
TempoRL: Learning When to Act

André Biedenkapp¹ Raghu Rajan¹ Frank Hutter^{1,2} Marius Lindauer³

Abstract

Reinforcement learning is a powerful approach to learn behaviour through interactions with an environment. However, behaviours are usually learned in a purely reactive fashion, where an appropriate action is selected based on an observation. In this form, it is challenging to learn *when* it is necessary to execute new decisions. This makes learning inefficient, especially in environments that need various degrees of fine and coarse control. To address this, we propose a proactive setting in which the agent not only selects an action in a state but also for how long to commit to that action. Our TempoRL approach introduces skip connections between states and learns a skip-policy for repeating the same action along these skips. We demonstrate the effectiveness of TempoRL on a variety of traditional and deep RL environments, showing that our approach is capable of learning successful policies up to an order of magnitude faster than vanilla Q -learning.

1. Introduction

Although reinforcement learning (RL) has celebrated many successes in the recent years (see e.g., Mnih et al., 2015; Lillicrap et al., 2016; Baker et al., 2020), in its classical form it is limited to learning policies in a mostly reactive fashion, i.e., observe a state and react to that state with an action. Guided by the reward signal, policies that are learned in such a way can decide which action is expected to yield maximal long-term rewards. However, these policies generally do not learn *when a new decision has to be made*. A more proactive way of learning, in which agents proactively commit to playing an action for multiple steps could further improve RL by (i) potentially providing better exploration compared to common one-step exploration; (ii) faster learning as proactive policies provide a form of

temporal abstraction by requiring fewer decisions; (iii) explainability as learned agents can indicate when they expect new decisions are required.

Temporal abstractions are a common way to simplify learning of policies with potentially long action sequences. Typically, the temporal abstraction is learned on the highest level of a hierarchy and the required behaviour on a lower level (see e.g. Sutton et al., 1999; Eysenbach et al., 2019). For example, on the highest level a *goal policy* learns which states are necessarily visited and on the lower level the *behaviour* to reach goals is learned. Spacing goals far apart still requires to learn complex behaviour policies whereas a narrow goal spacing requires to learn complex goal policies. Another form of temporal abstraction is to use actions that work at different time-scales (Precup et al., 1998). Take for example an agent that is tasked with moving an object. On the highest level the agent would follow a policy with abstract actions, such as *pick-up object*, *move object*, *put-down object*, whereas on the lower level actions could directly control actuators to perform the abstract actions.

Such hierarchical approaches are still reactive, but instead of reacting to an observation on only one level, reactions are learned on multiple levels. Though these approaches might allow us to learn *which* states are necessarily traversed in the environment, they do not enable us to learn *when* a new decision has to be made on the behaviour level.

In this work, we propose an alternative approach: a proactive view on learning policies that allows us to jointly learn a behaviour and how long to carry out that behaviour. To this end, we re-examine the relationship between agent and environment, and the dependency on time. This allows us to introduce *skip connections* for an environment. These skip connections do not change the optimal policy or state-action-values but allow us to propagate information much faster. We demonstrate the effectiveness of our method, which we dub TEMPORL with tabular and deep function approximation on a variety of environments with discrete and continuous action spaces. Our contributions are:

1. We propose a proactive alternative to classical RL.
2. We introduce skip-connections for MDPs by playing an action for several consecutive states, which leads to faster propagation of information about future rewards.

¹Department of Computer Science, University of Freiburg, Germany ²BCAI, Renningen, Germany ³Information Processing Institute (tnt), Leibniz University Hannover, Germany. Correspondence to: André Biedenkapp <biedenka@cs.uni-freiburg.de>.

3. We propose a mechanism based on a hierarchy for learning when to make new decisions through the use of skip-connections.
4. On classical and deep RL benchmarks, we show that TEMPORL outperforms plain DQN, DAR and FiGAR both in terms of learning speed and sometimes even by converging to better policies.

2. Related Work

A common framework for temporal abstraction in RL is the options framework (Precup et al., 1998; Sutton et al., 1999; Stolle & Precup, 2002; Bacon et al., 2017; Harutyunyan et al., 2018; Mankowitz et al., 2018; Khetarpal & Precup, 2019). Options are triples $\langle \mathcal{I}, \pi, \beta \rangle$ where \mathcal{I} is the set of admissible states that defines in which states the option can be played; π is the policy the option follows when it is played; and β is a random variable that determines when an option is terminated. In contrast to our work, options require a lot of prior knowledge about the environment to determine the set of admissible states, as well as the option policies themselves. However, Chaganty et al. (2012) proposed to learn options based on observed connectedness of states. Similarly, SoRB (Eysenbach et al., 2019) uses data from the replay buffer to build a connectedness graph, which allows to query sub-goals on long trajectories. Further work on discovering options paid attention to the termination criterion, learning persistent options (Harb et al., 2018) and meaningful termination criteria (Vezhnevets et al., 2016; Harutyunyan et al., 2019).

Similarly, in AI planning macro actions provide temporal abstractions. However, macro actions are not always applicable as some actions can be locked. Chrupa & Vallati (2019) propose to learn when macro actions become available again, allowing them to identify non-trivial activities. For various problem domains of AI planning, varieties of useful macro actions are known and selecting which macro actions to consider is not trivial. Vallati et al. (2019) propose a macro action selection mechanism that selects which macro actions should be considered for new problems. Further, Nasiriany et al. (2019) show that goal-conditioned policies learned with RL can be incorporated into planning. With complex state observations goal states are difficult to define.

An important element in DQN’s success in tackling various Atari games (Mnih et al., 2015) is due to the use of *frame skipping* (Bellemare et al., 2013). Thereby the agent skips over a few states, always playing the same action, before making a new decision. Without the use of frame skipping, the change between successive observations is small and would have required more observations to learn the same policy. Tuning the *skip-size* can additionally improve performance (Braylan et al., 2015; Khan et al., 2019). A similar line of research focuses on learning persistent poli-

cies which always act after a static, fixed time-step for one-dimensional (Metelli et al., 2020) and multi-dimensional actions (Lee et al., 2020). However, static skip-sizes might not be ideal. Dabney et al. (2020) demonstrated that temporally extended ϵ -greedy schedules improve exploration and thereby performance in sparse-reward environments while performing close to vanilla ϵ -greedy exploration on dense-reward environments.

Different techniques have been proposed to handle continuous time environments (Doya, 2000; Tiganj et al., 2017). Recently, Huang et al. (2019) proposed to use *Markov Jump Processes* (MJs). MJs are designed to study optimal control in MDPs where observations have an associated cost. The goal then is to balance the costs of observations and actions to act in an optimal manner with respect to total cost. Their analysis demonstrated that frequent observations are necessary in regions where an optimal action might change rapidly, while in areas of infrequent change, fewer observations are sufficient. In contrast to ours, this formalism strictly prohibits observations of the skipped transitions to save observation costs and thus losing a lot of information, which otherwise could be used to learn how to act while simultaneously learning when new decisions are required.

Schoknecht & Riedmiller (2002; 2003) demonstrated that learning with multi-step actions can significantly speed up RL. Relatedly, Lakshminarayanan et al. (2017) proposed *DAR*, a Q -network with multiple output heads per action to handle different repetition lengths, drastically increasing the action space but improving learning. In contrast to that, Sharma et al. (2017) proposed *FiGAR*, a framework that jointly learns an action policy and a second repetition policy that decides how often to repeat an action. Crucially, its repetition policy is not conditioned on the chosen action resulting in independent repetition and behaviour actions. The policies are learned together through a joint loss. Thus, counter to our work, the repetition policy only learns which repetition length works well on average for all actions. Further, FiGAR requires modification to the training method of a base agent to accommodate the repetition policy. When evaluating our method in the context of DQN, we compare against DAR and in the context of DDPG against FiGAR as they were originally developed and evaluated on these agent types. The appendix, code and experiment results are available at github.com/automl/TempoRL.

3. TempoRL

We begin this section by introducing skip connections into MDPs, propagating information about expected future rewards faster. We then introduce a novel learning mechanism that makes use of a hierarchy to learn a policy that is capable of not only learning which action to take, but also *when* a new action has to be chosen.




Figure 1. Example transitions with skip of length three (drawn with \cdots). At the same time we can also observe shorter skips of length two ($--$) and normal steps, i.e. skips of length one ($-$).

3.1. Temporal Abstraction through Skip MDPs

It is possible to make use of contextual information in MDPs (Hallak et al., 2015; Modi et al., 2018; Biedenkapp et al., 2020). To this end, we contextualize an existing MDP \mathcal{M} to allow for skip connections as $\mathcal{M}_c := \{\mathcal{M}_c\}_{c \in \mathcal{C}}$ with $\mathcal{M}_c := \langle \mathcal{S}, \mathcal{A}, \mathcal{P}_c, \mathcal{R}_c \rangle$. Akin to options, a skip-connection c is a triple $\langle s, a, j \rangle$, where s is the starting state for a skip transition (and not a set of states as in the options framework); a is the action that is executed when skipping through the MDP; and j is the skip-length, where $a \in \mathcal{A}$, $s \in \mathcal{S}$ and $j \in \mathcal{J} = \{1, \dots, J\}$. This context to the MDP induces different MDPs with shared state and action spaces $(\mathcal{S}, \mathcal{A})$, but different transitions \mathcal{P}_c and reward functions \mathcal{R}_c to account for the introduced skips.

In practice however, the transition and reward functions are unknown and do not allow to easily insert skips. Nevertheless, as we make use of action repetition, we can simulate a skip connection. A skip connects two states s and s' iff state s' is reachable from state s by repeating action a j -times. This gives us the following skip transition function:

$$\mathcal{P}_c(s, a, s') = \begin{cases} \prod_{k=0}^{j-1} \mathcal{P}_{s_k s_{k+1}}^a & \text{if reachable} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

with s_k and s_{k+1} the states traversed by playing action a for the k th time, and with $s_0 = s$ and $s_j = s'$. This change in the transition function is reflected accordingly in the reward:

$$\mathcal{R}_c(s, a, s') = \begin{cases} \sum_{k=0}^{j-1} \gamma^k \mathcal{R}_{s_k s_{k+1}}^a & \text{if reachable} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Thus, for skips of length 1 we recover the original transition function $\mathcal{P}_{\langle s, a, 1 \rangle}(s, a, s') = \mathcal{P}_{ss'}^a$ as well as the original reward function $\mathcal{R}_{\langle s, a, 1 \rangle}(s, a, s') = \mathcal{R}_{ss'}^a$. The goal with skip-MDPs is to find an optimal skip policy $\pi_J: \mathcal{S} \times \mathcal{A} \mapsto \mathcal{J}$, i.e., a policy that takes a state and a behaviour action as input and maps to a skip value that maximally reduces the total required number of decisions to reach the optimal reward. Thus, similar to skip-connections in neural networks, skip MDPs allow us to propagate information about future rewards much more quickly and enables us to determine *when* it becomes beneficial to switch actions.

3.2. Learning When to Make Decisions

In order to learn using skip connections we need a new mechanism that selects which skip connection to use. In

order to facilitate this, we propose using a hierarchy in which a *behaviour* policy determines the action a to be played given the current state s , and a *skip* policy determines how long to commit to this behaviour.

To learn the behaviour, we can make use of classical Q -learning, where the Q -function gives a mapping of expected future rewards when playing action a in state s_t at time t and continuing to follow the behaviour policy π thereafter.

$$Q^\pi(s, a) := \mathbb{E}[r_t + \gamma Q^\pi(s_{t+1}, a_{t+1}) | s = s_t, a] \quad (3)$$

To learn to skip, we first have to define a *skip-action space* that determines all possible lengths of skip-connections, e.g., $j \in \{1, 2, \dots, J\}$. To learn the value of a skip we can make use of n -step Q -learning with the condition that, at each step of the j steps, the action stays the same.

$$Q^{\pi_J}(s, j | a) := \mathbb{E} \left[\sum_{k=0}^{j-1} \gamma^k r_{t+k} + \gamma^j Q^\pi(s_{t+j}, a_{t+j}) | s = s_t, a, j \right] \quad (4)$$

We call this a flat hierarchy since the behaviour and the skip policy have to always make decisions at the same time-step; however, the behaviour policy has to be queried before the skip policy. Once we have determined both the action a and the skip-length j we want to perform, we execute this action for j steps. We can then use standard temporal difference updates to update the behaviour and skip Q -functions with all one-step observations and the overarching skip-observation. Note that the skip Q -function can also be conditioned on continuous actions if the behaviour policy can handle continuous action-spaces.

One interesting observation regarding this learning scheme is that, when playing skip action j , we are able to also observe all smaller skip transitions for all intermediate steps. Figure 1 gives a visual representation. Specifically, we can directly see that, when executing a skip of length j , we can observe and learn from $\frac{j \cdot (j+1)}{2}$ skip-transitions in total. As we observe all intermediate steps, we can use this trajectory of transitions to build a local connectedness graph (similar to Figure 1) from which we can look up all skip-connections. This allows us to efficiently learn the values of multiple skips, i.e. the action value at different time-resolutions. For pseudo-code and more details we refer to Appendix B.

3.3. Learning When to Make Decisions in Deep RL

When using deep function approximation for TEMPORL we have to carefully consider how we parameterize the skip policy. Commonly, in deep RL we do not only deal with featurized states but also with image-based ones. Depending on the state modality we can consider different architectures:

Concatenation The simplest parametrization of our skip-policy assumes that the state of the environment we are




Figure 2. Schematic representations of considered architectures for learning when to make decisions, where a_t is the action coming from a separate behaviour policy.

learning from is featurized, i.e., a state is a vector of individual informative features. In this setting, the skip-policy network can take any architecture deemed appropriate for the environment, where the input is a concatenation of the original state s_t and the chosen behaviour action a_t , i.e., $s'_t = (s_t, a_t)$, see Figure 2a. This allows the skip-policy network to directly learn features that take into account the chosen behaviour action. However, note that this concatenation assumes that the state is already featurized.

Contextualization In deep RL, we often have to learn to act directly from images. In this case, concatenation is not trivially possible. Instead we propose to use the behaviour action as context information further down-stream in the network. Feature learning via convolutions can then progress as normal and the learned high-level features can be concatenated with the action a_t and be used to learn the final skip-value, see Figure 2b.

Shared Weights Concatenation and contextualization learn individual policy networks for the behaviour and skip policies and do not share information between the two. To achieve this we can instead share parts of the networks, e.g., the part of learning higher-level features from images (see Figure 3). This allows us to learn the two policy networks with potentially fewer weights than two completely independently learned networks. In the forward and backward passes, only the shared feature representation with the corresponding output layers are active. Similar to the contextualization, the output layers for the skip-values require the selected action, i.e. the argmax of the action outputs, as additional input.




Figure 3. Architecture with shared feature representation for joint learning of when to make a decision and what action to take.

4. Experiments

We evaluated TEMPORL with tabular as well as deep Q -functions. We first give results for the tabular case. All code, the appendix and experiment data including trained policies are available at github.com/automl/TempoRL. For details on the used hardware see Appendix C.

4.1. Tabular TempoRL

In this subsection, we describe experiments for a tabular Q -learning implementation that we evaluated on various grid-worlds with sparse rewards (see Figure 4). We first evaluate our approach on the *cliff* environment (see Figure 4a) before evaluating the influence of the exploration schedule on both vanilla and TEMPORL Q -learning, which we refer to as Q and t - Q , respectively.




Figure 4. 6×10 Grid Worlds. Agents have to reach a fixed goal state from a fixed start state. Dots represent decision steps of vanilla and TEMPORL Q -learning policies.

Gridworlds All considered environments (see Figure 4) are discrete, deterministic, have sparse rewards and have size 6×10 . Falling off a cliff results in a negative reward (-1) and reaching a goal state results in a positive reward ($+1$). For a more detailed description of the gridworld environments we refer to Appendix D.

For this experiment, we limit our TEMPORL agent to a maximum skip length of $J = 7$; thus, a learned optimal policy requires 4 decision points instead of 3. For evaluations using larger skips we refer to Appendix E. Note that increasing the skip-length improves TEMPORL up to some point, at which it has too many irrelevant skip-actions at its disposal which slightly decreases the performance. We compare the learning speed, in terms of training policies, of our approach to a vanilla Q -learning agent. Both methods are trained for 10 000 episodes using the same ϵ -greedy strategy, where ϵ is linearly decayed from 1.0 to 0.0 over all episodes.

Figure 5a depicts the evaluation performance of both methods. TEMPORL is $13.6 \times$ faster than its vanilla counterpart to reach a reward of 0.5, and $12.4 \times$ faster to reach a reward of 1.0 (i.e., always reach the goal). Figure 5b shows the number of required steps in the environment, as well as the number of decision steps. TEMPORL is capable of finding a policy that reaches the goal much faster than vanilla Q -learning while requiring far fewer decision steps. Furthermore, TEMPORL recovers the optimal policy quicker than vanilla Q -learning. Lastly we can observe that after having trained for ≈ 6 000 episodes, TEMPORL starts to increase the number of decision points. This can be attributed to skip values of an action having converged to the same value and our implementation selecting a random skip as tie-breaker.

Table 1 summarizes the results on all environments in terms of normalized area under the reward curve and number of decisions for three different ϵ -greedy schedules. A reward AUC value closer to 1.0 indicates that the agent was capable




Figure 5. Evaluation performance of tabular Q -learning agents over 100 random seeds. (a) & (b): The agents were trained with a linearly-decaying ϵ -greedy policy on the cliff environment. (a) Achieved reward. (b) Length of executed policy (\cdots) and number of decisions (—) made by the policies. (c) Comparison to temporally extended ϵ -greedy exploration (te- ϵ -greedy Q in plot) on a 23×23 Gridworld (Dabney et al., 2020). t - Q is our proposed TEMPORL agent. The lines/shaded area represent the mean/standard deviation.

Table 1. Normalized AUC for reward and average number of decision steps. Both agents are trained with the same ϵ schedule.

(a) linearly decaying ϵ -schedule						
	Cliff		Bridge		ZigZag	
	Q	t - Q	Q	t - Q	Q	t - Q
Reward _{AUC}	0.92	0.99	0.75	0.97	0.57	0.92
Decisions	27.9	5.2	49.5	5.0	83.6	7.9
(b) logarithmically decaying ϵ -schedule						
Reward _{AUC}	0.96	0.99	0.94	0.98	0.90	0.96
Decisions	21.7	4.9	21.4	5.3	35.6	6.9
(c) constant $\epsilon = 0.1$						
Reward _{AUC}	0.99	0.99	0.98	0.99	0.95	0.99
Decisions	17.1	5.1	14.7	5.2	27.6	7.1

of learning to reach the goal quickly. A lower number of decisions is better as fewer decisions were required to reach the goal, making a policy easier to learn. In view of both metrics, TEMPORL readily outperforms the vanilla agent, learning much faster and requiring far fewer decisions.

Sensitivity to Exploration As the used exploration mechanism can have a dramatic impact on agent performance we evaluated the agents for three commonly used ϵ -greedy exploration schedules. In the cases of linearly and logarithmically decaying schedules, we decay ϵ over all 10 000 training episodes, starting from 1.0 and decaying it to 0 or 10^{-5} , respectively. In the constant case, we set $\epsilon = 0.1$.

As shown in Table 1, maybe not surprisingly, too much (linear) and too little (log) exploration are both detrimental to the agent’s performance. However, TEMPORL performs quite robustly even using suboptimal exploration strategies. TEMPORL outperforms its vanilla counterpart in all cases, showing the effectiveness of our proposed method.

Guiding Exploration To demonstrate TEMPORL not only benefits through better exploration but also learning *when* to act, we use the 23×23 Gridworld and agent hyperparameters as introduced by Dabney et al. (2020).

An agent starts in the top center and has to find a goal further down and to the left only getting a reward for reaching the goal within 1000 steps. Temporally-extended exploration (te- ϵ -greedy Q -learning; Dabney et al., 2020) is able to cover a space much better than 1-step exploration. However, it falls short in guiding the agent back to high reward areas. TEMPORL enables an agent to quickly find a successful policy that reach a goal while exploring around such a policy. Figure 5c shows TEMPORL reliably reaches the goal after ≈ 30 episodes. An agent using temporally-extended epsilon greedy exploration does not reliably reach the goal in this time-frame and on average requires twice as many steps.

4.2. Deep TempORL

In this section, we describe experiments for agents using deep function approximation implemented with PyTorch (Paszke et al., 2019) in version 1.4.0. We begin with experiments on featurized environments before evaluating on environments with image states. We evaluate TEMPORL for DQN with different architectures for the skip Q -function. We compare against *dynamic action repetition* (DAR; Lakshminarayanan et al., 2017) for the DQN experiments and against *Fine grained action repetition* (FiGAR Sharma et al., 2017) for experiments with DDPG.¹

4.2.1. ADVERSARIAL ENVIRONMENT – DDPG

Setup We chose to first evaluate on OpenAI gyms (Brockman et al., 2016) Pendulum-v0 as it is an adversarial setting where high action repetition is nearly guaranteed to overshoot the balancing point. Thus, agents using action repetition that make mistakes during training will have to spend additional time learning *when* it is necessary to be reactive; a challenge vanilla agents are not faced with. We trained all DDPG agents (Lillicrap et al., 2016) for a total of 3×10^4 training steps and evaluated the agents every 250 training steps. The first 10^3 steps follow a uniform random policy to generate the initial experience. We used Adam (Kingma & Ba, 2015) with PyTorchs default settings.

¹Neither DAR, nor FiGAR are publicly available and thus we used our own reimplementation available at github.com/automl/TempORL.

Table 2. Average normalized reward AUC for DDPG agents on Pendulum-v0. t-DDPG and FiGAR are evaluated over different maximal skip-lengths for 15 seeds. Corresponding learning curves are given in Appendix F

	t-DDPG							FiGAR						
DDPG	2	4	6	8	10	14	20	2	4	6	8	10	14	20
0.92	0.89	0.89	<i>0.90</i>	0.89	0.89	0.89	0.88	<i>0.76</i>	0.57	<i>0.39</i>	0.31	0.28	0.25	0.24

Agents All actor and critic networks of all **DDPG** agents consist of two hidden layers with 400 and 300 hidden units respectively. Following (Sharma et al., 2017), **FiGAR** introduces a second actor network that shares the input layer with the original actor network. The output layer is a softmax layer with J outputs, representing the probability of repeating the action for $j \in \{1, \dots, J\}$ time-steps. Both actor outputs are jointly input to the critic and gradients are directly propagated from the critic through both actors. **TEMPORL DDPG** (which we refer to as **t-DDPG** in the following) uses the concatenation architecture which takes the state with the action output of the DDPG actor as input and makes use of the critic’s Q -function when learning the skip Q -function. We evaluate t-DDPG and FiGAR on a grid of maximal skip lengths of $\{2, 4, 6, \dots, 20\}$. See Appendix F for implementation details and all used hyperparameters.

Pendulum Table 2 confirms that agents using action repetition indeed are slower in learning successful policies, as reflected by the normalized reward AUC. As **FiGAR** does not directly inform the skip policy about the chosen repetition value or vica versa, the agent tends to struggle quite a lot in this environment already with only two possible skip-values and is not capable of handling larger maximal skip values. In contrast to that, t-DDPG only slightly lags behind vanilla DDPG and readily adapts to larger skip lengths, by quickly learning to ignore irrelevant skip-values. Further, due to making use of n-step learning, t-DDPG starts out very conservative as large skip values appear to lead to larger negative rewards in the beginning. With more experience however, t-DDPG learns *when* switching between actions becomes advantageous, thereby approximately halving the required decisions (see Appendix F).

4.2.2. FEATURZIED ENVIRONMENTS – DQN

Setup We trained all agents for a total of 10^6 training steps using a constant ϵ -greedy exploration schedule with ϵ set to 0.1. We evaluated all agents every 200 training steps. We used the Adam with a learning rate of 10^{-3} and default parameters as given in PyTorch v1.4.0. For increased learning stability, we implemented all agents using double deep Q networks (van Hasselt et al., 2016). All agents used a replay buffer with size 10^6 and a discount factor γ of 0.99. The **TEMPORL** agents used an additional replay buffer of size 10^6 to store observed skip-transitions. We used the **MountainCar-v0** and **LunarLander-v2** environments. See

Appendix G for a detailed description of the environments.

Agents The basic **DQN** architecture consists of 3 layers, each with 50 hidden units and *ReLU* activation functions. The **DAR** baseline used the same architecture as the **DQN** agent but duplicated the output heads twice, each of which is associated with specific repetition values allowing for fine and coarse control. We evaluated possible coarse control values on the grid $\{2, 4, 6, 8, 10\}$, keeping the fine-control value fixed to 1 to allow for actions at every-time step.

For **TEMPORL** agents not using weight sharing we used the same **DQN** architecture for both Q -functions. The concatenation architecture used an additional input unit whereas the context architecture added the behaviour action as context at the third layer after using 10 additional hidden units to process the behaviour action. An agent using a weight-sharing architecture shared the first two layers of the **DQN** architecture and used the third layer of the **DQN** architecture to compute the behaviour Q -values. The skip-output used 10 hidden units to process the behaviour action and processed this output together with the hidden state of the 2nd layer in a 3rd layer with 60 hidden units. We refer to a **DQN** using **TEMPORL** as **t-DQN** in the following.

Influence of the Skip-Architecture We begin by evaluating the influence of architecture choice on our **t-DQN** on both environments, before giving a more in-depth analysis on the learning behaviour in the individual environments. To this end, we report the normalized reward AUC for all three proposed architectures and different maximal skip-lengths, see Table 3. Both the concat and context architectures behave similarly on both environments, which is to be expected as they differ very little in setup. Both architectures have an increase in AUC before reaching the best maximal skip-length for the respective environment. The shared architecture, mostly conceptualized for image-based environments, however shows more drastic reactions to choice of J , leading to the best result in the first and to the worst result in the other environment.

MountainCar Tables 3a & 4a depict the performance of the agents for different maximal skip lengths and Figure 6a shows the learning curves of the best **TEMPORL** architecture as well as the best found **DAR** agent. On **MountainCar** the **DQN** baseline struggles in learning a successful policy, resulting in a small AUC of 0.50 compared to the best result of t-DQN of 0.64. Furthermore, a well tuned **DAR** baseline,

Table 3. Average normalized reward AUC for different TEMPORL architectures and maximal skip-lengths over 50 seeds. All agents are trained with the same ϵ schedule. Bold faced values give the overall best AUC and cursive values the best per architecture.

(a) MountainCar-v0						(b) LunarLander-v2					
Max Skip	2	4	6	8	10	Max Skip	2	4	6	8	10
concat	0.469	0.523	0.602	0.626	<i>0.630</i>	concat	0.855	0.878	0.868	0.862	0.830
context	0.429	0.540	0.601	0.608	<i>0.620</i>	context	0.858	<i>0.876</i>	0.871	0.859	0.837
shared	0.440	0.464	0.592	0.561	0.644	shared	<i>0.851</i>	0.837	0.803	0.769	0.696

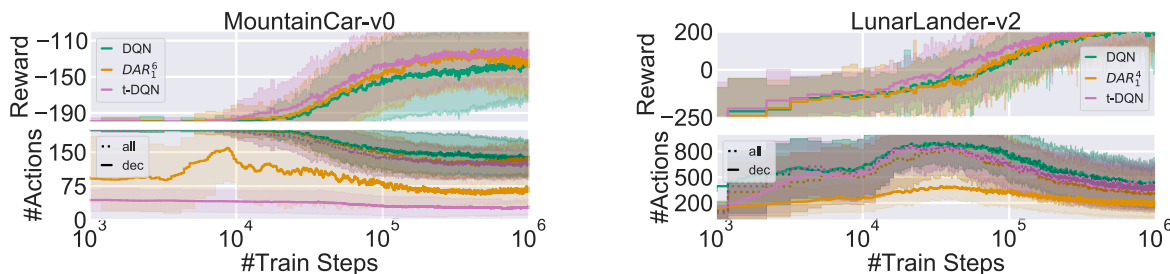


Figure 6. Evaluation performance of deep Q -learning agents on MountainCar-v0 and LunarLander-v2. Solid lines give the mean and the shaded area the standard deviation over 50 random seeds. The sub- and superscripts of **DAR** give the best found fine and coarse repetition values respectively. **t-DQN** is our proposed method using the best architecture as reported in Table 3. (top) Achieved rewards. (bottom) Length of executed policy (\cdots) and number of decisions ($—$) made by the policies.

Table 4. Average normalized reward AUC for maximal skip-length of 10 for MountainCar-v0 and 4 for LunarLander-v2 over 50 seeds. All agents are trained with the same ϵ schedule. We show varying $^{max}_{min}$ repetitions for **DAR** and the best **t-DQN** architecture (see Table 3). Bold faced values give the overall best AUC and cursive values the best per method which are plotted in Figure 6.

(a) MountainCar-v0						
		DAR				
DQN	t-DQN	$\frac{2}{1}$	$\frac{4}{1}$	$\frac{6}{1}$	$\frac{8}{1}$	$\frac{10}{1}$
<i>0.50</i>	0.64	0.43	0.45	<i>0.60</i>	0.56	0.56

(b) LunarLander-v2						
DQN	t-DQN	$\frac{2}{1}$	$\frac{4}{1}$	$\frac{6}{1}$	$\frac{8}{1}$	$\frac{10}{1}$
<i>0.83</i>	0.88	0.84	<i>0.85</i>	0.81	0.72	0.60

carefully trading off fine control and coarse control results in an AUC of 0.593. Figure 6a shows that DAR learns to trade off both coarse and fine-control. However, as DAR does not know that two output heads correspond to the same action, with different repetition values, DARs reward begins to drop in the end as it learns to overly rely on coarse control. During the whole training procedure the best t-DQN agent and the best DAR agent result in policies that require far fewer decisions, with t-DQN requiring only ≈ 50 decisions per episode reducing the number of decisions by a factor of ≈ 3 compared to vanilla DQN.

LunarLander For such a dense reward there is only a small improvement for t-DQN and a properly tuned DAR agent.

Again the t-DQN agent performs best, achieving a slightly higher AUC of 0.88 than the best tuned DAR agent (0.85), see Tables 3b & 4b. Further, Figure 6b shows that, in this setting, t-DQN agents quickly learn to be very reactive, acting nearly at every time-step. Again, DAR can not learn that some output heads apply the same behaviour action for multiple time-steps, preferring coarse over fine control.

4.2.3. ATARI ENVIRONMENTS

Setup We trained all agents for a total of 2.5×10^6 training steps (i.e. only 10 million frames) using a linearly ϵ -greedy exploration schedule over the first 200 000 time-steps with a final fixed ϵ set to 0.01. We evaluated all agents every 10 000 training steps and evaluated for 3 episodes. For increased learning stability we implemented all agents using double deep Q networks. For DQN we used the architecture of Mnih et al. (2015) which also serves as basis for our shared t-DQN and the DAR architecture. As maximal skip-value we chose 10. A detailed list of hyperparameters is given in Appendix H. Following (Bellemare et al., 2013), we used a frame-skip of 4 to allow for a fair comparison to the base DQN. We used *OpenAI Gym*'s We trained all agents on the games BEAMRIDER, FREEWAY, MSPACMAN, PONG and QBERT.

Learning When to Act in Atari Figure 7 depicts the learning curves as well as the number of steps and decisions. The training behaviour from our TEMPORL agents falls into one of three categories on all evaluated Atari games.




Figure 7. Evaluation performance on Atari environments. Solid lines give the mean and the shaded area the standard deviation over 15 random seeds. (top) Achieved rewards. (bottom) Length of executed policy (· · ·) and number of decisions (—) made by the policies. To make trends more visible, we smooth over a window of width 7.

(i) Our learned t-DQN exhibits a slight improvement in learning speed, on MSPACMAN and PONG² before being caught up by DQN, with both methods converging to the same final reward (see Figures 7a & H1a). Nevertheless, TEMPORL learns to make use of different degrees of fine and coarse control to achieve the same performance. For example, a trained proactive TEMPORL policy requires roughly 33% fewer decisions. DAR on the other hand, learns to overly rely on the coarse control, leading to far fewer decisions but also worse final performance.

(ii) On QBERT the learning performance of our t-DQN lags behind that of vanilla DQN over the first 10⁶ steps. Figure 7b (bottom) shows that in the first $\approx 0.5 \times 10^6$ steps, TEMPORL first has to learn which skip values are appropriate for Qbert. In the next $\approx 0.5 \times 10^6$ steps, our t-DQN begins to catch up in reward, while using its learned fine and coarse control, before starting to overtake its vanilla counterpart. As it was not immediately clear if this trend would continue after 2.5×10^6 training steps, we continued the experiment for twice as many steps. TEMPORL continues to outperform its vanilla counterpart, having learned to trade off different levels of coarse and fine control. The effect of over-reliance of DAR on the coarse control is further amplified on QBERT resulting in far worse policies than either vanilla DQN and TEMPORL.

(iii) In games such as FREEWAY and BEAMRIDER (Figures 7c & H1b), we see an immediate benefit to jointly learning *when* and *how* to act through TEMPORL. For these games, our t-DQNs begin to learn faster and achieve a better final reward than vanilla DQNs. An extreme example for this is FREEWAY, where the agents have to control a chicken to cross a busy multi-lane highway as often as possible within a fixed number of frames. To this end one action has to be played predominantly, whereas the other two possible actions are only needed to sometimes avoid an oncoming car. The vanilla DQN learns to nearly constantly play the predominant action, but does not learn proper avoidance strategies, leading to a reward of ≈ 25 (i.e successfully crossing the road 25 times). t-DQN on the other hand not




Figure 8. Example States in which TEMPORL makes new decisions. The agents are trained with a maximal skip-length of 10. (a) Example states in which TEMPORL learned when to make new decisions in MountainCar, starting slightly to the right of the valley. (b) Example states of Qbert. To make it easier to see where QBert is in the images we highlight him as a red square and indicate the taken trajectory with a blue arrow.

only learns faster to repeatedly play the predominant action, but also learns proper avoidance strategies by learning to anticipate *when* a new decision has to be made, resulting in an average reward close to the best possible reward of 34. Here, DAR profits from the use of a coarse control, learning faster than vanilla DQN. However, similarly to vanilla DQN, DAR learns a policy that can only achieve a reward of 25, not learning to properly avoid cars.

5. Analysis of TempoRL Policies

To analyze TEMPORL policies and the decisions *when* to act we selected trained agents and evaluated their policies on the environments they were trained on. Videos for all the behaviours we describe here are part of the supplementary³. In the tabular case we plot the key-states for an agent (see Figure 4) that can skip at most 7 steps ahead. On the Cliff environment TEMPORL learns to make a decision in the starting state, once it has cleared the cliff, once before reaching the other side (since it can not skip more than 7 states) and once to go down into the goal. Similar observations can be made for all grid-worlds. This shows that in this setting our TEMPORL agents are capable of skipping over unimportant states and learned *when* they are required to perform novel actions.

²Results for PONG and BEAMRIDER are given in Appendix H.

³Available at github.com/automl/TempoRL

Key-states in which TEMPORL decides to take new actions in the featurized MountainCar environment are shown in Figure 8a. Starting slightly to the right of the valley, the agent learns to gain momentum by making use of skips, repeating the left action⁴. As soon as TEMPORL considers the run-up to be sufficient to clear the hill on the right, it switches the action direction. From this point on TEMPORL sticks with this action and always selects the largest available skip-length (i.e. 10). Still, TEMPORL has to make many intermediate decisions, as the agent is limited by the maximal skip-length.

Finally, we evaluated TEMPORL’s skipping behaviour on Qbert. An example of key states in which TEMPORL decides to make new decisions are given in Figure 8b. Our TEMPORL agent learns to use large skip-values to reach the bottom of the left column, lighting up all platforms in between. After that the agent makes use of large skips to light up the second diagonal of platforms. Having lit up a large portion of the platform, TEMPORL starts to make fewer uses of skips. This behaviour is best observed in the video provided in the supplementary. Also, note that we included all trained networks in our supplementary such that readers can load the networks to study their behaviour.

This analysis confirms that TEMPORL is capable of not only reacting to states but also learning to anticipate *when* a switch to a different action becomes necessary. Thus, besides the benefit of improved training speed through better guided exploration, TEMPORL improves the interpretability of learned policies.

6. Conclusion

We introduced skip-connections into the existing MDP formulation to propagate information about future rewards much faster by repeating the same action several times. Based on skip-MDPs, we presented a learning mechanism that makes use of existing and well understood learning methods. We demonstrated that our new method, TEMPORL is capable of learning not only how to act in a state, but also *when* a new action has to be taken, without the need for prior knowledge. We empirically evaluated our method using tabular and deep function approximation and empirically evaluated the learning behaviour in an adversarial setting. We demonstrated that the improved learning speed not only comes from the ability of repeating actions but that the ability to learn which repetitions are helpful provided the basis of learning *when* to act. For both tabular and deep RL we demonstrated the high effectiveness of our approach and showed that even in environments requiring mostly fine-

⁴Note, in the particular example given in Figure 8a the agent first performs the left action twice, each for one time-step before it recognizes that it is gaining momentum and it can make use of large skips.

control our approach performs well. Further, we evaluated the influence of exploration strategies, architectural choices and maximum skip-values of our method and showed it to be robust.

As pointed out by Huang et al. (2019), observations might be costly. In such cases, we could make use of TEMPORL to learn how to behave and when new actions need to be taken; when using the learned policies, we could use the learned skip behaviour to only observe after having executed the longest skips possible. All in all, we believe that TEMPORL opens up new avenues for RL methods to be more sample efficient and to learn complex behaviours. As future work, we plan to study distributional TEMPORL as well as how to employ different exploration policies when learning the skip policies and behaviour policies.

Acknowledgements

The authors acknowledge support by the state of Baden-Württemberg through bwHPC, André Biedenkapp, Raghu Rajan and Frank Hutter by the German Research Foundation (DFG) through grant no INST 39/963-1 FUGG as well as funding by the Robert Bosch GmbH, and Marius Lindauer by the DFG through LI 2801/4-1. The authors would like to thank Will Dabney for providing valuable initial feedback as well as Fabio Ferreira and Steven Adriaensen for feedback on the first draft of the paper.

References

- Bacon, P., Harb, J., and Precup, D. The option-critic architecture. In S. Singh and Markovitch, S. (eds.), *Proceedings of the Conference on Artificial Intelligence (AAAI’17)*. AAAI Press, 2017.
- Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., and Mordatch, I. Emergent tool use from multi-agent autotutorials. In *Proceedings of the International Conference on Learning Representations (ICLR’20)*, 2020. Published online: [iclr.cc](https://arxiv.org/abs/2005.00991).
- Bellemare, M. G., Naddaf, Y., Veness, J., and Bowling, M. The arcade learning environment: An evaluation platform for general agents. *J. Artif. Intell. Res.*, 47:253–279, 2013.
- Biedenkapp, A., Bozkurt, H. F., Eimer, T., Hutter, F., and Lindauer, M. Dynamic Algorithm Configuration: Foundation of a New Meta-Algorithmic Framework. In Lang, J., Giacomo, G. D., Dilkina, B., and Milano, M. (eds.), *Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI’20)*, pp. 427–434, June 2020.
- Braylan, A., Hollenbeck, M., Meyerson, E., and Miikkulainen, R. Frame skip is a powerful parameter for learn-

- ing to play atari. In *Proceedings of the Workshops at Twenty-ninth National Conference on Artificial Intelligence (AAAI'15)*, 2015.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. OpenAI Gym. *arXiv:1606.01540 [cs.LG]*, 2016.
- Chaganty, A., Gaur, P., and Ravindran, B. Learning in a small world. In van der Hoek, W., Padgham, L., Conitzer, V., and Winikoff, M. (eds.), *International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2012*, pp. 391–397. IFAAMAS, 2012.
- Chrpa, L. and Vallati, M. Improving domain-independent planning via critical section macro-operators. In Hentenryck, P. V. and Zhou, Z. (eds.), *Proceedings of the Conference on Artificial Intelligence (AAAI'19)*, pp. 7546–7553. AAAI Press, 2019.
- Dabney, W., Ostrovski, G., and Barreto, A. Temporally-extended ϵ -greedy exploration. *arXiv:2006.01782 [cs.LG]*, 2020.
- Doya, K. Reinforcement learning in continuous time and space. *Neural Computation*, 12(1):219–245, 2000.
- Eysenbach, B., Salakhutdinov, R., and Levine, S. Search on the replay buffer: Bridging planning and reinforcement learning. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Proceedings of the 32nd International Conference on Advances in Neural Information Processing Systems (NeurIPS'19)*, pp. 15220–15231, 2019.
- Hallak, A., Castro, D. D., and Mannor, S. Contextual markov decision processes. *arXiv:1502.02259 [stat.ML]*, 2015.
- Harb, J., Bacon, P., Klissarov, M., and Precup, D. When waiting is not an option: Learning options with a deliberation cost. In McIlraith, S. and Weinberger, K. (eds.), *Proceedings of the Conference on Artificial Intelligence (AAAI'18)*, pp. 3165–3172. AAAI Press, 2018.
- Harutyunyan, A., Vrancx, P., Bacon, P., Precup, D., and Nowé, A. Learning with options that terminate off-policy. In McIlraith, S. and Weinberger, K. (eds.), *Proceedings of the Conference on Artificial Intelligence (AAAI'18)*, pp. 3173–3182. AAAI Press, 2018.
- Harutyunyan, A., Dabney, W., Borsa, D., Heess, N., Munos, R., and Precup, D. The termination critic. In Chaudhuri, K. and Sugiyama, M. (eds.), *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 89 of *Proceedings of Machine Learning Research*, pp. 2231–2240. PMLR, 2019.
- Huang, Y., Kavitha, V., and Zhu, Q. Continuous-time markov decision processes with controlled observations. In *Proceedings of the 57th Annual Allerton Conference on Communication, Control, and Computing*, pp. 32–39. IEEE, 2019.
- Khan, A., Feng, J., Liu, S., and Asghar, M. Z. Optimal skipping rates: training agents with fine-grained control using deep reinforcement learning. *Journal of Robotics*, 2019, 2019.
- Khetarpal, K. and Precup, D. Learning options with interest functions. In Hentenryck, P. V. and Zhou, Z. (eds.), *Proceedings of the Conference on Artificial Intelligence (AAAI'19)*, pp. 9955–9956. AAAI Press, 2019.
- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR'15)*, 2015. Published online: iclr.cc.
- Lakshminarayanan, A. S., Sharma, S., and Ravindran, B. Dynamic action repetition for deep reinforcement learning. In S. Singh and Markovitch, S. (eds.), *Proceedings of the Conference on Artificial Intelligence (AAAI'17)*, pp. 2133–2139. AAAI Press, 2017.
- Lee, J., Lee, B., and Kim, K. Reinforcement learning for control with multiple frequencies. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H. (eds.), *Proceedings of the 33rd International Conference on Advances in Neural Information Processing Systems (NeurIPS'20)*, volume 33, 2020.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. In *Proceedings of the International Conference on Learning Representations (ICLR'16)*, 2016. Published online: iclr.cc.
- Mankowitz, D. J., Mann, T. A., Bacon, P., Precup, D., and Mannor, S. Learning robust options. In McIlraith, S. and Weinberger, K. (eds.), *Proceedings of the Conference on Artificial Intelligence (AAAI'18)*, pp. 6409–6416. AAAI Press, 2018.
- Metelli, A., Mazzolini, F., Bisi, L., Sabbioni, L., and Restelli, M. Control frequency adaptation via action persistence in batch reinforcement learning. In III, H. D. and Singh, A. (eds.), *Proceedings of the 37th International Conference on Machine Learning (ICML'20)*. Proceedings of Machine Learning Research, 2020.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M. A., Fidjeland, A., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. Human-level control

- through deep reinforcement learning. *Nature*, 518(7540): 529–533, 2015.
- Modi, A., Jiang, N., Singh, S. P., and Tewari, A. Markov decision processes with continuous side information. In *Algorithmic Learning Theory (ALT'18)*, volume 83 of *Proceedings of Machine Learning Research*, pp. 597–618. PMLR, 2018.
- Moore, A. W. *Efficient memory-based learning for robot control*. PhD thesis, Trinity Hall, University of Cambridge, Cambridge, 1990.
- Nasiriany, S., Pong, V., Lin, S., and Levine, S. Planning with goal-conditioned policies. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Proceedings of the 32nd International Conference on Advances in Neural Information Processing Systems (NeurIPS'19)*, pp. 14814–14825. 2019.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. PyTorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Proceedings of the 32nd International Conference on Advances in Neural Information Processing Systems (NeurIPS'19)*, pp. 8024–8035, 2019.
- Precup, D., Sutton, R. S., and Singh, S. P. Theoretical results on reinforcement learning with temporally abstract options. In *Proceedings of the 10th European Conference on Machine Learning (ECML)'98*, pp. 382–393, 1998.
- Schoknecht, R. and Riedmiller, M. A. Speeding-up reinforcement learning with multi-step actions. In Dorronsoro, J. R. (ed.), *Proceedings of the International Conference on Artificial Neural Networks (ICANN'02)*, volume 2415 of *Lecture Notes in Computer Science*, pp. 813–818. Springer, 2002.
- Schoknecht, R. and Riedmiller, M. A. Reinforcement learning on explicitly specified time scales. *Neural Computing and Applications*, 12(2):61–80, 2003.
- Sharma, S., Lakshminarayanan, A. S., and Ravindran, B. Learning to repeat: Fine grained action repetition for deep reinforcement learning. In *Proceedings of the International Conference on Learning Representations (ICLR'17)*, 2017. Published online: [iclr.cc](https://arxiv.org/abs/1705.09962).
- Stolle, M. and Precup, D. Learning options in reinforcement learning. In *Proceedings of the 5th International Symposium on Abstraction, Reformulation and Approximation SARA'02*, volume 2371 of *Lecture Notes in Computer Science*, pp. 212–223. Springer, 2002.
- Sutton, R. S., Precup, D., and Singh, S. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2): 181–211, 1999.
- Tiganj, Z., Shankar, K. H., and Howard, M. W. Scale invariant value computation for reinforcement learning in continuous time. In *Proceedings of the AAAI Spring Symposia'17*, 2017.
- Vallati, M., Chrapa, L., and Serina, I. MEvo: a framework for effective macro sets evolution. *Journal of Experimental & Theoretical Artificial Intelligence*, 0(0):1–19, 2019.
- van Hasselt, H. Double q-learning. In Lafferty, J., Williams, C., Shawe-Taylor, J., Zemel, R., and Culotta, A. (eds.), *Proceedings of the 23rd International Conference on Advances in Neural Information Processing Systems (NeurIPS'10)*, pp. 2613–2621, 2010.
- van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double q-learning. In Schuurmans, D. and Wellman, M. (eds.), *Proceedings of the Thirtieth National Conference on Artificial Intelligence (AAAI'16)*, pp. 2094–2100. AAAI Press, 2016.
- Vezhnevets, A., Mnih, V., Osindero, S., Graves, A., Vinyals, O., and Agapiou, J. Strategic attentive writer for learning macro-actions. In Lee, D., Sugiyama, M., von Luxburg, U., Guyon, I., and Garnett, R. (eds.), *Proceedings of the 29th International Conference on Advances in Neural Information Processing Systems (NeurIPS'16)*, pp. 3486–3494, 2016.