

Curriculum-Based Imitation of Versatile Skills

Maximilian Xiling Li^{*†}, Onur Celik[†], Philipp Becker[†], Denis Blessing[†], Rudolf Lioutikov^{*}, Gerhard Neumann[†]

^{*} Intuitive Robots Lab, Karlsruhe Institute of Technology, Karlsruhe, Germany

[†] Autonomous Learning Robots, Karlsruhe Institute of Technology, Karlsruhe, Germany

Abstract—Learning skills by imitation is a promising concept for the intuitive teaching of robots. A common way to learn such skills is to learn a parametric model by maximizing the likelihood given the demonstrations. Yet, human demonstrations are often multi-modal, i.e., the same task is solved in multiple ways which is a major challenge for most imitation learning methods that are based on such a maximum likelihood (ML) objective. The ML objective forces the model to cover all data, it prevents specialization in the context space and can cause mode-averaging in the behavior space, leading to suboptimal or potentially catastrophic behavior. Here, we alleviate those issues by introducing a curriculum using a weight for each data point, allowing the model to specialize on data it can represent while incentivizing it to cover as much data as possible by an entropy bonus. We extend our algorithm to a Mixture of (linear) Experts (MoE) such that the single components can specialize on local context regions, while the MoE covers all data points. We evaluate our approach in complex simulated and real robot control tasks and show it learns from versatile human demonstrations and significantly outperforms current SOTA methods.¹

Index Terms—Imitation, Versatility, Curriculum Learning

I. INTRODUCTION

To increase the accessibility of robotic systems, we need intuitive ways to teach them new skills. One promising and well-studied approach to this problem is imitation learning. In this setting, the robot learns new skills from demonstrations provided by a domain expert, e.g., through teleoperation. Most imitation learning approaches then fit a simple probabilistic model to those demonstrations allowing them to copy the expert’s behavior by sampling from the model.

Learning a skill model from real world demonstrations commonly suffers from three significant problems, **a)** *outlier-sensitivity*, **b)** *locality-violations* in context space and **c)** *mode-averaging* in behavior space. Consider kinesthetic demonstrations of a robot hitting a table tennis ball. Commonly, some demonstrations will run into joint limits, however, these demonstrations might be few and significantly different from other demonstrations. In common approaches, such demonstrations have to be filtered out manually a priori or the learned model will include such degenerating demonstrations into its behavior, suffering from *outlier-sensitivity* [1]. Furthermore, the context points in such an example would cover a large area of the table. Covering disjoint or excessively large areas of the context space generally leads to bad behavior, especially if the skill parametrization is linear in the context, which is

a common assumption due to its computational simplicity. Avoiding this *locality-violation* allows each skill to specialize on a local subset of the contexts-space, performing better in that particular region than an overly general skill. Finally, significantly different strokes that return the ball to the same point are an example of the multi-modality generally found in real-world demonstrations [2]. Common approaches combine these different modes into one interpolated, suboptimal behavior that suffers from *mode-averaging*. If the number of components is less than the unknown number of modes present in the demonstrations, this problem even exists for mixture models that are optimized using common approaches such as EM and other likelihood-based methods.

We propose a novel approach to learning a Mixture of Experts (MoE) models that avoid these three problems. To this end, we introduce **a)** per-component weights on each data point that allow each component to select its own curriculum, i.e., the subset of data points it considers for its behavior. This weighting explicitly allows each component to ignore unfavored data points, avoiding *outlier-sensitivity*. **b)** a per-component Gaussian context distribution that prevents *locality-violations* in the context space and allows each component to learn specialized behavior for a region of the context space. And **c)** an optimization objective for updating the weights that prevent each component from *mode-averaging*, by learning the weights that optimize the likelihood and the context distribution simultaneously, while still incentivizing a larger coverage of data points through an entropy bonus.

We show our approaches’ advantages in three experiments and a detailed ablation study. For the experiments, we use a simple toy example, a sophisticated table-tennis simulation, and a highly multi-modal, contextual obstacle avoidance task using a Franka Emika Panda 7DoF Manipulator, as shown in Fig. 1. For the real robot task, we use versatile human demonstrations that exhibit the above-mentioned challenging properties. We compare our method to several baselines, i.e., Expectation Maximization [3], Expected Information Maximization [4], K-Nearest Neighbours, Normalizing Autoregressive Flows [5], and a Mixture Density Network-based approach [6]. We outperform the baselines in all tasks.

II. PRELIMINARIES

We consider a given set of demonstrations $\mathcal{T} = \{(\boldsymbol{\tau}_i, \mathbf{c}_i)\}_{i=1}^N$. A trajectory demonstration $\boldsymbol{\tau}_i = (\mathbf{s}_0, \mathbf{s}_1, \dots, \mathbf{s}_{T_i})$ consists of the recorded states, e.g., joint configurations, \mathbf{s}_t of the robot and the context \mathbf{c}_i defines task properties required

[†]Correspondence to maximilian.li@kit.edu

¹A reference implementation can be found at <https://github.com/intuitive-robots/ml-cur>




Fig. 1: **Left:** Virtual Table Tennis environment. The ball is launched above the red dot towards the robot and is supposed to hit on the green dot on return. **Center:** Data Collection for the obstacle avoidance task by a human demonstrator. **Right:** Virtual twin for the obstacle avoidance task. For each demonstration the obstacles, serving as contexts, are randomly placed.

to contextualize a given demonstration. The context vector \mathbf{c} encodes environmental states and desired goal conditions, such as obstacle or target positions. In contrast to prior work [7, 8, 9] the context vector is not required to be in the sensorimotor space of the robot nor does it specify a new coordinate system. The demonstration length T_i is different for each τ_i . We represent skills as movement primitives [10], specifically Probabilistic Movement Primitives (ProMPs) [11]. ProMPs use time-dependent radial basis functions Φ , that are centered at different points in time and linearly combined with a parameter vector ω to represent trajectories. This compactly represents each expert trajectory as a parameter vector ω which is fitted to the demonstration using linear regression

$$\omega_i = \left(\Phi^T \Phi + \lambda \mathbf{I} \right)^{-1} \Phi^T \tau_i. \quad (1)$$

Exact details are shown in [11]. Together with the context for each demonstration, we obtain a data set $\mathcal{D} = \{\mathbf{c}_i, \omega_i\}_{i=1}^N$ which we can use to fit our model. To capture the versatility, we employ a Mixture of Expert (MoE) model, given as

$$p(\omega|\mathbf{c}) = \sum_o p(\omega|\mathbf{c}, o)p(o|\mathbf{c}). \quad (2)$$

Here $p(o|\mathbf{c})$ is the softmax gating distribution and $p(\omega|\mathbf{c}, o)$ are linear conditional Gaussian experts.

III. RELATED WORK

Imitation Learning for Versatile Skills. In the simplest case, learning skills as movement primitives [10, 11] from human demonstrations amounts to regression. To capture versatility in the behavior we need to learn multi-modal models, e.g., mixture models. Common approaches that learn GMMs over MP parameters via EM, only use the resulting mixture models to find suitable MP parameters for novel contexts [12, 13]. Neither of these approaches capture versatile behavior for individual contexts. Other approaches learn a mixture over contextualized MPs using Gaussian Mixture Regression [14] and learn non-linear relations between context and MP parameters by using a mixture density network [6]. All of the above-mentioned approaches use maximum likelihood objectives and thus suffer from the mode-averaging problem. Expected Information Maximization (EIM)[4] computes the

information projection [15] based on samples that can be applied to skill-based imitation learning [16]. Yet, EIM relies on an intermediate density ratio estimation step [17], making the approach inefficient and hard to tune. EIM further suffers from *locality-violations*. Recent approaches use generative deep learning models [7, 18] or apply preference learning using additional human rankings [19]. However, these approaches do not leverage skill-based representations but instead operate on entire trajectories of state-action pairs.

Curriculum Learning. While curricula are a well-studied technique to improve the performance of supervised [20, 21] and reinforcement learning [22] approaches, we are not aware of any approaches using curricula for skill-based imitation learning. Our approach is inspired by a recent approach to versatile skill discovery in reinforcement learning [23]. They employ the same MoE parametrization, as a basis for a curriculum approach to versatile maximum entropy reinforcement learning. Yet, as they consider a fundamentally different setting they use a different formulation of the curriculum. Further, they only addressed *locality-violations* as reinforcement learning objectives naturally do not suffer from *mode-averaging*.

IV. CURRICULUM-BASED IMITATION LEARNING

Consider a given set of demonstrations $\mathcal{T} = \{(\tau_i, \mathbf{c}_i)\}_{i=1}^N$, where each tuple consists of a demonstrated trajectory τ_i and a corresponding context variable \mathbf{c}_i , e.g., a joint trajectory of a robot hitting a table tennis ball and the balls subsequent point of impact on the table. Our goal is now to learn a parameterized, generative contextual skill model $p_{\Theta}(\omega|\mathbf{c})$ over a motion primitive weighting ω given a particular context \mathbf{c} under the parameters Θ , e.g., predict the weights of a primitive that hits a table tennis ball such that it bounces of the table at a particular position. In this work we use the ProMP representation and, hence, apply the weight projection in Eq. 1 to transform the dataset $\mathcal{T} = \{(\tau_i, \mathbf{c}_i)\}_{i=1}^N$ into the corresponding dataset $\mathcal{D} = \{\omega_i, \mathbf{c}_i\}_{i=1}^N$. Note that such a projection is not unique to ProMPs and other primitive representations could be used just as well. We first motivate our objective for a single component in and then extend it to the Mixture of Experts (MoE) case. We call the resulting method *Maximum-Likelihood based Curriculum (ML-Cur)*.

Curriculum Imitation Learning for a Single Expert. An intuitive objective for fitting a model is to maximize the log-likelihood $\mathcal{L}(\boldsymbol{\theta}|\mathcal{D}) = \sum_i \log p_{\boldsymbol{\theta}}(\boldsymbol{\omega}_i|\mathbf{c}_i)$, with $\boldsymbol{\theta}$ being the parametrization of a single component, e.g., the mean and covariance of a ProMP weight distribution. While easy to optimize, this approach suffers from the discussed shortcomings *outlier-sensitivity*, *locality-violation* and *mode-averaging*. By introducing a weight ν_i for each demonstration the component is able to only consider a relevant subset of the given demonstrations, i.e. its curriculum, allowing the component to ignore outliers or data points from other modes. The weights ν sum up to 1 and are strictly positive. We further introduce a Gaussian context distribution $p_{\boldsymbol{\eta}}(\mathbf{c})$ restricting the component to a local subspace of the context and enables the component to learn more specialized behavior, addressing the *locality-violation* problem. The resulting objective

$$\arg \max_{\nu_{1:N}, \boldsymbol{\eta}, \boldsymbol{\theta}} \sum_i \nu_i (\log p_{\boldsymbol{\theta}}(\boldsymbol{\omega}_i|\mathbf{c}_i) + \log p_{\boldsymbol{\eta}}(\mathbf{c}_i)), \quad (3)$$

ensures that the maximum likelihood updates for $p_{\boldsymbol{\eta}}(\mathbf{c})$ and $p_{\boldsymbol{\theta}}(\boldsymbol{\omega}|\mathbf{c})$ are **i)** coupled, **ii)** only consider data points that $p(\boldsymbol{\omega}|\mathbf{c})$ can represent and **iii)** the considered data points are local in the context space. However, the objective in Eq. 3 would encourage greedy components with increasingly smaller curricula. We counter this behavior by introducing an entropy bonus over the data weights $\mathcal{H}[\boldsymbol{\nu}] = -\sum_i \nu_i \log \nu_i$, yielding the single component objective

$$\arg \max_{\nu_{1:N}, \boldsymbol{\eta}, \boldsymbol{\theta}} \sum_i \nu_i (\log p_{\boldsymbol{\theta}}(\boldsymbol{\omega}_i|\mathbf{c}_i) + \log p_{\boldsymbol{\eta}}(\mathbf{c}_i)) + \alpha \mathcal{H}(\boldsymbol{\nu}), \quad (4)$$

with scaling factor α . This objective yields an exciting new type of learning process which forms the core of our proposed method. At the beginning the component will focus on a small curriculum, i.e., a small set of data points, since the log-likelihood will dominate the objective. With increasing learning process the likelihood of more data points will increase and the entropy term will force the component to extend its curriculum to more data points until the ‘representational capacity’ of the component is reached and further data points can not be represented well any more. At the same time the context distribution ensures the locality of the curriculum in the context space.

Objective 4 can now be iteratively optimized for $p_{\boldsymbol{\theta}}(\boldsymbol{\omega}|\mathbf{c})$ and $p_{\boldsymbol{\eta}}(\mathbf{c})$ and subsequently for the data weights $\nu_{1:N}$. The updates for $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$ follow common derivations for MoE models. However, the closed form solution for the data weights

$$\nu_i^* \propto (p_{\boldsymbol{\theta}}(\boldsymbol{\omega}_i|\mathbf{c}_i) p_{\boldsymbol{\eta}}(\mathbf{c}_i))^{\frac{1}{\alpha}} \quad (5)$$

offers the valuable insight that the weight ν_i is high iff $p_{\boldsymbol{\theta}}(\boldsymbol{\omega}_i|\mathbf{c}_i) p_{\boldsymbol{\eta}}(\mathbf{c}_i)$ is high, allowing the component to adjust its curriculum by specializing on favored $(\boldsymbol{\omega}_i, \mathbf{c}_i)$ tuples. Note that this specialization is only possible as we iteratively optimize the individual terms. This weight update in combination with the locality provided by the Gaussian context distribution

further ensures that the weighted maximum likelihood updates for $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$ are not subject to mode averaging.

The scalar factor α in Eq. 4 trades off the importance of the likelihood terms and the entropy of the weight distribution. For high α values the updates for $p(\mathbf{c})$ and $p(\boldsymbol{\omega}|\mathbf{c})$ will approach the standard maximum likelihood update as the weights will get more and more uniform. In contrast, a too-small α can cause the expert to concentrate on a single data point only. Due to highly varying ranges of the log likelihoods and the entropy values during training, choosing a good value is not straightforward. We therefore additionally propose an extension that automatically tunes α .

Curriculum-Based Imitation Learning for MoE. We propose a multi-modal model that allows individual components to **a)** specialize on local context regions while avoiding spanning over the whole context space, **b)** ignore modes it cannot represent, and **c)** still cover as much as possible of the available data. Following [23], we define our mixture model as

$$p_{\Theta}(\boldsymbol{\omega}|\mathbf{c}) = \sum_o \frac{p_{\boldsymbol{\eta}}(\mathbf{c}|o) p_{\boldsymbol{\lambda}}(o)}{p(\mathbf{c})} p_{\boldsymbol{\theta}}(\boldsymbol{\omega}|\mathbf{c}, o), \quad (6)$$

where $p_{\boldsymbol{\theta}}(\boldsymbol{\omega}|\mathbf{c}, o)$ and $p_{\boldsymbol{\eta}}(\mathbf{c}|o)$ denote the component and the context distribution analogous to above, which are now additionally conditioned on the respective component. The K -variate categorical distribution $p_{\boldsymbol{\lambda}}(o)$ models a gating prior over the K components and $p(\mathbf{c}) = \sum_o p_{\boldsymbol{\eta}}(\mathbf{c}|o) p_{\boldsymbol{\lambda}}(o)$ denotes a prior over the context. The model in Eq. 6 differs from standard mixture of experts (Eq. 2) by modelling the component-wise context distribution $p_{\boldsymbol{\eta}}(\mathbf{c}|o)$ and $p_{\boldsymbol{\lambda}}(o)$ instead of the gating distribution $p(o|\mathbf{c})$. This formulation defines a mixture model where each component has its own curriculum while focusing on a local context region.

A naive approach to update the MoE model is to use the objective in Eq. 4 and let each component specialize independently on a different part of the data set. In this case however, components might adjust their curriculum similarly and thus, concentrate on the same data regions, as they are not aware on which regions the other components specialize. To allow each components to adjust its curriculum independently while coordinating the curriculum with other components, we define individual data weights $\nu_{o,i}$ for each component, yielding the per-datapoint weight $\nu_i = \frac{1}{K} \sum_o \nu_{o,i}$. In combination with the entropy bonus in Eq. 4 such weights force the components to focus on different parts of the context space. Following the described intuitions, we extend our objective (Eq. 4) to

$$\arg \max_{\Theta} \sum_i \sum_{o, \boldsymbol{\theta}, \boldsymbol{\eta}} \nu_{o,i} (\log p_{\boldsymbol{\theta}}(\boldsymbol{\omega}_i|\mathbf{c}_i, o) + \log p_{\boldsymbol{\eta}}(\mathbf{c}_i|o) + \log p_{\boldsymbol{\lambda}}(o)) + \alpha \mathcal{H}(\boldsymbol{\nu}), \quad (7)$$

with $\mathcal{H}(\boldsymbol{\nu}) = -\sum_i \frac{1}{K} \sum_o \nu_{o,i} \log \frac{1}{K} \sum_o \nu_{o,i}$ being the entropy over all per-component data weights and $\Theta = \{\nu_{1:K, 1:N}, \boldsymbol{\eta}_{1:K}, \boldsymbol{\theta}_{1:K}, \boldsymbol{\lambda}\}$ being the learned parameters of the model. This objective is very similar to the one in Eq. (4), but with the key distinction that the likelihood and weight terms are now conditioned on the mixture variable o .

Optimizing the MoE. As described and motivated above, we now seek to maximize Objective 7 w.r.t. $\nu_{1:K,1:N}, \eta_{1:K}$ and $\theta_{1:K}$. This optimization can be easily done with weighted maximum likelihood updates. However, the optimization w.r.t. $\nu_{o,i}$ is difficult since the sum over o appears inside the log of the entropy term, rendering a closed form solution for $\nu_{o,i}$ infeasible. We tackle this problem by introducing a (tight) variational lower bound similar to [24]

$$\arg \max_{\Theta} \sum_i \sum_{o, \theta, \eta} \nu_{o,i} (\log p_{\theta}(\omega_i | \mathbf{c}_i, o) + \log p_{\eta}(\mathbf{c}_i | o) + \log p_{\lambda}(o)) + \alpha \mathcal{H}(\nu_{o,i}), \quad (8)$$

with $\mathcal{H}(\nu_{o,i}) = -\frac{1}{K} \sum_i \sum_o \nu_{o,i} (\log \nu_{o,i} - \log \tilde{\nu}_{o|i})$ and $\tilde{\nu}_{o|i} = \nu_{o,i}^{old} / \sum_o \nu_{o,i}^{old}$ being the per component data entropy and component data responsibility respectively. Intuitively, the log of the data responsibility returns a high value, if only one component o is responsible for data point i , encouraging the components to divide the data points among each other and thus avoid overlapping support. Note that the old weight $\nu_{o,i}^{old}$ is simply fixed to the weight of the previous iteration. The lower bound in Eq. 8 does not have any effect on the weighted maximum likelihood update rules for $\eta_{1:K}, \theta_{1:K}$ and λ , but yields a closed-form solution for the component weight update

$$\nu_{o,i} \propto (p_{\theta}(\omega_i | \mathbf{c}_i, o) p_{\eta}(\mathbf{c}_i | o) p_{\lambda}(o))^{K/\alpha} \tilde{\nu}_{o|i}, \quad (9)$$

recovering the solution from Eq. 5 adjusted to mixture models.

Autotuned Entropy Scaling. Choosing the entropy scaling parameter α is not straightforward and might yield unsatisfying results, caused by highly varying ranges of the log likelihoods and the entropy values during training. To stabilize the optimization for $\nu_{o,i}$, we propose to reformulate our objective in Eq. 8 as a constrained optimization problem

$$\begin{aligned} \arg \max_{\Theta} \sum_i \sum_{o, \theta, \eta} \nu_{o,i} (\log p_{\theta}(\omega_i | \mathbf{c}_i, o) + \log p_{\eta}(\mathbf{c}_i | o) + \log p_{\lambda}(o) + \alpha_o \log \tilde{\nu}_{o|i}) \\ \text{s.t. } \mathcal{H}(\nu_{o,i}) \geq \mathcal{H}_{\min}, \end{aligned} \quad (10)$$

where we now consider a scaling α_o and a minimal entropy \mathcal{H}_{\min} per component o . Choosing \mathcal{H}_{\min} is easier and more intuitive compared to choosing the entropy scaling factor α as \mathcal{H}_{\min} ensures that each component covers a certain amount of samples, consequently preventing components from collapsing to a single sample. A good choice is to set $\mathcal{H}_{\min} = \log n_{\text{eff}}$, where n_{eff} is the desired number of effective samples per component. Furthermore, we are not restricted to choose one single α for all components anymore. The Lagrangian dual optimization yields the optimal solution for each weight

$$\nu_{o,i}^* \propto (p_{\theta}(\omega_i | \mathbf{c}_i, o) p_{\eta}(\mathbf{c}_i | o) p_{\lambda}(o))^{\frac{1}{\alpha_o}} \tilde{\nu}_{o|i},$$

with α_o being obtained by minimizing the dual function

$$\arg \min_{\alpha_o \geq 0} \alpha_o \left(\log \sum_i (p_{\theta}(\omega_i | \mathbf{c}_i, o) p_{\eta}(\mathbf{c}_i | o) p_{\lambda}(o))^{\frac{1}{\alpha_o}} \tilde{\nu}_{o|i} - \mathcal{H}_{\min} \right).$$

We optimize the convex dual using L-BFGS[25].




Fig. 2: **Planar Robot Reacher.** Trajectories in our planar reacher toy task. The target positions are marked in purple. Black lines represent the links of a 10DoF robotic arm with the end-effector highlighted in red. The environment obstacles are visualized as red boxes. The demonstrations used for training (left) show three ways of reaching the target. We fit MoEs with 2 components with both EM (center) and ML-Cur (right). ML-Cur avoids the mode-averaging problem that EM suffers from.

V. EXPERIMENTS

We show the advantages and strengths of our method on three experimental settings, including real-world robots and data. The baselines include Mixture of Expert (MoE) (Eq. 2) trained with both Expectation Maximization (EM) [3] and Expectation Information Maximization (EIM) [4]. While EM suffers from *outlier-sensitivity*, *locality-violations* and *mode-averaging*, EIM addresses *mode-averaging* yet still suffers from *locality-violations*. EM and ML-Cur only have a small number of hyper-parameters. EIM relies on an intermediate density ratio estimation [17], making the approach inefficient and hard to tune. We additionally compare to Mixture Density Networks (MDN)-based approach [6], Neural Autoregressive Flows (NFlow) [5] and K-nearest neighbor (KNN) [26].

Planar Robot Reacher. First, we consider a synthetic 2D robotic reaching task for illustrative purposes. A 10DoF planar robot must reach a specified target position while avoiding three obstacles. The task is contextualized by the 2D target position. The action space is a ten-dimensional joint angle position vector. Using 1780 collision-free training demonstrations, we fit MoE models with two components using EM and ML-Cur. As shown in Fig. 2 EM is unable to fit the training demonstrations as the maximum likelihood objective forces the model to average over the middle and lower mode in the data, which does not only prevents the robot from reaching the targets but also makes it crash into the obstacles. ML-Cur on the other hand simply ignores the mode with the least amount of samples and yields a successful and safe model.

Table Tennis. Next, we consider a virtual table tennis task modeled in MuJoCo [27]. We collect a data set using a highly versatile reinforcement learning agent [23] and randomly split it into 5,000 training and 500 test samples. A ball is launched from different start positions, and after bouncing on the robot's table side, the robot has to return the ball to a given goal position. The expert demonstrations entail varying movements, such as forehand and backhand strokes.

The context is given as a 4D vector, defining the XY ball launch position and the XY target position on the opponent's

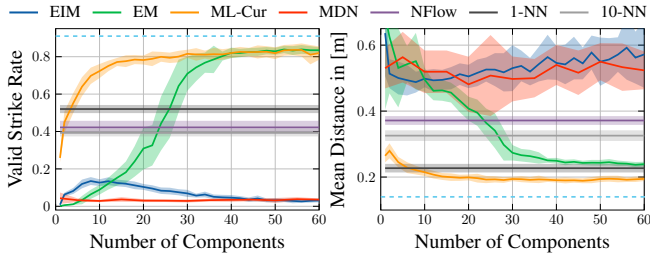


Fig. 3: **Table Tennis.** NFlow and KNN with $k = 1$ and $k = 10$ appear as straight lines, as these models don't have components. All models have been tested with 20 different random seeds and training-test selections of the available data. ML-Cur outperforms the baselines in both the valid strike rate and MDE. It most closely approaches the performance of the demonstrator, shown as a light blue dashed line.




Fig. 4: **Table Tennis.** Model performance with 15 components, for an increasing number of minimum effective samples over ten random seeds and training-test selections. For small numbers of effective samples, the model specializes each component to very few training samples. The performance suffers as the context space is not appropriately covered. The dashed line indicates the point where the minimum number of components times the number of effective samples per component becomes larger than the total sample size. Approaching this threshold leads to the discussed degenerating behavior.

side. We evaluated two metrics, the *valid strike rate*, i.e., how often the ball is successfully returned to the other side, and the *mean distance error (MDE)* to the target position for valid strikes. Fig. 3 shows the metrics for ML-Cur, EM, EIM, MDN, NFlow, KNN with one (1-NN) and ten (10-NN) nearest neighbors for a varying number of model components. ML-Cur generally outperforms EM, especially with lower numbers of components, while EIM and MDN do not seem to be able to solve the task as indicated by the low valid strike rate. Nflow, 1-NN, and 10-NN have a higher valid strike rate for a small number of components, but EM and especially ML-Cur outperform them with more components.

Additionally, for ML-Cur, we investigate the effect of changing the minimum number of effective samples to be covered by each component. We fix the number of MoE components at 15 and present the results in Fig. 4. For small numbers of effective samples, the model is free to specialize each component to only very few training samples. Thus, the performance suffers as the context space is not appropriately covered. For a large number of effective samples the curriculum weights become more uniform, leading to *outlier-*




Fig. 5: **Obstacle Avoidance.** Distribution of test errors over 20 trials. MDN and NFlow can not solve the task. KNN ($k = 1$) has the lowest collision rate, but the highest distance error. ML-Cur has the lowest distance error and second lowest collision rate, outperforming the other models.

sensitivity, locality-violations and even *mode-averaging*. Thus the MoE components are unable to specialize on samples and the overall model performance suffers.

Obstacle Avoidance. We conducted a real-world robot experiment with a Franka Emika Panda 7DoF Manipulator. We collect versatile human expert demonstrations using teleoperation with a virtual twin: Moving a physical robot by hand and mirroring its state in a MuJoCo [27]-based environment. We place three static, cylindrical objects in the simulation environment and move the robot end-effector around these obstacles to the desired target position. The task is contextualized by a 4D vector, describing the three obstacles and the goal's Y-axis position. The X-coordinate is fixed for all four objects. We train MoE models with twelve components using ML-Cur, EM, and EIM. As additional baselines, we include MDN, NFlow, and KNN with one (1-NN) and ten (10-NN) nearest neighbors. The models are trained on 228 collision-free demonstrations and tested on 50 random test contexts using MuJoCo. We repeated this experiment 20 times with different random seeds and train-test selections and show the models' error distribution on the test data in Figure 5. As metrics, we consider the rate of generated trajectories, which collide with at least one obstacle, and the mean distance error (MDE) of the collision-free trajectories to the target position. MDN is unable to generate valid trajectories, resulting in the lowest collision rate, but the highest distance to the goal position. KNN with $k = 1$ has a lower collision rate than ML-Cur, but slightly higher distance error. ML-Cur outperforms the remaining baselines.

Ablation Studies. We perform several ablation studies in order to show the relative importance of the different terms of the ML-Cur objective. To that end, we use the table tennis task with 20 components for a quantitative analysis and the planar reacher task for a qualitative analysis.

Data Weight Ablation. We compare ML-Cur as proposed with a version which is optimized ignoring the data weights $\nu_{o,i}$ as introduced in Eq. 8. Consequently, we treat every data point as equally important.

Ignoring the data weights lead to equal updates for all components and thus reduces the mixture of experts model




Fig. 6: The planar reacher configurations at the means of the context distributions. The six Gaussian context distributions are shown as colored uncertainty ellipses. 6a shows the result for ML-Cur. 6b demonstrates that removing the data weights reduces ML-Cur to a single expert model, averaging over the whole context space. 6c indicates that violating locality results in context components covering larger areas. 6d shows that removing the responsibilities from the optimization results in heavily overlapping components, yielding less versatile solutions.

TABLE I: We ablate key aspects of our objective on the table tennis task. As the ablation study shows all the aspects contribute to the superior performance of ML-Cur.

	Valid Strikes (\uparrow)	Distance Error (\downarrow)
ML-Cur	0.791 ± 0.034	0.201 ± 0.007
No Data Weights	0.002 ± 0.001	0.62 ± 0.218
With Locality Violation	0.653 ± 0.073	0.209 ± 0.013
Without Responsibilities	0.667 ± 0.077	0.216 ± 0.012

to a single component. As a result, the performance degrades drastically on all experiments as shown in Table I. In Fig. 6b we visualize how the model behaves on the planar reaching task. Due to the lack of the weights, it can not concentrate on specific data points and hence can not adjust its curriculum which lead to a poor performance.

Locality Violation Ablation. Next, we compare ML-Cur with a version which does not make use of the Gaussian context components $p_{\eta}(\mathbf{c}_i|o)$ during training. Without $p_{\eta}(\mathbf{c}_i|o)$ each expert might cover an arbitrary large or disjoint regions in the context space. After training, we fit the Gaussian context components in order to be able to use the mixture model for inference. As a result, mode averaging leads to performance losses as shown in Table I and visualized in Fig. 6c.

Responsibility Ablation. Lastly, we compare ML-Cur with a version where we do not maximize the data weight entropy $\mathcal{H}(\nu)$ and hence ignore the data weight responsibility $\tilde{\nu}_{o|i}$ but instead maximize the per component entropy $\mathcal{H}(\nu_o)$. As a consequence, components are not punished when covering equal regions in the context space, leading to performance losses as shown in Table I and visualized in Fig. 6d. Fig. 7 shows the expansion of context components throughout training.

VI. CONCLUSION

We introduced ML-Cur, a curriculum-based approach to versatile skill learning by imitation, using a re-parameterized MoE over movement primitives to model contextualized and versatile behavior. ML-Cur solves three common problems




Fig. 7: The expansion of the Gaussian context distributions in the table tennis task show the adjustment of the curriculum.

of maximum likelihood learning for MoEs: *outlier-sensitivity*, *locality-violation* in context space and *mode-averaging* in behavior space. We introduced a new objective that considers per-component datapoint weights and context distributions combined with an entropy bonus. The resulting objective yields expressive yet specialized components based on individually adjusted curricula. The variational lower bound allows for an efficient optimization while encouraging diversity between the components. The resulting method, ML-Cur, outperforms maximum likelihood approaches and even EIM in all conducted experiments, highlighting the advantages of our method. One of our experiments uses versatile real-world human data, which suffers from the initially described issues.

Limitations. A limitation of our current approach is the linearity of the experts. While using a mixture of movement primitives still allows us to represent complex behavior it also limits our choice of the context space. Hence, the experiments in this paper do not consider high dimensional context, e.g., images or point clouds. However, for higher dimensional contexts, the linear representation would not be sufficient, and we would need to work with nonlinear experts.

ACKNOWLEDGEMENTS

The research presented in this paper was funded by the Deutsche Forschungsgemeinschaft (German Research Foundation) – 448648559 and supported by the Carl Zeiss Foundation under the project JuBot (Jung Bleiben mit Robotern).

REFERENCES

- [1] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín, “What matters in learning from offline human demonstrations for robot manipulation,” in *5th Annual Conference on Robot Learning*, 2021. [Online]. Available: <https://openreview.net/forum?id=JrsfBJtDFdI>
- [2] M. Orsini, A. Raichuk, L. Hussenot, D. Vincent, R. Dadashi, S. Girgin, M. Geist, O. Bachem, O. Pietquin, and M. Andrychowicz, “What matters for adversarial imitation learning?” in *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34. Curran Associates, Inc., 2021, pp. 14 656–14 668. [Online]. Available: <https://proceedings.neurips.cc/paper/2021/file/7b647a7d88f4d6319bf0d600d168dbeb-Paper.pdf>
- [3] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum likelihood from incomplete data via the em algorithm,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 39, no. 1, pp. 1–22, 1977.
- [4] P. Becker, O. Arenz, and G. Neumann, “Expected information maximization: Using the i-projection for mixture density estimation,” in *International Conference on Learning Representations*, 2020.
- [5] C.-W. Huang, D. Krueger, A. Lacoste, and A. Courville, “Neural autoregressive flows,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 2078–2087.
- [6] 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.
- [7] M. Y. Seker, M. Imre, J. Piater, and E. Ugur, “Conditional neural movement primitives,” in *Proceedings of Robotics: Science and Systems*, Freiburg im Breisgau, Germany, June 2019.
- [8] S. Calinon, “Robot learning with task-parameterized generative models,” in *Proc. Intl Symp. on Robotics Research (ISRR)*, 2015.
- [9] A. T. Le, M. Guo, N. van Duijkeren, L. D. Rozo, R. Krug, A. G. Kupcsik, and M. Bürger, “Learning forceful manipulation skills from multi-modal human demonstrations,” *CoRR*, vol. abs/2109.04222, 2021. [Online]. Available: <https://arxiv.org/abs/2109.04222>
- [10] S. Schaal, J. Peters, J. Nakanishi, and A. Ijspeert, “Learning movement primitives,” in *Robotics research. the eleventh international symposium*. Springer, 2005, pp. 561–572.
- [11] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann, “Probabilistic movement primitives,” *Advances in neural information processing systems*, vol. 26, 2013.
- [12] K. Mülling, J. Kober, O. Kroemer, and J. Peters, “Learning to select and generalize striking movements in robot table tennis,” *The International Journal of Robotics Research*, vol. 32, no. 3, pp. 263–279, 2013.
- [13] M. Ewerton, G. Neumann, R. Lioutikov, H. B. Amor, J. Peters, and G. Maeda, “Learning multiple collaborative tasks with a mixture of interaction primitives,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1535–1542.
- [14] A. Pervez and D. Lee, “Learning task-parameterized dynamic movement primitives using mixture of gmms,” *Intelligent Service Robotics*, vol. 11, no. 1, pp. 61–78, 2018.
- [15] K. P. Murphy, *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [16] N. Freymuth, P. Becker, and G. Neumann, “Versatile inverse reinforcement learning via cumulative rewards,” in *NeurIPS 2021 Workshop on Robot Learning: Self-Supervised and Lifelong Learning, Virtual*, 2021.
- [17] M. Sugiyama, T. Suzuki, and T. Kanamori, *Density ratio estimation in machine learning*. Cambridge University Press, 2012.
- [18] K. Hausman, Y. Chebotar, S. Schaal, G. Sukhatme, and J. J. Lim, “Multi-modal imitation learning from unstructured demonstrations using generative adversarial nets,” *Advances in neural information processing systems*, vol. 30, 2017.
- [19] V. Myers, E. Biyik, N. Anari, and D. Sadigh, “Learning multimodal rewards from rankings,” in *Conference on Robot Learning*. PMLR, 2022, pp. 342–352.
- [20] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, “Curriculum learning,” in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 41–48.
- [21] P. Soviany, R. T. Ionescu, P. Rota, and N. Sebe, “Curriculum learning: A survey,” *arXiv preprint arXiv:2101.10382*, 2021.
- [22] S. Narvekar, B. Peng, M. Leonetti, J. Sinapov, M. E. Taylor, and P. Stone, “Curriculum learning for reinforcement learning domains: A framework and survey,” *Journal of Machine Learning Research*, vol. 21, pp. 1–50, 2020.
- [23] O. Celik, D. Zhou, G. Li, P. Becker, and G. Neumann, “Specializing versatile skill libraries using local mixture of experts,” in *Conference on Robot Learning*. PMLR, 2022, pp. 1423–1433.
- [24] O. Arenz, M. Zhong, and G. Neumann, “Trust-region variational inference with gaussian mixture models.” *J. Mach. Learn. Res.*, vol. 21, pp. 163–1, 2020.
- [25] D. C. Liu and J. Nocedal, “On the limited memory bfgs method for large scale optimization,” *Mathematical programming*, vol. 45, no. 1, pp. 503–528, 1989.
- [26] C. Bishop, “Pattern recognition and machine learning,” *Pattern Recognition and Machine Learning*, 2006.
- [27] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 5026–5033.