Active Learning based on Data Uncertainty and Model Sensitivity

Nutan Chen, Alexej Klushyn, Alexandros Paraschos, Djalel Benbouzid, Patrick van der Smagt

AI Research, Data:Lab, Volkswagen Group, Munich, Germany

{first name dot last name}@volkswagen.de

Abstract-Robots can rapidly acquire new skills from demonstrations. However, during generalisation of skills or transitioning across fundamentally different skills, it is unclear whether the robot has the necessary knowledge to perform the task. Failing to detect missing information often leads to abrupt movements or to collisions with the environment. Active learning can quantify the uncertainty of performing the task and, in general, locate regions of missing information. We introduce a novel algorithm for active learning and demonstrate its utility for generating smooth trajectories. Our approach is based on deep generative models and metric learning in latent spaces. It relies on the Jacobian of the likelihood to detect non-smooth transitions in the latent space, i.e., transitions that lead to abrupt changes in the movement of the robot. When non-smooth transitions are detected, our algorithm asks for an additional demonstration from that specific region. The newly acquired knowledge modifies the data manifold and allows for learning a latent representation for generating smooth movements. We demonstrate the efficacy of our approach on generalising elementary skills, transitioning across different skills, and implicitly avoiding collisions with the environment. For our experiments, we use a simulated pendulum where we observe its motion from images and a 7-DoF anthropomorphic arm.

I. INTRODUCTION

Learning is rarely random and it typically follows an intended, often greedy curriculum. Actively seeking to fill knowledge gaps is, in fact, tantamount to faster learning. This setup is especially advantageous when the data is scarce and expensive to acquire. To actively seek the missing knowledge, the learning algorithm is endowed with the ability to query an *oracle* for the next datum, or the next datum to label. The oracle is commonly a human labeller or a demonstrator. To this end, active learning often boils down to two components: a model that quantifies the learner's uncertainty and a *utility* function. The supremum of the utility function determines the next datum to be queried. Examples of utility functions include the Information Gain [1], Upper Confidence Bounds [2], or the Expected Improvement [3]. Active-learning approaches are efficient as, in general, require fewer data to achieve high performance [4]. Therefore, in robotics, where data acquisition is considered expensive, utilising active learning is crucial. In this work, we introduce an active-learning algorithm that is primarily aimed for motion planning. Our approach uses latent-variable models and proposes a novel utility function that operates in the latent space. It leverages recent advances in metric learning for latent-variable models [5], [6] and augments them with uncertainty estimation.



(a) Before active learning.

(b) After active learning.

Fig. 1: The FRANKA 7-DoF robotic arm performing a reaching movement. (a) Despite that the robot was uncertain about performing the movement, we executed it, and, as a result, the robot collided with the obstacles. (b) Our approach successfully detected that there is not enough information for executing the movement and asked for additional demonstrations. After acquiring the demonstrations, our approach was confident that no additional knowledge is needed and the robot successfully performed the reaching movement by avoiding the obstacles.

Current approaches in active learning for robotics, further discussed in Section II, mainly focus on acquiring demonstrations for learning new trajectories when data points are missing i.e., they can successfully quantify the missing information. However, they do not provide a data-driven interpolation approach between the already acquired data points, e.g., when generalising skills or transitioning between skills. In practice, a shortest path approach is used, but it suffers from limitations. First, the interpolation does not follow the shortest path in the data manifold and consequently, might drive the robot away from the known state-space regions that have been acquired from demonstrations, resulting, e.g., in collisions or in reaching joint limits. Second, the smoothness of the movement is also not taken into account. Yet, smooth and predicable motion is a key ingredient for safe interactions with robots.

We propose a novel active-learning algorithm that supports data-driven generalisation by allowing interpolations in the data manifold, while it is capable to simultaneously detect non-smooth, abrupt changes. Additionally, it quantifies the uncertainty during the interpolation and, hence, suggests new demonstrations to be provided by an oracle. Specifically, our approach leverages from latent-variable models in order to infer a meaningful representations of the motion trajectories and exploits the Jacobian of the data likelihood along the movement to discover abrupt motions. We demonstrate the benefits of our approach in a set of experiments by generating smooth generalisations of movements and, in addition, we demonstrate how the uncertainty of a movement can be used to implicitly avoid obstacles. For the experimental evaluation, we model the motion of a pendulum assuming that we have access only to image observations and we use a 7-DoF anthropomorphic arm to demonstrate our approach on avoiding obstacles.

II. RELATED WORK

Active learning is a common component in robotic systems, especially when it comes to efficiently acquiring new samples during the learning [7], [8], or to performing exploration in the action space [9], [10]. Interested by the problem of teaching robots by humans, the authors in [11] leverage the uncertainty of the hypothesis space in order to efficiently request demonstrations from a human operator. In [12], the robot is endowed with the ability to ask questions, in order to acquire new labels, new demonstrations, or new skill representations. The uncertainty estimation of Gaussian Processes is used in [13] in order to learn to broaden the robot reaching skills by querying new demonstrations whenever the uncertainty reaches a specified threshold. Active learning is also used to improve over random exploration for grasping tasks based on visual sensory input [9]. Also for grasping, active learning is combined with reactive control in order to explore interesting poses using an upper confidence bound (UCB) policy [14]. In [15], a goal-driven active learning approach is developed for learning skills in continuous sensorimotor spaces. In [16], the authors combine model-free and modelbased reinforcement learning methods-and the uncertainty thereof, in order to acquire robotic manipulation skills.

Our robot experiments involves a query-based learning system. The scenario is similar to [13], however, in the latter work the user has to manually choose the trigger.

More generally, active learning has been intensively studied in the machine learning literature [4]. In the latter survey, the author distinguishes three types of scenario, depending on how to select the query to be labelled by the oracle; namely membership queries synthesis, stream-based selective sampling, and pool-based sampling. In the membership queries synthesis scenario (e.g., [17], [18]), the learner first generates (or synthesises) the query to be annotated, instead of sampling it from an observed pool of data. In streambased selective sampling (e.g., [7]), the learner further filters the generated queries to be labelled and hence can decide to discard it based on a given "informativeness measure". Finally, pool-based sampling (e.g., [19], [20]) is motivated by applications wherein a large amount of unlabelled data can be collected but labelling by the oracle is costly. In our illustrative example in Section IV-A-the pendulum experimentpool-based sampling is used. The original problem of the experiment corresponds to an unsupervised learning task and does not require labels, however the setup allows us to select the most useful data from the pool and evaluate our method. In Section IV-B and IV-C however, stream-based selective sampling is used. These experiments imply robot interaction

and it remains expensive to obtain unlabelled data for such manipulations.

In order to model the uncertainty, kernel-based methods are commonly used in the active learning literature, such as Support Vector Machines and "margin-based uncertainty" [21]—when dealing with low-dimensional data. In [22], the authors combine Radial Basis Functions (RBF) kernels with information density. Bayesian Neural Networks are also used for active learning, as recently shown by [20] in the context of images classification. Following [23], we use an RBF network to model data uncertainty. Our method is an alternative for modelling defect detection by measuring model sensitivity.

III. APPLYING RIEMANNIAN GEOMETRY TO LATENT VARIABLE MODELS FOR ACTIVE LEARNING

Latent-variable models (LVM), defined by

$$p(\mathbf{x}) = \int p(\mathbf{x} \mid \mathbf{z}) \, p(\mathbf{z}) \, d\mathbf{z},\tag{1}$$

are widely used to find a representation of observable data $\mathbf{x} \in \mathbb{R}^{N_x}$ through latent variables $\mathbf{z} \in \mathbb{R}^{N_z}$ based on hidden, nonlinear regularities in \mathbf{x} .

We use latent variables for generating a sequence of observable data points, with the condition that every generated point has a high similarity to the previous one. However, the similarity between the successive data points depends on the information provided to the LVM. In case of a low similarity we apply active learning to get targeted the missing information.

Gauging the similarity of two data points in the latent space is one of the main topics in this paper. To solve this problem, we take the Jacobian of the likelihood into account by treating the latent space as a Riemannian manifold. The Riemannian metric defines a relationship based on the Jacobian of the likelihood, due to change of variables when moving from Riemannian (latent) to Euclidean (observation) space.

A *smooth* interpolation through our observable data can be obtained by following the geodesic, i.e. the lengthminimising curve between two points in the Riemannian space. Here, *smooth* refers to a strong similarity of successive data points.

However, even when following the geodesic, finding a smooth interpolation will fail under certain circumstances. For instance, when we are trying to interpolate between different classes. The distance between the data manifolds of the different classes in the observation space typically results in a high Jacobian value of the likelihood mean when interpolating from one class to the other. This implies that information is missing to provide a sequence of similar data points to connect the different manifolds smoothly. As a consequence, the variance of the likelihood changes as well.

Since the Jacobians of both the mean and the variance are taken into account by the Riemannian metric, this property can be turned to advantage when dealing with active learning. For instance, when trying to interpolate between different robot movements. Because missing data can be queried specifically if such boundaries are passed.

Building on that, the focus of our paper lies on applying Riemannian geometry to LVMs for active learning of robot movements.

A. Importance-weighted autoencoder

Since in most LVMs the integral in Eq. (1) is intractable, approximations are used which base on sampling [24] [25] or on variational inference [26] [27]. In the latter case, the problem is reformulated as the maximisation of the evidence lower bound (ELBO). The distribution $q(\mathbf{z})$ approximates the intractable posterior and $p_{\theta}(\mathbf{x}|\mathbf{z})$, defined as the generative model and parameterised by θ , approximates the likelihood. Let $\mathbf{X} = {\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)}}$ be observable data and $\mathbf{z}^{(i)}$ the corresponding latent variables. Then,

$$\ln p_{\theta}(\mathbf{X}) \geq \sum_{i=1}^{N} \mathbb{E}_{q(\mathbf{z}^{(i)})} \left[\ln \frac{p_{\theta}(\mathbf{x}^{(i)} | \mathbf{z}^{(i)}) p_{\theta}(\mathbf{z}^{(i)})}{q(\mathbf{z}^{(i)})} \right] = \mathcal{L}_{\text{ELBO}}.$$
(2)

Implementing $q(\mathbf{z}^{(i)}) = q_{\phi}(\mathbf{z}^{(i)}|\mathbf{x}^{(i)})$ with a neural network parameterised by ϕ , we obtain the variational autoencoder (VAE) introduced in [26], [27].

To overcome the limitations of ordinary VAEs and to achieve a tighter ELBO, we use importance-weighted autoencoders (IWAE) [28], [29] in our approach. IWAEs treat $q_{\phi}(\mathbf{z}|\mathbf{x})$ as a proposal distribution and obtain a tighter ELBO by using importance sampling:

$$\mathcal{L}_{\text{ELBO}} = \sum_{i=1}^{N} \mathbb{E}_{\mathbf{z}_{1}^{(i)}, \dots, \mathbf{z}_{K}^{(i)} \sim q_{\phi}(\mathbf{z}^{(i)} | \mathbf{x}^{(i)})} \Big[\ln \frac{1}{K} \sum_{k=1}^{K} w_{k}^{(i)} \Big],$$
(3)

with the importance weights

$$w_{k}^{(i)} = \frac{p_{\theta}(\mathbf{x}^{(i)} | \mathbf{z}_{k}^{(i)}) p_{\theta}(\mathbf{z}_{k}^{(i)})}{q_{\phi}(\mathbf{z}_{k}^{(i)} | \mathbf{x}^{(i)})}.$$
(4)

B. Riemannian geometry in latent variable models

Riemannian space is a differentiable manifold M which contains as an additional characteristic a metric to describe its geometric properties. The corresponding metric tensor **G** assigns to each point **z** in the latent space an inner product on the tangent space $T_z M$, defined by

$$\langle \mathbf{z}', \mathbf{z}' \rangle_{\mathbf{z}} := \mathbf{z}'^T \, \mathbf{G}(\mathbf{z}) \, \mathbf{z}',$$
 (5)

with $\mathbf{z}' \in T_{\mathbf{z}}M$ and $\mathbf{z} \in M$.

Let us assume we have a curve $\gamma : [0,1] \rightarrow \mathbb{R}^{N_z}$ in the Riemannian (latent) space that is transformed by a continuous function $f(\gamma(t))$ to an N_x -dimensional Euclidean (observation) space, where $\gamma(t) \in \mathbb{R}^{N_z}$. The length of this curve in the Euclidean space is defined as

$$L(\gamma) = \int_0^1 \sqrt{\left\langle \dot{\gamma}(t), \dot{\gamma}(t) \right\rangle_{\gamma(t)}} \mathrm{d}t, \tag{6}$$

with the metric tensor $\mathbf{G} = \mathbf{J}^T \mathbf{J}$, where \mathbf{J} is the Jacobian matrix of the likelihood and $\dot{\gamma}$ the time derivative of γ .

C. Using geodesics for trajectory generation

To approximate the geodesic we use a neural network that is optimised by minimising $L(\gamma)$. A singular-value decomposition of **G** ensures the geodesic is following the data manifold, as introduced in [5].

Although this method takes the sensitivity of the model into account, it does not capture data uncertainty. In other words: the high Jacobian values of the likelihood mean at the boundaries between different data manifolds are taken into account, but there is nothing that tells us where our model is uncertain due to missing data. The reason is a global variance of the generative model. To remedy that, the neural network of the generative model is extended by radial basis function (RBF) networks to be able to represent the likelihood variance too [6], [23].

In contrast to [6], we update the weights of the RBF networks during the training of the generative model and define a rule for an autonomous hyperparameter selection after the training is finished. The RBFs v and the precision $\psi(z)$ of the generative model are given by

$$\psi(\mathbf{z}) = \mathbf{W}\mathbf{v}(\mathbf{z}), \tag{7}$$
$$(\mathbf{z}) = \exp(-\lambda_k \|\mathbf{z} - \mathbf{c}_k\|_2), \text{ with } k = 1, \dots, K,$$

where K is the number of the radial basis functions. λ and c are variables representing bandwidth and centres, respectively. W are the weights to be optimised. The bandwidth is defined by

 v_k

$$\lambda_k = \alpha \left(\frac{\sum_i^K \|\mathbf{c}_k - \mathbf{c}_i\|_2}{k}\right)^{-2}, \qquad (8)$$

where α , a hyperparameter, denotes the curvature of the Riemannian metric. Since the variance is the reciprocal of the precision:

$$\boldsymbol{\sigma}^2(\mathbf{z}) = \boldsymbol{\psi}(\mathbf{z})^{-1},\tag{9}$$

it increases with the distance to the centres and the uncertainty of the model, respectively. It is not possible to directly compute the Jacobian of a sample $\mathbf{x} \sim p_{\theta}(\mathbf{x}|\mathbf{z})$. Hence, we reparameterise it by $\epsilon \sim \mathcal{N}(0, 1)$ [26], [30]:

$$\mathbf{J}(\mathbf{z}) = \mathbf{J}_{\mu}(\mathbf{z}) + \epsilon \ \mathbf{J}_{\sigma}(\mathbf{z}), \tag{10}$$

where J_{μ} and J_{σ} , the Jacobians of the mean and the standard deviation of the likelihood, represent the sensitivity and the data uncertainty of the model, respectively. The changes in the likelihood variance have influence on **G**, hence the equation introduced in [5] has to be updated. To simplify the calculation, we remove the stochasticity in **G** by taking the expectation [6]

$$\mathbb{E}_{p(\epsilon)}[\mathbf{G}(\mathbf{z})] = \mathbf{J}_{\mu}(\mathbf{z})^T \mathbf{J}_{\mu}(\mathbf{z}) + \mathbf{J}_{\sigma}(\mathbf{z})^T \mathbf{J}_{\sigma}(\mathbf{z}).$$
(11)

We differ from [6] in the optimisation procedure of the model: the centres **c** are computed by K-means and updated at every *n*-th iteration step during the training of the IWAE. For both the computation of the centres **c** and the RBFs **v**, the mean \mathbf{z}_{μ} of **z** is used. The weights **W** are optimised by



Fig. 2: Evaluations of the illustrative experiment based on a two-dimensional dataset. (a) Training and testing dataset with a sample sizes of $3.3 \cdot 10^3$ and $2 \cdot 10^3$ data points, respectively. The data-acquisition pool has the same distribution as the testing dataset and $2 \cdot 10^3$ data points. (b) Latent space of the trained model. The MF is represented by the grayscale. The blue points depict the mean of the training data. (c) Observation space of the acquired data after seven iterations when using different active-learning approaches, namely the MF, the Max Entropy, and a random acquisition strategy. The latter acquires data that is similarly distributed to the data in the acquisition pool, which has a large overlap with the training dataset and, therefore it does not provide an efficient learning approach. The Max Entropy does not take into account the MF of the model or acquires data points from the center of the latent space.

back-propagation. α is treated as a hyperparameter during the IWAE training. After the training is finished, α is updated to satisfy

$$\|\max[\mathbf{J}_{\mu}(\mathbf{z}_{\mu})] - \max[\mathbf{J}_{\sigma}(\mathbf{z}_{\mu})]\| < \epsilon, \text{ with } \epsilon \to 0.$$
 (12)

Satisfying Eq. (12) guarantees that the mean and the variance have a similar effect on the Riemannian metric tensor.

D. Active learning for robot trajectory generation

Active learning can be applied to targeted reduce the uncertainty of our model, which leads to smoother trajectory generations. In active learning an acquisition function a is used to detect where a model M is uncertain, so missing labels can be queried specifically:

$$\mathbf{x}^* = \operatorname*{arg\,max}_{\mathbf{x}\in D_{\text{pool}}} a(\mathbf{x}, M). \tag{13}$$

Our goal is to guarantee a smooth interpolation along the geodesic. This is only possible if there are no abrupt changes in the Jacobian of the likelihood, which is expressed by the determinant of the metric tensor **G** for a specific pair (\mathbf{x}, \mathbf{z}) . This leads to the following acquisition function:

$$a(\mathbf{x}, M) = \sqrt{\det \mathbf{G}(\mathbf{z})} \rightleftharpoons \mathrm{MF}(\mathbf{z}),$$
 (14)

also defined as the magnification factor (MF) [31]. The MF can be interpreted as the scaling factor when moving from the Riemannian (latent) to the Euclidean (observation) space, due to the change of variables.

In addition to the acquisition function, a threshold is necessary to tell the active learning algorithm whether a interpolation is smooth or not. The threshold is defined as

$$\tau(\Omega) = \frac{1}{N_{\Omega}} \sum_{i=1}^{N_{\Omega}} \omega_i + \sqrt{\operatorname{Var}(\Omega)}, \qquad (15)$$

where $\omega_i \in \Omega$, $\Omega = {MF(\mathbf{z}^{(1)}), \dots, MF(\mathbf{z}^{(N)})}$, and N_{Ω} is the cardinality of Ω .

When applying active learning to robot movements, we use a set of pairs of start and end points in the observation space $\Pi = \{(\mathbf{x}_0^{(1)}, \mathbf{x}_1^{(1)}), \dots, (\mathbf{x}_0^{(N_{\Pi})}, \mathbf{x}_1^{(N_{\Pi})})\}$. For each pair the geodesic $geo(\mathbf{x}_0, \mathbf{x}_1)$ is computed. To decide whether the movement (interpolation) between a start and an end point is smooth, only points along the geodesic $\gamma(t) \in geo(\mathbf{x}_0, \mathbf{x}_1)$ are taken into account. Thus, in contrast to the active learning approach described in Eq. (13), $D_{pool} = [\gamma(0), \gamma(1)]$ refers to the latent space. Based on whether the values of the magnification factor along the geodesic $geo(\mathbf{x}_0, \mathbf{x}_1)$ exceed $\tau(\Omega)$, the active learning algorithm decides if the trajectory between \mathbf{x}_0 and \mathbf{x}_1 is required to be demonstrated. In case of a required demonstration, retraining the model with the new data leads to low MFs along $geo(\mathbf{x}_0, \mathbf{x}_1)$.

Hence, the final result of our approach is a smooth movement or rather a smooth combination of movements of the robot—realised by reconstructing the latent variables along the geodesic.

IV. EXPERIMENTAL EVALUATION

We evaluate our approach in multiple scenarios. First, we use an artificial two-dimensional dataset to illustrate how our approach works. Then, we demonstrate that our



Fig. 3: The reconstruction error of the illustrative experiment when using the MF, Max Entropy, or random active-learning strategy. The acquisition functions acquire ten data points per iteration.

approach can work efficiently with high-dimensional data on simulated pendulum where the state is given by images. Finally, we present our results on controlling a 7-DoF robotic arm where smooth reaching movements are generated. When our approach detects that a trajectory would cross regions of the state space where not enough data have been acquired, it asks for additional demonstrations. Hence, it is used to implicitly avoid collisions and joint limits. The architectural design and the hyper-parameters used for our experiments are listed in the appendix.

A. Illustrative experiment

In the first experiment, we evaluate the efficiency of our approach in reducing the reconstruction error by actively acquiring data points from regions where the model does not have enough information. To better illustrate our approach, we generated an artificial two-dimensional dataset, where the IWAE maps the observation space to a two-dimensional latent space. The training dataset is depicted in Fig. 2a, whereas the resulting latent space and MF are shown in Fig. 2b.

Our algorithm asks for new data points from regions where the MF is high, and therefore it selects the central region. In contrast, the Max Entropy strategy assumes that enough information form the central region is present.

As a result, our MF active-learning approach can efficiently reduce the reconstruction error using fewer samples than the Max Entropy or the random acquisition approach. The results are shown in Fig. 3.

B. Trajectory planning for pendulum

We demonstrate the trajectory planning capabilities of our approach in a simulated 1-DoF pendulum system. The simulator provides a 16×16 -pixel image of the current state of the pendulum, which we use as input to our algorithm. We gathered an image dataset by collecting $T = 15 \cdot 10^3$ images for two different joint angle ranges, $R_1 = [0, 150)$ and $R_2 = [180, 330)$ degrees. Subsequently, we augmented the dataset by adding 0.05 Gaussian noise to each pixel, to avoid over fitting and to improve the coherence of the latent space.

After training, we generated four trajectories between the two datasets by following the geodesic. The generated trajectories are illustrated in Fig. 4a. The trajectories A_2A_3 and A_4A_1 move across regions of the state space where the MF exceeds a predetermined threshold, as not enough information has been collected from those regions. An illustration of the trajectories is shown in Fig. 4c.

Our approach requested for additional demonstrations from regions where the MF exceeds the threshold. Afterwards the model is retrained with the new data. As a result, the MF is reduced in these regions, as shown in Fig. 5a. The corresponding trajectories are significantly smoother after our active-learning approach was applied, as can be seen in Fig. 5b.

C. Generating robot trajectories with active learning

Deciding whether the robot is able to perform a task or a demonstration is required is not trivial. Therefore, we evaluate our approach in a robot trajectory generation setting, where the robot should consult the human operator to avoid collisions with the environment. Also, the generated trajectories should not have abrupt changes to enable the robot to precisely follow them. For this experiment, we used a Panda robot from FRANKA, a lightweight 7-DoF robotic arm with joint torque sensors.

For training our model, we provided demonstrations of reaching objects that were placed at two distinct locations. We used kinaesthetic teaching, i.e., the human demonstrator could freely move the robot to acquire a dataset of five demonstrations per reaching location. The setup is depicted in Fig. 1. During the demonstrations we recorded the joint angles of the robot at a rate of 1 kHz. Additionally, obstacles where placed in the workspace of the robot. Naively generating movements based on the demonstration dataset likely results in collisions.

We generated trajectories by computing an interpolation between the two distinct locations by following the geodesic, as shown in Fig. 6a. The geodesic trajectory crosses a region with high MF values.

Since the MF values along the proposed trajectory exceed the threshold, the algorithm asks the user for additional data. Thus, after collecting the data of the queried trajectory, we retrained our model on the new data and recomputed the geodesic. The updated latent space is shown in Fig. 6b, where the geodesic does not cross high-MF regions anymore. The end-effector trajectories before and after retraining are depicted in Fig. 6c. As a result, the robot arm moves close to the demonstrated path and avoids collisions with obstacles. A visualisation of the robot trajectory and its environment can be found in Fig. 7.

V. CONCLUSION

We introduced a new active-learning method, based on the model sensitivity in deep generative models. We showed that our method is suitable for efficiently learning new skills



Fig. 4: Evaluation of the pendulum experiment *before* active learning. (a) Latent space of the pendulum dataset after the initial training. The markers $\{A_1, A_2, A_3, A_4\}$ represent the position of the pendulum at $\{40, 120, 200, 280\}$ degrees, respectively. In addition, the colour encodes the pendulum rotation angles and the greyscale the MF values. (b) The trajectories $A_2 \rightarrow A_3$ and $A_4 \rightarrow A_1$ of the generated pendulum movements experience large MF values, leading to discontinuities. The discontinuities are marked by coloured lines. (c) The MF exceeds the threshold for two trajectories, $A_2 \rightarrow A_3$ and $A_4 \rightarrow A_1$.



Fig. 5: Evaluation of the pendulum experiment *after* active learning. (a) The resulting latent space of the pendulum dataset shows that the MF decreased along the planned trajectories. (b) The generated pendulum movements are smoother and they do not experience high MF values. (c) The MF values of all four generated pendulum trajectories are below the threshold.

from demonstrations while maintaining some smoothness between known motions. In addition to triggering a query for new demonstrations, the magnification factor also indicates whether the observed data contains unrealistic postures, sudden fast movements, or indicates previously unseen/untrained movements. Currently, the model is retrained when new data is acquired, in order to prevent the optimisation procedure to get stuck in local minima. We tackle this issue in future development and investigate alternative optimisation procedures that effectively allow for an online update.



- (a) Without active learning
- (b) After active learning

(c) The Cartesian trajectories of the end-effector

Fig. 6: Trajectory generation for reaching movements while the robot avoids obstacles. (a) The latent space after training with the initial demonstrations. In blue we depict the training data points. The generated trajectory crosses a high-MF region. (b) After providing the additional demonstration the resulting trajectory avoids collisions with the environment. (c) Cartesian trajectories of the end-effector, before (orange) and after (blue) active learning. The trajectories after active learning are smoother and follow a path close to the demonstrations.



Fig. 7: The FRANKA robot performing a novel reaching movement. (top) Our algorithm detected that an additional demonstration is required due to the lack of data. After preforming active learning the robot successfully avoids the obstacle. (bottom) Without the active-learning approach the robot collides with the environment.

ACKNOWLEDGMENT

We are very grateful to Justin Bayer for valuable suggestions concerning this work.

REFERENCES

- N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel, "Bayesian active learning for classification and preference learning," *arXiv preprint* arXiv:1112.5745, 2011.
- [2] R. Ganti and A. G. Gray, "Building bridges: viewing active learning from the multi-armed bandit lens," arXiv preprint arXiv:1309.6830, 2013.
- [3] E. Brochu, V. M. Cora, and N. De Freitas, "A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning," *arXiv preprint arXiv:1012.2599*, 2010.
- [4] B. Settles, "Active learning literature survey," Computer Science Technical Report, 2010.
- [5] N. Chen, A. Klushyn, R. Kurle, X. Jiang, J. Bayer, and P. van der Smagt, "Metrics for deep generative models," in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2018.
- [6] G. Arvanitidis, L. K. Hansen, and S. Hauberg, "Latent space oddity: on the curvature of deep generative models," in *International Conference* on Learning Representations (ICLR), 2018.
- [7] L. E. Atlas, D. A. Cohn, and R. E. Ladner, "Training connectionist networks with queries and selective sampling," in *Advances in neural information processing systems*, 1990, pp. 566–573.
- [8] D. A. Cohn, Z. Ghahramani, and M. I. Jordan, "Active learning with statistical models," *Journal of artificial intelligence research*, vol. 4, pp. 129–145, 1996.

- [9] M. Salganicoff, L. H. Ungar, and R. Bajcsy, "Active learning for vision-based robot grasping," *Machine Learning*, vol. 23, no. 2-3, pp. 251–278, 1996.
- [10] A. Morales, E. Chinellato, A. H. Fagg, and A. P. del Pobil, "An active learning approach for assessing robot grasp reliability," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1, 2004, pp. 485–490.
- [11] C. Chao, M. Cakmak, and A. L. Thomaz, "Transparent active learning for robots," in ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2010, pp. 317–324.
- [12] M. Cakmak and A. L. Thomaz, "Designing robot learners that ask good questions," in ACM/IEEE international conference on Human-Robot Interaction, 2012, pp. 17–24.
- [13] G. Maeda, M. Ewerton, T. Osa, B. Busch, and J. Peters, "Active incremental learning of robot movement primitives," in *Conference* on Robot Learning (CORL), 2017.
- [14] O. Kroemer, R. Detry, J. Piater, and J. Peters, "Combining active learning and reactive control for robot grasping," *Robotics and Autonomous* systems, vol. 58, no. 9, pp. 1105–1116, 2010.
- [15] A. Baranes and P.-Y. Oudeyer, "Active learning of inverse models with intrinsically motivated goal exploration in robots," *Robotics and Autonomous Systems*, vol. 61, no. 1, pp. 49–73, 2013.
- [16] S. Hangl, V. Dunjko, H. Briegel, and J. H. Piater, "Skill learning by autonomous robotic playing using active learning and creativity," *CoRR*.
- [17] D. Angluin, "Queries and concept learning," *Machine learning*, vol. 2, no. 4, pp. 319–342, 1988.
- [18] R. D. King, J. Rowland, S. G. Oliver, M. Young, W. Aubrey, E. Byrne, M. Liakata, M. Markham, P. Pir, L. N. Soldatova *et al.*, "The automation of science," *Science*, vol. 324, no. 5923, pp. 85–89, 2009.

- [19] D. D. Lewis and W. A. Gale, "A sequential algorithm for training text classifiers," in *Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*, 1994, pp. 3–12.
- [20] Y. Gal, R. Islam, and Z. Ghahramani, "Deep Bayesian Active Learning with Image Data," in *Proceedings of the 34th International Conference* on Machine Learning, 2017.
- [21] A. J. Joshi, F. Porikli, and N. Papanikolopoulos, "Multi-class active learning for image classification," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 2372–2379.
- [22] X. Li and Y. Guo, "Adaptive active learning for image classification," in *IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2013, pp. 859–866.
- [23] Q. Que and M. Belkin, "Back to the future: Radial basis function networks revisited," in *International Conference on Artificial Intelligence* and Statistics (AISTATS), vol. 51, 2016, pp. 1375–1383.
- [24] W. K. Hastings, "Monte carlo sampling methods using markov chains and their applications," 1970.
- [25] A. E. Gelfand and A. F. Smith, "Sampling-based approaches to calculating marginal densities," *Journal of the American statistical* association, vol. 85, no. 410, pp. 398–409, 1990.
- [26] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," *CoRR*, vol. abs/1312.6114, 2013.
- [27] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," pp. 1278–1286, 2014.
- [28] Y. Burda, R. B. Grosse, and R. Salakhutdinov, "Importance weighted autoencoders," *CoRR*, vol. abs/1509.00519, 2015.
- [29] C. Cremer, Q. Morris, and D. Duvenaud, "Reinterpreting importance-

weighted autoencoders," International Conference on Learning Representations Workshop Track, 2017.

- [30] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," in *Proceedings of the 31th International Conference on Machine Learning (ICML)*, 2014, pp. 1278–1286.
- [31] C. M. Bishop, M. Svens' en, and C. K. Williams, "Magnification factors for the SOM and GTM algorithms," in *Proceedings Workshop* on Self-Organizing Maps, 1997.
- [32] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization." *CoRR*, vol. abs/1412.6980, 2014.
- [33] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer* vision and pattern recognition, 2016, pp. 770–778.

Appendix

A. DETAILS OF THE TRAINING PROCEDURE

The Adam optimiser [32] was used for optimising the models of the three experiments. In the Tables I, II, and III, we provide the parameters we used during training. We abbreviate the fully connected layers by FC. Residual networks [33] are used for the pendulum and the FRANKA experiments. Increasing the depth of the generative model led to a more sensible and smoother magnification factor. With K, we refer to the number of importance-weighted samples.

TABLE I: Parameters of the illustrative experiment

recognition model	generative model	hyperparameters
Input $\in \mathbb{R}^2$ 2 tanh FC × 512 units linear FC output layer for means softplus FC output layer for variances	Input $\in \mathbb{R}^2$ 2 tanh FC × 512 units softplus FC output layer for means RBF for variances	learning rate = 2×10^{-3} K = 5 batch size = 150

TABLE II: Parameters for the pendulum experiment

recognition model	generative model	hyperparameters
Input $\in \mathbb{R}^{256}$ 2 tanh FC × 512 units linear FC output layer for means softplus FC output layer for variances	Input $\in \mathbb{R}^2$ 10 residual \times 128 units sigmoid FC output layer for means RBF for variances	learning rate = 10^{-4} K = 5 batch size = 32

TABLE	III:	Parameters	for	the	robot	experiment	t
-------	------	------------	-----	-----	-------	------------	---

recognition model	generative model	hyperparameters
Input $\in \mathbb{R}^7$ 2 tanh FC × 512 units linear FC output layer for means softplus FC output layer for variances	Input $\in \mathbb{R}^2$ 10 residual \times 64 units softplus FC output layer for means RBF for variances	learning rate = 5×10^{-4} K = 15 batch size = 150