

Robot Skill Adaptation via Soft Actor-Critic Gaussian Mixture Models

Iman Nematollahi*, Erick Rosete-Beas*, Adrian Röfer, Tim Welschehold, Abhinav Valada, Wolfram Burgard

Abstract—A core challenge for an autonomous agent acting in the real world is to adapt its repertoire of skills to cope with its noisy perception and dynamics. To scale learning of skills to long-horizon tasks, robots should be able to learn and later refine their skills in a structured manner through trajectories rather than making instantaneous decisions individually at each time step. To this end, we propose the Soft Actor-Critic Gaussian Mixture Model (SAC-GMM), a novel hybrid approach that learns robot skills through a dynamical system and adapts the learned skills in their own trajectory distribution space through interactions with the environment. Our approach combines classical robotics techniques of learning from demonstration with the deep reinforcement learning framework and exploits their complementary nature. We show that our method utilizes sensors solely available during the execution of preliminarily learned skills to extract relevant features that lead to faster skill refinement. Extensive evaluations in both simulation and real-world environments demonstrate the effectiveness of our method in refining robot skills by leveraging physical interactions, high-dimensional sensory data, and sparse task completion rewards. Videos, code, and pre-trained models are available at <http://sac-gmm.cs.uni-freiburg.de>.

I. INTRODUCTION

Thinking ahead is a hallmark of human intelligence. From early infancy, we form rich primitive object concepts through our physical interactions with the real world and apply this knowledge as an intuitive model of physics for reasoning about physically plausible trajectories and adapting them to suit our purposes [1]. This is at odds with most current deep imitation and reinforcement learning paradigms for robot sensorimotor control, which, despite recent progress [2]–[5], are typically trained to make isolated decisions at each time step of the trajectory. In fact, most existing methods for learning manipulation skills are end-to-end high-capacity models that map directly from pixels to actions [6]–[8]. However, although these approaches can capture complex relationships and are flexible to adapt in face of noisy perception, they require extensive amounts of data, and the trained agent is typically bound to take a distinct decision at every time step.

Learning from demonstration [9] is the classical paradigm to tackle the problem of representing skills with a trajectory-space policy. In this context, dynamical systems have shown to be a physically plausible motion generation mechanism that provides a high level of reactivity and robustness against perturbations in the environment [10]–[14]. Despite the great success of dynamical systems in affording flexible

*These authors contributed equally. All authors are with University of Freiburg, Germany. This work has been supported by German Federal Ministry of Education and Research under contract number 01IS18040B-OML and the BrainLinks-BrainTools center of the University of Freiburg.

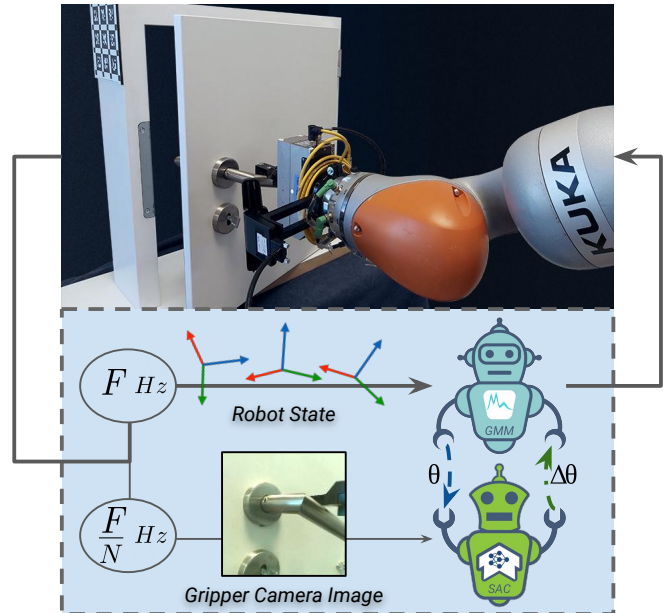


Fig. 1: Soft Actor-Critic Gaussian Mixture Models learn and refine real-world robot skills in the trajectory distribution space. While the Gaussian Mixture Model (GMM) agent learns a skill with few demonstrations and controls the robot with high-frequency F , the Soft Actor-Critic (SAC) agent affords to work with a lower frequency and refines the skill via adapting GMM parameters by leveraging physical interaction, high-dimensional sensory data, and sparse rewards.

robotic systems for industry, where a high-precision state of the environment is available, they are still of limited use in more complex real-world robotics scenarios. The main limitations of current dynamical systems in contrast to deep sensorimotor learning methods are their incompetence in handling raw high-dimensional sensory data such as images, and their susceptibility to noise in the perception pipeline.

In this paper, we advocate for hybrid models in learning robot skills: “Soft Actor-Critic Gaussian Mixture Models” (SAC-GMMs, see Fig 1). SAC-GMMs learn and refine robot skills in the real-world and present a hybrid model that combines dynamical systems and deep reinforcement learning in order to leverage their complementary nature. More precisely, SAC-GMMs learn a trajectory-based Gaussian mixture policy of skills from demonstrations and refine it by physical interactions of a soft actor-critic agent with the world. Our hybrid formulation allows the dynamical system to utilize high-dimensional observation spaces and cope with noise in demonstrations and sensory observations while maintaining a reactive and robust trajectory-based policy when interacting with dynamic environments. We argue

that maintaining this physically meaningful structure within the reinforcement learning refinement will yield enhanced performance and stability compared to residual corrections or direct learning of desired end-effector velocities. The method is simple, sample efficient and readily applicable in a variety of robotics scenarios. We exemplify this, by using our hybrid model for simulated peg insertion and power lever sliding skills, and a real-world door-opening skill. We demonstrate that SAC-GMM is able to successfully open a door in the real world after half an hour of physical interaction.

The main contributions of this paper are: 1) a hybrid model for learning and refining skills in trajectory distribution space, 2) exploiting high-dimensional sensory inputs obtainable solely during skill adaptation through physical interaction, such as tactile images, gripper camera images, and static camera depth maps to refine parameters of a dynamical system, 3) mitigating amount of robot exploration efforts for learning skills in sparse reward settings through a dynamical system model learned from few demonstrations, and 4) learning to refine two simulated and one real-world robot manipulation skills.

II. RELATED WORK

Learning adaptive skill models has been a long-standing goal in both robotics and machine learning. Our work builds on prior research in this domain which falls into model-based and data-driven paradigms and attempts to bridge the gap between them.

Learning from demonstration [9] is the paradigm in which the robot learns a new skill from demonstrations presented by humans. Dynamical systems [11], [12] use these demonstrations to provide an analytical description of robot motion over time. There are different underlying representations of dynamical systems, ranging from movement primitives [15] to probabilistic approaches [16]. There exists a large body of recent work proposing augmentations of dynamical systems, e.g., fitting of trajectories in a physically consistent fashion [17], ensuring kinematic feasibility of generated trajectories for a mobile robot [18], enabling dynamical systems to generate contact forces [19], integrating non-geometric features into action models [20], preventing mode collapse via mixture density networks [21], using diffeomorphisms to give rise to inherently stable dynamical system [22], and relaxing the point-to-point constraint on demonstrated trajectories [23]. Nonetheless, dynamical systems still suffer from the curse of dimensionality and are ineffective at learning from large or noisy datasets. In contrast, our skill model leverages high-dimensional sensory measurements to adapt its dynamical system and copes with noisy perception and dynamics.

An alternative approach to explicitly modeling robot skills via an analytical framework is to learn an implicit skill model using the robot's interaction data. These data-driven approaches are trained by optimizing directly for skill success. In particular, some recent works have proposed learning robot skills such as grasping [24], pick-and-stow [25], and part discovery [26] first in simulation, where interaction is

cheap and labeled, and then transferring the agent to the real world. Another transfer learning framework adoption is to learn a vision model from passive observations first and then to leverage the learned representations for learning manipulation skill models [27]. Further approaches attempt to learn the real-world dynamics in 2D [28] or 3D [29], and then use this model to perform specific skills such as pushing [30], peg insertion [31], or ball bearing [32]. Recent reinforcement learning (RL) works have proposed learning robot skills using images of the goal [33] or learned representations from unlabeled skill-related videos [34], [35] as the reward function. In general, these high-capacity deep neural networks are not sample-efficient and require a large amount of interaction data to enable a robot to learn complex skills. Moreover, they ignore low-level analytical descriptions of skills and are forced to make individual decisions at each time step. Compared to these approaches, our structured skill model reasons in trajectory space while still optimizing for skill success.

Our approach for learning adaptive skill models falls under a broader category of hybrid models [36], [37]. These approaches combine classical controllers or dynamical system frameworks with RL and learn robot skills with a fair amount of real-world interaction. More specifically, our hybrid skill model (see Fig. 1) leverages both 1) a dynamical system to provide an analytical description of a skill in the trajectory space and 2) a deep reinforcement learning framework to refine the skill via physical interactions and compensate for noisy perception. In contrast to prior work in this domain [38]–[40], our approach learns to adapt the skill in the trajectory space instead of predicting a residual action at each time step. We demonstrate that our approach mitigates the exploration time needed to refine the robot skill drastically. A further hybrid approach has been introduced by Rey *et al.* [41]. In contrast to their work, we use a higher dimensional observation for the reinforcement learning refinement of the skill, while their focus is on mathematical properties of the dynamical systems.

III. PROBLEM FORMULATION

We consider a motion to be driven by a dynamic system governed by a set of first order differential equations defining the dependency of the system velocity on the current state. Following this approach, robot skills are directly defined in trajectory space as a global map that specifies the required motion for reaching a target from arbitrary start poses. We are interested in using this structured model to enable a robot to learn and adapt skills quickly in the real world. Our goal is to first learn skills from a few demonstrations and later refine them through the robot's physical interactions in the world.

For this, we consider robot skills as a set of trajectories in 3D space. We denote a trajectory by $\Xi = \langle \xi_1, \dots, \xi_n \rangle$, where $\xi_i \in \mathbb{R}^d$ is the robot's geometric pose. Moreover, each execution of a robot skill induces a set of observations $\mathcal{O} = \langle \mathbf{o}_1, \dots, \mathbf{o}_n \rangle$ from the environment. Thus, a typical robot skill is described by its 6-DOF (Degree-of-Freedom) trajectory and co-occurring observations such as images,

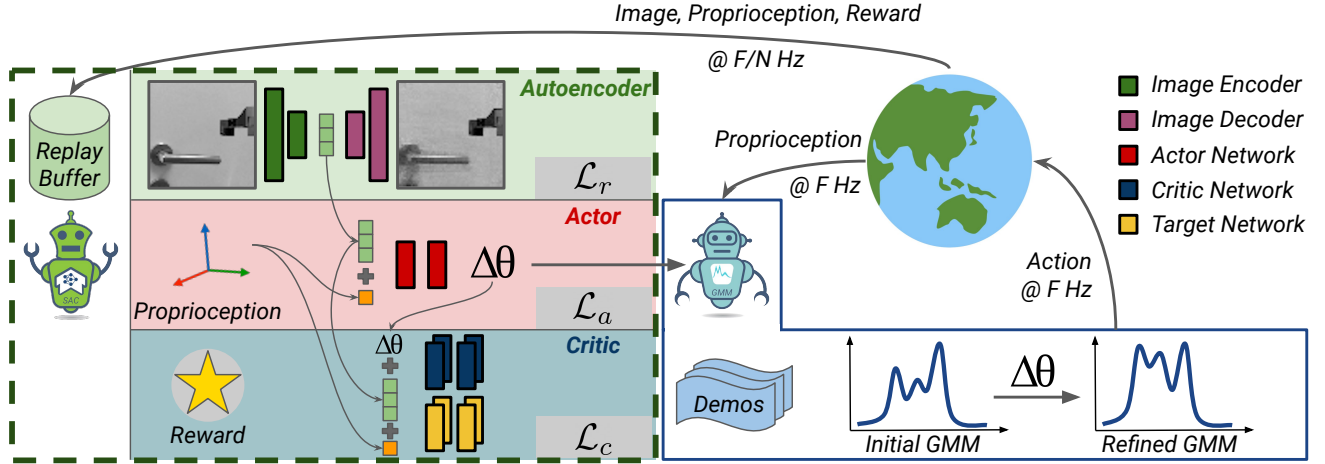


Fig. 2: Structure of SAC-GMM: Given a set of human demonstrations, the GMM agent learns an initial dynamical system parameterization θ that provides an analytical description of the robot's skill trajectory. After N interactions by the GMM agent with the environment (working with frequency F), the SAC agent receives a high-dimensional observation, robot state, and a sparse success reward from the environment. It then refines the initial GMM agent's trajectory parameters by $\Delta\theta$ for the next N interactions to optimize for skill success.

depth maps, or tactile measurements. The observations cannot be controlled directly by the robot but are reactions of the environment to the geometric course of the trajectory. Nevertheless, the robot has to adjust its geometric route in response to these observations.

To learn a robot skill, we require a set of demonstrations $\mathbf{X} = \langle \Xi_1, \dots, \Xi_j \rangle$, where each demonstration is a trajectory that exhibits the geometric route of the skill. Note that these demonstrations can be shown either by a human performing the skill or robot teleoperation. Our goal is to first learn a skill model $f_\theta(\xi) = \dot{\xi}$ that maps each robot pose to its first derivative (velocity), such that it fulfills the following desiderata: (1) it represents the skill in trajectory distribution space, and (2) f is a parametric function. The latter ensures that we can later refine our skill model by adapting its parameters θ . Note that this initial learning step does not include the high-dimensional observations as *i*) these might not be easy to capture in demonstrations and *ii*) they cannot be directly included in the parameterization.

We then aim to refine the previously learned skill directly in the parameterized trajectory space. The skill refinement policy $\pi_\phi(\xi, \mathbf{o}) = \Delta\theta$, builds on the robot pose and the robot's high-dimensional observations while performing the skill in a sparse task completion reward setup.

IV. SAC-GMM

We propose the hybrid approach Soft Actor-Critic Gaussian Mixture Models (SAC-GMM) consisting of two phases. In the first, we learn a dynamical system parameterization in form of a Gaussian mixture model from a few demonstrations. In the second, we refine this dynamical system with the soft actor-critic algorithm through physical interactions with the world. The architecture of our hybrid model is shown in Figure 2.

A. Dynamical System: Gaussian Mixture Model Agent

Dynamical systems afford an analytical representation of a motion's progression over time, and accordingly, they enable

the robot to generate trajectories while being robust in the face of perturbations. We formulate a robot skill as a control law driven by an autonomous dynamical system, defined by the robot pose ξ :

$$\dot{\xi} = f_\theta(\xi) + \epsilon, \quad (1)$$

where f_θ is the robot skill model, a parametric, non-linear, steady, and continuously differentiable function, and ϵ is a zero-mean additive Gaussian noise. From a machine learning perspective, learning the noise-free estimate of f from data is a regression problem and can be addressed by a mixture of Gaussians. Given a set of reference demonstrations \mathbf{X} for a robot skill, we parametrize Eq. (1) through Gaussian Mixture Regression (GMR) [16]. We first estimate the joint probability density $\mathcal{P}(\dot{\xi}, \xi)$ of the robot pose and its corresponding first-order derivative by a Gaussian mixture. Thereby, we parametrize the robot skill model f by $\theta = \{\pi_k, \mu_k, \Sigma_k\}_{k=1}^K$, where π_k are the priors (or mixing weights), μ_k the means and Σ_k the covariances of the k Gaussian functions.

By using this estimated joint probability density function, we employ Gaussian mixture regression (GMR) to retrieve $\dot{\xi}$ given ξ as the conditional distribution $\mathcal{P}(\dot{\xi} | \xi)$. This way our skill model can reproduce the demonstrated skill by estimating the next velocity at the current robot pose and thus generate a trajectory by updating the pose ξ with the generated velocity $\dot{\xi}$ scaled by a time step and proceeding iteratively. For detailed insights on using GMM encoded dynamic system for imitation learning we refer the reader to the extensively available literature [11], [12], [15]–[17].

B. Dynamical System Adaptation: Soft Actor-Critic Agent

Having learned the skill model f_θ , we can now leverage robot interactions with the world to explore and refine the initial model. We formulate this refinement as a reinforcement learning problem in which the agent has to modify the learned skill in the trajectory space and has only access to sparse rewards. In RL, the goal is to learn a policy π_ϕ in a partially observable Markov decision process, consisting of

an observation space \mathcal{O} , a state space \mathcal{S} and an action space \mathcal{A} . In our skill refinement scenario, the agent receives high-dimensional sensory measurements such as RGB images, tactile measurements, or depth maps which are encoded to a latent representation \mathbf{z} by an autoencoder. Together with the robot pose ξ , these form our continuous state space. The action space is also continuous and consists of the desired adaptation in the skill trajectory parameters $\Delta\theta$. Moreover, the environment emits a sparse reward only if the robot executes the skill effectively. Namely, if \mathbf{s}_t , \mathbf{a}_t and \mathbf{z}_t define the robot's state, action, and latent representation of observation respectively at time step t , then

$$\mathbf{s}_t := \{\xi_t, \mathbf{z}_t\}, \quad \mathbf{a}_t := \{\Delta\pi_k, \Delta\mu_k, \Delta\Sigma_k\}_{k=1}^K, \quad (2)$$

and consequently

$$\Delta\theta = \pi_\phi(\mathbf{a}_t | \mathbf{s}_t), \quad (3)$$

where π_ϕ is the robot skill refinement policy. We use $\rho_{\pi_\phi}(\mathbf{s}_t, \mathbf{a}_t)$ to denote the state-action marginal of the trajectory distribution induced by the policy π_ϕ . The robot has to learn this policy from its interactions with the world, such that it maximizes the expected total reward of the refined skill trajectory. Our particular choice for the reinforcement learning framework to learn the skill refinement policy is the soft actor-critic (SAC) algorithm [5]. SAC is an off-policy actor-critic method that, in addition to maximizing the expected total reward, aims to maximize the entropy of the stochastic policy. In doing so, SAC encourages robot exploration and avoids converging to non-optimal deterministic policies. Moreover, SAC exhibits good sample efficiency and stability, necessary ingredients for quick skill adaptation in the real world. Therefore, the objective for an optimal skill refinement policy is

$$J(\pi_\phi) = \sum_{t=0}^T \gamma^t \mathbb{E}_{(\mathbf{s}_t, \mathbf{a}_t) \sim \rho_{\pi_\phi}} [r(\mathbf{s}_t, \mathbf{a}_t) + \alpha \mathcal{H}(\pi_\phi(\cdot | \mathbf{s}_t))],$$

where α is the temperature parameter that determines the relative importance of the entropy term against the reward and regulates the stochasticity of the policy. Our SAC agent stores a collection of $\{\mathbf{s}_t, \mathbf{o}_t, \mathbf{a}_t, r_t, \mathbf{s}_{t+1}, \mathbf{o}_{t+1}\}_{t=1}^T$ transition tuples in a replay buffer \mathcal{D} , and concurrently learns an autoencoder AE_ω , a policy π_ϕ and two Q-functions Q_{ψ_1} and Q_{ψ_2} (to prevent overly optimistic value estimates) and their target networks. More concretely, we use the autoencoder AE_ω to learn a low-dimensional latent representation of the robot's high-dimensional observations. We adopt a similar strategy to the Contractive Autoencoders [42] and enforce an L2 penalty on the learned representation to encourage robust features. For training the autoencoder, actor, and critic networks, SAC randomly samples a batch from the replay buffer and performs stochastic gradient descent on minimizing the following loss objectives for autoencoder, critic, and actor respectively:

$$\begin{aligned} \mathcal{L}_r(\omega, \mathcal{D}) &= \mathbb{E}_{\mathbf{o}_t \sim \mathcal{D}} [\log p_\omega(\mathbf{o}_t | \mathbf{z}_t) + \lambda_z \|\mathbf{z}_t\|^2], \\ \mathcal{L}_c(\psi_i, \mathcal{D}) &= \mathbb{E}_{(\mathbf{s}_t, \mathbf{a}_t) \sim \mathcal{D}} [(Q_{\psi_i}(\mathbf{s}_t, \mathbf{a}_t) - Q^{\text{tar}}(\mathbf{s}_t, \mathbf{a}_t))^2], \\ \mathcal{L}_a(\phi, \mathcal{D}) &= \mathbb{E}_{\substack{\mathbf{s}_t \sim \mathcal{D} \\ \mathbf{a}_t \sim \pi_\phi}} [\alpha \log \pi_\phi(\mathbf{a}_t | \mathbf{s}_t) - \min_{i=1,2} Q_{\psi_i}(\mathbf{s}_t, \mathbf{a}_t)], \end{aligned}$$

where Q^{tar} is the target for the Q functions and is computed using the immediate reward, the value estimate of target Q network and an entropy regularization term.

C. Full Model

Figure 2 shows how our hybrid model learns and refines a robot skill. The GMM agent is fitted on the provided demonstrations and represents a dynamical system, controlling the motion in the trajectory space. After each N interaction steps with the world driven by the GMM encoded dynamics, the SAC agent receives the current state \mathbf{s}_t consisting of the latent observation \mathbf{z}_t and the robot state ξ_t , and additionally a reward r_t for the previous step. It then generates an action $\mathbf{a}_t := \{\Delta\pi_k, \Delta\mu_k, \Delta\Sigma_k\}_{k=1}^K$ according to the current \mathbf{s}_t . Note that with the latent observation \mathbf{z}_t the state \mathbf{s}_t contains more information than the data the original GMM agent was trained on. The SAC agent's action output $\{\Delta\pi_k, \Delta\mu_k, \Delta\Sigma_k\}_{k=1}^K$ is then used to adapt the original GMM for the next N interactions with the environment. The SAC agent optimizes $\pi_\phi(\mathbf{a}_t | \mathbf{s}_t)$ for skill success using ADAM to optimize the autoencoder, critics, and actor networks. The learning rate for the autoencoder is 3×10^{-5} , and for critics and actor is 3×10^{-4} . The observations that the autoencoder receives are 64×64 pixels. In all our experiments, we use $N = 32$ and $K = 3$. While updating GMM parameters, we ensure that mixing weights sum to one and covariance matrices stay positive semi-definite [43]. At inference time, the procedure remains the same. The robot interacts with the environment based on the GMM dynamic system encoding, while after each N steps, the SAC agent receives information about the current state and skill progress and adapts the original GMM agent accordingly.

V. EXPERIMENTAL EVALUATION

We evaluate SAC-GMM for learning robot skills in both simulated and real-world environments. The goals of these experiments are to investigate: (i) whether our hybrid model is effective in performing skills in realistic noisy environments, (ii) if exploiting high-dimensional data boosts the dynamical system adaptation, and (iii) how refining robot skills in trajectory space compares with alternative exploration policies in terms of accuracy and exploration budget.

A. Experimental Setup

We evaluate our approach in both simulated and real-world environments. We investigate two robot skills in simulation: Peg Insertion, and Power Lever Sliding. The environments are simulated with PyBullet and are shown in Figure 3. For the peg insertion skill, we design a cylindrical peg and its corresponding circular hole with 50 mm and 56 mm diameters, respectively. The robot receives a success reward whenever the peg is inserted correctly in the hole. We recognize the importance of the sense of touch in fitting a peg smoothly in a hole. Hence, during the skill adaptation phase, we provide the robot with vision-based tactile sensors [44], [45] on both fingers. For the power lever sliding skill, we design a lever with a relative height of 50 mm to the base. The robot is rewarded whenever it grasps the lever accurately and slides




Fig. 3: Left: Simulated peg insertion scenario; Right: Simulated power lever sliding scenario

it to the end, such that the light is turned on. For this skill, we provide the robot with static depth perception during the skill refinement to facilitate accurate lever pose estimation. In each demonstration as well as training and test episodes, we randomly change the relative position of the robot to the hole and lever. Moreover, to simulate realistic real-world conditions, we also consider noise regarding the detected position of the hole and lever. Namely, in both environments, we add Gaussian noise with a standard deviation of 1 cm in the target pose’s x , y , and z dimensions. For both skills, we collect 20 demonstrations by teleoperating the robot in the simulation.

For the real-world experiment, we investigate a door opening skill, and place our KUKA iiwa manipulator in front of a miniature door, see Figure 4. We attach an ArUco marker on the backside of the door to detect when it was successfully opened and grant a reward accordingly. Moreover, to enable our robot to run autonomously without any human intervention, we equip our door with a door closing mechanism. Thus, the door shuts when the robot releases the door handle and starts a new interaction episode. We provide our robot with an Intel SR300 camera mounted on the gripper for an RGB eye-in-hand view during the skill adaptation phase. We use OpenPose [46] to track the human hand and collect 5 human demonstrations of the door opening skill. Naturally, due to the difference in its geometry, it is not ideal for the robot to exactly reproduce the observed human hand poses and a task dependent offset to correct this difference would be required for ideal performance of the GMM. However, our approach can also cope with this additional source of noise. For the real-world experiments the initial GMM model is fitted using LPV-DS [17] while in simulation we rely on SEDS [11]. While these dynamic system encodings guarantee stability and convergence, our performed modification does not mathematically maintain these guarantees. However, since the adaptations are small, we did not face any issues with diverging or unstable trajectories in unseen starting configurations. All reported failure cases stem from lacking precision of the motions.

B. Evaluation Protocol

We compare our skill model against the following models: **GMM**: This baseline corresponds to the same dynamical system that we learn with the provided demonstrations in the first step of our approach. However, it is not able to explore or leverage high-dimensional observations to refine its performance.




Fig. 4: Real-world door opening. While skill refinement, the robot gets a reward when it opens the door. Wrist-mounted camera provides high-dimensional observations during task execution.

SAC: We employ the soft actor-critic agent [5] to explore and learn the skills. We initialize the replay buffer of the SAC agent with the demonstrations of skills. This baseline employs the same SAC structure that we use in SAC-GMM, including the autoencoder for high-dimensional observations and the network architecture. However, this baseline does not have access to a dynamical system, and consequently, does not reason on the trajectory level.

Res-GMM: Analogous to our approach, this baseline first learns a GMM agent using the demonstrations and then employs a SAC agent for skill refinement. In contrast to our approach, this SAC agent acts at each time step (instead of each N step), and instead of predicting change in trajectory parameters, it predicts a residual velocity which is summed up with the GMM agent’s predicted velocity. This baseline is inspired by recent approaches in the Residual RL domain [38]–[40]. We do not replicate these approaches directly as we want to evaluate the policy refinement and not the choice of underlying dynamical system in which these recent approaches differ.

For the quantitative evaluation of skill models, we employ them to perform the skill 10 times at each evaluation step. We report the average accuracy of each skill model over these episodes in our plots.

C. Experiments in Simulation

We start by evaluating our method on the peg insertion and power lever sliding skills in simulation. Quantitative results of the average success rate of each skill model over 100 trials per five different random seeds are reported in Table I. Our SAC-GMM successfully performs the peg insertion and lever sliding skills with an average final success rate of 86% and 81% respectively. It achieves significantly higher success rates than the GMM baseline, proving its effectiveness in refining the robot skills through physical interaction. Fig. 5a shows the success rate of all skill models during training for

Model \ Task	Peg Insertion	Lever Sliding
GMM	20%	54%
SAC	0%	0%
Res-GMM	30%	62%
SAC-GMM	86%	81%

TABLE I: Our SAC-GMM outperforms baseline skill models significantly in both robot skills by refining the initial dynamical systems (GMM) substantially.




Fig. 5: (a) SAC-GMM vs. baseline models for the peg insertion skill, (b) SAC-GMM exploits high-dimensional data (tactile measurements here) to cope better with noisy perception (c) SAC-GMM learns to open a real door after half an hour of physical interactions.

the peg insertion skill. The SAC baseline which uses skill demonstrations as initialization for its replay buffer, performs poorly and cannot learn the robot skills in the sparse reward setting. Though the Res-GMM skill model which adds a residual velocity to the GMM at each time step achieves a notable success rate of 30%, it is much slower in skill adaptation than SAC-GMM and suffers from a higher variance. This is caused by Res-GMM’s inability to reason about skill refinement on the trajectory level and its obligation to take a residual corrective action at each time step of the trajectory.

To analyze the influence of noise in the perception pipeline and the efficacy of SAC-GMM in exploiting high-dimensional sensory data, we conduct several experiments on the simulated environments (see Table II). We observe that SAC-GMM can fully leverage high-dimensional data such as tactile measurements and depth maps in all scenarios to achieve a superior skill success rate. Fig. 5b shows the success rate of SAC-GMM with and without having access to the tactile sensors [44] during refinement of the peg insertion skill in the noisy setup. We find that SAC-GMM utilizes the high-dimensional observation to deal better with noise and learn the skill faster. We observe that the pure SAC model which does not utilize dynamical systems and accordingly does not reason in the trajectory space cannot learn robot skills in sparse reward settings. We further analyze the SAC skill model in the shaped reward settings. The results in Table III show how SAC-GMM stands out in terms of performance and exploration budget. SAC-GMM outperforms SAC even though it only requires sparse rewards and a third of exploratory interaction steps in the world.

D. Real-World Door Opening

Fig. 5c reports the results for the door opening skill in the real world. We find that, although the initial dynamical system (the GMM agent fitted on human demonstrations) enables the robot to reach the door handle, the robot misses

Model \ Task	Peg Insertion				Lever Sliding			
	No Noise		With Noise		No Noise		With Noise	
	No Tactile	With Tactile	No Tactile	With Tactile	No Depth	With Depth	No Depth	With Depth
GMM	20%	x	10%	x	54%	x	39%	x
SAC	0%	0%	0%	0%	0%	0%	0%	0%
Res-GMM	24%	30%	26%	23%	49%	62%	29%	38%
SAC-GMM	44%	86%	33%	56%	68%	81%	42%	52%

TABLE II: The average success rate of skill models over 100 trials per five different random seeds, under various noise and sensors settings for the simulated robot skills.

Model	Reward	Interactions	Peg Insertion	Lever Sliding
SAC	Dense	60K	60%	30%
SAC-GMM	Sparse	20K	86%	81%

TABLE III: SAC-GMM significantly outperforms SAC, even though SAC has access to shaped rewards and takes three times more exploratory interaction steps in the world.

the proper position to apply its force and can only open the door with a 10% success rate. This failure is due to the robot’s noisy perception and dynamics. Our SAC-GMM exploits the wrist-mounted camera RGB images and sparse door opening rewards and achieves a 90% success rate after only half an hour of physical interactions (~100 episodes) with the door. The SAC baseline fails to learn the skill and the Res-GMM model performs poorly, as adding residual velocities at each time step results in non-smooth trajectories. Videos of these experiments are available at <http://sac-gmm.cs.uni-freiburg.de>.

VI. CONCLUSIONS

In this work, we present “Soft Actor-Critic Gaussian Mixture Models” as a new framework for learning robot skills. This hybrid model leverages reinforcement learning to refine robot skills represented via dynamical systems in their trajectory distribution space and exploits the natural synergy between data-driven and analytical frameworks. Extensive experiments carried out in both simulation and real-world settings, demonstrate that our proposed skill model: 1) learns to refine robot skills through physical interactions in realistic noisy environments, 2) exploits high-dimensional sensory inputs available during skill refinement to cope better with noise, and 3) performs robot skills significantly better than comparable alternatives considering the performance accuracy and exploration costs.

An interesting extension of our work is to investigate transferring skills learned in one environment to new environments, where the visual appearance of the overall setting will differ considerably while the goals of the skills remain similar (e. g., transfer the skill of door opening in one room to other rooms), and refining skills in new environments. While for the initial dynamic system stability and convergence guarantees can be formulated, the modifications on the model performed online by our SAC agent do not formally abide by these guarantees. Even though we did not experience any issues with divergence or instabilities, further adjustments as suggested by Rey *et al.* [41] could be included to guarantee convergence and stability.

REFERENCES

- [1] B. M. Lake, T. D. Ullman, J. B. Tenenbaum, and S. J. Gershman, "Building machines that learn and think like people," *Behavioral and brain sciences*, vol. 40, 2017.
- [2] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in Neural Information Processing Systems*, vol. 29, pp. 4565–4573, 2016.
- [3] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [4] S. Fujimoto, H. Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *Int. Conf. on Machine Learning*, 2018, pp. 1587–1596.
- [5] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *Int. Conf. on Machine Learning*, 2018, pp. 1861–1870.
- [6] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine, "One-shot visual imitation learning via meta-learning," in *Conf. on Robot Learning*, 2017, pp. 357–368.
- [7] A. Zeng, S. Song, K.-T. Yu, E. Donlon, F. R. Hogan, M. Bauza, D. Ma, O. Taylor, M. Liu, E. Romo, *et al.*, "Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching," in *Int. Conf. on Robotics & Automation*, 2018.
- [8] L. Berscheid, P. Meißner, and T. Kröger, "Self-supervised learning for precise pick-and-place without object model," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4828–4835, 2020.
- [9] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, "Survey: Robot programming by demonstration," *Springer Handbook of Robotics*, pp. 1371–1394, 2008.
- [10] S. Schaal, S. Kotosaka, and D. Sternad, "Nonlinear dynamical systems as movement primitives," *International Journal of Humanoid Robotics*.
- [11] S. M. Khansari-Zadeh and A. Billard, "Learning stable nonlinear dynamical systems with gaussian mixture models," *IEEE Transactions on Robotics*, vol. 27, no. 5, pp. 943–957, 2011.
- [12] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, "Dynamical movement primitives: learning attractor models for motor behaviors," *Neural computation*, vol. 25, no. 2, pp. 328–373, 2013.
- [13] J. Silvério, L. Roza, S. Calinon, and D. G. Caldwell, "Learning bimanual end-effector poses from demonstrations using task-parameterized dynamical systems," in *Int. Conf. on Intelligent Robots and Systems*, 2015, pp. 464–470.
- [14] S. Manschitz, M. Gienger, J. Kober, and J. Peters, "Mixture of attractors: A novel movement primitive representation for learning motor skills from demonstrations," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 926–933, 2018.
- [15] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, "Learning and generalization of motor skills by learning from demonstration," in *Int. Conf. on Robotics & Automation*, 2009, pp. 763–768.
- [16] S. Calinon, F. D'halluin, E. L. Sauser, D. G. Caldwell, and A. G. Billard, "Learning and reproduction of gestures by imitation," *IEEE Robotics & Automation Magazine*, vol. 17, no. 2, pp. 44–54, 2010.
- [17] N. Figueroa and A. Billard, "A physically-consistent bayesian non-parametric mixture model for dynamical system learning," in *Conf. on Robot Learning*, 2018, pp. 927–946.
- [18] T. Welschehold, C. Dornhege, F. Paus, T. Asfour, and W. Burgard, "Coupling mobile base and end-effector motion in task space," in *Int. Conf. on Intelligent Robots and Systems*.
- [19] W. Amanhoud, M. Khoramshahi, and A. Billard, "A dynamical system approach to motion and force generation in contact tasks," in *Robotics: Science and Systems*, 2019.
- [20] I. Nematollahi, D. Kuhner, T. Welschehold, and W. Burgard, "Augmenting action model learning by non-geometric features," in *Int. Conf. on Robotics & Automation*, 2019, pp. 7769–7775.
- [21] Y. Zhou, J. Gao, and T. Asfour, "Movement primitive learning and generalization: Using mixture density networks," *IEEE Robotics & Automation Magazine*, vol. 27, no. 2, pp. 22–32, 2020.
- [22] M. A. Rana, A. Li, D. Fox, B. Boots, F. Ramos, and N. Ratliff, "Euclideanizing flows: Diffeomorphic reduction for learning stable dynamical systems," in *Learning for Dynamics and Control*, 2020.
- [23] J. Uraim, M. Ginesi, D. Tateo, and J. Peters, "Imitationflow: Learning deep stable stochastic dynamic systems by normalizing flows," in *Int. Conf. on Intelligent Robots and Systems*, 2020, pp. 5231–5237.
- [24] A. Mousavian, C. Eppner, and D. Fox, "6-dof grasnet: Variational grasp generation for object manipulation," in *Int. Conf. on Computer Vision*, 2019, pp. 2901–2910.
- [25] L. Hermann, M. Argus, A. Eitel, A. Amiranashvili, W. Burgard, and T. Brox, "Adaptive curriculum generation from demonstrations for sim-to-real visuomotor control," in *Int. Conf. on Robotics & Automation*, 2020, pp. 6498–6505.
- [26] S. Y. Gadre, K. Ehsani, and S. Song, "Act the part: Learning interaction strategies for articulated object part discovery," *arXiv preprint arXiv:2105.01047*, 2021.
- [27] L. Yen-Chen, A. Zeng, S. Song, P. Isola, and T.-Y. Lin, "Learning to see before learning to act: Visual pre-training for manipulation," in *Int. Conf. on Robotics & Automation*, 2020, pp. 7286–7293.
- [28] C. Finn, I. Goodfellow, and S. Levine, "Unsupervised learning for physical interaction through video prediction," *Advances in Neural Information Processing Systems*, vol. 29, pp. 64–72, 2016.
- [29] I. Nematollahi, O. Mees, L. Hermann, and W. Burgard, "Hindsight for foresight: Unsupervised structured dynamics models from physical interaction," in *Int. Conf. on Intelligent Robots and Systems*, 2020.
- [30] A. Kloss, M. Bauza, J. Wu, J. B. Tenenbaum, A. Rodriguez, and J. Bohg, "Accurate vision-based manipulation through contact reasoning," in *Int. Conf. on Robotics & Automation*, 2020, pp. 6738–6744.
- [31] M. A. Lee, Y. Zhu, P. Zachares, M. Tan, K. Srinivasan, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, "Making sense of vision and touch: Learning multimodal representations for contact-rich tasks," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 582–596, 2020.
- [32] S. Tian, F. Ebert, D. Jayaraman, M. Mudigonda, C. Finn, R. Calandra, and S. Levine, "Manipulation by feel: Touch-based control with deep predictive models," in *Int. Conf. on Robotics & Automation*, 2019.
- [33] A. Singh, L. Yang, K. Hartikainen, C. Finn, and S. Levine, "End-to-end robotic reinforcement learning without reward engineering," *arXiv preprint arXiv:1904.07854*, 2019.
- [34] P. Sermanet, C. Lynch, Y. Chebotar, J. Hsu, E. Jang, S. Schaal, S. Levine, and G. Brain, "Time-contrastive networks: Self-supervised learning from video," in *Int. Conf. on Robotics & Automation*, 2018.
- [35] O. Mees, M. Merklinger, G. Kalweit, and W. Burgard, "Adversarial skill networks: Unsupervised robot skill learning from video," in *Int. Conf. on Robotics & Automation*, 2020, pp. 4188–4194.
- [36] S. Bahl, M. Mukadam, A. Gupta, and D. Pathak, "Neural dynamic policies for end-to-end sensorimotor learning," in *Advances in Neural Information Processing Systems*, 2020.
- [37] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser, "Tossing-bot: Learning to throw arbitrary objects with residual physics," *IEEE Transactions on Robotics*, vol. 36, no. 4, pp. 1307–1319, 2020.
- [38] G. Schoettler, A. Nair, J. Luo, S. Bahl, J. A. Ojea, E. Solowjow, and S. Levine, "Deep reinforcement learning for industrial insertion tasks with visual inputs and natural rewards," *arXiv preprint arXiv:1906.05841*, 2019.
- [39] T. Davchev, K. S. Luck, M. Burke, F. Meier, S. Schaal, and S. Ramamoorthy, "Residual learning from demonstration: adapting dynamic movement primitives for contact-rich insertion tasks," *arXiv preprint arXiv:2008.07682*, 2020.
- [40] A. Ranjbar, N. A. Vien, H. Ziesche, J. Boedecker, and G. Neumann, "Residual feedback learning for contact-rich manipulation tasks with uncertainty," *arXiv preprint arXiv:2106.04306*, 2021.
- [41] J. Rey, K. Kronander, F. Farshidian, J. Buchli, and A. Billard, "Learning motions from demonstrations and rewards with time-invariant dynamical systems based policies," *Autonomous Robots*, vol. 42, no. 1, pp. 45–64, 2017.
- [42] S. Rifai, P. Vincent, X. Muller, X. Glorot, and Y. Bengio, "Contractive auto-encoders: Explicit invariance during feature extraction," in *Int. Conf. on Machine Learning*, 2011.
- [43] N. J. Higham, "Computing a nearest symmetric positive semidefinite matrix," *Linear algebra and its applications*, 1988.
- [44] S. Wang, M. Lambeta, P.-W. Chou, and R. Calandra, "Tacto: A fast, flexible and open-source simulator for high-resolution vision-based tactile sensors," *arXiv preprint arXiv:2012.08456*, 2020.
- [45] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, *et al.*, "Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 3838–3845, 2020.
- [46] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh, "Openpose: realtime multi-person 2d pose estimation using part affinity fields," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 43, no. 1, pp. 172–186, 2019.