

Conceptual Reinforcement Learning for Language-Conditioned Tasks

Shaohui Peng^{1, 2, 3}, Xing Hu¹, Rui Zhang^{1, 3}, Jiaming Guo^{1, 2, 3}, Qi Yi^{1, 3, 4},
Ruizhi Chen^{2, 5}, Zidong Du^{1, 3}, Ling Li^{2, 5}, Qi Guo¹, Yunji Chen^{1, 2, †}

¹ SKL of Processors, Institute of Computing Technology, CAS

² University of Chinese Academy of Sciences

³ Cambricon Technologies

⁴ University of Science and Technology of China

⁵ SKL of Computer Science, Institute of Software, CAS

{pengshaohui18z, huxing, zhangrui, guojiaming18s, duzidong, guoqi, cyj}@ict.ac.cn,
yiqi@mail.ustc.edu.cn, {ruizhi, liling}@iscas.ac.cn

Abstract

Despite the broad application of deep reinforcement learning (RL), transferring and adapting the policy to unseen but similar environments is still a significant challenge. Recently, the language-conditioned policy is proposed to facilitate policy transfer through learning the joint representation of observation and text that catches the compact and invariant information across environments. Existing studies of language-conditioned RL methods often learn the joint representation as a simple latent layer for the given *instances* (episode-specific observation and text), which inevitably includes noisy or irrelevant information and cause spurious correlations that are dependent on instances, thus hurting generalization performance and training efficiency. To address this issue, we propose a conceptual reinforcement learning (CRL) framework to learn the *concept-like* joint representation for language-conditioned policy. The key insight is that concepts are compact and invariant representations in human cognition through extracting similarities from numerous instances in real-world. In CRL, we propose a multi-level attention encoder and two mutual information constraints for learning compact and invariant concepts. Verified in two challenging environments, RTFM and Messenger, CRL significantly improves the training efficiency (up to 70%) and generalization ability (up to 30%) to the new environment dynamics.

Introduction

Deep reinforcement learning has been successfully applied in various areas such as video games (Mnih et al. 2013) and robot control (Lillicrap et al. 2016). However, it is still a significant challenge to transfer and adapt the policy to unseen but similar environments (Kirk et al. 2021). Recently, researchers propose language-conditioned policy that adopts the language as an intermediate information channel to facilitate policy transfer (Narasimhan et al. 2018; Oh et al. 2017). Based on extra-textual descriptions that specify environment dynamics, the language-conditioned policy can learn the connection between training and testing environments to promote generalization ability.

The key challenge of language-conditioned policy is how to learn the joint representation of observation and textual

description. The joint representation should catch the compact and invariant information about the actual cause of the reward that is shared in various training environments, so that the policy can generalize to unseen environments. Existing studies of language-conditioned reinforcement learning methods often implicitly learn the joint representation as a simple latent layer for the given *instances* (episode-specific observation and text), which cannot guarantee to catch the invariant factors across different instances (Zhong et al. 2020; Hanjie et al. 2021; Branavan et al. 2011; Narasimhan et al. 2018). Such a strategy inevitably includes noisy or irrelevant information and causes spurious correlations that are dependent on instances, thus hurting generalization performance and training efficiency (Zhong et al. 2021).

Inspired by the human activities that extract the similarities across numerous instances in the real-world to form “concept” for better understanding the world from the abstract perspective and solving similar tasks efficiently, we propose to explicitly learn the *concept-like* representation for language-conditioned policy. Concretely, we mainly consider two important features for such representation, *invariance* and *compactness*, which boost both transfer performance and training efficiency. Based on the invariance of concepts, policies can be quickly transferred between similar environments. Meanwhile, based on the abstraction of concepts, the learned joint representation is more compact, and thus the complexity of policies can be reduced to improve training efficiency. For example, when playing video games, we can treat various monster instances with different names or appearances as the same concept “enemy”, and handle them in a similar policy such as attacking them to get drops, thereby adapting to an unseen environment quickly.

Based on the above analyses, we propose a conceptual reinforcement learning (CRL) framework which learns concept-based representation (referred as concepts) for language-conditioned policy to promote sample efficiency and generalization performance in unseen testing environments. In CRL, we propose a multi-level attention concept encoder and two mutual information constraints for learning concepts. Specifically, in the concept encoder, each level of the attention module generates one concept, and the generated concepts participate in subsequent levels of the attention module. All outputs of the concept encoder are con-

catenated as the final concept-based representation. Such a mechanism simulates the multi-step reasoning process and guarantees that the encoder can extract concept-related information from text and observation. More importantly, we augment two mutual information constraints to guarantee the invariance and compactness of the learned concepts, which is the major difference from latent representation learned implicitly by existing methods. We minimize mutual information between concepts and entities of observation in different environments, which guides the concepts to catch the shared properties and keep them invariant across different tasks. Besides, we leverage an information bottleneck between the textual description and concepts to ensure the compactness of the concepts. Finally, we train the policy upon the generated concept-based representation to drive the learning of concepts by the RL objective. In conclusion, the concept-based representation learned by CRL has three evident advantages, *invariance*, *compactness*, and *better interpretability*. First, the invariance significantly promotes the generalization performance of the policy in unseen environments. Second, the compactness makes the training process of the policy more efficient. Third, the language-conditioned policy constructed upon concepts has higher interpretability to make the policy easier to employ and understand.

We verify CRL on two challenge environments with textual descriptions, RTFM (Zhong et al. 2020) and Messenger (Hanjie et al. 2021), both of which are benchmarks to evaluate the generalization ability of language-conditioned policy to new environment dynamics. The former focus on multi-step reasoning across observation and textual description, while the latter focus on Out-Of-Distribution (OOD) settings of test environments. Compared with state-of-the-art methods, the proposed CRL gains significant generalization performance (up to 30%) and training efficiency promotion (up to 70%) on both RTFM and Messenger.

Background

In this paper, we focused on conceptual reinforcement learning in the language-conditioned policy that applied to the environment with textual description. The text in such environments specifies the relation of entities, goals, and some environment dynamics, and it is regenerated every episode to initialize different but similar environments. For example, Read To Fight Monsters (RTFM) is a well-known environment where the agent needs reading and multi-step relation reasoning to understand the relation between tools and monsters. Compared with other RL environments, the input of the policy in such environments includes two parts, the world observation o and the additional textual description t :

- **World observation** $o \in \mathbb{R}^{h \times w \times d}$: h and w represent the height and width of the grid world. A d -word symbolic name represents the content in the grid. Some non-overlapping entities $e \in \mathbb{R}^{n \times d}$ are scattered in the grid world, including the agent and other interactive objects. The entities in o except the agent are sampled from an invisible large candidate set and thus may change in every episode. In different episodes, the property of the same entities is unfixed and needs to be inferred from the tex-

tual description of the current episode. Figure 2 shows an example of world observation. There are five entities in o , including the agent, monsters, and tools.

- **Text** $t \in \mathbb{R}^{num_{sent} \times l_{sent} \times d_t}$: The textual description of the environment contains less than num_{sent} sentences, and each sentence consists of less than l_{sent} tokens. d_t is the embedding dimension of each token. The text keeps fixed in an episode but changes between episodes. t describes the goal, the entities’ property, and the relation between entities. Figure 2 shows an example of textual description in the RTFM environment. “The Star Alliance team is made up of beetle, jackal, and shaman. Defeat the Star Alliance.” explains that the entity “shaman” is the enemy, and the goal of the agent is to defeat it.

The language-conditioned policy $\pi(a|o, t)$ takes world observation o and text t as inputs and outputs an action a . As the text t described some dynamics of the environment and changes every episode, the reward function $r_\pi(o, t, a, o')$ also takes text t as input. The target is to find an optimal policy π to maximize the accumulated rewards $R_\pi(o, t) = E[\sum_{n=0}^{\infty} \gamma^n r(o_{t+n}, t, a_t, o_{t+n+1}) | o_t = o]$. Existing methods implicitly learn the joint representation as a simple latent layer for the given instances (episode-specific observation and text), which cannot guarantee catching the invariant factors across different instances. Our method (CRL) explicitly learned invariant compact concepts for policy to promote the policy’s generalization performance and training efficiency.

The Conceptual Reinforcement Learning (CRL) Framework

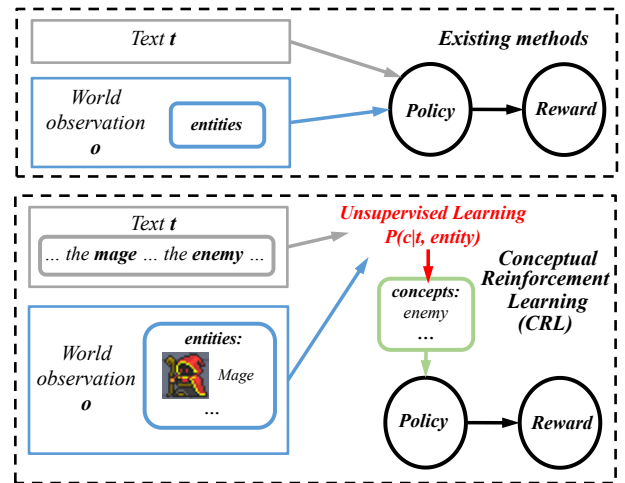


Figure 1: Different modeling of CRL and existing language-conditioned policy. CRL additionally learns the concept encoder $P(c|t, entity)$ compared with existing methods. Taking a certain episode as an example, the encoder computes the concept of the entity “Mage” as “enemy” through reasoning in the text. The concepts are shared properties of different entities in various episodes and are fed into the language-conditioned policy as a joint representation.

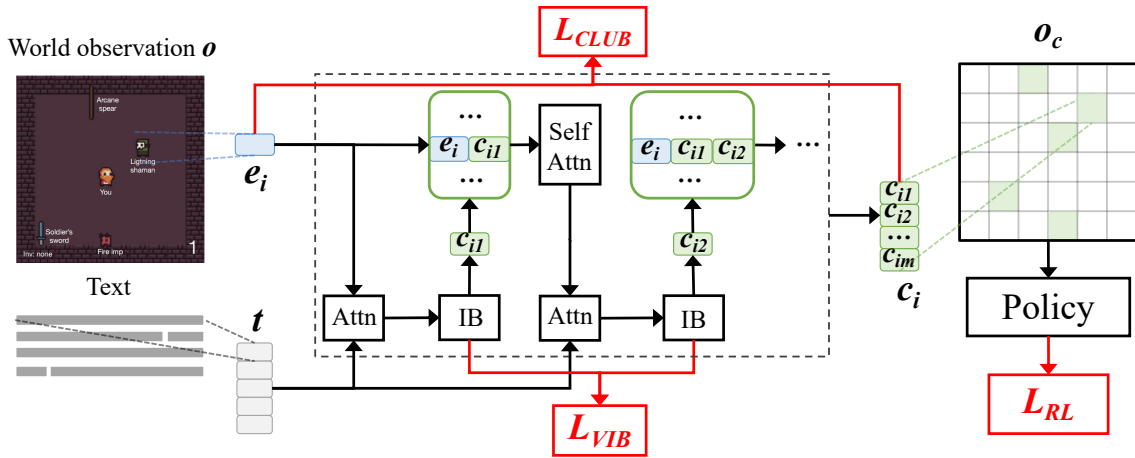


Figure 2: Overview of Conceptual Reinforcement Learning (CRL) framework. Take an example from RTFM. CRL gets inputs including the world observation o and the text t . The multi-level concept encoder of CRL leverage entities e_i to conduct multiple attention over the text t to get e_i 's concepts $c_i = \{c_{i1}, \dots, c_{im}\}$. CRL leverages two mutual information constraints to guarantee the two advantages of produced concepts, where L_{CLUB} for invariance and L_{VIB} for compactness. All entities' concepts are combined to form the new concept-based representation o_c for the language-conditioned policy.

In this section, we introduce the Conceptual Reinforcement Learning (CRL) framework, aiming to learn the concept-like representation with both advantages of *invariance* and *compactness* across different scenarios for better generalization. *Invariance* enables the concepts can generalized to unseen similar environments, and *compactness* promotes the training efficiency of the policy. To achieve these two goals, CRL includes the multi-level attention module as the concept encoder, which simulates the multi-step reasoning process to extract concept-related information from textual description and observation. Meanwhile, CRL leverages two mutual information constraints to guarantee the invariance and compactness of the learned concepts.

Concept Encoder

Motivation We are inspired that humans can extract similarities from numerous instances in the real-world to form concepts, which are compact and invariant representations. We propose introducing concepts as the representation for language-conditioned policy to make the policy be trained efficiently and generalize to similar scenarios. The analysis of the modeling for the language-conditioned policy of the existing methods and our method that introduce concepts as intermediate representation is as follows.

Existing methods: Figure 1 shows the modeling of existing methods optimizing policy for higher reward relies on all entities and text. However, entities and the text change every episode and may have distribution shifts on unseen similar scenarios. For example, the role assignments to the same entity are not overlapped between training and test games in RTFM and Messenger. The implicitly learned joint representation inevitably catches irrelevant information about entities and further causes spurious correlations.

Concept Learning: As shown in Figure 1, the CRL framework introduces an explicitly learned representation,

called concepts, between inputs and the policy. Concepts (c) are abstractions of entities' shared properties inferred from the interactive information of text and entities. For example, there is a role concept in Messenger. Every entity should be a sender, receiver, or decoy regardless of its name or movement patterns. We model concepts $P(c|t, entity)$ through an unsupervised learning encoder and replace the world observation o as a concept-based world representation consisting of learned concepts. As a result, the language-conditioned policy only relies on concepts, which are shared properties between entities that the agent needs to understand and exploit for solving tasks. Concepts have two evident advantages, *invariance* and *compactness*. Invariance means that concepts are an invariant representation for entities not only across episodes but also in train or test scenarios, making the policy generalize well in similar unseen environments. Besides, compared with the large entities candidate set and complex textual description, the concept is a compact representation, which makes the data utilization of policy training much more efficient.

Definitions In each game episode, the environment offers the world observation $o \in \mathbb{R}^{h \times w \times d}$, which includes n entities $e = \{e_1, \dots, e_n\} \in \mathbb{R}^{n \times d}$, and a textual description $t \in \mathbb{R}^{num_{sent} \times l_{sent} \times d_t}$. We assume that there are m kinds of invisible concept variables $C = \{C_1, C_2, \dots, C_m\}$ representing the shared properties of entities. Each entity e_i can be labeled as concatenation of concept values $c_i = \{c_{i1}, \dots, c_{im}\} \in \mathbb{R}^{m \times d_c}$, where d_c is the embedding dimension of a single concept. The concept encoder is to model the distribution of each entity's concepts $p(c_i|e_i, t)$ given the entity e_i and text t . The concepts generated by the concept encoder may not be identical to the ground truth concept since we learn the encoder in an unsupervised way, but the generated concepts have the same advantages as actual concepts, which are invariant and compact.

Architecture The overall design of the CRL framework is replacing e_i in the world observation o as derived concept values c_i to construct the concepts-based observation for the policy training. As the most important part of the CRL framework, the multi-level attention concept encoder $f_\theta(e_i, t)$ takes in the entity embedding $e_i \in \mathbb{R}^{d_e}$ and text t as input and output concepts c_i . The dotted box in Figure 2 shows the concept encoder. The concept encoder contains m Attention modules, and each module outputs a single concept $c_{ij} \in \mathbb{R}^{d_c}$. These modules produce both basic concepts inferred by raw text information and complex concepts derived from the multi-step reasoning process for the underlying information in the text and basic concepts.

Specifically, the first Attention module could generate concepts inferred by direct information in the text. For example, in Messenger, different entities have different roles, and the text claims this information explicitly. “the thing that is not able to move is the mage who possesses the enemy that is deadly” points out that the role concept of entity “mage” is “enemy”. In this Attention module, we compute query q_1 from the entity embedding e_i and the key k_1 and value v_1 from the text t as follows:

$$\begin{aligned} q_1 &= MLP(e_i) \\ k_1 &= GRU_{key_1}(t) \\ v_1 &= GRU_{val_1}(t). \end{aligned} \quad (1)$$

And then, we get the direct information of the concept c_{i1} through the standard Attention module as shown in Figure 2. The following IB module is a Variational Information Bottleneck encoder for the concept c_{i1} , aiming to improve the compactness of concepts and will be explained in the next section.

The subsequent Attention modules have an extra Self-Attention model to simulate the multi-step reasoning process to generate complex concepts, considering the acquisition of some complex concepts needs not only raw text information but also the preceding concepts. For example, in RTFM, “Fire monsters are defeated by fanatical and shimmering weapons.” demonstrates the relation between “fire” entities and “fanatical” entities. However, to reason whether the entity “fanatical sword” is the concept “useful tool”, we need to know the concept of the “fire” entity is an “enemy” or not. So we need to build a multi-level Attention mechanism that concatenates entities with discovered shallow concepts as the query to discover the deep concepts. As shown in Figure 2, the query tensors in the following Attention module are obtained through a Self-Attention mechanism on the concatenation of all entity embeddings and their preceding concepts:

$$\begin{aligned} q_j &= SelfAttn(Concat(e_i, c_{i1}, \dots, c_{i,j-1})) \\ k_j &= GRU_{key_j}(t) \\ v_j &= GRU_{val_j}(t), \text{ where } 1 < j \leq m \end{aligned} \quad (2)$$

Finally, we concatenate all concepts together to get c_i , and replace all $e_i, i \in \{0, \dots, n\}$ in the world observation o with

their concepts to construct a new concept-based world representation $o_c \in \mathbb{R}^{h \times w \times d_c}$. We use o_c rather than the original world observation o and text t as the representation of the environment state to train the policy $\pi(a|o_c)$.

Mutual Information Constraints

The model architecture of CRL introduced in the last section can obtain an effective concept representation during concept-based policy training. To guarantee the two important features, invariance and compactness, of produced concepts, we add extra constraints on the concept encoder f_θ during policy training. Specifically, we consider the invariance and compactness properties from the view of information theory, detailed as follows.

Invariance Constraint As the entities change every episode and the properties of the same entity inferred from the text of different episodes may differ, the concepts should only reflect the shared properties across episodes and not contain information about the episode-specific entities. Therefore, although the concepts are inferred from entities and the text, we need to make the marginal distributions of produced concepts and entities independent to ensure invariance. As a result, the policy constructed upon the concepts would focus on abstract properties rather than information about episode-specific entities, so it can eliminate spurious correlations. We minimize the mutual information (MI), a fundamental measure of the dependence between two random variables, to force the learned concept independent of the entity. Mathematically, the MI between two variables X and Y is:

$$I(X; Y) = \mathbb{E}_{p(x,y)} \left[\log \frac{p(x,y)}{p(x)p(y)} \right] \quad (3)$$

We adopt CLUB (Cheng et al. 2020) to minimize the upper bound of mutual information $I(e; c)$ between the entities e and its concept c as shown in Figure 2. With N samples $\{e_i, c_i\}_{i=1 \rightarrow N}$, the upper bound of MI estimated by CLUB is:

$$\begin{aligned} I_{CLUB}(e; c) &= \\ &= \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N [\log q(c_i|e_i) - \log q(c_j|e_i)], \end{aligned} \quad (4)$$

where $q(c_i|e_i)$ is a multi-layer MLP predictor. We can minimize mutual information between entities and their concepts through the following extra loss:

$$\begin{aligned} L_{CLUB}(\theta) &= \sum_{i=1}^n I_{CLUB}(c_i; e_i) \\ &= \sum_{i=1}^n I_{CLUB}(f_\theta(e_i, t); e_i), \end{aligned} \quad (5)$$

where n is the number of entities in the world observation o .

Compactness Constraint The RL training process implicitly ensures the effectiveness of concepts since the agent solves tasks using concept-based observation. We propose an additional compactness constraint to eliminate noisy or

irrelevant information during concept encoding, so that only necessary information for the task is included in the concepts. We adopt an information bottleneck (IB) to compress the concepts, a technique in the information theory to compress the noise information from the input variable X and retain the representative information for the output variable Y in the middle variable Z . The mutual information objective of IB is as follows:

$$\min_{p(z|x)} I(X; Z) - \beta I(Z; Y) \quad (6)$$

In our work, as the target of retaining useful information is implicitly done through the RL objective, we only utilize the first term in equation 6 to compress irrelevant information when encoding text t to the concept c . We adopt Deep Variational Information Bottleneck (Deep VIB) to minimize the upper bound of $I(c; t)$ as shown in Figure 2. In practice, we implement the encoder head (IB in Figure 2) of f_θ with the form of $\mathcal{N}(c_{ij} | \mu^{d_c}(t'_{ij}), \sigma^{d_c}(t'_{ij}))$, representing the Gaussian distribution of the j -th concept of e_i . And t'_{ij} is the output of the j -th Attention module containing the information of the concept, μ^{d_c} and σ^{d_c} are MLP layers to output distribution parameters with d_c dimensions. Then we can sample a noise ϵ from $\mathcal{N}(0, 1)$ to compute the concept c_{ij} using the reparameterization trick. We minimize the Kullback-Leibler divergence between encoder head $p_\theta(c_{ij}|t'_{ij})$ and Normal distribution $\mathcal{N}(0, 1)$ to compress the noise and irrelevant information:

$$L_{VIB}(\theta) = \sum_{i=1}^n \sum_{j=1}^m \mathbb{E}_{\epsilon \sim p(\epsilon)} [KL(p_\theta(c_{ij}|t'_{ij}), \mathcal{N}(0, 1))] \quad (7)$$

In summary, to keep the invariance and compactness of the concept, we add extra mutual information constraints L_{CLUB} and L_{VIB} on the concept encoder when training policy. Concretely, these two constraints can be easily integrated into the standard RL method. The new objective is:

$$L_{CRL}(\theta) = L_{RL}(\theta) + \alpha_1 L_{CLUB}(\theta) + \alpha_2 L_{VIB}(\theta), \quad (8)$$

where $L_{RL}(\theta)$ is the original RL objective, and coefficients α_1, α_2 are hyperparameters. (Details are in Appendix B).

Experimental Setup

Environments

To verify the performance of CRL, we evaluate the framework on two challenging benchmarks, RTFM and Messenger. **RTFM** (Zhong et al. 2020) is a game that the agent needs to **Read Text** to obtain the correct tool in the grid world and then **Fight** the correct **Monster**. The key challenge of RTFM is that the agent needs multi-step reasoning to interpret the text. RTFM has a train set of environment dynamics (including entities and role assignments) and an independently identically distribution (i.i.d.) held-out test set. The environment also has four kinds of complexity settings, including base 6×6 grids setting (6×6), describing entities in the form of many-to-one group assignments to make disambiguation more difficult (`group`), allowing moving monsters that hunt down the player (`dyna`) and using natural

language descriptions (`n1`). Because of the high complexity, the environment also offers four curriculum stages. **Messenger** (Hanjie et al. 2021) is also a grid environment (10×10). The agent must extract the role information of the entities from the text manual to acquire a message from the sender and deliver it to the receiver. Unlike RTFM, Messenger’s key challenge is the out-of-distribution (OOD) problem. The entities and role assignments are no-repeat and have distribution shifts between train and test scenarios, which may lead to spurious correlations. The Messenger offers three difficulty stages, including only message acquiring or delivering (S1), both acquiring and delivering (S2), and adding decoy entities and irrelevant descriptions (S3).

The details of the environment and the CRL implementation are shown in Appendix A and B.

Baselines

We compare CRL with two state-of-art methods, **txt2 π** on RTFM and **EMMA** on Messenger, respectively.

txt2 π (Zhong et al. 2020) exploits the prior knowledge of the observation structure in RTFM, which builds the representation by capturing the three-way interactions between the goal description, the environment dynamic description, and the world observation. **txt2 π** outperforms the language-conditioned CNNs and the Feature-wise linear modulation (FiLM) on RTFM. **EMMA (Entity Mapper with Multimodal Attention)** (Hanjie et al. 2021) uses the entity-conditioned attention module to select relevant information in the text to generate the representation. It achieves the best performance on Messenger compared with other baselines.

Results

We first compare the performance of CRL and the state-of-the-art in RTFM and Messenger to show its effectiveness. Then, we conduct ablation studies to show the influential impact of the mutual information constraints and pre-defined concepts number during concept generation. At last, we visualize the concept embeddings in both environments to intuitively show what CRL learns.

Transfer in RTFM

We compare the transfer performance of CRL and **txt2 π** in the RTFM environment, as shown in Table 1. CRL outperforms **txt2 π** in terms of both the transfer performance and training efficiency. 1). *Transfer performance*: CRL outperforms **txt2 π** on the base 6×6 environment and all variant environments when transferred from the base 6×6 . The results show that after training on the 6×6 environment, CRL can successfully transfer the learned policy to all variant environments and get up to 46% promotion on win rate. In comparison, **txt2 π** needs four curriculum stages to achieve a similar win rate (details in Appendix A). This is because CRL learned invariant concepts in the base environment that can be reused in other environments quickly. 2). *Training efficiency*: Besides, CRL also gets large training efficiency promotions. CRL saves 70% training steps (100 million to 30 million) on 6×6 and 40% on other transfer environments (50 million to 30 million). The learning curves in Figure 3

Transfer from	Method	Transfer to								Training Steps
		6 × 6	6 × 6 dyna	6 × 6 groups	6 × 6 nl	6 × 6 dyna groups	6 × 6 groups nl	6 × 6 dyna nl	6 × 6 dyna groups nl	
random	txt2π	84 ± 20	26 ± 7	25 ± 3	45 ± 6	23 ± 2	25 ± 3	23 ± 2	23 ± 2	100M
	CRL	93 ± 3	36 ± 10	33 ± 10	38 ± 4	17 ± 2	20 ± 2	17 ± 3	16 ± 3	30M(↓ 70%)
6 × 6	txt2π		85 ± 9	82 ± 19	78 ± 24	64 ± 12	52 ± 13	53 ± 18	40 ± 8	50M
	CRL		89 ± 9	96 ± 3	97 ± 2	85 ± 1	87 ± 2	85 ± 4	86 ± 3	30M(↓ 40%)

Table 1: RTFM results in different settings. All results get from 5 random seeds.

	Method	Train	Test (OOD)	Training Steps
S1	EMMA	98 ± 2.1	85 ± 1.4	30M
	CRL	98 ± 1	88 ± 2.5	30M
S2	EMMA	96 ± 2	45 ± 12	30M
	CRL	98 ± 2	76 ± 5	30M
S3	EMMA	19 ± 2.9	10 ± 0.8	30M
	CRL	43 ± 5	32 ± 1.9	30M

Table 2: Messenger results in different stages. All results get from 5 random seeds.

show the efficiency comparison intuitively. The reason for efficiency promotion is the compactness of the concept. In conclusion, benefiting from the invariance and compactness of the concept, the policy learned through CRL can transfer to similar new environments efficiently.

Zero-Shot Generalization in Messenger

Different from the i.i.d setting of RTFM, there are out-of-distribution (OOD) problems between the training and test environments in Messenger. We compare the training performance and zero-shot generalization performance on unseen test games. 1). *Zero-short generalization performance*: Table 2 shows that CRL outperforms EMMA on the more challenging stages (S2 and S3), including the higher training win rate (up to 20% in S3) and lower generalization gap (up to 30% in S2). The results show that the invariant concepts can mitigate spurious correlations in OOD generalization scenarios. 2). *Training efficiency*: CRL also promotes the training efficiency on Messenger (see Figure 4), but the lift is less than that in RTFM.

Ablation Study

Mutual information (MI) constraints: To show the significant effect of the MI constraints when learning concepts, we verified the performance of the ablated variant CRL w/o MI. CRL w/o MI shares the same architecture with CRL, but without the MI constraints. Figures 3 and 4 show the comparison of learning curves. The results show that the performance of CRL with MI constraints outperforms the previous method apparently, especially in generalization ability. In contrast, the ablation variant degenerates to other baselines that implicitly learn the latent joint representation. The ab-




Figure 3: Ablation results in RTFM




Figure 4: Ablation results in Messenger. The solid and dotted curves represent training and test results.

lation experiments proved the importance of MI constraints when learning compact and invariant concepts (The ablation study of separate MI constraints refers to Appendix C).

Pre-defined concept number: We also investigate the influence of the pre-defined concept number m on CRL. Considering the environment RTFM, whose ground truth m is 2, we test $m = 1$ to $m = 4$, and compare the performance in RTFM-base and transfer performance on the hardest RTFM-final respectively. Figure 5 shows that when the pre-defined m is larger than the ground truth value, the efficiency is slightly influenced, while the asymptotic performance is stable. The reason is that the encoder model has more parameters and becomes harder to train, but the MI constraints can guarantee the learned concepts are compact and invariant. When the pre-defined m is smaller than the ground truth, which makes the encoder cannot exact concept-related information, the performance is affected obviously. The above results inspired the user to set a larger concept number m .




Figure 5: Ablation study of pre-defined concept number (m)

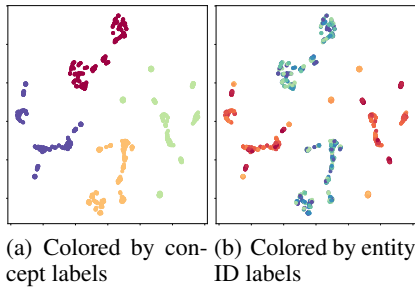


Figure 6: Visual results of **concept vectors** in RTFM

Visual Results of the Concept

To intuitively show concepts learned by CRL, we represent the TSNE (t-Distributed Stochastic Neighbor Embedding) results of learned concept vectors in Figures 6. Specifically, we randomly initialize the environments and save the concept vectors of the entities generated by CRL, then compute the TSNE results of 1000 concept vectors. RTFM has two concepts from the human view, and each has two values: goal/decoy monster and useful/useless tools. Figure 6 shows the TSNE results colored by concept labels (unavailable for the policy) and entity IDs, respectively. Results show that the learned concept vectors separate into 4 clusters corresponding to four concept values. Furthermore, the same concept of different entities aggregates to the same cluster. The phenomenon implies that using entity straightforwardly as the observation for policy training may introduce irrelevant information and cause spurious correlations, while produced concepts by CRL can keep invariant across different entities of various scenarios to address this issue.

Related Works

Language-conditioned Reinforcement Learning: Recently, there have been many works combining reinforcement learning with language and learning language-conditioned policy. The most related field is **reading to act**, for example, RTFM (Zhong et al. 2020) and Messenger (Hanjie et al. 2021). The agent is additionally provided information of the environment dynamics through text. Previous methods often implicitly learn the joint representation of observation and text when optimizing the policy (Hanjie et al. 2021; Narasimhan et al. 2018; Zhong et al. 2020, 2021; Branavan et al. 2011). The implicit joint representation will

inevitably include noisy or irrelevant information and cause spurious correlations. Our method explicitly learns the invariant compact representation to address this issue.

Besides reading to act, there are two kinds of scenarios, instruction following and text game. **Instruction following** means that agents need to solve different tasks specified by high-level instructions in the environment (Hermann et al. 2017; Chaplot et al. 2018; Bahdanau et al. 2019). Different instruction indicates different reward functions in the same environment dynamic, which limits the generalization ability of the agent across different dynamics. Traditional methods require prior knowledge to model the relation between instructions and observations (Andreas et al. 2015). More recently, some works directly embed both the observation and instruction as the joint input of the policy and train the policy under the reinforcement learning paradigm (Mei et al. 2016; Misra et al. 2017; Goyal et al. 2019). **Text game** (Osborne et al. 2021) environments are transformed from Interactive Fiction Games where the agent must interact with the environment through textual observation and action, such as TextWorld (Côté et al. 2018) and Jericho (Hausknecht et al. 2020). The language-conditioned RL methods in the text game focus on exploration (Madotto et al. 2020; Ammanabrolu et al. 2020), action space reduction (Guo et al. 2020; Jain et al. 2020; Yao et al. 2020), and reward shaping (Murugesan et al. 2021; Adolphs et al. 2020).

Invariant Representation Learning Learning invariant representation that can generalize across environments is critical to the application of RL algorithms (Arjovsky et al. 2019; Zhang et al. 2020). Some works leverage inductive biases, for example, object-oriented architecture (Kansky et al. 2017; Yi et al. 2022) and disentangled representation (Higgins et al. 2017; Peng et al. 2022). Some methods want to eliminate misleading information like temporal information (Raileanu et al. 2021; Guo et al. 2021) or background (Wang et al. 2021). Other methods define the metric of invariance for optimizing the representation, like Bisimulation metrics (Zhang et al. 2021) and policy similarity (Agarwal et al. 2021). The above invariant representation Learning methods cannot work for language-condition policy. Our method combines the attention-based reasoning module and mutual information constraints to learn an invariant and compact representation for language-conditioned policy.

Conclusion

The language-conditioned policy is proposed to facilitate policy transfer through learning the joint representation of observation and text that catches the compact and invariant information across various environments. We propose a conceptual reinforcement learning (CRL) framework to learn the joint representation with both advantages of *invariance* and *compactness* for language-conditioned policy. Verified in two challenging environments, RTFM and Messenger, CRL significantly improves the training efficiency and generalization ability to the new environment dynamics. As concept-like representation is valuable for general policy, we hope to broaden the application scope of CRL in the future.

Acknowledgements

This work is partially supported by the National Key Research and Development Program of China (under Grant 2018AAA0103300), the NSF of China (under Grants 61925208, 62002338, 62102399, U19B2019, 61732020), Beijing Academy of Artificial Intelligence (BAAI), CAS Project for Young Scientists in Basic Research (YSBR-029), Youth Innovation Promotion Association CAS and Xplore Prize.

References

- Adolphs, L.; et al. 2020. LeDeepChef Deep Reinforcement Learning Agent for Families of Text-Based Games. In *AAAI*, 7342–7349.
- Agarwal, R.; Machado, M. C.; Castro, P. S.; and Bellemare, M. G. 2021. Contrastive Behavioral Similarity Embeddings for Generalization in Reinforcement Learning. In *ICLR*.
- Ammanabrolu, P.; Tien, E.; Hausknecht, M. J.; and Riedl, M. O. 2020. How to Avoid Being Eaten by a Grue: Structured Exploration Strategies for Textual Worlds. *CoRR*, abs/2006.07409.
- Andreas, J.; et al. 2015. Alignment-Based Compositional Semantics for Instruction Following. In *EMNLP*, 1165–1174.
- Arjovsky, M.; Bottou, L.; Gulrajani, I.; and Lopez-Paz, D. 2019. Invariant Risk Minimization. *CoRR*, abs/1907.02893.
- Bahdanau, D.; Hill, F.; Leike, J.; Hughes, E.; Hosseini, S. A.; Kohli, P.; and Grefenstette, E. 2019. Learning to Understand Goal Specifications by Modelling Reward. In *ICLR (Poster)*.
- Branavan, S. R. K.; et al. 2011. Learning to Win by Reading Manuals in a Monte-Carlo Framework. In *ACL*, 268–277.
- Chaplot, D. S.; Sathyendra, K. M.; Pasumarthi, R. K.; Rajagopal, D.; and Salakhutdinov, R. 2018. Gated-Attention Architectures for Task-Oriented Language Grounding. In *AAAI*, 2819–2826.
- Cheng, P.; Hao, W.; Dai, S.; Liu, J.; Gan, Z.; and Carin, L. 2020. CLUB: A Contrastive Log-ratio Upper Bound of Mutual Information. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, 1779–1788.
- Côté, M.-A.; Kádár, Á.; Yuan, X.; Kybartas, B. A.; Barnes, T.; Fine, E.; Moore, J.; Hausknecht, M. J.; Asri, L. E.; Adada, M.; Tay, W.; and Trischler, A. 2018. TextWorld: A Learning Environment for Text-based Games. In *CGW@IJCAI*.
- Goyal, P.; et al. 2019. Using Natural Language for Reward Shaping in Reinforcement Learning. In *IJCAI*, 2385–2391.
- Guo, J.; Zhang, R.; Zhang, X.; Peng, S.; Yi, Q.; Du, Z.; Hu, X.; Guo, Q.; and Chen, Y. 2021. Hindsight Value Function for Variance Reduction in Stochastic Dynamic Environment. In *IJCAI*, 2476–2482. ijcai.org.
- Guo, X.; Yu, M.; Gao, Y.; Gan, C.; Campbell, M.; and Chang, S. 2020. Interactive Fiction Game Playing as Multi-Paragraph Reading Comprehension with Reinforcement Learning. In *EMNLP (1)*, 7755–7765.
- Hanjie, A. W.; et al. 2021. Grounding Language to Entities and Dynamics for Generalization in Reinforcement Learning. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, 4051–4062.
- Hausknecht, M. J.; Ammanabrolu, P.; Côté, M.; and Yuan, X. 2020. Interactive Fiction Games: A Colossal Adventure. In *AAAI*, 7903–7910.
- Hermann, K. M.; Hill, F.; Green, S.; Wang, F.; Faulkner, R.; Soyer, H.; Szepesvari, D.; Czarnecki, W. M.; Jaderberg, M.; Teplyashin, D.; Wainwright, M.; Apps, C.; Hassabis, D.; and Blunsom, P. 2017. Grounded Language Learning in a Simulated 3D World. *CoRR*, abs/1706.06551.
- Higgins, I.; Pal, A.; Rusu, A. A.; Matthey, L.; Burgess, C. P.; Pritzel, A.; Botvinick, M. M.; Blundell, C.; and Lerchner, A. 2017. DARLA: Improving Zero-Shot Transfer in Reinforcement Learning. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, 1480–1490.
- Jain, V.; Fedus, W.; Larochelle, H.; Precup, D.; and Bellemare, M. G. 2020. Algorithmic Improvements for Deep Reinforcement Learning Applied to Interactive Fiction. In *AAAI*, 4328–4336.
- Kansky, K.; Silver, T.; Mély, D. A.; Eldawy, M.; Lázaro-Gredilla, M.; Lou, X.; Dorfman, N.; Sidor, S.; Phoenix, D. S.; and George, D. 2017. Schema Networks: Zero-shot Transfer with a Generative Causal Model of Intuitive Physics. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, 1809–1818.
- Kirk, R.; Zhang, A.; Grefenstette, E.; and Rocktäschel, T. 2021. A Survey of Generalisation in Deep Reinforcement Learning. *CoRR*, abs/2111.09794.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2016. Continuous control with deep reinforcement learning. In *ICLR (Poster)*.
- Madotto, A.; Namazifar, M.; Huizinga, J.; Molino, P.; Ecoffet, A.; Zheng, H.; Papangelis, A.; Yu, D.; Khatri, C.; and Tür, G. 2020. Exploration Based Language Learning for Text-Based Games. In *IJCAI*, 1488–1494.
- Mei, H.; et al. 2016. Listen, Attend, and Walk: Neural Mapping of Navigational Instructions to Action Sequences. In *AAAI*, 2772–2778.
- Misra, D. K.; et al. 2017. Mapping Instructions and Visual Observations to Actions with Reinforcement Learning. In *EMNLP*, 1004–1015.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. A. 2013. Playing Atari with Deep Reinforcement Learning. *CoRR*, abs/1312.5602.
- Murugesan, K.; Atzeni, M.; Kapanipathi, P.; Shukla, P.; Kumaravel, S.; Tesauro, G.; Talamadupula, K.; Sachan, M.; and Campbell, M. 2021. Text-based RL Agents with Commonsense Knowledge: New Challenges, Environments and Baselines. In *AAAI*, 9018–9027.
- Narasimhan, K.; et al. 2018. Grounding Language for Transfer in Deep Reinforcement Learning. *J. Artif. Intell. Res.*, 63: 849–874.

Oh, J.; Singh, S.; Lee, H.; and Kohli, P. 2017. Zero-Shot Task Generalization with Multi-Task Deep Reinforcement Learning. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, 2661–2670.

Osborne, P.; et al. 2021. A Survey of Text Games for Reinforcement Learning informed by Natural Language. *CoRR*, abs/2109.09478.

Peng, S.; Hu, X.; Zhang, R.; Tang, K.; Guo, J.; Yi, Q.; Chen, R.; Zhang, X.; Du, Z.; Li, L.; Guo, Q.; and Chen, Y. 2022. Causality-driven Hierarchical Structure Discovery for Reinforcement Learning. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.

Raileanu, R.; et al. 2021. Decoupling Value and Policy for Generalization in Reinforcement Learning. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, 8787–8798.

Wang, X.; et al. 2021. Unsupervised Visual Attention and Invariance for Reinforcement Learning. In *CVPR*, 6677–6687.

Yao, S.; Rao, R.; Hausknecht, M. J.; and Narasimhan, K. 2020. Keep CALM and Explore: Language Models for Action Generation in Text-based Games. In *EMNLP (1)*, 8736–8754.

Yi, Q.; Zhang, R.; Peng, S.; Guo, J.; Hu, X.; Du, Z.; Zhang, X.; Guo, Q.; and Chen, Y. 2022. Object-Category Aware Reinforcement Learning. In Oh, A. H.; Agarwal, A.; Belgrave, D.; and Cho, K., eds., *Advances in Neural Information Processing Systems*.

Zhang, A.; Lyle, C.; Sodhani, S.; Filos, A.; Kwiatkowska, M.; Pineau, J.; Gal, Y.; and Precup, D. 2020. Invariant Causal Prediction for Block MDPs. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, 11214–11224.

Zhang, A.; McAllister, R. T.; Calandra, R.; Gal, Y.; and Levine, S. 2021. Learning Invariant Representations for Reinforcement Learning without Reconstruction. In *ICLR*.

Zhong, V.; Hanjie, A. W.; Wang, S.; Narasimhan, K.; and Zettlemoyer, L. 2021. SILG: The Multi-environment Symbolic Interactive Language Grounding Benchmark. *ArXiv*, abs/2110.10661.

Zhong, V.; et al. 2020. RTFM: Generalising to Novel Environment Dynamics via Reading. *ArXiv*, abs/1910.08210.