

# Deep Reinforcement Learning for Event-Driven Multi-Agent Decision Processes

Kunal Menda<sup>1</sup>, Yi-Chun Chen, Justin Grana, James W. Bono, Brendan D. Tracey,  
Mykel J. Kochenderfer<sup>2</sup>, and David Wolpert, *Fellow, IEEE*

**Abstract**—The incorporation of macro-actions (temporally extended actions) into multi-agent decision problems has the potential to address the curse of dimensionality associated with such decision problems. Since macro-actions last for stochastic durations, multiple agents executing decentralized policies in cooperative environments must act asynchronously. We present an algorithm that modifies generalized advantage estimation for temporally extended actions, allowing a state-of-the-art policy optimization algorithm to optimize policies in Dec-POMDPs in which agents act asynchronously. We show that our algorithm is capable of learning optimal policies in two cooperative domains, one involving real-time bus holding control and one involving wildfire fighting with unmanned aircraft. Our algorithm works by framing problems as “event-driven decision processes,” which are scenarios in which the sequence and timing of actions and events are random and governed by an underlying stochastic process. In addition to optimizing policies with continuous state and action spaces, our algorithm also facilitates the use of event-driven simulators, which do not require time to be discretized into time-steps. We demonstrate the benefit of using event-driven simulation in the context of multiple agents taking asynchronous actions. We show that fixed time-step simulation risks obfuscating the sequence in which closely separated events occur, adversely affecting the policies learned. In addition, we show that arbitrarily shrinking the time-step scales poorly with the number of agents.

**Index Terms**—Artificial intelligence, autonomous vehicles, discrete event simulation, distributed decision-making, neural networks, multi-agent systems.

## I. INTRODUCTION

IN COOPERATIVE multi-agent environments, policies become difficult to optimize using reinforcement learning due to the curse of dimensionality. Additionally, the delay between accrued rewards and responsible actions is a significant problem in cooperative multi-agent systems, where one agent may receive undeserved reward from favorable actions performed by other agents. A possible way to address

these problems is to have agents learn multi-step *macro-actions*, rather than the low-level *primitive actions* typically considered in reinforcement learning algorithms.

Suppose we were to use reinforcement learning to learn a policy to fly an unmanned aircraft to a nearby forest. Learning the precise actuator commands from the single reward of whether or not we got there would be intractable. Macro-actions are typically defined by a policy that maps observations to primitive actions and a set of conditions specifying when the macro-action terminates [1]. We can pre-train macro-actions to perform higher level tasks such as way-point tracking, and have a policy learn to use them to navigate to our goal.

When considering cooperative multi-agent environments, the problem of mapping a single reward for successful cooperation to each individual agent’s primitive actions becomes exponentially more difficult as we scale the number of agents. For this reason, the use of macro-actions becomes substantially more necessary. For example, consider a problem in which many unmanned aircraft must cooperate to extinguish multiple fires. By restricting our action space to only high-level macro-actions (such as *fly to a specified fire*, or *attempt to extinguish a fire*), we can rely on traditional controllers to execute the low-level control policies required by these macro-actions. With macro-actions, we can substantially reduce the size of the state space (as many of the degrees-of-freedom pertaining to the pose of any given aircraft may not be relevant at this level of abstraction), as well as significantly reduce reward-delays and ease credit assignment. We can now attribute the rewards of extinguishing a large fire to the fact that we attempted to do so alongside a supporting agent, as opposed to a long sequence of low-level observations and actuator inputs. In general, using macro-actions allows us to shrink the state-space to only the subset of states relevant to deciding between the macro-actions. Additionally, using them enables us to substantially reduce the effective length of a typical episode since we need only consider time-steps in which a new macro-action is selected [2], [3].

Even when considering only *decentralized* policies in which each agent makes an independent decision based on their partial observation of the environment, a challenge arises when extending the use of macro-actions to cooperative multi-agent domains. Since macro-actions are temporally extended, agents executing decentralized policies must act asynchronously. However, few off-the-shelf reinforcement learning algorithms consider the problem of multiple-agents acting asynchronously. To address asynchrony, we view such a multi-agent decision process as event-driven, where agents choose

Manuscript received July 24, 2017; revised April 13, 2018; accepted June 1, 2018. Date of publication July 25, 2018; date of current version March 29, 2019. This work was supported in part by NASA under Grant NNX14A111G and in part by ARO under Grant W911NF-15-1-0127. The Associate Editor for this paper was Q. Zhang. (*Corresponding author: Kunal Menda.*)

K. Menda and M. J. Kochenderfer are with the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305 USA (e-mail: kmenda@stanford.edu).

Y.-C. Chen is with the Anderson School of Management, University of California at Los Angeles, Los Angeles, CA 90095 USA.

J. Grana and B. D. Tracey are with the Santa Fe Institute, Santa Fe, NM 87501 USA.

J. W. Bono is with Economists Inc., San Francisco, CA 94105 USA.

D. Wolpert is with the Resident Faculty, Santa Fe Institute, Santa Fe, NM 87501 USA.

Digital Object Identifier 10.1109/TITS.2018.2848264

a new macro-action when prompted by an event in some set of events occurring in the environment (such as completion of the currently executed macro-action, or the availability of new information). This article presents a method for learning optimal macro-action policies for event-driven decision processes.

In many real world problems, the durations over which macro-actions are active are drawn from continuous distributions. However, current methods for planning and learning with macro-actions in multi-agent settings simulate the environment by discretizing time into fixed time-steps, as well as using discretized state and action spaces [4]–[9]. When using macro-actions in multi-agent environments, temporal discretization poses a unique problem because assuming a fixed time step may group more than one temporally distinct event within a single time step. We refer to this as a *race condition*. In these race conditions, information about the relative timing of the events is lost, and learning with such a simulation environment can lead to optimized policies that transfer poorly to the real world.

As an example in which event-sequences can affect the policy learned, consider two aircraft which can each fly to one of two locations or remain in place. If they choose the same location to fly to, they both accrue a large negative reward (as they may collide and crash). If they choose different locations, they accrue a small positive reward. Given this, we may have that one agent chooses a location first, and the other observes that choice and decides to fly to the other location. However, if the time-step was large enough such that both agents always had to choose their action simultaneously, they may be forced to learn a strategy in which they both always choose the safe option of remaining in place. This example demonstrates an instance in which the choice of simulation time-step can adversely affect the policy learned by an arbitrary reinforcement learning algorithm—a problem which is unique to multi-agent domains with temporally-extended and asynchronously executed actions.

At the cost of computation time, one could mitigate the risk of race conditions by arbitrarily shrinking the time-step. However, we will demonstrate that in order to maintain a chosen threshold for the probability of a race condition, one must decrease the time-step quadratically with the number of agents. Since the computation time for a simulation is typically inversely related to the time-step, the computation time also scales quadratically with the number of agents.

Representing policies as deep neural networks has shown success in the domain of decentralized multi-agent decision making [10], and allows for the representation of continuous state-spaces when trained with policy gradient methods. This paper presents an extension of the PS-TRPO algorithm [10] to accommodate temporally extended actions, which both allows for continuous state-space representation and does not assume a fixed time-step. By using such an algorithm, we can simulate event-driven processes using event-driven simulators, which step from event to event as opposed to from time-step to time-step. By making use of such simulators, we can eliminate race conditions and scale the number of agents while only linearly scaling the computation time of the simulation.

This article presents two key contributions. The first contribution is a modification to the PS-TRPO algorithm that allows it to optimize macro-action policies. In addition to interfacing with event-driven simulation and continuous state/action spaces, the algorithm does not require any expert demonstrations for policy optimization, unlike a state-of-the-art algorithm in this domain [5]. We demonstrate that our algorithm is able to learn optimal policies in cooperative multi-agent environments, including real-time bus holding control and wildfire fighting with unmanned aircraft. The second contribution is a pair of experiments on the wildfire domain that demonstrate the utility of moving from fixed time-step simulations to event-driven simulations. The first of these experiments shows that large time-steps can result in race conditions that cause poor policy transfer between simulation and the real world, and the second demonstrates that mitigating the risk of race conditions forces a quadratic decrease in time-step with the number of agents.

To our knowledge, in addition to being able to interface with event-driven simulators, our algorithm will be the first that uses deep reinforcement learning to optimize decentralized macro-action policies in multi-agent environments. While other research on macro-actions allow for planning over multiple levels of hierarchy, this paper will focus on the case in which learning is over macro-actions that exist at a single level of hierarchy, leaving the extension to full hierarchical learning to future work.

## II. PRELIMINARIES AND RELATED WORK

### A. Generalized Advantage Estimation and TRPO

Policy gradient methods are widely used for optimizing policies through reinforcement learning. In these methods, we define a policy as a mapping from a history of observations  $o_{0:k}$  to a distribution over actions at the  $k$ th decision-instant. A policy  $\psi_{\theta}(o_{0:k})$  is parameterized using some parameter vector  $\theta$ . Policy gradient methods update  $\theta$  by estimating the gradient direction that improves performance the most. Generalized Advantage Estimation (GAE) is a method for computing approximate policy gradients from simulation trajectories [11]. Previous policy-gradient methods, though providing unbiased estimates of the policy gradient, result in gradient estimates with high variance that worsen with long time-horizons. Further, it is argued that GAE, which is a method specified by two hyper-parameters  $\gamma$  and  $\lambda$ , reduces this variance while maintaining a tolerable level of bias [11].

GAE defines the *advantage function*  $A(s_k, a_k)$  to be the difference between  $V(s_k)$  and  $Q(s_k, a_k)$ , where  $V(s_k)$  is the value being in state  $s$  at the decision-instant  $k$  and  $Q(s_k, a_k)$  is the value of taking action  $a$  from state  $s$  at the decision-instant  $k \in \mathbb{N}$ . They incorporate the hyper-parameter  $\gamma$  as follows:

$$V^{\gamma}(s_k) := \mathbb{E}_{s_{k+1:\infty}, a_{k:\infty}} \left[ \sum_{l=0}^{\infty} \gamma^l r_{k+l} \right] \quad (1)$$

$$Q^{\gamma}(s_k, a_k) := \mathbb{E}_{s_{k+1:\infty}, a_{k+1:\infty}} \left[ \sum_{l=0}^{\infty} \gamma^l r_{k+l} \right] \quad (2)$$

$$A^{\gamma}(s_k, a_k) := Q^{\gamma}(s_k, a_k) - V^{\gamma}(s_k) \quad (3)$$

where  $r_k$  is the reward received from taking action  $a_k$  from state  $s_k$ . The discounted approximation to the policy-gradient is then defined as:

$$g^\gamma := \mathbb{E}_{s_0:\infty, a_0:\infty} \left[ \sum_{k=0}^{\infty} \gamma^k A^\gamma(s_k, a_k) \nabla_{\theta} \log \psi_{\theta}(a_k | s_k) \right] \quad (4)$$

When a policy is used to simulate a batch of episodes in one training epoch, the GAE advantages are estimated by first computing  $\delta_k^V$ , which is an estimate of the immediate advantage of action of  $a_k$ , and then the GAE advantage estimate  $A_k^{GAE(\gamma, \lambda)}$  as follows:

$$\delta_k^V := r_k + \gamma V(s_{k+1}) - V(s_k) \quad (5)$$

$$A_k^{GAE(\gamma, \lambda)} = \sum_{l=0}^{\infty} (\gamma \lambda)^l \delta_{k+l}^V \quad (6)$$

Here,  $V(s_k)$  is the approximate value function, referred to as a *baseline*, and  $\lambda$  is a hyper-parameter that is used in addition to  $\gamma$  to control the bias-variance trade-off in gradient estimation. The baseline  $V(s_k)$  is also a function approximator (such as a linear mapping or neural network) parameterized by some vector  $\phi$ . Once the advantages are computed, an algorithm such as Trust-Region Policy Optimization (TRPO) updates the parameters  $\theta$  by constraining the KL-divergence between the previous and new function approximations, and updates  $\phi$  by training the baseline model with supervised learning to map states to their average discounted returns in the set of trajectories [12].

### B. Dec-POMDPs and PS-TRPO

Initial efforts to extend the partially observable Markov decision process (POMDP) framework to multi-agent settings attempted centralized control over the joint state and action spaces of all agents, making planning intractable. To address this problem, efforts have been dedicated to solving *decentralized* versions of the same problems, in which each agent has access to only some local observation of the state space (which may include additional information communicated by other agents), and must choose their action based on only that observation. PS-TRPO is an extension to the TRPO algorithm to cooperative multi-agent domains modeled as decentralized POMDPs (Dec-POMDPs), in which multiple agents act in a single environment with decentralized execution of the same policy  $\psi_{\theta}$  [10]. Here agent  $i$ 's  $k$ th action is given by  $a_{i,k} \sim \psi_{\theta}(o_{i,0:k})$ , where  $o_{i,0:k}$  is the  $i$ th agent's observation history. The parameters  $\theta$  of  $\psi_{\theta}$  are then updated by the TRPO algorithm using advantages computed from all agents' trajectories. They assume a reward structure in which all rewards are shared jointly by agents.

### C. MacDec-POMDPs and the PoEM Algorithm

A policy in a Dec-POMDP is a mapping from an agent's local observations to their local action. However, we may want to optimize a policy in an environment in which the observation spaces high-dimensional. There often exists a natural hierarchical abstraction over the policy space, in which high-level controllers issue commands to low-level controllers, which lead to the development of the *options* framework [1].

An agent  $i$ 's option, or macro-action space  $M_i$  is defined as containing macro-action tuples  $m_i = \langle I_i^m, \beta_i^m, \pi_i^m \rangle$ , where for each macro-action  $m_i \in M_i$ ,  $I_i^m$  specifies a set of states from which the macro-action can be initiated,  $\beta_i^m$  specifies the probability of the action terminating in any given state after having been initiated, and  $\pi_i^m$  specifies the low-level controller corresponding to that macro-action, mapping the agent's observation to a primitive (non-macro) action.

Planning over macro-actions in Dec-POMDPs can be formalized as a *MacDec-POMDP* [13]. Since macro-actions can extend over arbitrary time horizons, we must treat our decision process as a semi-MDP and find an optimal policy  $\psi$  that maps an agent's observation history to a macro-action. The current state-of-the-art model-free reinforcement learning algorithm for optimizing policies in the MacDec-POMDP setting is the PoEM algorithm [5]. This algorithm aims to optimize a finite state controller (FSC) that represents a macro-action policy using expectation-maximization (EM). The algorithm is shown to scale linearly with the number of agents, but quadratically with the number of *nodes* in the FSC, a number which is proportional to the discretization of the observation and action space. Additionally, the algorithm generates experience histories by requiring demonstrations from an expert heuristic controller. In this paper, we will extend the PS-TRPO algorithm to optimize policies in MacDec-POMDPs, preserving the algorithms' ability to optimize over continuous state and action spaces and requiring no expert demonstrations for policy optimization.

## III. EXTENDING PS-TRPO TO MACDEC-POMDPs

To extend the PS-TRPO algorithm to optimize macro-action policies, we must address the fact that actions are asynchronous and temporally extended. We begin by framing a MacDec-POMDP as an event-driven process, allowing us to associate decision instants to discrete events. With this framework, we can modify GAE to allow PS-TRPO to optimize a policy with trajectories generated from MacDec-POMDPs.

### A. MacDec-POMDPs as Event-Driven Processes

Similar to the definition of the option presented in Section II-C, we define the macro-action space as the tuple  $\langle I_i^m, E_i^m, \pi_i^m \rangle$ . For each macro-action  $m$ ,  $I_i^m$  and  $\pi_i^m$  are the set of states from which agent  $i$  can take the macro-action and the lower-level controller the macro-action specifies, respectively. However,  $E_i^m$  now specifies a set of events, the triggering of any of which terminates the macro-action, prompting the agent to select another macro-action given a new observation. In our example where unmanned aircraft coordinate to extinguish fires, the macro-action chosen by one aircraft may be to fly to a particular fire. A lower-level policy tuned to optimally carry out this order would take over control until a relevant event occurs, such as receiving information that the fire being flown to was just extinguished. When such an event occurs, the agent would be given a new observation and the opportunity to choose a new action.

Unlike conventional MDPs, agents may receive rewards at any instant in time, not just when an event prompts them to act. For example, if rewards were given to the whole



team for accomplishments, one agent may accrue a reward for the team without triggering an event that prompts the another agent to choose a new action. Hence, we adopt a convention by which all rewards accrued by an agent over the duration  $\Delta t_{i,k}$ , in which macro-action  $m_{i,k}$  is being carried out, can be collected into a single reward to associate with that macro-action. We let:

$$r_{i,k} = \sum_{j=0}^{J_{i,k}} e^{-\gamma \Delta \tau_j} \rho_j + \int_0^{\Delta t_{i,k}} e^{-\gamma t} c_{i,k}(T_{i,k} + t) dt \quad (7)$$

$$T_{i,k} = \sum_{j=0}^{k-1} \Delta t_{i,j} \quad (8)$$

Here,  $\rho_{0:J_{i,k}}$  are the  $J_{i,k}$  discrete rewards accrued by agent  $i$  at times  $T_{i,k} + \Delta \tau_{0:J_{i,k}}$ ,  $c_{i,k}(t)$  is a continuous-time reward generator over the interval that the macro-action is being executed, and  $\gamma$  is the continuous-time discount rate. The rewards  $\rho_{0:J_{i,k}}$  are accrued on events occurring in the environment during the time when an agent's  $k$ th macro-action is active.

### B. Adapting GAE for Temporally Extended Actions

A simulation episode will now generate a trajectory  $\langle \dots, o_{i,k}, m_{i,k}, \Delta t_{i,k}, r_{i,k}, \dots \rangle$  for each agent  $i$ . With these trajectories, we can modify Equations 5 and 6 in GAE to account for the temporal extension of actions. We will refer to GAE modified for macro-actions as M-GAE. The modified equations are:

$$\delta_{i,k}^V = r_{i,k} + e^{-\gamma \Delta t_{i,k}} V(o_{i,k+1}) - V(o_{i,k}) \quad (9)$$

$$A_{i,k}^{M-GAE(\gamma, \lambda)} = \sum_{l=0}^{\infty} e^{-\gamma \lambda (T_{i,l} - T_{i,k})} \delta_{k+l}^V \quad (10)$$

The modifications now factor in the duration over which macro-actions are active into the discounted sum of rewards that is associated to the taking of each macro-action, as opposed to assuming every action lasts a single time-step.

Using M-GAE, the PS-TRPO algorithm can now optimize policies in MacDec-POMDPs framed as event-driven processes.

## IV. EVENT-DRIVEN SIMULATORS

Many methods for learning policies in multi-agent settings assume that simulations are performed with some fixed step-size. However, in reality, the timings of the various agents' decisions may be drawn from continuous, stochastic distributions. As mentioned earlier, using simulators with fixed time-steps leaves the possibility open that more than one relevant event can occur in the same time-step. Consequently, a fixed time-step simulation obfuscates the relative timing of these events, and may result in policies being learned on the simulator that transfer poorly to the real world, where the sequences may be important.

A key advantage of framing MacDec-POMDPs as event-driven processes is that it allows us to use event-driven simulation. Event-driven simulators, such as the Python module *SimPy* [14], step the environment from event to event as opposed to from one time-step to the next. Underlying the simulation is the assumption that all entities cycle through first




Fig. 1. An example bus corridor.

performing some instantaneous processing, and then *yielding* until some event they specify occurs, which could be the lapse of some specified amount of time.

Even though not every real-world environment is event-driven, modeling one as an event-driven process does not necessarily sacrifice much simulation fidelity. For example, suppose that we have two aircraft, intending to fly to two different fires. We can simulate their individual flights by sampling a duration for each aircraft from a distribution dependent on their desired flight path and average velocities. We can then assert that after those durations have elapsed, each aircraft will have reached their respective fires. By doing so, we can treat the initiation and termination of each agent's flight as distinct events, allowing the simulator to skip simulating anything in-between. In this example, we do, however, lose the ability to simulate the influence flight paths have on each other. This simplification may be benign because we want to learn high-level policies that may not need to take into account these details.

To motivate the use of event-driven simulation, in the latter part of the next section we will demonstrate that the choice of step-size can have adverse effects on the quality of policy learned in simulation. Though we can always shrink the step-size by a sufficient amount to ensure high-fidelity simulation, we will show through a set of Monte-Carlo experiments that doing so causes the computational cost of simulation to scale poorly with the number of agents compared to an event-driven simulation.

## V. PERFORMANCE EXPERIMENTS

The experiments in this section validate the claim that the PS-TRPO algorithm with M-GAE is able to optimize policies in multi-agent environments where agents select from a space of temporally extended actions. The real-time bus holding problem will be used to validate this claim. This section also shows the benefit of using event-driven simulators over fixed time-step simulators in the wildfire fighting problem.

### A. Real-Time Bus Holding

Consider a fleet of buses servicing the bus-corridor, shown in Figure 1. Buses sequentially visit a set of stops numbered 1 through  $K$ , returning to Stop 1 after the loop completes, unloading all passengers, and repeating the cycle. A phenomenon called *bus bunching* can occur, in which buses queue to visit a set of stops, leaving no time between their arrivals, after which a long delay occurs until the bunch cycles back to the stop. This phenomenon leads to unreliable and sub-optimal service, and the research community has explored a strategy called *bus holding* to mitigate the issue. In this control strategy, a bus is instructed to *hold* at its current stop for some specified amount

of time, in addition to the time it waits for passenger boarding and alighting. The problem has previously been framed as a Dec-POMDP, and effective decentralized policies have been discovered through reinforcement learning [15]. We will follow the problem formulation described by Chen *et al.* [15], and show that PS-TRPO with M-GAE generates policies that match their results in performance.

1) *Problem Specification*: Our environment consists of a bus corridor of  $N$  buses and  $K$  bus stops. All agents are initialized in a queue for access to Stop 1, with Bus 1 being the first in the queue and Bus  $N$  being the last. Upon arrival at Stop  $k$ ,  $U_{i,k} = \lfloor q_k \cdot L_{i,k} \rfloor$  passengers will attempt to alight the bus, where  $q_k \in [0, 1]$  and  $L_{i,k} \in [0, L_{\max}]$  is the load of the bus when reaching Stop  $k$ . Here,  $\lfloor \cdot \rfloor$  is used to indicate rounding to the nearest integer. Additionally,  $B_{i,k} = \min(\lfloor \nu_k \cdot h_{i,k} \rfloor, L_{\max} - L_{i,k} + U_{i,k})$  passengers will attempt to board the bus, where  $\nu_k \in [0, \infty)$  is the passenger arrival rate, and  $h_{i,k}$  is the *headway* between Bus  $i$  and Bus  $i - 1$  ahead of it, defined as the time elapsed between the departure of Bus  $i - 1$  and the arrival of Bus  $i$  at Stop  $k$ . The bus will wait a nominal amount of time of  $S_{i,k} = \max(t_a U_{i,k}, t_b B_{i,k})$ , where  $t_a$  and  $t_b$  are the rate at which passengers alight and board, respectively. The bus will then wait an additional amount of time chosen from its action space before departing from the stop. The bus will then travel for time  $rt_{k+1}$  to Stop  $k + 1$ , and join the queue for arrival at that stop.

Upon arriving at a stop, Bus  $i$  will receive an observation of  $[z_1, z_2, z_3, z_4]$ , where  $z_1$  is the current stop's index,  $z_2$  is the headway  $h_{i,k}$ ,  $z_3$  is the current load  $L_{i,k}$ , and  $z_4$  is the sojourn-time elapsed between its current and previous observations. The bus then selects a holding time as some multiple  $a_{i,k} \in \{0, 1, 2, 3\}$  of the parameter  $T_{\text{hold}}$ . The goal of policy optimization is for each bus to choose actions in such a manner that all headways  $h_{i,k}$  are as close as possible to the planned headway  $H$ . Hence, whenever a bus  $i$  arrives at a stop  $k$ , a reward of  $-\nu_k |h_{i,k} - H|^2$  will be accrued by all agents.

The parameters for the bus corridor environment used in our experiments are summarized in Table I. All policies are trained on maximum episode lengths of three simulated hours, which corresponds to four to five cycles of all buses through the corridor. The policy trained using PS-TRPO with M-GAE is represented by a Multi-Layer Perceptron (MLP). The parameters to represent the policy and train it using PS-TRPO with M-GAE are summarized in Table II.<sup>1</sup>

As benchmarks for comparison, we compare the trained MLP policy to the performance of not holding at all, which we refer to as *No Holding*, and a policy that we call *Optimized Thresholds*. In this policy, a set of three ordered thresholds  $T_1 > T_2 > T_3$  are given, and the action chosen is:

$$a_{i,k}(h_{i,k}) = \begin{cases} 0 & \text{if } h_{i,k} > T_1 \\ 1 & \text{if } T_2 < h_{i,k} < T_1 \\ 2 & \text{if } T_3 < h_{i,k} < T_2 \\ 3 & \text{if } h_{i,k} < T_3 \end{cases} \quad (11)$$

<sup>1</sup>The sensitivity of results to enlarging the hidden layers, adding hidden layers, or perturbing the TRPO Max. Step, was not found to be substantial in this experiment.

TABLE I  
ENVIRONMENT PARAMETERS USED IN REAL-TIME  
BUS HOLDING EXPERIMENTS

Parameter	Value	Unit
$N$	6	Buses
$K$	10	Stops
$L_{\max}$	75	pax
$rt_{k=\{1,\dots,10\}}$	180	s
$t_a$	1.8	s
$t_b$	3.0	s
$H$	6.0	min
$T_{\text{hold}}$	30.0	s
Stop	$q_k$	$\nu_k \left[ \frac{\text{pax}}{\text{min}} \right]$
1	1.0	1.5
2	0	2.25
3	0.1	1.4
4	0.25	4.5
5	0.25	2.55
6	0.5	1.8
7	0.5	1.43
8	0.1	1.05
9	0.75	0.75
10	0.1	0.45

TABLE II  
PARAMETERS USED FOR MLP POLICY REPRESENTATION AND  
TRAINING IN REAL-TIME BUS HOLDING EXPERIMENTS

Parameter	Value	Unit
MLP Hidden Size	32	neurons
$\gamma$	$10^{-5}$	$s^{-1}$
$\lambda$	2.0	
TRPO Max Step	0.01	
Batch Size	540	$\frac{\text{episodes}}{\text{epoch}}$




Fig. 2. Training curves for MLP policy on the Real-Time Bus Holding Environment, for various values of the parameter  $\lambda$ .

The values  $T_{1:3}$  are optimized using the SciPy implementation of differential evolution [16], [17]. The optimization target was average return under the same reward function used to optimize the MLP policy with the same episode length.

2) *Results*: As shown in Figures 2 and 3, the optimized MLP policy is able to achieve an average total return per episode that is substantially greater than those of the No Holding and Optimized Thresholds policies, though the Optimized Thresholds policy is itself a substantial improvement over

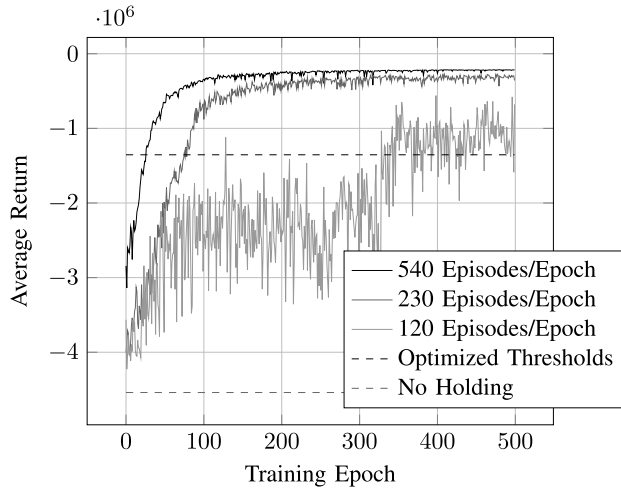


Fig. 3. Training curves for MLP policy on the Real-Time Bus Holding Environment, trained with various batch sizes.

No Holding. In Figure 2, we see that  $\lambda$  slightly influences the total return, but the tested values produce policies that are relatively similar in performance when compared to the Optimized Thresholds. In Figure 3, we do, however, observe that decreasing the batch size degrades convergence. This degradation is due to the fact that the estimated gradients are noisier with smaller batches.

Figures 4 and 5 show that the No Holding policy results in the fleet of buses moving together as a single bunch, as expected. The entire load is borne by the leading bus, which saturates to being full, and the buses behind it remain almost empty. On the other hand, the Optimized Thresholds policy has similar steady-state performance as the optimized MLP policy, but takes approximately 200 minutes from the start of the episode to achieve uniformity in bus arrivals at stops, while the optimized MLP policy requires only around 100 minutes. Neither of these two policies saturate the load of any bus in steady-state operation, but the optimized MLP policy maintains lower variance in the distribution of loads over buses when they reach any given stop. Though the variance in the distribution of loads over buses is not explicitly optimized for, the variance occurs as a consequence of non-uniformity in arrival times. Hence, smaller variance in this distribution of loads suggests that the optimized MLP policy is mitigating the adverse effects of bus-bunching more effectively than Optimized Thresholds can. This experiment provides evidence that the PS-TRPO algorithm using M-GAE is well suited to optimize policies with continuous observation spaces in event-driven, multi-agent environments.

### B. Wildfire Fighting Problem

The goal of this experiment is to demonstrate the utility of using event-driven simulations, which step from event-to-event, over fixed time-step simulations. We introduce a wildfire fighting environment. We will learn policies on both an event-driven simulation of this environment, which makes no assumption regarding time-step, as well as fixed time-step simulations of the environment, where events that occur

TABLE III  
ENVIRONMENT PARAMETERS USED IN THE WILDFIRE FIGHTING PROBLEM EXPERIMENTS

Parameter	Value ( $\pm\sigma$ )	Unit
$N$	3	Aircraft
$K$	9	Fires
$v_{uav}$	0.015	m/s
$T_{hold}$	3.0 ( $\pm 0.3$ )	s
$T_{health}$	3.0	s
$T_{max}$	265	s
$r_{ext}$	1.0	points
$r_{pen}$	20.0	points

within the same time-step are considered to have occurred simultaneously. We will then test all learned policies on the event-driven simulation to examine whether approximation errors introduced by temporal discretization can degrade the performance of the policy in the event-driven environment. As the time-step goes to zero, the simulation should be effectively identical to the event-driven simulation. Thus, it is expected that a sufficiently small time-step should lead to minimal to no degradation in the performance of the policy, when tested on the event-driven simulator.

1) *Problem Specification*: There are  $N$  unmanned aircraft that are tasked with extinguishing  $K$  fires. At a decision-instant, an aircraft can choose any of the five closest fires to move straight toward at velocity  $v_{uav}$ , or to hold its current location for a time  $T_{hold}$ . If an aircraft is currently at a fire, and it chooses to hold, it will attempt to extinguish the fire. Each fire is given a *health*  $T_{health}$ , which is the amount of time aircraft must collectively attempt to extinguish the fire before it is extinguished. When any fire is extinguished, all agents receive a reward of  $r_{ext}$ . If an aircraft is moving toward a fire or currently attacking it, it is considered *interested* in the fire. However, if an aircraft attempts to extinguish a fire that another aircraft is currently attempting to extinguish, all agents receive a penalty of  $r_{pen}$  at the instant it makes that decision. This penalty is meant to simulate an adverse effect of aircraft crowding a space, but also makes the learning problem more interesting by increasing the conditioning of optimal actions on the status of other agents.

At any decision instant  $k$ , aircraft  $i$  receives an observation  $o_{i,k}$  consisting of its own  $x$  and  $y$  location, its distance to the five closest fires, the interest in each of those fires, whether each of those fires has already been extinguished, and the remaining health of each of the fires. A simulation episode lasts a maximum simulated duration of  $T_{max}$ . The parameters used by our simulation environment are summarized in Table III. The nine fires form three clusters of three, with the cluster centroids evenly spaced 0.99 m from the origin, and with the fires evenly spaced 0.01 m from their centroids. The initial locations of the aircraft are randomized uniformly within the bounds of  $[-1, 1]$  m at the start of each episode. The environment parameters have been selected to illustrate the adverse effects of simulation step-size on a reinforcement learning problem with multiple agents acting asynchronously, as opposed to most accurately depict real-world settings. Once again, an MLP is used to represent the decentralized policy.

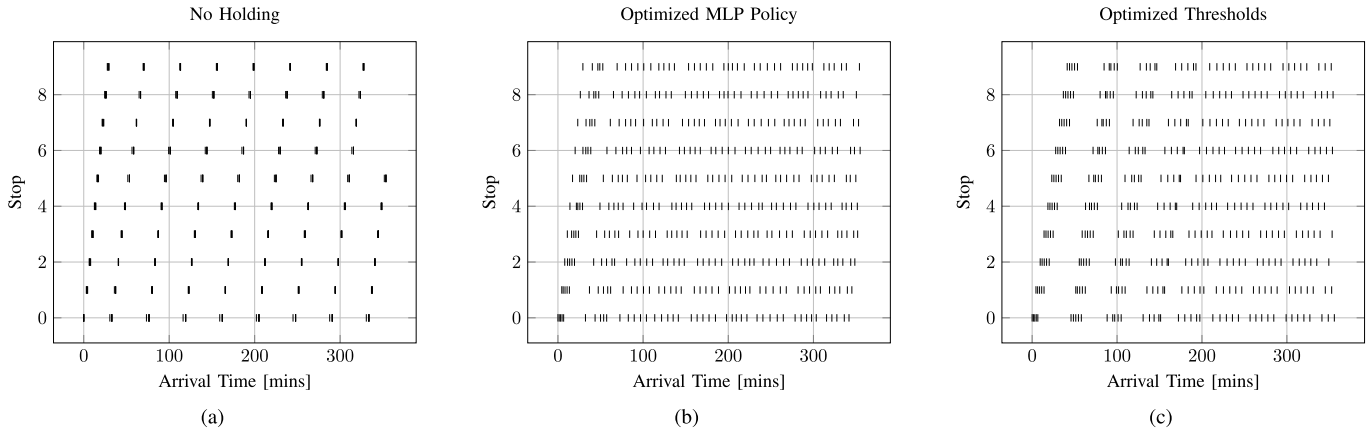


Fig. 4. Arrival times of the six buses at the ten stops, under (a) a policy that uses no holding, (b) a policy learned via PS-TRPO using M-GAE, and (c) a policy that optimizes three headway thresholds for selecting between the actions.

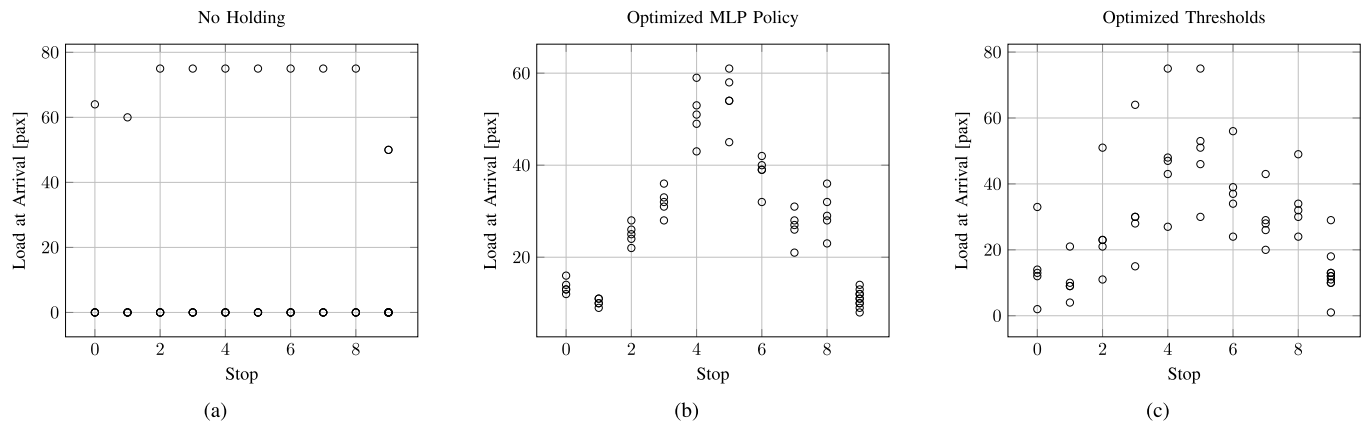


Fig. 5. Load at arrival of the six buses as they arrive at the 10 stops, under (a) a policy that uses no holding, (b) a policy learned via PS-TRPO using M-GAE, and (c) a policy that optimizes three headway thresholds for selecting between the actions. Results are displayed for the buses' fourth cycle through the stops to show the steady-state behavior.

TABLE IV

PARAMETERS USED FOR MLP POLICY REPRESENTATION AND TRAINING IN THE WILDFIRE FIGHTING PROBLEM EXPERIMENTS

Parameter	Value	Unit
MLP Hidden Size	32	neurons
$\gamma$	0.02	$s^{-1}$
$\lambda$	1.0	
TRPO Max Step	0.01	
Batch Size	600~1200	$\frac{\text{episodes}}{\text{epoch}}$

The parameters used to train policies using PS-TRPO with M-GAE are summarized in Table IV.<sup>2</sup>

2) *Results*: To ensure that intelligent policies can be learned in this problem domain, we first trained a policy on the event-driven simulation for an excessive 3000 epochs. The learned behavior was then examined to see if actions taken

<sup>2</sup>We experimented with enlarging the hidden layers, adding hidden layers, and adjusting the TRPO maximum step size. We found that performance was not significantly sensitive to these adjustments. The value of  $\gamma$  is chosen to be non-zero to discourage time-wasting actions, and  $\lambda$  is chosen to impose to extra credit assignment to recent actions. The batch size was varied to achieve acceptable gradient variance while minimizing training time.

by the policy match intuition. Since there is a large penalty for attempting to hold at a fire that another agent is currently holding at, intuition would suggest that doing so when we observe another agent's interest in the fire should be avoided. Meanwhile, it would be sensible to hold at a fire if it does not have another agent interested in it. We find that agents only hold at a fire that has another agent interested in it 27% of the time, while they hold at a fire that has no other agent interested in it 84% of the time. If the agent chooses another fire to target, we find that it chooses the closest and second closest fires 15% of the time, respectively, while choosing the third closest fire 68% of the time. Since fires are arranged in clusters of three, and the separation between clusters is much larger than the separation between the fires within them, agents appear to choose the closest cluster of fires 98% of the time. It is important to note that if an agent is currently at a fire, then moving to the closest fire is equivalent to staying where it is and immediately making a new decision. Though it is peculiar that the *third* closest fire is being chosen with the highest probability, there is little cost to moving to another fire when already at a cluster.

We encode a policy with the behaviors listed in the previous paragraph as follows. If the agent is at a live fire, it will hold



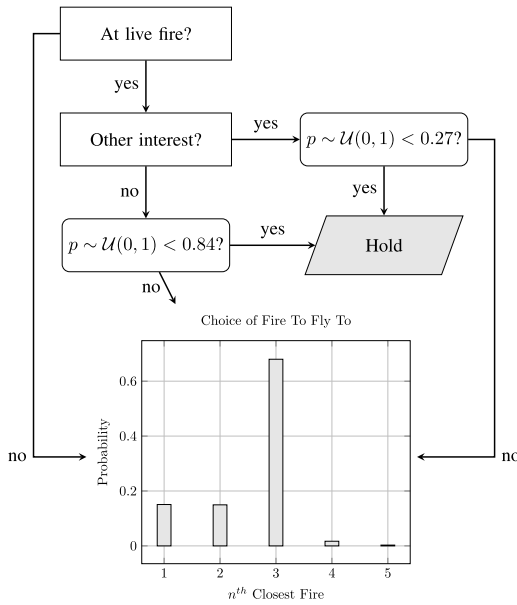


Fig. 6. Approximated behavior of the learned MLP policy in the wildfire fighting problem.

with 84% probability if there is no other interest. If there is other interest, it will hold with 27% probability. If it does not hold, will fly to other fires with the probabilities listed above. A flowchart representing this policy is shown in Figure 6. When testing this simpler approximation of the learned policy, we find that it performs 28% worse than the learned policy when compared on average discounted return. Thus, although this simpler policy seems to approximate the learned policy in important ways, there appear to be subtleties in the learned policy that improve performance.

With confidence that policies training in this problem domain learn intelligent behaviors, we compare the effects of varying the fixed time-step of simulation on the policy learned. All policies were trained for 300 epochs, after which all training curves had reached asymptotes as policies settled to local optima. For each simulation environment, five separate policies were trained from scratch, to account for variation in the local optima on which they may settle. Each of these policies were then tested on the event-driven simulation environment, and the average discounted returns for each of policies were averaged to estimate the performance of a policy learned on that simulation environment and tested on the event-driven environment.

Figure 7 shows performance statistics for fixed time-step simulation environments with time-steps  $\Delta T$  ranging from 0.1s to 10.0s. We see that when the time-step is less than one second, the learned policies perform as well as those learned on the event-driven simulation. However, if the time-step is greater than one second, the performance is substantially degraded. Since the nominal value for  $T_{\text{hold}}$  is  $3.0 \pm 0.03\text{s}$ , and the value for  $T_{\text{health}}$  is 3.0s, half of agents' attempts to extinguish fires would require them to make an additional attempt at doing so, since the fire would not have been extinguished in the single attempt. However, if the simulation time-step is greater than one second, the time held

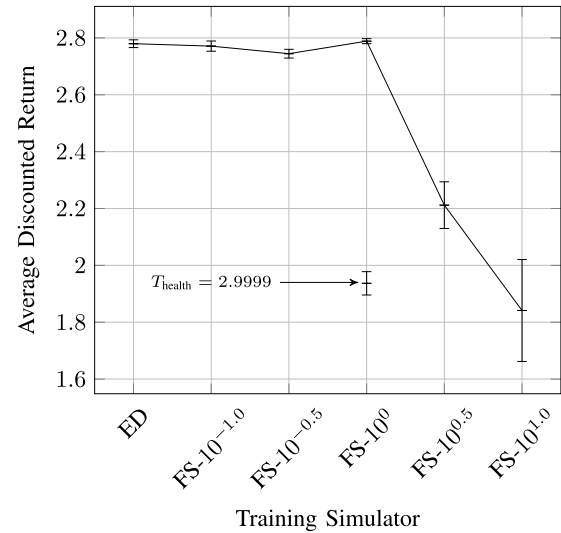


Fig. 7. Performance when tested on the Event-Driven (ED) simulator of policies learned on the (ED) simulator and Fixed Step-size (FS) simulations of various step sizes. Performance drops substantially when the FS-10<sup>0</sup> simulator has  $T_{\text{health}} = 2.9999$  instead of  $T_{\text{health}} = 3.0$ .

at a fire by any agent is rounded to a value greater than 3.0s, resulting in a policy learned expecting that any agents' attempt to extinguish a fire will result in it being extinguished. This expectation causes the performance of policies learned on simulations with these time-steps to degrade when tested on the event-driven simulator.

To confirm that this is the cause of policy degradation, we modify the value of  $T_{\text{health}}$  to be 2.9999. Previously, the approximation made by a time-step of 1s computed the fire to still be alive when it was supposed to be. However, with the modified value of  $T_{\text{health}}$ , this is no longer the case. That is, with this change, aircraft attempting to extinguish fires in the  $\Delta T = 1\text{s}$  simulator will always see fires extinguished in a single attempt. As expected, this causes the performance of the policy to degrade when tested on the event-driven environment.<sup>3</sup>

This experiment provides an example of how the assumed time-step can force a policy to be learned that transfers poorly to an event-driven simulation, or the real-world, where temporal approximations are not made. While this experiment was tailored to demonstrate the problem, we can generally expect the choice of time-step to cause poor policy transfer if it is large enough to obfuscate the sequence in which events occur. In the Wildfire Fighting Problem, the performance degradation occurs because the order between the event of a hold-action ending and the fire being extinguished are obfuscated. If restricted to using fixed time-step simulation, we would thus hope to maintain a time-step that is much smaller than the typically duration that separates important events.

We note that the use of event-driven simulation is appropriate in this domain because we can assume that the operation

<sup>3</sup>Note that using a time-step of 10<sup>-0.5</sup>s appears to produce better results than using a time-step of 1.0s. Though it is statistically possible that they are not different, a possible explanation for why they might be is that hold actions occasionally take  $2 \cdot 10^{-0.5}\text{s}$ , encouraging more careful use of the action.



TABLE V  
RELATIONSHIP BETWEEN  $\Delta T$  AND  $N$ , FOR MAINTAINING  
CONSTANT  $\delta$ , AS THE TIMING OF EVENTS ARE DRAWN  
FROM VARIOUS DISTRIBUTIONS

Distribution	Parameters	$\beta$ (Std. Error)
Exponential	$\lambda = 1$	2.216 (0.055)
Chi-Squared	$k = 5$	2.213 (0.045)
Gamma	$k = 5, \theta = 1$	2.219 (0.037)

of macro-actions, which in this experiment are simple way-point trackers, is independent of the behavior of other agents. However, by using event-driven simulation, we do compromise some simulation fidelity, such as the possibility that two agents' flight paths may intersect and result in expensive evasive maneuvers. Hence, if the dynamics of agents cannot be decoupled during the operation of macro-actions, then a simulation with a sufficiently small time-step is preferable. In the next experiment, we characterize how rapidly the time-step must be shrunk as the number of agents in the environment increases, illustrating the advantage of opting for event-driven simulation.

### C. Scaling the Time-Step to avoid Race-Conditions

Above, we defined a *race condition* to be a time-step in which more than one event occurs. For an  $N$ -agent problem, we can assume that the time at which events pertaining to the  $N$  agents occur are drawn from some distribution  $p(t)$  for  $t \geq 0$ . Hence, for  $N$  samples from  $p(t)$ , a race condition occurs if any two samples fall into the same interval of size  $\Delta T$ . We can always shrink  $\Delta T$  arbitrarily so that the probability of a race condition is less than some threshold  $\delta$ . However, since this comes at the cost of computation time, it is of interest to examine how  $\Delta T$  must shrink as we increase  $N$ . We do so by selecting some example distribution, sampling 10,000 sets of  $N$  samples from the distribution, and for each set of  $N$  samples, we check for a race condition by seeing if more than two of the  $N$  samples fall in the same interval of some chosen  $\Delta T$ . We then vary  $\Delta T$  until the probability of a race condition is less than some threshold value of  $\delta = 0.1$ . To determine a relationship for each distribution tested, we vary  $N$  from 2 to 20, solve for the corresponding  $\Delta T$ , and fit the functional form:

$$\Delta T = \alpha N^\beta \quad (12)$$

Table V shows that for various example distributions,  $\Delta T$  varies inversely with approximately  $N^{2.2}$ . For a fixed duration simulation, computation time scales inversely with the time-step. Hence, we find that computation time in fixed time-step simulations of event-driven process scales approximately as  $N^{2.2}$  with the number of agents  $N$ , assuming we maintain a fixed low-probability of race conditions. Though the underlying distribution  $p(t)$  of an arbitrary environment is unlikely to exactly be any of the example distributions chosen, this experiment suggests that an arbitrary distribution will also have  $\beta \approx 2.2$ , or at the very least,  $\beta > 1$ . Event-driven simulators, on the other hand, will process a number of events

which are in many cases linearly related to the number of agents, providing a substantial computational advantage when simulating large-scale multi-agent problems.

## VI. CONCLUSION

This paper presented an algorithm for learning neural network policies for control in cooperative multi-agent environments with asynchronous and temporally extended actions. The novelty of this contribution is that it extends an algorithm called PS-TRPO that does not require discretization of the observation space. This contrasts with existing algorithms tackling the same class of problem. Section V-A showed that our algorithm is able to learn an optimal policy for real-time bus holding, and is able to achieve better qualitative and quantitative performance than a sensible baseline policy. Since the majority of decision problems can benefit from temporal abstraction, this work is relevant the many multi-agent decision problems currently too intractable to solve when restricted only to primitive actions.

The approach to extending PS-TRPO involves framing the decentralized multi-agent decision problem as event-driven, allowing us to in many circumstances model our environment using event-driven simulation. Using such a simulator does not require assuming a fixed time-step, allowing us to eliminate *race conditions*. Section V-B showed that these artifacts from temporal discretization obfuscated the event sequences and can result in learning policies that are sensitive to the time-step. Further, Section V-C showed that arbitrarily shrinking the time-step to avoid these artifacts scales poorly with the number of agents, motivating the use of event-driven simulation. More generally, this work demonstrates the utility of using event-driven simulation in modeling decision problems in which the sequence of events is critical to simulation fidelity, and the duration between events is drawn from a continuous distributions with semi-infinite support. The source code for this work can be found at <https://github.com/sisl/event-driven-rlab>. The algorithm presented in this paper is built as a modification to the TRPO framework in `rllib` [18].

The work presented here borrows inspiration from hierarchical reinforcement learning, which attempts to simplify learning problems by stratifying decisions in levels of abstraction. However, the problem of simultaneously optimizing high and low-level policies is an open area of research. Our framework has assumed that we are only attempting to optimize a high-level policy, but the logical next step is to relax this assumption. Another avenue for future works lies in the use of multi-fidelity simulation, in which the event-driven simulation may be used to balance simulation fidelity and accuracy.

## ACKNOWLEDGMENT

The authors acknowledge the support of Guillaume P. Brat from the NASA Ames Research Center, as well as the patient help from Jayesh K. Gupta. The authors would also like to thank the anonymous reviewers for their helpful comments.

## REFERENCES

- [1] R. S. Sutton, D. Precup, and S. Singh, "Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning," *Artif. Intell.*, vol. 112, nos. 1–2, pp. 181–211, 1999.

- [2] A. G. Barto and S. Mahadevan, "Recent advances in hierarchical reinforcement learning," *Discrete Event Dyn. Syst.*, vol. 13, no. 4, pp. 341–379, 2003.
- [3] R. He, E. Brunskill, and N. Roy, "Efficient planning under uncertainty with macro-actions," *J. Artif. Intell. Res.*, vol. 40, pp. 523–570, Mar. 2011.
- [4] M. Ghavamzadeh, S. Mahadevan, and R. Makar, "Hierarchical multi-agent reinforcement learning," *Auton. Agents Multi-Agent Syst.*, vol. 13, no. 2, pp. 197–229, 2006.
- [5] M. Liu, C. Amato, E. P. Anesta, J. D. Griffith, and J. P. How, "Learning for decentralized control of multiagent systems in large, partially-observable stochastic environments," in *Proc. AAAI Conf. Artif. Intell. (AAAI)*, 2016, pp. 2523–2529.
- [6] J. Shen, G. Gu, and H. Liu, "Multi-agent hierarchical reinforcement learning by integrating options into MAXQ," in *Proc. Int. Multi-Symp. Comput. Comput. Sci. (IMSCCS)*, vol. 1, Jun. 2006, pp. 676–682.
- [7] D. M. Clement and M. Huber, "Temporal and agent abstractions in multiagent reinforcement learning," in *Proc. IEEE Int. Conf. Syst., Man, Cybern. (SMC)*, Oct. 2016, pp. 2190–2195.
- [8] S. Omidshafiei, A. Agha-Mohammadi, C. Amato, S. Liu, J. P. How, and J. L. Vian, "Health-aware multi-UAV planning using decentralized partially observable semi-Markov decision processes," in *Proc. AIAA Infotech Aerosp.*, 2016, p. 1407.
- [9] C. Amato, G. Konidaris, A. Anders, G. Cruz, J. P. How, and L. P. Kaelbling, "Policy search for multi-robot coordination under uncertainty," *Int. J. Robot. Res.*, vol. 35, no. 14, pp. 1760–1778, 2016.
- [10] J. K. Gupta, M. Egorov, and M. J. Kochenderfer, "Cooperative multi-agent control using deep reinforcement learning," in *Proc. Adapt. Learn. Agents Workshop (ALA)*, 2017, pp. 66–83.
- [11] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2016, pp. 1–14.
- [12] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2015, pp. 1889–1897.
- [13] C. Amato, G. D. Konidaris, and L. P. Kaelbling, "Planning with macro-actions in decentralized POMDPs," in *Proc. Int. Conf. Auton. Agents Multiagent Syst. (AAMAS)*, 2014, pp. 1273–1280.
- [14] K. Müller and T. Vignaux, (2002). *SimPy: Simulating Systems in Python*. [Online]. Available: <https://pypi.python.org/pypi/simpy>
- [15] W. Chen, K. Zhou, and C. Chen, "Real-time bus holding control on a transit corridor based on multi-agent reinforcement learning," in *Proc. IEEE Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2016, pp. 100–106.
- [16] R. Storn and K. Price, "Differential evolution—A simple and efficient heuristic for global optimization over continuous spaces," *J. Global Optim.*, vol. 11, no. 4, pp. 341–359, 1997.
- [17] E. Jones, T. Oliphant, and P. Peterson. (2001). *SciPy: Open Source Scientific Tools for Python*. [Online]. Available: <http://www.scipy.org/>
- [18] Y. Duan, X. Chen, R. Houthoofd, J. Schulman, and P. Abbeel, "Benchmarking deep reinforcement learning for continuous control," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2016, pp. 1329–1338.



**Yi-Chun Chen** was born in Taipei, Taiwan, in 1991. He received the B.S. degree in physics from National Taiwan University in 2013 and the M.S. degree from the Institute for Computational and Mathematical Engineering, Stanford University, in 2017. He is currently pursuing the Ph.D. degree with the UCLA Anderson School of Management, with research focus on operations research. He loves mathematics, algorithms, and artificial intelligence.



**Justin Grana** received the Ph.D. degree in economics from American University. He is currently a Post-Doctoral Researcher with the Santa Fe Institute. He is interested in integrating game theory and continuous-time stochastic processes. He has worked on applying such a combination to computer network security, collusive cartel formulation, and air traffic control.



**James W. Bono** received the B.A. degree in international studies from American University, Washington, DC, USA, and the Ph.D. degree in economics from University of California at Irvine, Irvine, CA, USA. He is currently a Senior Economist with Economists Inc., specializing in microeconomics and game theory.



**Brendan D. Tracey** received the B.S. degree in mechanical engineering from University of Rochester, and the M.S. and Ph.D. degrees in aeronautics and astronautics from Stanford University. He is currently a Post-Doctoral Researcher with the Santa Fe Institute and Massachusetts Institute of Technology. He researches algorithmic developments for machine learning, uncertainty quantification, and the design of engineering systems.



**Mykel J. Kochenderfer** received the B.S. and M.S. degrees in computer science from Stanford University, Stanford, CA, USA, and the Ph.D. degree from The University of Edinburgh, Edinburgh, U.K. He is currently an Assistant Professor of aeronautics and astronautics with Stanford University. He is also the Director of the Stanford Intelligent Systems Laboratory, conducting research on advanced algorithms and analytical methods for the design of robust decision-making systems.



**David Wolpert** (F'16) received the bachelor's degree from Princeton University and the M.S. and Ph.D. degrees from UCSB, all in physics. He has been with a data mining startup, IBM, NASA, and Los Alamos. He is currently a Professor with the Santa Fe Institute (with affiliations at MIT and ASU). His work ranges from game theory, to modeling the dynamics of social organizations, to machine learning, to optimization theory, and to the thermodynamics of computation.



**Kunal Menda** received the B.S. degree in engineering from Harvey Mudd College, Claremont, CA, USA. He is currently pursuing the M.S. and Ph.D. degrees in aeronautics and astronautics with Stanford University, Stanford, CA, USA. He is also a Research Assistant with the Stanford Intelligent Systems Laboratory, conducting research on deep reinforcement learning using hierarchical abstractions.