimination strategy. To summarize, our contributions are as follows:

- Our main contribution is to provide a learning-based perspective to utilize self-assessment results to learn a failure-aware policy for self-assessable robotics tasks.
- We propose two effective architectures for the failure-aware policy. One tends to degenerate actions similar to failed ones, while another uses recurrent network to implicitly represent failure memory of the trial sequence.
- We evaluate our method with three typical self-assessable robotics tasks, including sequential image classification, object reorientation and localization. Both simulated and real-world experiments validate the effectiveness of our policy, and the two architectures present different advantages according to the task properties. Moreover, when the actions are correlated, learning a failure-aware policy can achieve better performance than the process-of-elimination strategy.

II. RELATED WORKS

Robotic Self-assessment. Recently, robotic self-assessment has become a topic of interest in human-robot interaction. [10] highlights the importance of online Competence Assessment (CA) for safe real-world operation of robots. [11] further extends the term to Proficiency Self-Assessment (PSA), which shows the ability of a robot to predict, estimate, or measure its performance given a context or environment before action execution. Actually, this term can be extended to all robotics tasks, and a self-assessable robotics task means that the robot has some PSA metrics to evaluate its performance. For example, lots of works discard unsafe selected actions with prior knowledge of the global environment [2], [3], [4], [5], [6], [7], or with estimated environment dynamics [12], [13], [14], or by pre-acting with visual imagination [15], [16], [17], [18]. These works either simply use self-assessment metrics to filter actions and re-choose the action with the next-highest affordance or handcrafted safer policy, or consume a large amount of data to build the environment or imagination module, which brings another problem of estimation bias. In contrast, our work directly uses the result representation from self-assessment during the training process, thus easily integrating self-assessment into our policy distribution.

Failure-aware Policy Learning. In this paper, we define the failure-aware policy to be aware of the previous action

failures, and utilize the failed trials to predict more reliable actions. However, there are few studies under this definition. Hence, we review works that study policies predicting the success probability of the current action [13], [15], [16], [17], [18], [19], [20], or estimating the novelty and uncertainty of current observation [21], [22], [23], [24] which can be referred to as studies of failure-prediction policy. Other works like [14] predict the error of current action execution and propose an error-aware policy which takes as input the predicted future state error, and generates the corrected action. Similarly, [25] conducts an error detector by checking the reconstruction of the current state. However, most of these policies only measure the success probability of the current actions, without awareness of previous failed trials.

Robotic Exploration. Robotic exploration is a more general domain of our work, which can be regarded as exploration conditioned on the previous failures. Traditional works [26], [27], [28] design algorithms to explore states with less visiting times (*i.e.* with larger entropy). Also, many recent works follow the similar idea to set bonus to states deemed to be interesting or novel [29], [30], [31], [32], [33], [34]. Other works like [1] propose to train a set of policies to overlap a group of contexts with a disagreement penalty. Another view is to decouple the exploration policy from the exploitation policy, thus eliminating the inductive bias from task reward and stabilizing the policy training [35].

III. PROBLEM FORMULATION

In this work, we define self-assessable robotics tasks as those where a robot has some self-assessment rules to evaluate its performance before actually executing actions, and the evaluation results enable action re-choosing. Such self-assessment rules often serve as a safe module in real applications by predicting collision with simulation, which can provide a relatively accurate failure awareness. As a result, in this paper, we assume that self-assessment rules that correctly distinguish failure. Since there is no action execution, the observation stays invariant. Following the general framework in Fig. 1(a), in this section, we formulate and compare the process-of-elimination strategy and our method. Given the invariant observation o , a discrete action set \mathcal{A} , and a self-assessment module **SA**, we can formulate the same part of the process-of-elimination strategy and our method as follows:

$$\begin{aligned} a_0 &= \pi_0(o)|_{a_0 \in \mathcal{A}} \\ f_t(a_t|o) &:= \mathbf{SA}(o, a_t)|_{t \geq 0} = \begin{cases} 1, & \text{if successful} \\ 0, & \text{otherwise} \end{cases} \quad (1) \\ a_{t+1} &= \pi_{\text{FA}}(f_t|o)|_{a_{t+1} \in \mathcal{A}} \end{aligned}$$

where π_0 is the learned policy, a_t is the selected action at step t , f_t is a binary distribution defined by the **SA** results up to step t , and π_{FA} is a failure-aware policy conditioned on invariant observation o . Note that in this work, π_0 is assumed as a differentiable policy that contains an observation encoder and a decoder predicting the action affordances.

Process-of-elimination. This is an intuitive way to design π_{FA} (Fig. 1(b)) which chooses the action with the next-highest affordance after figuring out that the selected action

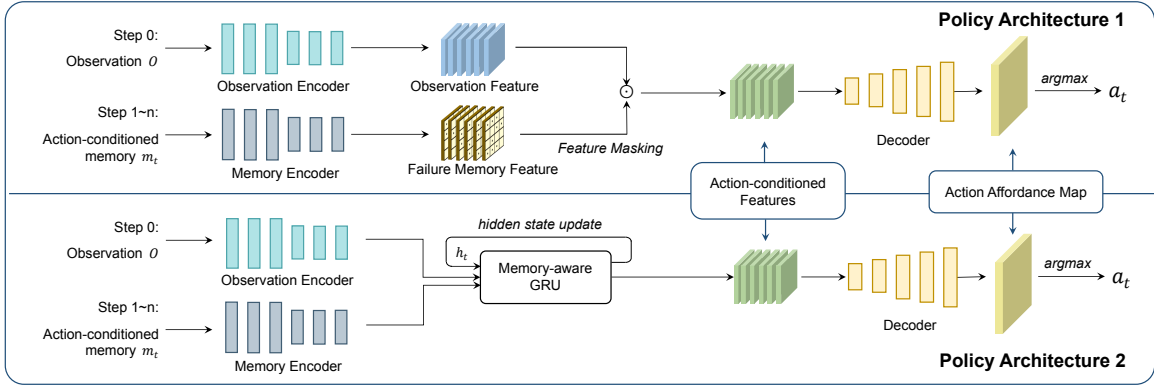


Fig. 2. Two architectures for the failure-aware policy π_{FA} . Note that the observation encoder and the decoder are components of the learned policy π_0 , which accelerates the failure-aware policy training process.

is unreliable, and can be formulated as follows:

$$a_{t+1} = \pi_{FA}(f_t|o)|_{a_{t+1} \in \mathcal{A}} = \pi_0(o) * f_t(a_t|o)|_{a_{t+1} \in \mathcal{A}} \quad (2)$$

In this formulation, π_{FA} is a handcrafted sorting policy, which simply multiplies the learned policy distribution by f_t . In this way, the previous action failures do not influence the affordance distribution of the remaining untried actions.

Failure-aware Learning Policy. In this paper, we propose to integrate the self-assessment results into the failure-aware policy training, and use the result representation m_t from f_t as the variable state (Fig. 1(c)). Our framework can be represented as follows:

$$\begin{aligned} m_t &\sim f_t(a_t|o) \\ a_{t+1} &= \pi_{FA}(m_t|o; \theta)|_{a_{t+1} \in \mathcal{A}} \end{aligned} \quad (3)$$

where π_{FA} is a learnable failure-aware policy, m_t is a representation of f_t with the same size as the action set \mathcal{A} , and θ represents the learnable parameters of neural network.

Self-assessment Representation. In our paper, m_t is used to represent the results of self-assessment. Concretely, it is a binary or normalized matrix of which each element represents the trial memory of the corresponding action, thus named action-conditioned memory. m_t is initialized at the beginning of every episode. For a binary m_t , it is initialized as an all-one matrix with the same size as the action set \mathcal{A} , while for a normalized representation, it is initialized as the normalized affordance map predicted by the learned policy π_0 . If an action is assessed as failed during the trial process, then the corresponding element is set to zero, thus updating m_t during the whole episode.

IV. METHODS

A. System Overview

Fig. 1(c) shows our pipeline. At the first step, the learned policy π_0 takes as input the observation o to get an initial action a_0 . If a_0 fails according to the self-assessment module SA, a self-assessment result representation m_t will be constructed and fed into the failure-aware policy π_{FA} to rechoose another action a_t . Note that the network parameters of π_0 and π_{FA} are partially shared. Compared to the process-of-elimination strategy, we propose to investigate the dependency between the self-assessment results and the remaining untried actions by learning the failure-aware policy.

B. Failure-aware Policy Architecture

Considering that the memory of previous failure can be utilized either explicitly or implicitly, we propose two architectures for the failure-aware policy π_{FA} , which are shown in Fig. 2. Both of them contain an observation encoder and a decoder, which are the shared components of the learned policy π_0 , which accelerates the training process.

Policy Architecture 1. The first proposed architecture is to encode the self-assessment result representation m_t into the same shape of the embedding feature generated from the observation encoder, and conduct an element-wise product to mask the observation feature by the failure memory feature embedding, which generates an action-conditioned feature. In this way, failure memory is explicitly considered into the feature embedding with a mask-like operation, thus affecting a shift on the action distribution.

$$\begin{aligned} e &= \mathbf{E}_o(o) \odot \mathbf{E}_m(m_t) \\ a_t &= \operatorname{argmax}_{a_t \in \mathcal{A}} \mathbf{D}(e) \end{aligned} \quad (4)$$

where \mathbf{E}_o , \mathbf{E}_m and \mathbf{D} symbolize the observation encoder, the memory encoder and the decoder respectively. a_t is selected from the action affordance map generated from \mathbf{D} . Note that the memory encoder can be simplified as a replica transform to the shape of the observation feature, or an identity transform in specific implementations.

Policy Architecture 2. The second architecture aims to bring the failure memory as a recurrent form across the episode using a memory-aware module. In this architecture, the feature embedding comes from the observation o at the first step, and from the updated m_t in the following steps. In this way, the failure memory is implicitly delivered across the decision process as a latent embedding.

$$e = \begin{cases} \mathbf{E}_o(o) & t = 0 \\ \mathbf{E}_m(m_t) & t > 0 \end{cases} \quad (5)$$

where \mathbf{E}_o and \mathbf{E}_m are of the same definitions in Eq. 4. With the feature embedding, the memory-aware module obtains awareness of the observation at the first step and implicitly represents it by the hidden vector, then produces new action distributions with recurrent memory in the following steps.

$$a_t = \operatorname{argmax}_{a_t \in \mathcal{A}} \mathbf{D}(\text{GRU}(e)) \quad (6)$$

where \mathbf{D} is of the same definition in Eq. 4, and **GRU** corresponds to the memory-aware module.

C. Policy Learning

We implement behavior cloning to train the learned policy π_0 . Note that there is no sequential decision making for this stage. For each step, the policy generates an action under the observation and gets feedback from self-assessment. Also, since parameters of π_0 and π_{FA} are partially shared, it can be seen as pre-training for the failure-aware policy π_{FA} .

To train the failure-aware policy, we apply value-based RL algorithms. For each episode, the policy is provided with the observation at the first step and chooses an action. If the feedback from the self-assessment module **SA** is positive, then the episode ends. Otherwise, the failure-aware policy π_{FA} will take as input the memory representation m_t and re-choose an action until receiving positive feedback from **SA**. During the training process, network parameters shared with π_0 will be fixed.

V. EXPERIMENTAL RESULTS

In this section, we will conduct experiments in three self-assessable robotics tasks to: 1) evaluate the effectiveness and advantages of our failure-aware policy compared with other methods; 2) show the different performances of the two failure-aware policy architectures; 3) investigate what kind of policy is optimal for the sequential decision making problem under invariant observation.

A. Experimental Setup

We consider three typical self-assessable robotics tasks (Fig. 3) for evaluation. The first task is sequential image classification on ImageNet [36] motivated by [1]. In this task, the robot observes an image at the beginning of an episode, and identifies a label for this image. After choosing a label, the self-assessment module will indicate whether the choice is correct or not. The second task is object reorientation, where a robot is supposed to choose a reorientation object pose to achieve a feasible pick-reorient-place process (*i.e.* successful path planning of the whole manipulation) with path planning cost as less as possible [13]. For this task, the policy is trained in SAPIEN [37] with a UR5 arm, tested with unseen samples, and evaluated in real world. Self-assessment is conducted by path planning algorithms, which guarantee the execution success. And the last is localization on synthetic dataset [38] and real-world dataset UPO [39] and Bicocca [40], [41], which predicts the position of the robot given a global map and an observed scan, and gets the assessment of the localization accuracy. In this task, if the predicted position is at the $k \times k$ neighborhood of the ground truth position, then the action is regarded as successful. Note that the robot will re-choose action after getting the self-assessment results until evaluated as successful or up to the limited trial times. In real applications, we can use registration algorithms as the self-assessment module which measures localization accuracy. Details of the self-assessment module of these tasks, the learned policy, and implementations of our two architectures




Fig. 3. Three tasks for evaluation.

TABLE I
TESTING PERFORMANCE OF SEQUENTIAL IMAGE CLASSIFICATION

Method	tsr/%	tns
RE	0.54	2.93
LPRE	70.06	1.01
SP	89.05	1.39
FMP-1	89.05	1.39
FMP-2	89.07	1.39

as well as the training and testing settings can be found in Appendix [42].

B. Metrics and Baselines

In this work, we aim to achieve a reliable action as soon as possible, since online self-assessment costs time and energy. Thus, we limit the sequential trial number to $t = 5$ times, and measure the algorithms with the following metrics.

- **Task Success Rate (tsr)**: Average task success rate across all testing samples. If the policy passes the self-assessment within 5 trials, then the corresponding testing sample is regarded as successful.
- **Trial Number to Success (tns)**: Average trial number to get a success feedback from self-assessment of all success samples.

All the tasks are measured with **tsr** and **tns**, which demonstrate the effectiveness and efficiency of action re-choosing based on previous failed trials. Also, an additional metric is tested for the object reorientation task:

- **Planning Cost (pc)**: Average path planning cost across all testing samples. It is a unique metric for object reorientation task, where robot is supposed to choose a reorientation pose with path planning cost as less as possible.

We compare the performance of our system to the following baseline approaches:

Random Exploration (RE). A policy which selects actions uniformly at random from the candidate action set \mathcal{A} .

Learned Policy with Random Exploration (LPRE). A policy that uses the learned policy π_0 for the first step, and if failed, then use the **RE** policy among the remaining actions.

Sorting Policy (SP). A policy that uses the process-of-elimination strategy.

Also, we name our policies of two architecture as **Failure-aware Feature Masking Policy (FMP-1)** and **Failure-aware Recurrent Memory Policy (FMP-2)** respectively.

C. Results

Comparisons to Baselines. First, we compare our method with baselines in three tasks. For sequential image classification, we evaluate each method with the validation dataset of ImageNet [36]. Note that in this task we regard the output of π_0 as the observation feature embedding (more details can be found in Appendix [42]). That is, **SP** and **FMP-1** have the same settings for this task, thus reporting the

TABLE II
TESTING PERFORMANCE OF OBJECT REORIENTATION

Method	tsr/%	100/pc	tns
RE	77.50	3.07	2.18
LPRE	79.33	3.48	1.90
SP	81.64	3.25	2.05
FMP-1	86.96	3.91	1.77
FMP-2	89.37	4.55	1.61

TABLE III
TESTING PERFORMANCE OF LOCALIZATION IN SYNTHETIC ENVIRONMENTS.

Method	tsr/%	tns
LPRE	83.95	1.01
SP	84.23	1.01
FMP-1	94.53	1.47
FMP-2	85.02	1.04

same performances. We can see from Table I that except for **RE**, **LPRE** shows the worst performance, while other three methods demonstrate similar performances. Referring to the analysis in [1], which proves that **SP** is an optimal policy for the sequential image classification task, our experimental results further figure out that **FMP-1** and **FMP-2** can also achieve optimal performance for this task.

For the object reorientation experiments, we present an additional metric “pc” in the reciprocal form, since the path planning cost will be infinite if the planning fails. Each method is evaluated with 207 unseen samples. Results in Table II show that **FMP-2** outperforms other methods across all metrics, followed by **FMP-1**, which demonstrates that integrating the previous failure memory into policy training endows better policy tune-up during online testing and better generalization performance. Also, **FMP-2** reports better performance than **FMP-1** across all metrics. This might be due to the fact that, in this task, there exist some candidate reorientation poses which are similar to each other, thus leading to a similar point cloud feature. However, similar poses do not mean similar task assessment results. For example, flipping an object will fail due to the collision between the gripper and the table. But if the object is with a similar pose which leaves small space for the gripper, such manipulation might succeed. For **FMP-1**, it utilizes previous failures by conducting feature masking, which might hinder some possible successful actions. Moreover, randomly choosing actions among the remaining ones (**LPRE**), or applying memory as a mask (**SP**) neglects the dependency between the failure trials and the remaining actions, which shows lower performance compared to our policies. Besides, we can find that the performance **RE** is not too bad due to the small action set of this task.

As for the localization task, the testing localization samples include three testing sequences in synthetic environments [38] and two testing sequences in real environments. Evaluation results in synthetic environments are shown in Table III ($k = 15$, more detailed results and ablation studies on k can be found in Appendix [42]). It is obvious that **FMP-1** achieves the best performance in the episode success rate with an average trial number less than 2. This large performance margin might come from the feature masking process, which




Fig. 4. Testing Performance of Localization in Real Environments.

hinders the similar feature of failed actions and encourages the policy to jump from the previous choices. **FMP-2** also shows better task success rate compared to **SP**. Note that the trial numbers to success of **LPRE** and **SP** are both close to 1, which indicates that these two policies cannot conduct effective adjustment by the self-assessment results. Fig. 4 demonstrates the testing results in real environments, which shows that our policies have better task success rate than **SP** and **LPRE** in all real environments. Overall, **FMP-2** shows better episode success rate, while **FMP-1** costs less trials to locate the position. This might be because the complex geometry of real-world maps calls for exploration in a small scope, which is the advantage of **FMP-2**, while **FMP-1** might hinder these similar positions if one of them fails. However, by jumping from the previous failure zone, **FMP-1** is able to achieve success in less trials. Note that we use $k = 15$ because the complex geometry of real-world maps leads to multi-modal predictions, and we choose the action with maximum affordance in our experiments.

Case Studies. Fig. 5 presents some testing cases in the localization task of three policies. Since **SP** does not change its original distribution, all its decisions depend on the distribution generated from π_0 . Instead, **FMP-1** concerns more on the feature correlation. When aware of a failed action, it can hinder the similar feature, thus jumping out of the previous failure zone. Also, by leveraging the recurrent implicit memory, **FMP-2** is also capable of adjusting the decisions, but shows a more conservative exploration process. Also, we can see the normalized probability changes of all feasible reorientation poses in an object reorientation case in Fig. 6. In this task, the feasible poses are not unique, and the similar poses (with near pose indexes) do not mean similar feasibility. In this task, **FMP-2** performs better because it tends to explore the near pose first to confirm its feasibility. More case studies are shown in Appendix [42].

D. Discussion

Considering the architecture designs of our models and all the experimental results, we further analyze the advantages of our method, and how to choose an optimal policy in a specific tasks:

What are the advantages of two policy architectures?

The same advantage of these two policy architectures is learnable. By integrating the self-assessment results into the policy training, **FMP-1** and **FMP-2** acquire the awareness of the previous failures, and have the capability of trimming the policy distribution according to these failures. However, these two architectures show different properties. **FMP-1** concerns more on the feature correlation, and tends to hinder the actions with similar features to that of the previously failed choices.

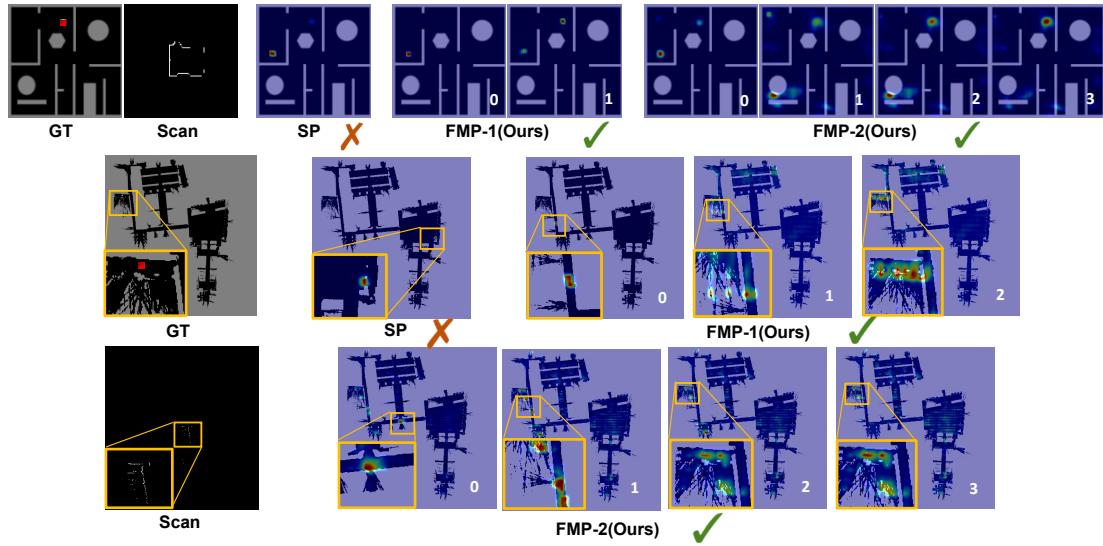


Fig. 5. Testing cases in localization task of three policies. The top row is a case in a synthetic environment. The left two columns show the global map with the ground truth position labeled as a red point, and the scan observation. Other columns show the prediction process and the distributions of three policies. The remaining row shows a case in the UPO dataset, where the keys are zoomed in with yellow boxes. The distribution is reflected by the color, where the value comes larger as the color comes closer to red. ✓ means that the policy successfully find the right position, while ✗ means a failure.




Fig. 6. A Testing case in reorientation task of three policies. The setting of this case is: magic clean, initial pose: $(-0.16, 0.16, 0.03, 1.59, 0.01, -3.14)$, target pose: $(-0.15, 0.09, 0.13, 0.0, 0.0, 0.0)$, feasible reorientation poses indexes: (4, 13, 18, 23). We plot the normalized probability distributions of feasible poses the selected pose at all trial steps, and the stars label successful trials. ✓ means that the policy successfully finds the right pose, while ✗ means a failure. The decision sequences of three policies are **SP**: 12→7→2→17→8, **FMP-1**: 22→15→3→2→23, **FMP-2**: 24→17→13.

Thus, **FMP-1** can achieve better performance in tasks where similar features (*e.g.* pose, geometry, visual attribute, and etc) lead to similar self-assessment results. Instead, **FMP-2** pays more attention to the recurrent memory, which encodes the observation feature and the previous trials. Consequently, **FMP-2** conducts a more conservative exploration process than **FMP-1**, and demonstrates better performances in tasks requiring an adjustment in a small scope.

What kind of policies is optimal? For a task where similar features lead to similar self-assessment results, **FMP-1** can help jump out of local minimum, while for a task which needs to adjust in a small scope of the initial action, **FMP-2** shows better exploration strategy. Also, analyzing the results of the three tasks, we figure out that, when the actions are correlated, our method outperforms the process-of-elimination strategy. That is, *when the actions are correlated, the equality is broken between the action with the next-highest affordance and the action with the highest affordance conditioned on previous failures*. For sequential image classification, there is little correlation among different class choices. And the training data is adequate for fitting the distributions of all classes. Therefore, a simple process-of-elimination strategy can achieve optimal performance. Instead, as the action correlation increases, integrating action correlation into the

policy learning endows better performance. Hence, we can achieve better performances in the localization task, and show the biggest advantage in the object reorientation task, whose actions have the highest correlation.

E. Conclusion and Limitation

In this paper, we propose to integrate the self-assessment results to learn a failure-aware policy, and propose two policy architectures. Experiments in three self-assessable robotics tasks demonstrate that our method outperforms other methods with higher task success rate with less trials. Moreover, we find that the action correlation has a large impact on the effect of our algorithm. The main limitation of our method lies in the assumption of the discrete action set. This limitation comes from the construction of the representation m_t . In our paper, it is represented as a finite matrix, of which each element corresponds an action. In future work, more general representations of m_t for continuous actions can be studied. Also, in this paper, we assume that the self-assessment module can accurately predict failure, and π_0 is a learned differentiable policy with a feature bottleneck layer. Further works can involve the assessment uncertainty in more real-world applications, and extend to more general π_0 . More analysis and experiments can be assessed in Appendix [42].

REFERENCES

- [1] D. Ghosh, J. Rahme, A. Kumar, A. Zhang, R. P. Adams, and S. Levine, "Why generalization in rl is difficult: Epistemic pomdps and implicit partial observability," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [2] D. Isele, A. Nakhaei, and K. Fujimura, "Safe reinforcement learning on autonomous vehicles," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1–6.
- [3] K. Srinivasan, B. Eysenbach, S. Ha, J. Tan, and C. Finn, "Learning to be safe: Deep rl with a safety critic," *arXiv preprint arXiv:2010.14603*, 2020.
- [4] H. Krasowski, X. Wang, and M. Althoff, "Safe reinforcement learning for autonomous lane changing using set-based prediction," in *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2020, pp. 1–7.
- [5] K. Mokhtari and A. R. Wagner, "Safe deep q-network for autonomous vehicles at unsignalized intersection," *arXiv preprint arXiv:2106.04561*, 2021.
- [6] K. Lin, R. Zhao, Z. Xu, and J. Zhou, "Efficient large-scale fleet management via multi-agent deep reinforcement learning," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2018, pp. 1774–1783.
- [7] D. Chen, Z. Li, Y. Wang, L. Jiang, and Y. Wang, "Deep multi-agent reinforcement learning for highway on-ramp merging in mixed traffic," *arXiv preprint arXiv:2105.05701*, 2021.
- [8] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [9] K. Cho, B. Merriënboer, C. Gulcehre, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," in *EMNLP*, 2014.
- [10] G. J. Burghouts, A. Huizing, and M. A. Neerincx, "Robotic self-assessment of competence," *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2020.
- [11] A. Norton, H. Admoni, J. Crandall, T. Fitzgerald, A. Gautam, M. Goodrich, A. Saretzky, M. Scheutz, R. Simmons, A. Steinfeld *et al.*, "Metrics for robot proficiency self-assessment and communication of proficiency in human-robot teams," *ACM Transactions on Human-Robot Interaction*, 2022.
- [12] P. Deptula, H.-Y. Chen, R. A. Licitra, J. A. Rosenfeld, and W. E. Dixon, "Approximate optimal motion planning to avoid unknown moving avoidance regions," *IEEE Transactions on Robotics*, vol. 36, no. 2, pp. 414–430, 2019.
- [13] K. Xu, H. Yu, R. Huang, D. Guo, Y. Wang, and R. Xiong, "Efficient object manipulation to an arbitrary goal pose: Learning-based anytime prioritized planning," *2022 IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [14] V. Kumar, S. Ha, and C. K. Liu, "Error-aware policy learning: Zero-shot generalization in partially observable dynamic environments," *Robotics: Science and Systems (RSS)*, 2021.
- [15] C. Finn and S. Levine, "Deep visual foresight for planning robot motion," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 2786–2793.
- [16] F. Ebert, C. Finn, S. Dasari, A. Xie, A. Lee, and S. Levine, "Visual foresight: Model-based deep reinforcement learning for vision-based robotic control," *arXiv preprint arXiv:1812.00568*, 2018.
- [17] A. Wang, T. Kurutach, K. Liu, P. Abbeel, and A. Tamar, "Learning robotic manipulation through visual planning and acting," in *Robotics: science and systems*, 2019.
- [18] N. Di Palo and E. Johns, "Safari: Safe and active robot imitation learning with imagination," *arXiv preprint arXiv:2011.09586*, 2020.
- [19] K. Xu, H. Yu, Q. Lai, Y. Wang, and R. Xiong, "Efficient learning of goal-oriented push-grasping synergy in clutter," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6337–6344, 2021.
- [20] Z. Xu, Z. He, J. Wu, and S. Song, "Learning 3d dynamic scene representations for robot manipulation," in *Conference on Robot Learning*. PMLR, 2021, pp. 126–142.
- [21] C. Richter and N. Roy, "Safe visual navigation via deep learning and novelty detection," 2017.
- [22] L. Wellhausen, R. Ranftl, and M. Hutter, "Safe robot navigation via multi-modal anomaly detection," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1326–1333, 2020.
- [23] G. Kahn, A. Villafior, V. Pong, P. Abbeel, and S. Levine, "Uncertainty-aware reinforcement learning for collision avoidance," *arXiv preprint arXiv:1702.01182*, 2017.
- [24] B. Lütjens, M. Everett, and J. P. How, "Safe reinforcement learning with model uncertainty estimates," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8662–8668.
- [25] J. Wong, A. Tung, A. Kurenkov, A. Mandlekar, L. Fei-Fei, S. Savarese, and R. Martín-Martín, "Error-aware imitation learning from teleoperation data for mobile manipulation," in *Conference on Robot Learning*. PMLR, 2022, pp. 1367–1378.
- [26] R. Martínez-Cantin, N. de Freitas, A. Doucet, and J. A. Castellanos, "Active policy learning for robot planning and exploration under uncertainty," in *Robotics: Science and systems*, vol. 3, 2007, pp. 321–328.
- [27] J. Peters, K. Mulling, and Y. Altun, "Relative entropy policy search," in *Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.
- [28] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [29] R. Houthoofd, X. Chen, Y. Duan, J. Schulman, F. De Turck, and P. Abbeel, "Vime: Variational information maximizing exploration," *Advances in neural information processing systems*, vol. 29, 2016.
- [30] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos, "Unifying count-based exploration and intrinsic motivation," *Advances in neural information processing systems*, vol. 29, 2016.
- [31] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, "Curiosity-driven exploration by self-supervised prediction," in *International conference on machine learning*. PMLR, 2017, pp. 2778–2787.
- [32] H. Tang, R. Houthoofd, D. Foote, A. Stooke, O. Xi Chen, Y. Duan, J. Schulman, F. DeTurck, and P. Abbeel, "# exploration: A study of count-based exploration for deep reinforcement learning," *Advances in neural information processing systems*, vol. 30, 2017.
- [33] Y. Burda, H. Edwards, A. Storkey, and O. Klimov, "Exploration by random network distillation," *arXiv preprint arXiv:1810.12894*, 2018.
- [34] M. C. Machado, M. G. Bellemare, and M. Bowling, "Count-based exploration with the successor representation," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 04, 2020, pp. 5125–5133.
- [35] W. F. Whitney, M. Bloesch, J. T. Springenberg, A. Abdolmaleki, K. Cho, and M. Riedmiller, "Decoupled exploration and exploitation policies for sample-efficient reinforcement learning," *arXiv preprint arXiv:2101.09458*, 2021.
- [36] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [37] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang *et al.*, "Sapien: A simulated part-based interactive environment," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 097–11 107.
- [38] R. Chen, H. Yin, Y. Jiao, G. Dissanayake, Y. Wang, and R. Xiong, "Deep samplable observation model for global localization and kidnapping," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2296–2303, 2021.
- [39] R. Ramón-Vigo, J. Pérez, F. Caballero, and L. Merino, "Navigating among people in crowded environment: Datasets for localization and human robot interaction," in *Proceedings of the Workshop on Robots in Clutter: Perception and Interaction in Clutter, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Citeseer, 2014.
- [40] A. Bonarini, W. Burgard, G. A. E. Fontana, M. Matteucci, D. Sorrenti, and J. Tardos, "Rawseeds: Robotics advancement through web-publishing of sensorial and elaborated extensive data sets," in *Workshop on Benchmarks in Robotics Research at IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2006)*, 2006, pp. 1–5.
- [41] S. Ceriani, G. Fontana, A. Giusti, D. Marzorati, M. Matteucci, D. Migliore, D. Rizzi, D. G. Sorrenti, and P. Taddei, "Rawseeds ground truth collection systems for indoor self-localization and mapping," *Autonomous Robots*, vol. 27, no. 4, pp. 353–371, 2009.
- [42] K. Xu, "Paper with appendix," <https://arxiv.org/abs/2302.13024>, 2021.
- [43] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [44] R. Wightman, "Pytorch image models," <https://github.com/rwightman/pytorch-image-models>, 2019.
- [45] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.