Emergent Real-World Robotic Skills via Unsupervised Off-Policy Reinforcement Learning

Archit Sharma $^{\alpha}$, Michael Ahn, Sergey Levine, Vikash Kumar, Karol Hausman $^{\beta},$ Shixiang Gu $^{\beta}$ Google Research

Abstract-Reinforcement learning provides a general framework for learning robotic skills while minimizing engineering effort. However, most reinforcement learning algorithms assume that a well-designed reward function is provided, and learn a single behavior for that single reward function. Such reward functions can be difficult to design in practice. Can we instead develop efficient reinforcement learning methods that acquire diverse skills without any reward function, and then re-purpose these skills for downstream tasks? In this paper, we demonstrate that a recently proposed unsupervised skill discovery algorithm can be extended into an efficient off-policy method, making it suitable for performing unsupervised reinforcement learning in the real world. Firstly, we show that our proposed algorithm provides substantial improvement in learning efficiency, making rewardfree real-world training feasible. Secondly, we move beyond the simulation environments and evaluate the algorithm on real physical hardware. On quadrupeds, we observe that locomotion skills with diverse gaits and different orientations emerge without any rewards or demonstrations. We also demonstrate that the learned skills can be composed using model predictive control for goal-oriented navigation, without any additional training.

I. INTRODUCTION

Reinforcement learning (RL) has the potential of enabling autonomous agents to exhibit intricate behaviors and solve complex tasks from high-dimensional sensory input without hand-engineered policies or features [57, 51, 40, 36, 16]. These properties make this family of algorithms particularly applicable to the field of robotics where hand-engineering features and control policies have proven to be challenging and difficult to scale [30, 28, 52, 29, 17, 26]. However, applying RL to real-world robotic problems has not fully delivered on its promise. One of the reasons for this is that the assumptions that are required in a standard RL formulation are not fully compatible with the requirements of real-world robotics systems. One of these assumptions is the existence of a ground truth reward signal, provided as part of the task. While this is easy in simulation, in the real world this often requires special instrumentation of the setup, as well as the ability to reset the environment after every learning episode, which often requires tailored reset mechanisms or manual labor. If we could relax some of these assumptions, we may be able to fully utilize the potential of RL algorithms in realworld robotic problems.

In this context for robotics, the recent work in unsupervised learning becomes relevant — we can learn skills without any

Contact: architsh@google.com



Fig. 1. A 12 degree of freedom quadruped (D'Kitty) discovers diverse locomotion skills without any rewards or demonstrations. We develop off-DADS, an asynchronous and off-policy version of Dynamics-Aware Discovery of Skills (DADS) [55], that enables sample-efficient skill discovery based on mutual-information based optimization.

external reward supervision and then re-purpose those skills to solve downstream tasks using only a limited amount of interaction. Of course, when learning skills without any reward supervision, we have limited control over the kinds of skills that emerge. Therefore, it is critical for unsupervised skill learning frameworks to optimize for diversity, so as to produce a large enough repertoire of skills such that potentially useful skills are likely to be part of this repertoire. In addition, a framework like this needs to offer the user some degree of control over the dimensions along which the algorithm explores. Prior works in unsupervised reinforcement learning [46, 41, 48, 2, 13, 14, 55] have demonstrated that interesting behaviors can emerge from reward-free interaction between the agent and environment. In particular, [13, 14, 55] demonstrate that the skills learned from such unsupervised interaction can be harnessed to solve downstream tasks. However, due to their sample-inefficiency, these prior works in unsupervised skill learning have been restricted to simulation environments (with a few exceptions such as Baranes and Oudeyer [5], Pong et al. [49], Lee et al. [34]) and their feasibility of executing on real robots remains unexplored.

In this paper, we address the limiting sample-inefficiency challenges of previous reward-free, mutual-information-based learning methods and demonstrate that it is indeed feasible to carry out unsupervised reinforcement learning for acquisition

 $[\]alpha$ Work done in Google AI residency, β equal contribution

Paper appendix is available here: https://arxiv.org/abs/2004.12974

of robotic skills. To this end, we build on the work of Sharma et al. [55] and derive a sample-efficient, off-policy version of a mutual-information-based, reward-free RL algorithm, Dynamics-Aware Discovery of Skills (DADS), which we refer to as off-DADS. Our method uses a mutual-information-based objective for diversity of skills and specification of taskrelevant dimensions (such as x-y location of the robot) to specify where to explore. Moreover, we extend off-DADS to be able to efficiently collect data on multiple robots, which together with the efficient, off-policy nature of the algorithm, makes reward-free real-world robotics training feasible. We evaluate the asynchronous off-DADS method on D'Kitty, a compact cost-effective quadruped, from the ROBEL robotic benchmark suite 3. We demonstrate that diverse skills with different gaits and navigational properties can emerge, without any reward or demonstration. We present simulation experiments that indicate that our off-policy algorithm is up to 4x more efficient than its predecessor. In addition, we conduct real-world experiments showing that the learned skills can be harnessed to solve downstream tasks using model-based control, as presented in [55].

II. RELATED WORK

Our work builds on a number of recent works [17] 33 26, 37, 19 44 that study end-to-end reinforcement learning of neural policies on real-world robot hardware, which poses significant challenges such as sample-efficiency, reward engineering and measurements, resets, and safety [17] 11, 60]. Gu et al. [17], Kalashnikov et al. [26], Haarnoja et al. [19], Nagabandi et al. [44] demonstrate that existing off-policy and model-based algorithms are sample efficient enough for real world training of simple manipulation and locomotion skills given reasonable task rewards. Eysenbach et al. [12], Zhu et al. [60] propose reset-free continual learning algorithms and demonstrate initial successes in simulated and real environments. To enable efficient reward-free discovery of skills, our work aims to address the sample-efficiency and reward-free learning jointly through a novel off-policy learning framework.

Reward engineering has been a major bottleneck not only in robotics, but also in general RL domains. There are two kinds of approaches to alleviate this problem. The first kind involves recovering a task-specific reward function with alternative forms of specifications, such as inverse RL [45, 1], 61, 22] or preference feedback [8]; however, these approaches still require non-trivial human effort. The second kind proposes an intrinsic motivation reward that can be applied to different MDPs to discover useful policies, such as curiosity for novelty [53, 46, 54, 6, 47, 9, 31], entropy maximization [23, 49, 34, 15], and mutual information [27, 25, 10, 38, 39, 14, 13, 41, 55]. Ours extends the dynamics-based mutualinformation objective from Sharma et al. [55] to sampleefficient off-policy learning.

Off-policy extension to DADS [55] poses challenges beyond those in standard RL [50, 24, 59, 42]. Since we learn an action abstraction that can be related to a low-level policy in hierarchical RL (HRL) [58, 56, 32, 4, 43], we encounter similar difficulties as in off-policy HRL [43] [35]. We took inspirations from the techniques introduced in [43] for stable off-policy learning of a high-level policy; however, on top of the non-stationarity in policy, we also need to deal with the non-stationarity in reward function as our DADS rewards are continually updated during policy learning. We successfully derive a novel off-policy variant of DADS that exhibits stable and sample-efficient learning.

III. BACKGROUND

In this section, we setup the notation to formally introduce the reinforcement learning problem and the algorithmic foundations of our proposed approach. We work in a Markov decision process (MDP) $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p, r)$, where \mathcal{S} denotes the state space of the agent, \mathcal{A} denotes the action space of the agent, $p: \mathcal{S} \times \mathcal{S} \times \mathcal{A} \to [0, \infty)$ denotes the underlying (stochastic) dynamics of the agent-environment which can be sampled starting from the initial state distribution $p_0: \mathcal{S} \to [0, \infty)$, and a reward function $r: \mathcal{S} \times \mathcal{A} \to [0, \infty)$. The goal of the optimization problem is to learn a controller $\pi(a_t \mid s_t)$ which maximizes $\mathbb{E}[\sum_t \gamma^t r(s_t, a_t)]$ for a discount factor $\gamma \in [0, 1)$.

Within deep reinforcement learning, there are several methods to optimize this objective. In particular, offpolicy methods centered around learning a Q-function [36, 20, 16, 26 are known to be suitable for reinforcement learning on robots. At a high level, algorithms estimate $Q^{\pi}(s_t, a_t) = \mathbb{E}[\sum_{i>t} \gamma^{i-t} r(s_i, a_i)]$, where the expectation is taken over state-action trajectories generated by the executing policy π in the MDP \mathcal{M} after taking action a_t in the state s_t . Crucially, Q^{π} can be estimated using data collected from arbitrary policies using the temporal-difference learning (hence off-policy learning). For continuous or large discrete action spaces, a parametric policy can be updated to $\pi'(a \mid s) \leftarrow \operatorname{argmax}_{a} Q^{\pi}(s, a)$, which can be done approximately using stochastic gradient descent when Q is differentiable with respect to a [36, 21, 18]. While the off-policy methods differ in specifics of each step, they alternate between estimating Q^{π} and updating π using the Q^{π} till convergence. The ability to use trajectories sampled from arbitrary policies enables these algorithms to be sample efficient.

A. Unsupervised Reinforcement Learning

In an unsupervised learning setup, we assume a MDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, p)$ without any reward function r, retaining the previous definitions and notations. The objective is to systematically acquire diverse set of behaviors using autonomous exploration, which can subsequently be used to solve downstream tasks efficiently. To this end, a skill space \mathcal{Z} is defined such that a behavior $z \in \mathcal{Z}$ is defined by the policy $\pi(a|s, z)$. To learn these behaviors in a reward-free setting, the information theoretic concept of mutual information is generally employed. Intuitively, mutual information $\mathcal{I}(x, y)$ between two random variables x, y is high when given x, the uncertainty in value of y is low and vice-versa. Formally,

$$\mathcal{I}(x,y) = \int p(x,y) \log \frac{p(x,y)}{p(x)p(y)} dxdy$$

Dynamics-aware Discovery of Skills (DADS) [55] uses the concept of mutual information to encourage skill discovery with predictable consequences. It uses the following conditional mutual information formulation to motivate the algorithm: $\mathcal{I}(s', z \mid s)$ where s' denotes the next state observed after executing the behavior z from the state s. The joint distribution can be factorized as follows: $p(z, s, s') = p(z)p(s \mid z)p(s' \mid s, z)$, where p(z) denotes the prior distribution over \mathcal{Z} , $p(s \mid z)$ denotes the stationary distribution induced by $\pi(a \mid s, z)$ under the MDP \mathcal{M} and $p(s' \mid s, z) = \int p(s' \mid s, a)\pi(a \mid s, z)da$ denotes the transition dynamics. The conditional mutual information can be written as

$$\mathcal{I}(s', z \mid s) = \int p(z, s, s') \log \frac{p(s' \mid s, z)}{p(s' \mid s)} dz ds ds'$$

At a high level, optimizing for $\mathcal{I}(s', z \mid s)$ encourages π to generate trajectories such that s' can be determined from s, z (*predictability*) and simultaneously encourages π to generate trajectories where s' cannot be determined well from s without z (*diversity*). Note, computing \mathcal{I} is intractable due to intractability of $p(s' \mid s, z)$ and $p(s' \mid s)$. However, one can motivate the following reinforcement learning maximization for π using variational inequalities and approximations as discussed in [55]:

$$J(\pi) = \mathbb{E}_{z,s,s' \sim p(z,s,s')} [r(s,z,s')]$$
$$r(s,z,s') = \frac{q_{\phi}(s' \mid s, z)}{\sum_{i=1}^{L} q_{\phi}(s' \mid s, z_i)} + \log L$$

for $\{z_i\}_{i=1}^L \sim p(z)$ where q_{ϕ} maximizes

$$J(q_{\phi}) = \mathbb{E}_{z,s,s' \sim p(z,s,s')}[\log q_{\phi}(s' \mid s, z)]$$

Sharma et al. [55] propose an on-policy alternating optimization: At iteration t, collect a batch $\mathcal{B}^{(t)}$ of trajectories from the current policy $\pi^{(t)}$ to simulate samples from p(z, s, s'), update $q_{\phi}^{(t)} \rightarrow q_{\phi}^{(t+1)}$ on $\mathcal{B}^{(t)}$ using stochastic gradient descent to approximately maximize $J(q_{\phi})$, label the transitions with reward $r^{(t+1)}(s,z,s')$ and update $\pi^{(t)} \rightarrow \pi^{(t+1)}$ on $\mathcal{B}^{(t)}$ using any reinforcement learning algorithm to approximately maximize $J(\pi)$. Note, the optimization encourages the policy π to produce behaviors predictable under $q_{\phi}(s'|s,z)$, while rewarding the policy for producing diverse behaviors for different $z \in \mathcal{Z}$. This can be seen from the definition of r(s, z, s'): The numerator will be high when the transition $s \rightarrow s'$ has a high log probability under the current skill z (high $q_{\phi}(s' \mid s, z)$ implies high predictability), while the denominator will be lower if the transition has low probability under z_i (low $q_{\phi}(s' \mid s, z_i)$ implies q_{ϕ} is expecting a different transition under the skill z_i).

Interestingly, the variational approximation $q_{\phi}(s' \mid s, z)$, called skill dynamics, can be used for model-predictive control. Given a reward function at test-time, the sequence of skill $z \in Z$ can be determined online using model-predictive control by simulating trajectories using skill dynamics q_{ϕ} .

IV. TOWARDS REAL-WORLD UNSUPERVISED LEARNING

The broad goal of this section is to motivate and present the algorithmic choices required for accomplishing rewardfree reinforcement learning in the real-world. We address the issue of sample-efficiency of learning algorithms, which is the main bottleneck to running the current unsupervised learning algorithms in the real-world. In the same vein, an asynchronous data-collection setup with multiple actors can substantially accelerate the real-world execution. We exploit the off-policy learning enterprise to demonstrate unsupervised learning in the real world, which allows for both sampleefficient and asynchronous data collection through multiple actors [26].

A. Off-Policy Training of DADS

We develop the off-policy variant of DADS, which we call off-DADS. For clarity, we can restate $J(\pi)$ in the more conventional form of expected discounted sum of rewards. Using the definition of the stationary distribution $p(s \mid z) = \sum_{t=0}^{T} \gamma^t p(s_t = s \mid z)$ for a γ -discounted episodic setting of horizon T, we can write:

$$J(\pi) = \mathbb{E}\left[\sum_{t=0}^{T-1} \gamma^t r(s_t, z, s_{t+1})\right]$$

where the expectation has been taken with respect to trajectories generated by $\pi(a \mid s, z)$ for $z \sim p(z)$. This has been explicitly shown in Appendix A Now, we can write the Q-value function as

$$Q^{\pi}(s_t, a_t) = \mathbb{E}[\sum_{i \ge t} \gamma^{i-t} r(s_t, z, s_{t+1})]$$

For problems with a fixed reward function, we can use offthe-shelf off-policy reinforcement learning algorithms like soft actor-critic [20, 21] or deep deterministic policy gradient [36]. At a high level, we use the current policy $\pi^{(t)}$ to sample a sequence of transitions from the environment and add it to the replay buffer \mathcal{R} . We uniformly sample a batch of transitions $\mathcal{B}^{(t)} = \{(s_i, z_i, a_i, s'_i)\}_{i=0}^B$ from \mathcal{R} and use it to update $\pi^{(t)}$ and $Q^{\pi^{(t)}}$.

However, in this setup: (a) the reward is non-stationary as r(s, z, s') depends upon q_{ϕ} , which is learned simultaneously to π, Q^{π} and (b) learning q_{ϕ} involves maximizing $J(q_{\phi})$ which implicitly relies on the current policy π and the induced stationary distribution $p(s \mid z)$. For (a), we recompute the reward r(s, z, s') for the batch $\mathcal{B}^{(t)}$ using the current iterate $q_{\phi}^{(t)}$. For (b), we propose two alternative methods:

- We use samples from current policy $\pi^{(t)}$ to maximize $J(q_{\phi})$. While this does not introduce any additional bias, it does not take advantage of the off-policy data available in the replay buffer.
- Reuse off-policy data while maximizing $J(q_{\phi})$.

To re-use off policy data for learning q_{ϕ} , we have to consider importance sampling corrections, as the data has been sampled from a different distribution. While we can derive an unbiased gradient estimator, as discussed in Appendix **B**, we motivate Algorithm 1: Asynchronous off-DADS

Initialize parameters $\pi(a \mid s, z), q_{\phi}(s' \mid s, z);$ Initialize replay buffer \mathcal{R} ; // collector threads while not done do Sync $\pi_c \leftarrow \pi$; Sample $z \sim p(z)$; Collect episode using π_c ; // store $\pi_c(a \mid s, z)$ end // training thread $s \leftarrow 0, n \leftarrow 0;$ while not done do while n < s + newsteps do Sync \mathcal{B} ; $n + = size(\mathcal{B})$; // get data from actors $\mathcal{R} \leftarrow \mathcal{R} \cup \mathcal{B};$ end for $i \leftarrow 1$ to T_q do Sample $\{s_j, z_j, a_j, s'_j, \pi_c(a_j \mid s_j, z_j)\}_{j=1}^{B_q} \sim \mathcal{R};$ $w_j \leftarrow clip(\frac{\pi(a_j \mid s_j, z_j)}{\pi_c(a_j \mid s_j, z_j)}, \frac{1}{\alpha}, \alpha);$ Update q_{ϕ} using $\{s_j, z_j, s'_j, w_j\}_{j=1}^{B_q}$; end for $i \leftarrow 1$ to T_{π} do Sample $\{s_j, z_j, a_j, s'_j, \pi_c(a_j \mid s_j, z_j)\}_{j=1}^{B_{\pi}} \sim \mathcal{R};$ $r_j \leftarrow r(s_j, z_j, s'_j); // \text{DADS reward}$ Update π using $\{s_j, z_j, a_j, s'_j, r_j)\}_{j=1}^{B_{\pi}};$ end $s \leftarrow n;$ end

an alternate estimator which is simpler and more stable numerically, albeit biased. Consider the definition of $J(q_{\phi})$:

$$J(q_{\phi}) = \mathbb{E}_{z,s,s'\sim p} \Big[\log q_{\phi}(s' \mid s, z) \Big]$$

= $\int p(z)p(s \mid z)p(s' \mid s, z) \log q_{\phi}(s' \mid s, z) dz ds ds'$
= $\int p(z)p(s \mid z)\pi(a \mid s, z)p(s' \mid s, a)$
 $\log q_{\phi}(s' \mid s, z) dz ds da ds'$

where we have used the fact that $p(s' | s, z) = \int p(s' | s, a)\pi(a | s, z)da$. Now, consider that the samples have been generated by a behavior policy $\pi_c(a | s, z)$. The corresponding generating distribution can be written as: $p^{\pi_c}(z, s, s') = \int p(z)p_c(s | z)\pi(a | s, z)p(s' | s, a)da$, where the prior p(z) over Z and the dynamics p(s' | s, a) are shared across all policies, and $p_c(s | z)$ denotes the stationary state distribution induced by π_c . We can rewrite $J(q_{\phi})$ as

$$J(q_{\phi}) = \int p(z)p_c(s \mid z)\pi_c(a \mid s, z)p(s' \mid s, a)$$
$$\frac{p(s \mid z)\pi(a \mid s, z)}{p_c(s \mid z)\pi_c(a \mid s, z)}\log q_{\phi}(s' \mid s, z)dzdsdads'$$

which is equivalent to

$$J(q_{\phi}) = \mathbb{E}_{z,s,a,s' \sim p^{\pi_c}} \left[\frac{p(s \mid z)\pi(a \mid s, z)}{p_c(s \mid z)\pi_c(a \mid s, z)} \log q_{\phi}(s' \mid s, z) \right]$$

Thus, the gradient for $J(q_{\phi})$ with respect to ϕ can be written as:

$$\nabla_{\phi} J(q_{\phi}) = \mathbb{E} \Big[\frac{p(s \mid z)\pi(a \mid s, z)}{p_c(s \mid z)\pi_c(a \mid s, z)} \nabla_{\phi} \log q_{\phi}(s' \mid s, z) \Big]$$
$$\approx \frac{1}{B_q} \sum_{i=1}^{B_q} \Big[\frac{\pi(a_i \mid s_i, z_i)}{\pi_c(a_i \mid s_i, z_i)} \nabla_{\phi} \log q_{\phi}(s'_i \mid s_i, z_i) \Big]$$

The estimator is biased because we compute the importance sampling correction as $w_i = clip(\frac{\pi(a_i|s_i,z_i)}{\pi_c(a_i|s_i,z_i)}, \frac{1}{\alpha}, \alpha)$ which ignores the intractable state-distribution correction $\frac{p(s|z)}{p_c(s|z)}$ [18]. This considerably simplifies the estimator while keeping the estimator numerically stable (enhanced by clipping) as compared to the unbiased estimator derived in Appendix B. In context of off-policy learning, the bias due to state-distribution shift can be reduced using a shorter replay buffer.

Our final proposed algorithm is summarized in the Algorithm []. At a high level, we use n actors in the environment which use the latest copy of the policy to collect episodic data. The centralized training script keeps adding new episodes to the shared replay buffer \mathcal{R} . When a certain threshold of new experience has been added to \mathcal{R} , the buffer is uniformly sampled to train q_{ϕ} to maximize $J(q_{\phi})$. To update π , we sample the buffer uniformly again and compute r(s, z, s') for all the transitions using the latest q_{ϕ} . The labelled transitions can then be passed to any off-the-shelf off-policy reinforcement learning algorithm to update π and Q^{π} .

V. EXPERIMENTS

In this section, we experimentally evaluate our robotic learning method, off-DADS, for unsupervised skill discovery. First, we evaluate the off-DADS algorithm itself in isolation, on a set of standard benchmark tasks, to understand the gains in sample efficiency when compared to DADS proposed in [55], while ablating the role of hyperparameters and variants of off-DADS. Then, we evaluate our robotic learning method on D'Kitty from ROBEL [3], a real-world robotic benchmark suite. We also provide preliminary results on D'Claw from ROBEL, a manipulation oriented robotic setup in Appendix D.

A. Benchmarking off-DADS

We benchmark off-DADS and its variants on continuous control environments from OpenAI gym [7], similar to [55]. We use the HalfCheetah, Ant, and Humanoid environments, with state-dimensionality 18, 29, and 47 respectively. We also consider the setting where the skill-dynamics only observes the global x, y coordinates of the Ant. This encourages the agent to discover skills which diversify in the x, y space, yielding skills which are more suited for locomotion [13, 55].

To evaluate the performance of off-DADS and the role of hyperparameters, we consider the following variantions:

• *Replay Buffer Size*: We consider two sizes for the replay buffer \mathcal{R} : 10,000 (s) and 1,000,000 (l). As alluded to,



Fig. 2. Sample efficiency comparison of off-DADS with DADS (red). We control the effect of state-distribution shift using length of replay buffers (s implies short and l implies long replay buffer) and importance sampling corrections (1 and 10 being the values of the clipping parameter). We observe that all variants of off-DADS outperforms DADS in terms of sample efficiency, and using a short replay buffer with importance sampling clip parameter set to 10 consistently gives the best performance.

this controls how *on-policy* the algorithm is. A smaller replay buffer will have lower bias due to state-distribution shift, but can lose sample efficiency as it discards samples faster [18].

• Importance Sampling: We consider two settings for the clipping parameter in the importance sampling correction: $\alpha = 1$ and $\alpha = 10$. The former implies that there is no correction as all the weights are clipped to 1. This helps evaluate whether the suggested importance sampling correction gives any gains in terms of sample efficiency.

This gives us four variants abbreviated as s1, s10, l1 and l10. We also evaluate against the off-DADS variant where the skill-dynamics is trained on on-policy samples from the current policy. This helps us evaluate whether training skill-dynamics on off-policy data can benefit the sample efficiency of off-DADS. Note, while this ablation helps us understand the algorithm, this scheme would be wasteful of samples in asynchronous off-policy real world training, where the data from different actors could potentially be coming from different (older) policies. Finally, we benchmark against the baseline DADS, as formulated in [55]. The exact hyperparameters for each of the variants are listed in Appendix C. We record curves for five random seeds for the average intrinsic reward r(s, z, s') as a function of samples from the environment and report the average curves in Figure [2].

We observe that all variants of off-DADS consistently outperform the on-policy baseline DADS on all the environments. The gain in sample efficiency can be as high as four times, as is the case for Ant (x-y) environment where DADS takes 16 million samples to converge to the same levels as shown



Fig. 3. (Left) D'Kitty robot from the ROBEL benchmark. (Right) D'Kitty with the LED configuration for PhaseSpace tracking.

for off-DADS (about 0.8 average intrinsic reward). We also note that irrespective of the size of the replay buffer, the importance sampling correction with $\alpha = 10$ outperforms or matches $\alpha = 1$ on all environments. This positively indicates that the devised importance sampling correction makes a better bias-variance trade-off than no importance sampling. The best performing variant on every environment except Ant (x-y) is the s10. While training skill-dynamics on-policy provides a competitive baseline, the short replay buffer and the clipped importance sampling counteract the distribution shift enough to benefit the overall sample efficiency of the algorithm. Interestingly on Ant (x-y), the best performing variant is l10. The long replay buffer variants are slower than the short replay buffer variants but reach a higher average intrinsic reward. This can be attributed to the smaller state-space for skill-dynamics (only 2-dimensional) and thus, the state-distribution correction required is potentially negligible but at the same time the offpolicy data is helping learn better policies.

B. Real-world Training

We now demonstrate the off-DADS can be deployed for real world reward-free reinforcement learning. To this end, we choose the ROBEL benchmark [3]. In particular, we deploy off-DADS on D'Kitty shown in the Figure [3]. D'Kitty is a 12 DOF compact quadruped capable of executing diverse gaits. We also provide preliminary results for D'Claw, a manipulation-oriented setup from ROBEL in Appendix [D].

C. D'Kitty Experimental Setup

To run real-world training, we constructed a walled $4m \times 4m$ cordoned area, shown in Figure 4. The area is equipped with 5 PhaseSpace Impulse X2 cameras that are equidistantly mounted along two bounding orthogonal walls. These cameras are connected to a PhaseSpace Impulse X2E motion capture system which performs 6 DOF rigid body tracking of the D'Kitty robots' chassis at 480Hz. We use two D'Kitty robots for data collection and training in experiment. Each D'Kitty, we attach one PhaseSpace LED controller which controls 8 active LED markers that are attached to the top surface of the D'Kitty chassis as shown in Figure 3] Each D'Kitty is tethered



Fig. 4. Two quadrupeds in a cordoned area. The LEDs allow the robots and the object to be tracked using the PhaseSpace cameras.

via 3 cables: USB serial to the computer running off-DADS, 12V/20A power to the D'Kitty robot, and USB power to the LED controller. To reduce wire entanglement with the robot, we also have an overhead support for the tethering wires.

Dynamixel Property	Value
Model	XM-W210
Control Mode	Position Control
Baudrate	1 Mbps
PWM Limit	450 (50.85%)
Voltage Range	9.5V to 16V

Fig. 5. Dynamixel motor configuration for the D'Kitty robots.

D. Algorithmic Details

We first test the off-DADS algorithm variants in simulation. For the D'Kitty observation space, we use the Cartesian position and Euler orientation (3 + 3), joint angles and velocities (12 + 12), the last action executed by the D'Kitty (12) and the upright (1), which is the cosine of the orientation with global z-axis. The concatenated observation space is 43-dimensional. Hyperparameter details for off-DADS (common to all variants) are as follows: The skill space \mathcal{Z} is 2D with support over $[-1,1]^2$. We use a uniform prior p(z) over \mathcal{Z} . We parameterize $\pi(a \mid s, z), Q^{\pi}(s, a, z)$ and $q_{\phi}(s' \mid s, z)$ using neural networks with two hidden layers of size 512. The output of $\pi(a \mid s, z)$ is parameterized by a normal distribution $\mathcal{N}(\mu, \Sigma)$ with a diagonal covariance which is scaled to [-1,1] using tanh transformation. For q_{ϕ} , we reduce the observation space to the D'Kitty co-ordinates (x, y). This encourages skill-discovery for locomotion behaviors [55], [13]. We parameterize q_{ϕ} to predict $\Delta s = s' - s$, a general trick in model-based control which does not cause any loss in representational power as the next state can be recovered by adding the prediction to the current state. We use soft-actor critic [20] to optimize π, Q^{π} . To learn q_{ϕ} , we sample batches of size 256 and use the Adam optimizer with a fixed learning rate of 0.0003 for $T_q = 8$ steps. For soft-actor critic, we again use Adam optimizer with a fixed learning rate of 0.0003 while sampling batches of size



Fig. 6. (Left) Training curves for D'Kitty in both simulation and real-world. We find the off-DADS with a short replay and importance sampling clipping parameter $\alpha = 10$ to be the most suitable for the real-world learning. We find the real-world learning curve closely follows the simulation learning curve.

256 for 128 steps. Discount factor $\gamma = 0.99$, with a fixed entropy coefficient of 0.1. For computing the DADS reward r(s, z, s'), we set L = 100 samples from the prior p(z). We set the episode length to be 200, which terminates prematurely if the upright coefficient falls below 0.9 (that is the D'Kitty is tilting more than 25 degrees from the global z-axis).

In terms of off-DADS variants, we evaluate the four variants discussed in the previous section. For all the variants, we collect at least 500 steps in the simulation before updating q_{ϕ} and π . The observations for the variants resemble those of the Ant (x-y) environment. We observe that the variants with a replay buffer of size 10,000 are much faster to learn than the replay buffer of size 1,000,000. Asymptotically, we observe the long replay buffer outperforms the short replay buffer though. We also observe setting $\alpha = 10$ benefits the cause of sample efficiency.

For the real robotic experiment, we choose the hyperparameters \mathcal{R} to be of size 10,000 and we set $\alpha = 10$. While asymptotically better performance is nice, we prioritized sample efficiency. For the real experiment, we slightly modify the collection condition. For every update of q_{ϕ} and π , we ensure there are 200 new steps and at least 3 new episodes in the replay buffer \mathcal{R} .

E. Emergent Locomotion Behaviors

With the setup and hyperparameters described in the previous sections, we run the real-world experiment. The experiment was ran over 3 days, with the effective training time on the robot being 20 hours (including time spent in maintaining the hardware). We collected around 300,000 samples in total as shown in the learning curve in Figure 6. We capture the emergence of locomotion skills in our video supplement. Figure 1 and Figure 7 show some of the diversity which emerges in skills learned by D'Kitty using off-DADS, in terms of orientation and gaits.



Fig. 7. Diverse gaits learned by the D'Kitty in our real world experiments.

Broadly, the learning occurs in the following steps: (a) D'Kitty first tries to learn how to stay upright to prolong the length of the episode. This happens within the first hour of the episode. (b) It spends the next few hours trying to move around while trying to stay upright. These few hours, the movements are most random and the intrinsic reward is relatively low as they do not correlate well with z. (c) About 5-6 hours into the training, it starts showing a systematic gait which it uses to move in relatively random directions. This is when the intrinsic reward starts to rise. (d) A few more hours of training and this gait is exploited to predictably move in specific directions. At this point the reward starts rising rapidly as it starts diversifying the directions the agent can move in predictably. Interestingly, D'Kitty uses two different gaits to capture and grow in two different directions of motion, which can be seen in the video supplement. (e) At about 16 hours of training, it can reliably move in different directions and it is trying to further increase the directions it can move in predictably. Supplementary videos are available here: https://sites.google.com/view/dads-skill

One interesting difference from simulation where the D'Kitty is unconstrained, is that the harnesses and tethering despite best attempts restrain the movement of the real robot to some extent. This encourages the agent to invest in multiple gaits and use simpler, more reliable motions to move in different directions.

F. Challenges in real-world training

We discuss some of the challenges encountered during real-world reinforcement learning, particularly in context of locomotive agents.

• Reset & Autonomous operation: A good initial state distribution is necessary for the exploration to proceed towards the desirable state distribution. In context of locomotion, a good reset comprises of being in an upright position and relocating away from the extremities of the area. For the former, we tried two reset mechanisms: (a) scripted mechanism, which is shown in the supplemen-

tary video and (b) reset detector which would continue training if the D'Kitty was upright (based on height and tilt with z-axis), else would wait (for human to reset). However, (a) being programmatic is not robust and does not necessarily succeed in every configuration, in addition to being slow. (b) can be really fast considering that D'Kitty is reasonably compact, but requires human oversight. Despite human oversight, the reset detector can falsely assume the reset is complete and initiate the episode, which requires an episode filter to be written. Relocating from the extremities back to the center is a harder challenge. It is important because the tracker becomes noisy in those regions while also curbing the exploration of the policy. However, this problem only arises when the agent shows significant skills to navigate. There are other challenges besides reset which mandate human oversight into the operation. Primarily, random exploration can be tough on the robot, requiring maintenance in terms of tightening of screws and sometimes, replacing motors. We found latter can be avoided by keeping the motor PWMs low (450 is good). While we make progress towards reward-free learning in this work, we leave it to future work to resolve problems on the way to fully autonomous learning.

Constrained space: For the full diversity of skills to emerge, an ideal operation would have unconstrained space with accurate tracking. However, realistically that is infeasible. Moreover, real operation adds unmodelled constraints over simulation environments. For example, the use of harness to reduce wire entanglement with the robot adds tension depending on the location. When operating multiple D'Kitties in the same space, there can be collisions during the operation. Likelihood of such an event progressively grows through the training. About two-fifths through the training, we started collecting with only one D'Kitty in the space to avoid future collisions. Halfway through the training, we decided to expand the area to its limits and re-calibrate our tracking system for the larger area. Despite the expansion, we were still short on space for the operation of just one D'Kitty. To remedy the situation, we started decreasing the episode length. We went from 200 steps to 100, then to 80 and then to 60. However, we observed that short episodes started affecting the training curve adversarially (at about 300k samples). Constrained by the physical limits, we finished the training. While a reasonable skill diversity already emerges in terms of gaits and orientations within the training we conduct, as shown in Figure 1, more skills should be discovered with more training (as suggested by the simulation results as well as the fact the reward curve has not converged). Nonetheless, we made progress towards the conveying the larger point of reward-free learning being realized in real-world.

Distance Travelled	2.17 ± 0.59
Percentage of Falls	5%

Fig. 8. Distance (in m) travelled by skills sampled randomly from the prior in 100 steps, which is 10s of execution time. We also report the number of times D'Kitty falls before completing 100 steps. The results have been averaged over 20 trials.

G. Harnessing Learned Skills

Qualitatively, we see that a diverse set of locomotive skills can emerge from reward-free training. However, as has been discussed in [55], these skills can be harnessed for downstream tasks using model-based control on the learned skill-dynamics $q_{\phi}(s' \mid s, z)$. First, we partly quantify the learned skills in Figure 8 We execute skills randomly sampled from the prior and collect statistics for these runs. In particular, we find that despite limited training, the skills are relatively robust and fall in only 5% of the runs, despite being proficient in covering distance. Interestingly, the learned skills can also be harnessed for model-based control as shown in Figure 9. The details for model-predictive control follow directly from [55], which elucidates on how to to do so using skill-dynamics q_{ϕ} and π . We have included video supplements showing modelpredictive control in the learned skill space for goal navigation.



Fig. 9. Navigation via model-predictive-control over skills learned by off-DADS. Using the skill-dynamics learned in the unsupervised training, the planner composes skills to move towards the goal. The trajectories visualized show movement of D'Kitty towards the goal box, marked by a set of LED markers.

VI. CONCLUSION

In this work, we derived off-DADS, a novel off-policy variant to mutual-information-based reward-free reinforcement learning framework. The improved sample-efficiency from off-policy learning enabled the algorithm to be applied on a real hardware, a quadruped with 12 DoFs, to learn various locomotion gaits under 20 hours without human-design reward functions or hard-coded primitives. Given our dynamics-based formulation from [55], we further demonstrate those acquired skills are directly useful for solving downstream tasks such as navigation using online planning with no further learning. We detail the successes and challenges encountered in our experiments, and hope our work could offer an important foundation toward the goal of unsupervised, continual reinforcement learning of robots in the real world for many days with zero human intervention.

REFERENCES

- [1] Pieter Abbeel and Andrew Y Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings* of the twenty-first international conference on Machine learning, page 1, 2004.
- [2] Joshua Achiam, Harrison Edwards, Dario Amodei, and Pieter Abbeel. Variational option discovery algorithms. arXiv preprint arXiv:1807.10299, 2018.
- [3] Michael Ahn, Henry Zhu, Kristian Hartikainen, Hugo Ponte, Abhishek Gupta, Sergey Levine, and Vikash Kumar. Robel: Robotics benchmarks for learning with lowcost robots. arXiv preprint arXiv:1909.11639, 2019.
- [4] Pierre-Luc Bacon, Jean Harb, and Doina Precup. The option-critic architecture. In *Thirty-First AAAI Confer*ence on Artificial Intelligence, 2017.
- [5] Adrien Baranes and Pierre-Yves Oudeyer. Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems*, 61(1):49–73, 2013.
- [6] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. In Advances in neural information processing systems, pages 1471–1479, 2016.
- [7] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.
- [8] Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In Advances in Neural Information Processing Systems, pages 4299–4307, 2017.
- [9] Cédric Colas, Pierre Fournier, Olivier Sigaud, Mohamed Chetouani, and Pierre-Yves Oudeyer. Curious: intrinsically motivated modular multi-goal reinforcement learning. arXiv preprint arXiv:1810.06284, 2018.
- [10] Christian Daniel, Gerhard Neumann, and Jan Peters. Hierarchical relative entropy policy search. In Artificial Intelligence and Statistics, pages 273–281, 2012.

- [11] Gabriel Dulac-Arnold, Daniel Mankowitz, and Todd Hester. Challenges of real-world reinforcement learning. arXiv preprint arXiv:1904.12901, 2019.
- [12] Benjamin Eysenbach, Shixiang Gu, Julian Ibarz, and Sergey Levine. Leave no trace: Learning to reset for safe and autonomous reinforcement learning. *International Conference on Learning Representations (ICLR)*, 2018.
- [13] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.
- [14] Carlos Florensa, Yan Duan, and Pieter Abbeel. Stochastic neural networks for hierarchical reinforcement learning. *arXiv preprint arXiv:1704.03012*, 2017.
- [15] Seyed Kamyar Seyed Ghasemipour, Richard Zemel, and Shixiang Gu. A divergence minimization perspective on imitation learning methods. *Conference on Robot Learning (CoRL)*, 2019.
- [16] Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous deep q-learning with modelbased acceleration. In *International Conference on Machine Learning*, pages 2829–2838, 2016.
- [17] Shixiang Gu, Ethan Holly, Timothy Lillicrap, and Sergey Levine. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In 2017 IEEE international conference on robotics and automation (ICRA), pages 3389–3396. IEEE, 2017.
- [18] Shixiang Shane Gu, Timothy Lillicrap, Richard E Turner, Zoubin Ghahramani, Bernhard Schölkopf, and Sergey Levine. Interpolated policy gradient: Merging on-policy and off-policy gradient estimation for deep reinforcement learning. In Advances in neural information processing systems, pages 3846–3855, 2017.
- [19] Tuomas Haarnoja, Sehoon Ha, Aurick Zhou, Jie Tan, George Tucker, and Sergey Levine. Learning to walk via deep reinforcement learning. *arXiv preprint arXiv:1812.11103*, 2018.
- [20] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- [21] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. arXiv preprint arXiv:1812.05905, 2018.
- [22] Dylan Hadfield-Menell, Smitha Milli, Pieter Abbeel, Stuart J Russell, and Anca Dragan. Inverse reward design. In Advances in neural information processing systems, pages 6765–6774, 2017.
- [23] Elad Hazan, Sham M Kakade, Karan Singh, and Abby Van Soest. Provably efficient maximum entropy exploration. arXiv preprint arXiv:1812.02690, 2018.
- [24] Nan Jiang and Lihong Li. Doubly robust off-policy value evaluation for reinforcement learning. arXiv preprint arXiv:1511.03722, 2015.

- [25] Tobias Jung, Daniel Polani, and Peter Stone. Empowerment for continuous agent—environment systems. *Adaptive Behavior*, 19(1):16–39, 2011.
- [26] Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, et al. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. arXiv preprint arXiv:1806.10293, 2018.
- [27] Alexander S Klyubin, Daniel Polani, and Chrystopher L Nehaniv. All else being equal be empowered. In *European Conference on Artificial Life*, pages 744–753. Springer, 2005.
- [28] Jens Kober and Jan R Peters. Policy search for motor primitives in robotics. In Advances in neural information processing systems, pages 849–856, 2009.
- [29] Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- [30] Nate Kohl and Peter Stone. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *IEEE International Conference on Robotics and Automation*, 2004. Proceedings. ICRA'04. 2004, volume 3, pages 2619–2624. IEEE, 2004.
- [31] Varun Raj Kompella, Marijn Stollenga, Matthew Luciw, and Juergen Schmidhuber. Continual curiosity-driven skill acquisition from high-dimensional video inputs for humanoid robots. *Artificial Intelligence*, 247:313–335, 2017.
- [32] George Konidaris, Scott Kuindersma, Roderic Grupen, and Andrew Barto. Autonomous skill acquisition on a mobile manipulator. In *Twenty-Fifth AAAI Conference* on Artificial Intelligence, 2011.
- [33] Vikash Kumar, Emanuel Todorov, and Sergey Levine. Optimal control with learned local models: Application to dexterous manipulation. In 2016 IEEE International Conference on Robotics and Automation (ICRA), pages 378–383. IEEE, 2016.
- [34] Lisa Lee, Benjamin Eysenbach, Emilio Parisotto, Eric Xing, Sergey Levine, and Ruslan Salakhutdinov. Efficient exploration via state marginal matching. arXiv preprint arXiv:1906.05274, 2019.
- [35] Andrew Levy, George Konidaris, Robert Platt, and Kate Saenko. Learning multi-level hierarchies with hindsight. *arXiv preprint arXiv:1712.00948*, 2017.
- [36] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971, 2015.
- [37] A Rupam Mahmood, Dmytro Korenkevych, Gautham Vasan, William Ma, and James Bergstra. Benchmarking reinforcement learning algorithms on real-world robots. *arXiv preprint arXiv:1809.07731*, 2018.
- [38] Georg Martius, Ralf Der, and Nihat Ay. Information driven self-organization of complex robotic behaviors. *PloS one*, 8(5), 2013.

- [39] Georg Martius, Luisa Jahn, Helmut Hauser, and Verena V. Hafner. Self-exploration of the stumpy robot with predictive information maximization. In AngelP. del Pobil, Eris Chinellato, Ester Martinez-Martin, John Hallam, Enric Cervera, and Antonio Morales, editors, *Proc. From Animals to Animats, SAB 2014*, volume 8575 of *LNCS*, pages 32–42. Springer, 2014.
- [40] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602, 2013.
- [41] Shakir Mohamed and Danilo Jimenez Rezende. Variational information maximisation for intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pages 2125–2133, 2015.
- [42] Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc Bellemare. Safe and efficient off-policy reinforcement learning. In Advances in Neural Information Processing Systems, pages 1054–1062, 2016.
- [43] Ofir Nachum, Shixiang Shane Gu, Honglak Lee, and Sergey Levine. Data-efficient hierarchical reinforcement learning. In Advances in Neural Information Processing Systems, pages 3303–3313, 2018.
- [44] Anusha Nagabandi, Kurt Konoglie, Sergey Levine, and Vikash Kumar. Deep dynamics models for learning dexterous manipulation. arXiv preprint arXiv:1909.11652, 2019.
- [45] Andrew Y Ng, Stuart J Russell, et al. Algorithms for inverse reinforcement learning. In *Icml*, volume 1, pages 663–670, 2000.
- [46] Pierre-Yves Oudeyer and Frederic Kaplan. What is intrinsic motivation? a typology of computational approaches. *Frontiers in neurorobotics*, 1:6, 2009.
- [47] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by selfsupervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 16–17, 2017.
- [48] Deepak Pathak, Pulkit Agrawal, Alexei A. Efros, and Trevor Darrell. Curiosity-driven exploration by selfsupervised prediction. In *ICML*, 2017.
- [49] Vitchyr H Pong, Murtaza Dalal, Steven Lin, Ashvin Nair, Shikhar Bahl, and Sergey Levine. Skew-fit: Statecovering self-supervised reinforcement learning. arXiv preprint arXiv:1903.03698, 2019.
- [50] Doina Precup. Eligibility traces for off-policy policy evaluation. *Computer Science Department Faculty Publication Series*, page 80, 2000.
- [51] Martin Riedmiller. Neural fitted q iteration-first experiences with a data efficient neural reinforcement learning method. In *European Conference on Machine Learning*, pages 317–328. Springer, 2005.
- [52] Martin Riedmiller, Thomas Gabel, Roland Hafner, and Sascha Lange. Reinforcement learning for robot soccer. *Autonomous Robots*, 27(1):55–73, 2009.
- [53] Jürgen Schmidhuber. Curious model-building control

systems. In Proc. international joint conference on neural networks, pages 1458–1463, 1991.

- [54] Jürgen Schmidhuber. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230– 247, 2010.
- [55] Archit Sharma, Shixiang Gu, Sergey Levine, Vikash Kumar, and Karol Hausman. Dynamics-aware unsupervised discovery of skills. *International Conference on Learning Representations (ICLR)*, 2020.
- [56] Martin Stolle and Doina Precup. Learning options in reinforcement learning. In *International Symposium* on abstraction, reformulation, and approximation, pages 212–223. Springer, 2002.
- [57] Richard S Sutton, Andrew G Barto, et al. *Introduction* to reinforcement learning, volume 135. MIT press Cambridge, 1998.
- [58] Richard S Sutton, Doina Precup, and Satinder Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial intelligence*, 112(1-2):181–211, 1999.
- [59] Philip Thomas and Emma Brunskill. Data-efficient offpolicy policy evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 2139–2148, 2016.
- [60] Henry Zhu, Justin Yu, Abhishek Gupta, Dhruv Shah, Kristian Hartikainen, Avi Singh, Vikash Kumar, and Sergey Levine. The ingredients of real world robotic reinforcement learning. *International Conference on Representation Learning (ICLR)*, 2020.
- [61] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *Aaai*, volume 8, pages 1433–1438. Chicago, IL, USA, 2008.