Disentangling Transfer in Continual Reinforcement Learning

Anonymous Author(s)
Affiliation
Address
email

Abstract

The ability of continual learning systems to transfer knowledge from previously seen tasks in order to maximize forward transfer is a significant challenge for the field, limiting the applicability of continual learning solutions to realistic scenarios. Consequently, this study aims to broaden our understanding of transfer and its driving forces in the specific case of continual reinforcement learning. We adopt SAC as the underlying RL algorithm and Continual World as a suite of continuous control tasks. We systematically study how different components of SAC (the actor and the critic, exploration, and data) affect transfer efficacy, and we provide recommendations regarding various modeling options. The best set of choices, dubbed `ClonEx-SAC`, is evaluated on the recent Continual World benchmark. `ClonEx-SAC` achieves 87% final success rate compared to 80% of PackNet, the best method in the benchmark. Moreover, the transfer grows from 0.18 to 0.54 in the metric provided by Continual World.

1 Introduction

The ability of continual learning (CL) systems ([16, 21]) to utilize knowledge from previously seen tasks in order to maximize transfer on the current task is a significant challenge for the field. Achieving progress in this area would bring benefits both for real-life applications and multiple machine learning domains [23, 17, 44, 9], including reinforcement learning (RL), as advocated in [45]. In particular, it would constitute a critical step towards making efficient lifelong learning agents a reality.

The goal of this paper is to expand our understanding of transfer and its driving factors in continual reinforcement learning (CRL). As the underlying RL algorithm we assume soft actor-critic (SAC), see [15], and use Continual World [45] as the suite of continuous control environments. We systematically study the critic and actor networks, the key components of SAC, with regard to their influence on transfer. Similarly, we measure the impact of various choices for exploration and using buffer data. The low-level mechanisms of transfer are not yet fully understood even in the supervised learning case [27]. To the best of our knowledge, our work is the first one that undertakes a comprehensive study of this important topic in RL. To this end, we proceed in two stages: exploring a two-task setting and a full continual learning scenario.

We start by exploring a simplified two-tasks setting in Section 4. This allows us to leave out the impact of forgetting, as well as limit the choices regarding exploration and data handling. We use 100 pairs of robotic tasks from the Continual World benchmark. Here, we render two key observations: 1) the role of the critic is the most important for transfer, while exploration and actor play smaller,

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

but non-negligible, parts; 2) contributions of the individual components are mostly independent. Additionally, we show that the concept of feature reuse which is often used to explain supervised transfer learning [27, 32] might not be directly applicable in RL.

In Section 5, we aim to understand new effects which emerge for the full continual learning scenario, typically in longer sequences. In CL context, we need to take into account forgetting, being mindful of the fact that existing CL methods often favor mitigating forgetting at the expense of transfer, see [45]. Main results include: 1) reusing policies from previous tasks for exploration considerably improves performance; 2) behavioral cloning to rehearse past tasks is beneficial for both average performance and forward transfer, outperforming other considered methods; 3) regularizing the critic typically does not help for the performance of CL methods.

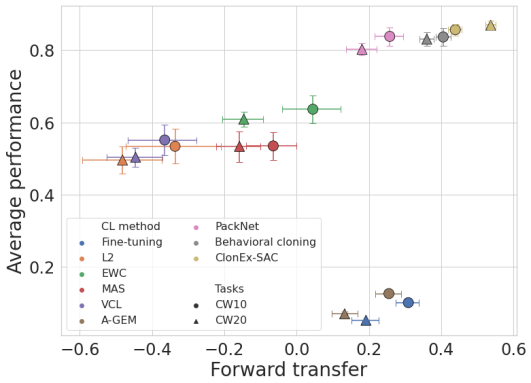

Figure 1: Performance of the `ClonEx-SAC` method compared with competitive baselines, on CW10 and CW20 task sequences. Average performance and forward transfer are shown, together with 90% bootstrap confidence intervals.

The result of our comprehensive analysis is a set of general recommendations. We also determine the combination which outperforms all other options, dubbed `ClonEx-SAC`. This method utilizes behavioral *cloning* to mitigate catastrophic forgetting for the actor. Moreover, at the beginning of each task, `ClonEx-SAC` queries all previous policies, the best of which generates initial *exploration* data. `ClonEx-SAC` achieves 87% final success rate compared to 80% of PackNet, the best method in the Continual World benchmark, see Figure 1. Importantly, we observe a sharp transfer increase from 0.18 to 0.54 in the metric provided in the benchmark. Notably, the value of forward transfer closely matches the reference forward transfer adjusted for exploration, which is a soft upper bound for transfer, as introduced in [45].

2 Related work

Continual learning algorithms are often categorized into three classes: regularization-based e.g. [1, 22, 30], parameter isolation e.g. [25] and rehearsal methods e.g. [5, 6]; see also CL survey papers [7, 16, 31]. [42, 3] advocate the need to develop CL methods suitable for reinforcement learning training as a necessary step towards learning artificial intelligent agents to operate in open-ended and changing environments. [21] provides a detailed review of this combination and a taxonomy of possible setups. [45] proposes a sequence of robotic tasks as a benchmark, comparing popular CL methods adapted to RL and advocating for putting more emphasis on transfer. The authors of [19] show how the synaptic Benna-Fusi model can be added on top of value-based RL methods to mitigate forgetting at both intra- and inter-task scales. A simple approach to cloning policies from previous tasks is employed in [43], and a similar replay strategy has been used in [34]. [20] tackles the case when task boundaries are not provided. Although most of the research is concerned with model-free continual reinforcement learning, an approach to model-based continual RL was presented in [18].

Transfer learning, which focuses on the reuse of machine learning models, has been extremely successful recently. In computer vision, convolutional neural networks [23, 17] and vision transformers [10] pre-trained on large datasets can be repurposed and fine-tuned on the target task. Modern transformer-based models [44, 9] train on large natural language corpora turned out to be very flexible and can be adapted to diverse downstream tasks, sometimes quite surprisingly [33, 24]. General surveys of transfer learning techniques are provided in [50, 39]. Interestingly, recent research [28, 48] suggests that there are still some gaps in our understanding of transfer learning. [28] analyses the low-level reasons for transfer, exhibiting surprising phenomena such as transfer between datasets

with permuted images. [48] performs large-scale experiments investigating representation transfer in a wide variety of visual tasks.

In reinforcement learning scenarios, the structure of the underlying MDP can be exploited to facilitate the transfer. [40, 49, 41, 38] present methods on how to find and use mappings between different domains. [29, 4] apply reward function reshaping. [26, 37] achieve transfer by means of high-level skills and hierarchical RL. Other lines of work exploits the model structure [35, 12] or enforce modularity [2, 8]. In this work, we aim to complement these studies, by focusing on the benefits of reusing neural network parameters, and other choices that exploit the RL structure, like exploration and data reuse.

3 Background

3.1 Continual learning and reinforcement learning

Continual learning tackles the problem of learning in non-stationary settings [7]. Typically, the solution is expected to perform well on all encountered tasks, although various metrics expressing different requirements are formulated. The popular CL desiderata include reducing the *forgetting* on previous tasks and increasing the *forward transfer*, i.e. speeding up the learning by reusing knowledge from previous tasks [11, 45]. Other desiderata focus on limiting resources, such as the number of samples, computation time, model size, or additional memory size. These requirements are often conflicting, so usually some trade-offs have to be made [16, 45, 31].

Combining CL with RL adds another layer of complexity. In this work, we focus on the SAC algorithm [15], which is often considered to be the method of choice for continuous control RL [47, 46, 36]. As an actor-critic algorithm, it is based on the interplay between its two parts, see Section 3.2. This is a fairly complicated algorithmic setup, which presents a number of challenges when using jointly with CL. In particular, since optimization of the critic and actor networks is intertwined, it is hard to understand and decouple the impact of individual components. Additionally, because of this interplay, training biases get easily exacerbated, often leading to inferior performance or even a collapse. Another complication is that the training objectives for the actor and the critic are different. The critic minimizes the Bellman error which is known to be a fragile objective [13] susceptible to training biases, and might correlate poorly with the value error (which we would like to minimize). As the actor optimizes over predictions of the critic, it might also suffer from these problems, even if less directly. Finally, since the policy and the data we see change during the training, there is an inherent distribution shift present, even within a single task.

3.2 SAC

In our study, we focus on soft actor-critic (SAC) [14], an off-policy actor-critic RL algorithm, based on the maximum entropy principle. **The critic** strives to approximate the entropy-corrected Q-function under the current policy, optimizing the Bellman error. **The actor** tries to find actions that maximize the Q-function. **The replay buffer** holds the seen experience and provides data for the actor and critic updates at each learning step. **The exploration policy** is used to gather data at the beginning of each task for a set number of K steps. By default, in most SAC implementations, this means sampling actions uniformly from the action space.

3.3 Continual World

We perform our experiments on the Continual World [45] benchmark. It contains a set of realistic robotic tasks, where a simulated Sawyer robot manipulates everyday objects. The structure of the observation and action spaces remains the same between the tasks; an observation is a 12-dimensional vector describing the coordinates of the robot's gripper and relevant objects. The 4-dimensional action space describes the gripper movement. In training, a dense reward function is used to make the tasks solvable; in evaluation, the binary success metric is used to indicate whether the desired goal has been reached. The tasks are arranged in sequences, training in each task lasts for 1M steps. CW10

sequence contains 10 different tasks arranged in a fixed order. CW20 consists of CW10 repeated twice, allowing to measure how much knowledge can be transferred in case of task repetitions. We use both CW10 and CW20 in our evaluations, as well as shorter sequences containing pairs of tasks from CW10.

3.4 Metrics

Following standard practice in continual learning literature, we report average performance and forgetting metrics. We also measure transfer as defined in [45]. Below we briefly recall these three metrics. Assume $p_i(t) \in [0, 1]$ to be the performance (success rate) of task i at time t, and that each of the N tasks is trained for Δ steps, so the total number of steps is $T = N \cdot \Delta$.

Average performance The average performance at time t is defined as $\mathrm{P}(t) := \frac{1}{N}\sum_{i=1}^{N} p_i(t)$. Its final value, $P(T)$, is a scalar summary of the performance and is presented in the result tables.

Forward transfer The forward transfer is computed as a normalized area between the training curve of the measured run and the training curve of a reference curve from training from scratch. Let us denote by $p_i^b \in [0, 1]$ the reference performance. Then the forward transfer on task i, FT_i, is defined as

$$\mathrm{FT}_i := \frac{\mathrm{AUC}_i - \mathrm{AUC}_i^b}{1 - \mathrm{AUC}_i^b}, \quad \mathrm{AUC}_i := \frac{1}{\Delta}\int_{(i-1)\cdot\Delta}^{i\cdot\Delta} p_i(t)\mathrm{d}t, \quad \mathrm{AUC}_i^b := \frac{1}{\Delta}\int_{0}^{\Delta} p_i^b(t)\mathrm{d}t.$$

The average forward transfer for all tasks, FT, is defined as $\mathrm{FT} = \frac{1}{N}\sum_{i=1}^{N} \mathrm{FT}_i$.

Forgetting For the task i, one can measure drop in performance after the end of learning on this task as $F_i = p_i(i \cdot \Delta) - p_i(T)$. Forgetting metric is defined as $F = \frac{1}{N}\sum_{i=1}^{N} F_i$.

3.5 Experimental setup

We follow the experimental setup from [45]. The actor and the critic are implemented as two separate MLP networks, each with 4 hidden layers of 256 neurons. We refer to the 4 hidden layers as the *backbone* and the last output layer as the *head*. By default, we assume the *multi-head* (MH) setting, where each task has its separate output head, but we also consider the *single-head* (SH) setting, where only a single head is used for all tasks. The SAC exploration phase takes $K = 10$k steps. All experiments in this paper were performed with 10 different seeds unless noted otherwise. We compute 90% confidence intervals through bootstrapping. More details on the experimental setup can be found in Appendix A.

4 Transfer in isolation

In this section, we study *what enables transfer between RL tasks*. We assume a two-task setting, where we measure the forward transfer from the first to the second task, disregarding continual learning specific issues (e.g. forgetting), which we defer to the next section. We utilize all 100 pairs of CW10 tasks, see Section 3.3, to evaluate the impact of *critic, actor, and exploration* given by SAC.

We will say that the actor or the critic are *carried over* (from the previous tasks) if their parameters are reused as the initialization in the next task; otherwise, the parameters are re-initialized. We also refer to the exploration policy as being carried over, if we use the policy from the previous task (or tasks) to gather the data during the first 10k steps of the SAC exploration phase (see Section 3.2); otherwise, a uniform random exploration policy is being used. We use both multi-head (MH) and single-head (SH) settings, with the former being default.

Figure 2 illustrates the impact of the individual components on transfer for each pair. The (i, j)-th entry in the matrix contains the forward transfer value when carrying over components from task i to task j. Table 1 presents the aggregated statistics from the matrices given in Figure 2: the average FT

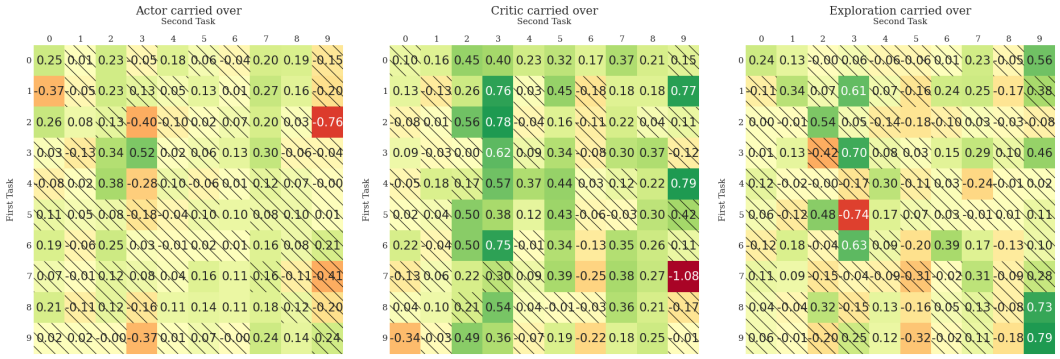

(a) Actor carried over (b) Critic carried over (c) Exploration carried over

Figure 2: The effect of carrying over different components on the performance on pairs of tasks from CW10. We shade an entry if the 90% confidence interval contains 0, indicating that we cannot be sure whether the component which was carried over makes a difference.

(including and excluding the diagonal), and the number of pairs with positive, negative and neutral FT[1]. Table 2 reports the transfer properties for all possible combinations of components present in Table 1, omitting single-head critic (since it performs worse in Table 1).

4.1 Carrying over SAC components

This section has two key observations. First, the role of the critic is the most important for FT, while exploration and actor play smaller, but non-negligible, parts. Second, the components are "transfer independent", in the sense that the transfer of the combination of the components is close to the sum of transfers yielded by each component alone.

The evidence for the first finding is presented in detail for each pair in Figure 2 and summarized in Table 1. More precisely, the average forward transfer across all pairs attributed to carrying over of the critic equals 0.2 (resp. 0.15) for MH (resp. SH) setup. This separates the critic from the actor and exploration, which yield (for the default MH setup) 0.06 and 0.09, respectively.

The importance of the critic is further emphasized by showing that restraining its learning capabilities, even when the weights are initialized to the parameters learned in the previous task, negatively impacts FT. This is shown in the last row of Table 1, where only the critic's head is allowed to train, while the body of the network is kept frozen and carried over from the previous task. On the one hand, this result could be surprising, since feature reuse is a common technique in supervised learning (e.g. in vision [27, 32]). On the other hand, the deterioration in FT can be explained by RL-specific factors. Namely, this procedure can hinder both the policy training (since the mechanics of SAC intertwines actor and critic) and the critic training (due to inflated Bellman errors).

As to the second finding, i.e. the "transfer independence" of the components, the results of the underlying analysis can be found in Table 2. We observe that the reported FT for the combination of components follows closely the sum of FTs for individual components (reported in Table 1). Furthermore, we observe that including all the components results in the highest transfer of 0.35.

There is a couple of remaining interesting observations. First, Figure 2 exhibits several vertical patterns, meaning that transfer depends more on the second task. Second, the effect on transfer increases on the diagonal, when the exploration is carried over. This seems reasonable since the policy in the new task is initialized to the already learned policy on the same task. Finally, resetting the head (MH setup) is beneficial in the case of the critic, while it hurts the actor.

[1] We say that a pair has positive (resp. negative) FT if the corresponding confidence interval is above (resp. below) 0. Otherwise, we mark it as neutral.

Table 1: Summary of the transfer statistics from the transfer matrices when transferring only a single component. FT and FT (no diag) represent average forward transfer across all pairs with and without considering the diagonal (transfer from a task to the same task), respectively. Subsequent columns denote the number of pairs with positive, negative and neutral transfer.

name	FT	FT (no diag)	# pos.	# neg.	# neutral
Actor (MH)	0.06 [0.03, 0.10]	0.05 [0.01, 0.09]	30	5	65
Critic (MH)	0.20 [0.17, 0.23]	0.19 [0.16, 0.23]	54	5	41
Exploration	0.09 [0.06, 0.13]	0.06 [0.03, 0.10]	28	9	63
Actor (SH)	0.12 [0.09, 0.15]	0.12 [0.09, 0.15]	37	1	62
Critic (SH)	0.15 [0.12, 0.18]	0.13 [0.10, 0.16]	41	19	40
Critic (train only head)	-1.29 [-1.33, -1.25]	-1.30 [-1.35, -1.26]	0	100	0

Table 2: Summary of transfer statistics when multiple components are carried over. We observe that impact of each component is largely independent from other components. That is, FT when carrying over multiple components is close to then sum of FT when carrying over each of them separately.

name	FT	FT (no diag)	# pos.	# neg.	# neutral
Actor (MH) + Critic (MH)	0.27 [0.23, 0.31]	0.25 [0.21, 0.29]	58	4	38
Actor (SH) + Critic (MH)	0.29 [0.25, 0.32]	0.28 [0.24, 0.32]	59	2	39
Actor (MH) + Exp.	0.16 [0.12, 0.20]	0.14 [0.10, 0.18]	39	3	58
Actor (SH) + Exp.	0.21 [0.17, 0.25]	0.18 [0.14, 0.22]	53	0	47
Critic (MH) + Exp.	0.30 [0.27, 0.34]	0.28 [0.24, 0.32]	64	2	34
Actor (SH) + Critic (MH) + Exp.	0.36 [0.32, 0.39]	0.33 [0.29, 0.36]	68	0	32
Actor (MH) + Critic (MH) + Exp.	0.35 [0.31, 0.38]	0.32 [0.29, 0.36]	70	1	29

5 Transfer in continual learning

In Section 4.1, we focused on direct transfer in the two-task setting. Now, we move to the full continual learning scenario, which brings two main differences: 1) we measure the performance of all tasks in the sequence, so forgetting now plays a significant role; 2) typically, we consider longer sequences of tasks of length 10 and 20 (CW10 and CW20, respectively). For longer sequences, forgetting and transfer may have complex mutual interactions. To reduce forgetting, CL methods usually regularize neural network weights, which in turn may be harmful to transfer. On the other hand, transfer benefits from accumulated knowledge – if forgetting is not mitigated, there might be nothing to transfer from.

We will investigate three main themes. The first one is reusing previous policies for exploration. For long sequences, there are multiple design choices available compared to the two-task scenarios. Secondly, we investigate CL with data reuse, an approach successful in supervised learning. We show that the CRL setup is more complex and requires careful investigation. Finally, given the importance of the critic for transfer (see Section 4), we study whether the critic should be regularized or not, and conclude that typically, the answer is negative.

We study these issues in conjunction with various CL methods: Fine-tuning, Perfect memory, EWC, PackNet, L2, A-GEM, MAS, and VCL. These are standard CL approaches adopted and tested in the RL setting [45], see details in Appendix B. We note that CL methods used here are mostly successful in mitigating forgetting; in this section, we report average performance and forward transfer, deferring forgetting to the Appendices C and E.

5.1 Exploration

When using the SAC algorithm, at the beginning of each task, there is a short period of exploration with a random policy, see Section 3.2. The experiments in Section 4.1 showed that the transfer increases if the policy from the previous task is used instead.

Now we pass from two-task scenarios to longer ones, and analyze the following options for choosing exploration policy, which we call: *random, preceding, uniform-previous*, and *best-return*, and define them as follows. In the first task, we always use a random policy, and assume that the tasks are numbered from 1 to N. Consider now $i \in \{2, \dots, N\}$. For the *random* strategy, we randomly sample from the action space, which is a default choice for SAC. For other strategies, at the beginning of each exploration episode, we choose a previous actor head to generate data instead of the random policy. In the case of *preceding* strategy, we use the $(i-1)$-th actor's head. For *uniform-previous* policy, we use the j-th actor's head, where $j := \text{RANDOM_UNIFORM}(\{1, \dots, i-1\})$. Finally, in *best-return* strategy, we first try every possible head, and then act using the $j_{\max}$-th actor's head, where $j_{\max} := \text{argmax}_{j \in \{1,\dots,i-1\}} R^i_j$; R^i_j is the return of the j-th head policy on the i-th task.

We evaluate how these strategies interact with various CL methods. We pick Fine-tuning, Behavioral cloning, L2, EWC, and PackNet. The results for two well-performing methods, EWC and PackNet, are presented in Table 3, with the rest being deferred to Appendix E. For EWC, choosing any non-random policy significantly improves upon the baseline random strategy. This is particularly visible in the CW20 sequence, which contains repeated tasks, and arguably can benefit more from an informed strategy like best-return. Interestingly, the results for the rather simple uniform-previous approach are quite competitive. We observe increased performance also for other methods except for PackNet, for which effects are negligible.

Table 3: Average performance and forward transfer for different exploration strategies on CW10 and CW20 sequences. Strategies are added on top of EWC and PackNet methods.

Method, exploration	CW10 perf.	CW10 f. transfer	CW20 perf.	CW20 f. transfer
EWC, random	0.63 [0.60, 0.66]	0.03 [-0.04, 0.09]	0.60 [0.59, 0.62]	-0.14 [-0.19, -0.09]
EWC, preceding	**0.70** [0.67, 0.73]	0.09 [0.03, 0.15]	0.61 [0.59, 0.64]	-0.14 [-0.19, -0.09]
EWC, uniform-previous	**0.72** [0.69, 0.75]	**0.24** [0.19, 0.28]	**0.70** [0.68, 0.73]	0.21 [0.17, 0.25]
EWC, best-return	**0.70** [0.68, 0.73]	**0.25** [0.21, 0.28]	**0.71** [0.69, 0.73]	**0.28** [0.25, 0.31]
PackNet, random	0.84 [0.81, 0.86]	0.26 [0.22, 0.29]	0.80 [0.79, 0.82]	0.18 [0.14, 0.22]
PackNet, preceding	0.84 [0.82, 0.85]	0.24 [0.20, 0.27]	0.81 [0.80, 0.83]	0.20 [0.16, 0.24]
PackNet, uniform-previous	0.84 [0.81, 0.86]	0.21 [0.15, 0.26]	0.80 [0.78, 0.82]	0.23 [0.18, 0.27]
PackNet, best-return	0.85 [0.83, 0.86]	0.23 [0.20, 0.26]	0.82 [0.81, 0.83]	0.23 [0.21, 0.25]

5.2 Data rehearsal

Rehearsal techniques work very well in supervised continual learning [6]. In RL, two major approaches to utilizing previous data have been studied: applying them as offline data using SAC loss, and behavioral cloning of the previous policies. The former, dubbed Perfect memory, was reported to perform poorly [45]. Behavioral cloning achieves more promising results [43, 34]. We study these two approaches with an emphasis on the effects on transfer.

In Perfect memory, all the experiences are kept in the buffer. SAC training is applied to data from the current task and offline data from the previous ones. For Behavioral cloning, an additional small buffer is filled at the end of training on each task. We annotate a subset of samples from the main SAC buffer using the trained actor and critic networks. When training the new task, we sample data from expert buffers and apply auxiliary losses (with different weights), minimizing the KL divergence between current and saved outputs for the actor and L2 distance for the critic; see details in Appendix B.

Firstly, we study the effect of rehearsal on transfer in the two-task scenario, using 100 task pairs from CW10, as in Section 4.1. We observe that using Perfect memory has a detrimental effect on transfer, and cloning the actor only has a rather neutral effect; we report results for these and more setups in Appendix E.

Secondly, we perform experiments on longer sequences, CW10 and CW20; see Table 4. For reference, we include two methods tested in [45], Fine-tuning and PackNet. Fine-tuning achieves the highest

Table 4: Average performance and forward transfer for Perfect memory and Behavioral cloning methods, as described in Section 5.2. Fine-tuning and PackNet are included for reference.

method	CW10 perf.	CW10 f. transfer	CW20 perf.	CW20 f. transfer
Perfect memory	0.27 [0.24, 0.30]	-1.13 [-1.23, -1.04]	0.09 [0.06, 0.12]	-1.32 [-1.41, -1.24]
Behavioral cloning	**0.84** [0.81, 0.86]	**0.41** [0.38, 0.43]	**0.83** [0.81, 0.85]	**0.36** [0.34, 0.38]
Fine-tuning	0.10 [0.10, 0.10]	0.31 [0.27, 0.34]	0.05 [0.05, 0.05]	0.19 [0.15, 0.23]
PackNet	**0.84** [0.81, 0.86]	0.26 [0.22, 0.29]	0.80 [0.79, 0.82]	0.18 [0.14, 0.22]

transfer and PackNet the highest overall performance out of the methods tested in [45]. Behavioral cloning performs very well. In terms of the average performance, it is on par with PackNet on CW10 and better on CW20. Importantly, it significantly outperforms the baselines in terms of transfer. We can see that Perfect memory works poorly, in line with existing literature. In Appendix C, we present the results for five other CL methods benchmarked in [45].

We observe an interesting phenomenon. While transfer was not improved in two-task scenario, it has increased in the case of longer sequences. This hints that perhaps the learner accumulates knowledge of the previous tasks and, thus, can reuse the most relevant one to improve the training of the current task. Additionally, perhaps behavioral cloning loss acts as a regularizer and helps shaping more general features, thus further improving transfer.

5.3 Regularizing the critic

This section is devoted to the study of critic regularization in CRL methods. Since the primary objective of CRL is the final performance of the actor, we have some flexibility in how we treat the critic. We can even completely ignore forgetting in the critic, as recommended in [45]. Other works suggest that regularization might be beneficial [20].

To understand this better, we carefully measure the performance for various critic regularization weights for EWC, L2, and Behavioral cloning. We first find a good value for the actor weight, with the one corresponding to the critic being set to 0. Then, with this value, we perform the sweep over the critic weight. For Behavioral cloning, we consider a range of values from 0.0001 to 100 (the latter being the coefficient for the actor). First, we perform experiments for the pairs of tasks from CW10. Small weight values exhibit similar performance to non-regularized critics i.e. with the weight 0. For larger, we see a performance drop. Then, we study a longer CW10 sequence, for which there is a sharp drop, even for the smallest weights: 0.82 (critic coef. = 0.0001), 0.77 (critic coef. = 0.001); compared to 0.84 for the non-regularized critic. The complete results are presented in Appendix E.

This confirms the practical recommendation from [45] to regularize only the actor. One possible explanation is that TD-learning used for the critic is highly sensitive to biases introduced by regularization.

6 Combining the improvements: `ClonEx-SAC`

Based on the experimental findings presented so far, we propose to combine the discovered enhancements in a simple method for continual reinforcement learning. This method significantly improves the performance in the Continual World benchmark [45]. In particular, we observe a sharp transfer increase to a value that matches a soft upper bound for transfer introduced in [45].

We incorporate the following choices in the method:

- We use behavioral cloning for the actor, which, as we showed in Section 5.2, effectively mitigates forgetting and increases transfer.
- We use best-return exploration, as described in Section 5.1, which efficiently reuses old policy heads for faster exploration.

- As indicated in Section 5.3, we do not use any CL regularization for the critic.
- We use multiple output heads for both actor and critic to profit from transferred representations without introducing too much bias in the new tasks, as discussed in Section 4.1.

We dub the method `ClonEx-SAC` to reflect the usage of the behavioral **clon**ing, improved **ex**ploration, and SAC algorithm.

We compare `ClonEx-SAC` with the behavioral cloning and 7 methods considered in [45], on the CW10 and CW20 sequences. We present results in Figure 1 (see Introduction) and Appendix C. `ClonEx-SAC` achieves 87% final performance compared to 80% of PackNet, the best method in [45]. The forward transfer of `ClonEx-SAC`, improves sharply from 0.19, the best previous result, to 0.54. Notably, `ClonEx-SAC`'s result closely matches the reference forward transfer, see below. We conjecture that this excellent transfer is an important factor in the final performance. We also notice that improvements brought separately by the behavioral cloning and the best-return exploration strategy work well together.

Reference forward transfer (RT) was introduced in [45] as a soft upper bound for transfer. For a sequence of tasks $t_1, \ldots, t_N$, it is defined defined as $\text{RT} := \frac{1}{N}\sum_{i=2}^{N} \max_{j<i} \text{FT}(t_j, t_i)$, where $F(t_j, t_i)$ denotes the forward transfer for the pair of tasks t_j, t_i. Intuitively, RT estimates the level of forward transfer, which could be achieved when a method is able to remember and transfer all meaningful aspects of previously seen tasks. Note that in principle, higher values of RT could still be achievable if the knowledge from the previous tasks is composed. In our setup, the values of RT are 0.44 for CW10 and 0.55 for CW20. In both cases, they are closely matched by the forward transfer of `ClonEx-SAC`. We note that our RT values are higher than the one reported in [45], since their work does not take into account the effects of improved exploration.

7 Limitations

We are fully aware that our analyses do not cover the entire spectrum of problems that one might be interested in when studying transfer in CRL. Here we summarize a few limitations of our work:

- We build on top of the SAC algorithm. There is a risk that some of the conclusions from this paper would differ for another choice of the underlying RL method.
- We focus on the Continual World suite. There is a possibility that some of the results from this paper would differ in environments from other domains or with different, potentially structured state spaces.
- `ClonEx-SAC` requires retaining data from previous tasks, which may not always be feasible (e.g., due to privacy concerns).

8 Conclusions

In this work, we identify and study some of the key factors contributing to transfer in continual reinforcement learning. In the first part of the study, we focus on the transfer alone, disregarding other CL desiderata, and analyze how much different components of the SAC algorithm (actor, critic, exploration) contribute to it. We identify the critic as the leading component. In the second part, we study further effects that are relevant to the full continual learning setup with long task sequences. In particular, we show that behavioral cloning and reusing previous policies for exploration significantly improve both transfer and the final performance. This leads to a new method, `ClonEx-SAC`, which outperforms considered baselines. We believe that this work constitutes the first step toward understanding the mechanisms behind transfer in continual reinforcement learning. There are still important issues to be resolved, e.g., pinpointing the exact role of feature reuse or the interplay between transfer and forgetting. We hope that these will be addressed by the community in the future.

References

[1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part III*, volume 11207 of *Lecture Notes in Computer Science*, pages 144–161. Springer, 2018.

[2] Jacob Andreas, Dan Klein, and Sergey Levine. Modular multitask reinforcement learning with policy sketches. In *International Conference on Machine Learning*, pages 166–175. PMLR, 2017.

[3] Bowen Baker, Ingmar Kanitscheider, Todor M. Markov, Yi Wu, Glenn Powell, Bob McGrew, and Igor Mordatch. Emergent tool use from multi-agent autocurricula. *CoRR*, abs/1909.07528, 2019.

[4] Tim Brys, Anna Harutyunyan, Matthew E Taylor, and Ann Nowé. Policy transfer using reward shaping. In *AAMAS*, pages 181–188, 2015.

[5] Arslan Chaudhry, Marc'Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with A-GEM. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.

[6] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet Kumar Dokania, Philip H. S. Torr, and Marc'Aurelio Ranzato. Continual learning with tiny episodic memories. *CoRR*, abs/1902.10486, 2019.

[7] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *arXiv preprint arXiv:1909.08383*, 2019.

[8] Coline Devin, Abhishek Gupta, Trevor Darrell, Pieter Abbeel, and Sergey Levine. Learning modular neural network policies for multi-task and multi-robot transfer. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 2169–2176. IEEE, 2017.

[9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina N. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. 2018.

[10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

[11] Sebastian Farquhar and Yarin Gal. Towards robust evaluations of continual learning. *CoRR*, abs/1805.09733, 2018.

[12] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017.

[13] Scott Fujimoto, David Meger, Doina Precup, Ofir Nachum, and Shixiang Shane Gu. Why should i trust you, bellman? the bellman error is a poor replacement for value error, 2022.

[14] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861–1870. PMLR, 2018.

[15] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Soft actor-critic algorithms and applications. *CoRR*, abs/1812.05905, 2018.

[16] Raia Hadsell, Dushyant Rao, Andrei A. Rusu, and Razvan Pascanu. Embracing change: Continual learning in deep neural networks. *Trends in Cognitive Sciences*, 24(12):1028 – 1040, 2020.

[17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[18] Yizhou Huang, Kevin Xie, Homanga Bharadhwaj, and Florian Shkurti. Continual model-based reinforcement learning with hypernetworks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 799–805. IEEE, 2021.

[19] Christos Kaplanis, Murray Shanahan, and Claudia Clopath. Continual reinforcement learning with complex synapses. In *International Conference on Machine Learning*, pages 2497–2506. PMLR, 2018.

[20] Samuel Kessler, Jack Parker-Holder, Philip J. Ball, Stefan Zohren, and Stephen J. Roberts. Same state, different task: Continual reinforcement learning without interference. *CoRR*, abs/2106.02940, 2021.

[21] Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. Towards continual reinforcement learning: A review and perspectives, 2020.

[22] James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *CoRR*, abs/1612.00796, 2016.

[23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In Peter L. Bartlett, Fernando C. N. Pereira, Christopher J. C. Burges, Léon Bottou, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States*, pages 1106–1114, 2012.

[24] Kevin Lu, Aditya Grover, Pieter Abbeel, and Igor Mordatch. Pretrained transformers as universal computation engines. *arXiv preprint arXiv:2103.05247*, 2021.

[25] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 7765–7773. IEEE Computer Society, 2018.

[26] Neville Mehta, Sriraam Natarajan, Prasad Tadepalli, and Alan Fern. Transfer in variable-reward hierarchical reinforcement learning. *Machine Learning*, 73(3):289–312, 2008.

[27] Behnam Neyshabur, Hanie Sedghi, and Chiyuan Zhang. What is being transferred in transfer learning? *Advances in neural information processing systems*, 33:512–523, 2020.

[28] Behnam Neyshabur, Hanie Sedghi, and Chiyuan Zhang. What is being transferred in transfer learning? In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 512–523. Curran Associates, Inc., 2020.

[29] Andrew Y Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *Icml*, volume 99, pages 278–287, 1999.

[30] Cuong V. Nguyen, Yingzhen Li, Thang D. Bui, and Richard E. Turner. Variational continual learning. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018.

[31] German Ignacio Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019.

[32] Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse? towards understanding the effectiveness of maml. *arXiv preprint arXiv:1909.09157*, 2019.

[33] Machel Reid, Yutaro Yamada, and Shixiang Shane Gu. Can wikipedia help offline reinforcement learning? *CoRR*, abs/2201.12122, 2022.

[34] David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy Lillicrap, and Gregory Wayne. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019.

[35] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.

[36] Kun Shao, Zhentao Tang, Yuanheng Zhu, Nannan Li, and Dongbin Zhao. A survey of deep reinforcement learning in video games. *arXiv preprint arXiv:1912.10944*, 2019.

[37] Lorenzo Steccanella, Simone Totaro, Damien Allonsius, and Anders Jonsson. Hierarchical reinforcement learning for efficient exploration and transfer. *arXiv preprint arXiv:2011.06335*, 2020.

[38] Erik Talvitie and Satinder P Singh. An experts algorithm for transfer learning. In *IJCAI*, pages 1065–1070, 2007.

[39] Chuanqi Tan, Fuchun Sun, Tao Kong, Wenchang Zhang, Chao Yang, and Chunfang Liu. A survey on deep transfer learning. In Věra Kůrková, Yannis Manolopoulos, Barbara Hammer, Lazaros Iliadis, and Ilias Maglogiannis, editors, *Artificial Neural Networks and Machine Learning – ICANN 2018*, pages 270–279, Cham, 2018. Springer International Publishing.

[40] Matthew E Taylor and Peter Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(7), 2009.

[41] Matthew E Taylor, Peter Stone, and Yaxin Liu. Transfer learning via inter-task mappings for temporal difference learning. *Journal of Machine Learning Research*, 8(9), 2007.

[42] Open Ended Learning Team, Adam Stooke, Anuj Mahajan, Catarina Barros, Charlie Deck, Jakob Bauer, Jakub Sygnowski, Maja Trebacz, Max Jaderberg, Michaël Mathieu, Nat McAleese, Nathalie Bradley-Schmieg, Nathaniel Wong, Nicolas Porcel, Roberta Raileanu, Steph Hughes-Fitt, Valentin Dalibard, and Wojciech Marian Czarnecki. Open-ended learning leads to generally capable agents. *CoRR*, abs/2107.12808, 2021.

[43] René Traoré, Hugo Caselles-Dupré, Timothée Lesort, Te Sun, Guanghang Cai, Natalia Díaz Rodríguez, and David Filliat. Discorl: Continual reinforcement learning via policy distillation. *CoRR*, abs/1907.05855, 2019.

[44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017.

[45] Maciej Wołczyk, Michał Zając, Razvan Pascanu, Lukasz Kucinski, and Piotr Miłoś. Continual world: A robotic benchmark for continual reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 2021.

[46] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 33:5824–5836, 2020.

[47] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura, editors, *3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings*, volume 100 of *Proceedings of Machine Learning Research*, pages 1094–1100. PMLR, 2019.

[48] Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruyssen, Carlos Riquelme, Mario Lucic, Josip Djolonga, André Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, Lucas Beyer, Olivier Bachem, Michael Tschannen, Marcin Michalski, Olivier Bousquet, Sylvain Gelly, and Neil Houlsby. The visual task adaptation benchmark. *CoRR*, abs/1910.04867, 2019.

[49] Zhuangdi Zhu, Kaixiang Lin, and Jiayu Zhou. Transfer learning in deep reinforcement learning: A survey. *arXiv preprint arXiv:2009.07888*, 2020.

[50] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning, 2021.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Section **??**.
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] In Section 7
   (c) Did you discuss any potential negative societal impacts of your work? [No] We do not see this work as having a significant societal impact.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]

3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code, including the scripts used to run the experiments from the paper, are in the supplementary materials.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We conduct each experiments with multiple seeds (at least 10)
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] We describe these details in Appendix F.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]