

Generalizing Robot Imitation Learning with Invariant Hidden Semi-Markov Models

Ajay Kumar Tanwani^{¶,§}, Jonathan Lee[§], Brijen Thananjeyan[§], Michael Laskey[§],
Sanjay Krishnan[§], Roy Fox[§], Ken Goldberg[§], Sylvain Calinon[¶]

Abstract. Generalizing manipulation skills to new situations requires extracting invariant patterns from demonstrations. For example, the robot needs to understand the demonstrations at a higher level while being invariant to the appearance of the objects, geometric aspects of objects such as its position, size, orientation and viewpoint of the observer in the demonstrations. In this paper, we propose an algorithm that learns a joint probability density function of the demonstrations with invariant formulations of hidden semi-Markov models to extract invariant segments (also termed as sub-goals or options), and smoothly follow the generated sequence of states with a linear quadratic tracking controller. The algorithm takes as input the demonstrations with respect to different coordinate systems describing virtual landmarks or objects of interest with a task-parameterized formulation, and adapt the segments according to the environmental changes in a systematic manner. We present variants of this algorithm in latent space with low-rank covariance decompositions, semi-tied covariances, and non-parametric online estimation of model parameters under small variance asymptotics; yielding considerably low sample and model complexity for acquiring new manipulation skills. The algorithm allows a Baxter robot to learn a pick-and-place task while avoiding a movable obstacle based on only 4 kinesthetic demonstrations.

Keywords: hidden Markov models, imitation learning, adaptive systems

1 Introduction

Generative models are widely used in robot imitation learning to estimate the distribution of the data for regenerating samples from the model [1]. Common applications include probability density function estimation, image regeneration, dimensionality reduction and so on. The parameters of the model encode the task structure which is inferred from the demonstrations. In contrast to direct trajectory learning from demonstrations, many problems arise in robotic applications that require higher contextual level understanding of the environment. This requires learning invariant mappings in the demonstrations that can generalize across different environmental situations such as size, position, orientation of objects, and viewpoint of the observer. Recent trend in imitation learning is forgoing such a task structure for end-to-end supervised learning which requires a large amount of training demonstrations.

[§]University of California, Berkeley.

[¶]Idiap Research Institute, Switzerland.

Corresponding author: ajay.tanwani@berkeley.edu




Fig. 1: Conceptual illustration of hidden semi-Markov model (HSMM) for imitation learning: (left) 3-dimensional Z-shaped demonstrations composed of 5 equally spaced trajectory samples, (middle) demonstrations are encoded with a 3 state HMM represented by Gaussians (shown as ellipsoids) that represent the blue, green and red segments respectively. The transition graph shows a duration model (Gaussian) next to each node, (right) the generative model is combined with linear quadratic tracking (LQT) to synthesize motion in performing robot manipulation tasks from 5 different initial conditions marked with orange squares (see also Fig. 2).

The focus of this paper is to learn the joint probability density function of the human demonstrations with a family of **Hidden Markov Models (HMMs)** in an **unsupervised** manner [22]. We combine tools from statistical machine learning and optimal control to segment the demonstrations into different components or sub-goals that are sequenced together to perform manipulation tasks in a smooth manner. We first present a simple algorithm for imitation learning that combines the decoded state sequence of a hidden semi-Markov model [22,35] with a linear quadratic tracking controller to follow the demonstrated movement [2](see Fig. 1). We then augment the model with a task-parameterized formulation such that it can be systematically adapted to changing situations such as pose/size of the objects in the environment [4,26,31]. We present latent space formulations of our approach to exploit the task structure using: 1) mixture of factor analyzers decomposition of the covariance matrix [16], 2) semi-tied covariance matrices of the mixture model [26], and 3) Bayesian non-parametric formulation of the model with Hierarchical Dirichlet process (HDP) for online learning under small variance asymptotics [27]. The paper unifies and extends our previous work on encoding manipulation skills in a task-adaptive manner [25,26,27]. Our objective is to reduce the number of demonstrations required for learning a new task, while ensuring effective generalization in new environmental situations.

1.1 Related Work

Imitation learning provides a promising approach to facilitate robot learning in the most ‘natural’ way. The main challenges in imitation learning include [18]: 1) **what-to-learn** – acquiring meaningful data to represent the important features of the task from demonstrations, and 2) **how-to-learn** – learning a control policy from the features to reproduce the demonstrated behaviour. Imitation learning algorithms typically fall into **behaviour cloning** or **inverse reinforcement learning (IRL)** approaches. IRL aims to recover the unknown reward function that is being optimized in the demonstrations, while be-

haviour cloning approaches directly learn from human demonstrations in a supervised manner. Prominent approaches to imitation learning include Dynamic Movement Primitives [10], Generative Adversarial Imitation Learning [9], one-shot imitation learning [5] and so on [20].

This paper emphasizes learning manipulation skills from human demonstrations in an unsupervised manner using a family of hidden Markov models by sequencing the atomic movement segments or primitives. HMMs have been typically used for recognition and generation of movement skills in robotics [11,15,23,34]. Other related application contexts in imitation learning include options framework [7,12], sequencing primitives [17], and neural task programs [33].

A number of variants of HMMs have been proposed to address some of its shortcomings, including: 1) how to bias learning towards models with longer self-dwelling states, 2) how to robustly estimate the parameters with high-dimensional noisy data, 3) how to adapt the model with newly observed data, and 4) how to estimate the number of states that the model should possess. For example, [13] used HMMs to incrementally group whole-body motions based on their relative distance in HMM space. [15] presented an iterative motion primitive refinement approach with HMMs. [19] used the Beta Process Autoregressive HMM for learning from unstructured demonstrations. Figueroa et al. used the transformation invariant covariance matrix for encoding tasks with a Bayesian non-parametric HMM [6].

In this paper, we address these shortcomings with an algorithm that learns a hidden semi-Markov model [22,35] from a few human demonstrations for segmentation, recognition, and synthesis of robot manipulation tasks (see Sec. 2). The algorithm observes the demonstrations with respect to different coordinate systems describing virtual landmarks or objects of interest, and adapts the model according to the environmental changes in a systematic manner in Sec. 3. Capturing such invariant representations allows us to compactly encode the task variations than using a standard regression problem. We present variants of the algorithm in latent space to exploit the task structure in Sec. 4. In Sec. 5, we show the application of our approach to learning a pick-and-place task from a few demonstrations, with an outlook to our future work.

2 Hidden Markov Models

Hidden Markov models (HMMs) encapsulate the spatio-temporal information by augmenting a mixture model with latent states that sequentially evolve over time in the demonstrations [22]. HMM is thus defined as a doubly stochastic process, one with sequence of hidden states and another with sequence of observations/emissions. Spatio-temporal encoding with HMMs can handle movements with variable durations, recurring patterns, options in the movement, or partial/unaligned demonstrations. Without loss of generality, we will present our formulation with semi-Markov models for the remainder of the paper. Semi-Markov models relax the Markovian structure of state transitions by relying not only upon the current state but also on the duration/elapsed time in the current state, i.e., the underlying process is defined by a *semi-Markov* chain with a variable duration time for each state. The state duration stay is a random integer variable that assumes values in the set $\{1, 2, \dots, s^{\max}\}$. The value corresponds to the

number of observations produced in a given state, before transitioning to the next state. **Hidden Semi-Markov Models** (HSMMs) associate an observable output distribution with each state in a semi-Markov chain [35], similar to how we associated a sequence of observations with a Markov chain in a HMM.

Let $\{\xi_t\}_{t=1}^T$ denote the sequence of observations with $\xi_t \in \mathbb{R}^D$ collected while demonstrating a manipulation task. The state may represent the visual observation, kinesthetic data such as the pose and the velocities of the end-effector of the human arm, haptic information, or any arbitrary features defining the task variables of the environment. The observation sequence is associated with a hidden state sequence $\{z_t\}_{t=1}^T$ with $z_t \in \{1 \dots K\}$ belonging to the discrete set of K cluster indices. The cluster indices correspond to different segments of the task such as reach, grasp, move etc. We want to learn the joint probability density of the observation sequence and the hidden state sequence. The transition between one segment i to another segment j is denoted by the transition matrix $\mathbf{a} \in \mathbb{R}^{K \times K}$ with $a_{i,j} \triangleq P(z_t = j | z_{t-1} = i)$. The parameters $\{\mu_j^S, \Sigma_j^S\}$ represent the mean and the standard deviation of staying s consecutive time steps in state j as $p(s)$ estimated by a Gaussian $\mathcal{N}(s | \mu_j^S, \Sigma_j^S)$. The hidden state follows a categorical distribution with $z_t \sim \text{Cat}(\pi_{z_{t-1}})$ where $\pi_{z_{t-1}} \in \mathbb{R}^K$ is the next state transition distribution over state z_{t-1} with Π_i as the initial probability, and the observation ξ_t is drawn from the output distribution of state j , described by a multivariate Gaussian with parameters $\{\mu_j, \Sigma_j\}$. The overall parameter set for an HSMM is defined by $\left\{ \Pi_i, \{a_{i,m}\}_{m=1}^K, \mu_i, \Sigma_i, \mu_i^S, \Sigma_i^S \right\}_{i=1}^K$.

2.1 Encoding with HSMM

For learning and inference in a HMM [22], we make use of the intermediary variables as: 1) **forward variable**, $\alpha_{t,i}^{\text{HMM}} \triangleq P(z_t = i, \xi_1 \dots \xi_t | \theta)$: probability of a datapoint ξ_t to be in state i at time step t given the partial observation sequence $\{\xi_1, \dots, \xi_t\}$, 2) **backward variable**, $\beta_{t,i}^{\text{HMM}} \triangleq P(\xi_{t+1} \dots \xi_T | z_t = i, \theta)$: probability of the partial observation sequence $\{\xi_{t+1}, \dots, \xi_T\}$ given that we are in the i -th state at time step t , 3) **smoothed node marginal** $\gamma_{t,i}^{\text{HMM}} \triangleq P(z_t = i | \xi_1 \dots \xi_T, \theta)$: probability of ξ_t to be in state i at time step t given the full observation sequence ξ , and 4) **smoothed edge marginal** $\zeta_{t,i,j}^{\text{HMM}} \triangleq P(z_t = i, z_{t+1} = j | \xi_1 \dots \xi_T, \theta)$: probability of ξ_t to be in state i at time step t and in state j at time step $t + 1$ given the full observation sequence ξ . Parameters $\left\{ \Pi_i, \{a_{i,m}\}_{m=1}^K, \mu_i, \Sigma_i \right\}_{i=1}^K$ are estimated using the EM algorithm for HMMs, and the duration parameters $\{\mu_i^S, \Sigma_i^S\}_{i=1}^K$ are estimated empirically from the data after training using the most likely hidden state sequence $z_t = \{z_1 \dots z_T\}$ (see App. 7.1 for details).

2.2 Decoding from HSMM

Given the learned model parameters, the probability of the observed sequence $\{\xi_1 \dots \xi_t\}$ to be in a hidden state $z_t = i$ at the end of the sequence (also known as *filtering* prob-

lem) is computed with the help of the forward variable as

$$P(z_t \mid \xi_1, \dots, \xi_t) = h_{t,i}^{\text{HMM}} = \frac{\alpha_{t,i}^{\text{HMM}}}{\sum_{k=1}^K \alpha_{t,k}^{\text{HMM}}} = \frac{\pi_i \mathcal{N}(\xi_t \mid \mu_i, \Sigma_i)}{\sum_{k=1}^K \pi_k \mathcal{N}(\xi_t \mid \mu_k, \Sigma_k)}. \quad (1)$$

Sampling from the model for predicting the sequence of states over the next time horizon $P(z_t, z_{t+1}, \dots, z_{T_p} \mid \xi_1, \dots, \xi_t)$ can be done in two ways: **1) stochastic sampling:** the sequence of states is sampled in a probabilistic manner given the state duration and the state transition probabilities. By stochastic sampling, motions that contain different options and do not evolve only on a single path can also be represented. Starting from the initial state $z_t = i$, the s duration steps are sampled from $\{\mu_i^S, \Sigma_i^S\}$, after which the next transition state is sampled $z_{t+s+1} \sim \pi_{z_{t+s}}$. The procedure is repeated for the given time horizon in a receding horizon manner; **2) deterministic sampling:** the most likely sequence of states is sampled and remains unchanged in successive sampling trials. We use the forward variable of HSMM for deterministic sampling from the model. The forward variable $\alpha_{t,i}^{\text{HSMM}} \triangleq P(z_t = i, \xi_1 \dots \xi_t \mid \theta)$ requires marginalizing over the duration steps along with all possible state sequences. The probability of a datapoint ξ_t to be in state i at time step t given the partial observation sequence $\{\xi_1, \dots, \xi_t\}$ is now specified as [35]

$$\alpha_{t,i}^{\text{HSMM}} = \sum_{s=1}^{\min(s^{\max}, t-1)} \sum_{j=1}^K \alpha_{t-s,j}^{\text{HSMM}} a_{j,i} \mathcal{N}(s \mid \mu_i^S, \Sigma_i^S) \prod_{c=t-s+1}^t \mathcal{N}(\xi_c \mid \mu_i, \Sigma_i), \quad (2)$$

where the initialization is given by $\alpha_{1,i}^{\text{HSMM}} = \Pi_i \mathcal{N}(1 \mid \mu_i^S, \Sigma_i^S) \mathcal{N}(\xi_1 \mid \mu_i, \Sigma_i)$, and the output distribution in state i is conditionally independent for the s duration steps given as $\prod_{c=t-s+1}^t \mathcal{N}(\xi_c \mid \mu_i, \Sigma_i)$. Note that for $t < s^{\max}$, the sum over duration steps is computed for $t - 1$ steps, instead of $s^{\max}$. Without the observation sequence for the next time steps, the forward variable simplifies to

$$\alpha_{t,i}^{\text{HSMM}} = \sum_{s=1}^{\min(s^{\max}, t-1)} \sum_{j=1}^K \alpha_{t-s,j}^{\text{HSMM}} a_{j,i} \mathcal{N}(s \mid \mu_i^S, \Sigma_i^S). \quad (3)$$

The forward variable is used to plan the movement sequence for the next T_p steps with $t = t + 1 \dots T_p$. During prediction, we only use the transition matrix and the duration model to plan the future evolution of the initial/current state and omit the influence of the spatial data that we cannot observe, i.e., $\mathcal{N}(\xi_t \mid \mu_i, \Sigma_i) = 1$ for $t > 1$. This is used to retrieve a step-wise reference trajectory $\mathcal{N}(\hat{\mu}_t, \hat{\Sigma}_t)$ from a given state sequence z_t computed from the forward variable with,

$$z_t = \{z_t, \dots, z_{T_p}\} = \arg \max_i \alpha_{t,i}^{\text{HSMM}}, \quad \hat{\mu}_t = \mu_{z_t}, \quad \hat{\Sigma}_t = \Sigma_{z_t}. \quad (4)$$

Fig. 2 shows a conceptual representation of the step-wise sequence of states generated by deterministically sampling from HSMM encoding of the Z-shaped data. In the next section, we show how to synthesise robot movement from this step-wise sequence of states in a smooth manner.




Fig. 2: Sampling from HSMM from an unseen initial state ξ_0 over the next time horizon and tracking the step-wise desired sequence of states $\mathcal{N}(\hat{\mu}_t, \hat{\Sigma}_t)$ with a linear quadratic tracking controller. Note that this converges although ξ_0 was not previously encountered.

2.3 Motion Generation with Linear Quadratic Tracking

We formulate the motion generation problem given the step-wise desired sequence of states $\{\mathcal{N}(\hat{\mu}_t, \hat{\Sigma}_t)\}_{t=1}^{T_p}$ as sequential optimization of a scalar cost function with a linear quadratic tracker (LQT) [2]. The control policy $\mathbf{u}_t$ at each time step is obtained by minimizing the cost function over the **finite time horizon** T_p ,

$$c_t(\xi_t, \mathbf{u}_t) = \sum_{t=1}^{T_p} (\xi_t - \hat{\mu}_t)^\top \mathbf{Q}_t (\xi_t - \hat{\mu}_t) + \mathbf{u}_t^\top \mathbf{R}_t \mathbf{u}_t, \quad (5)$$

$$\text{s.t. } \xi_{t+1} = \mathbf{A}_d \xi_t + \mathbf{B}_d \mathbf{u}_t,$$

starting from the initial state ξ_1 and following the discrete linear dynamical system specified by $\mathbf{A}_d$ and $\mathbf{B}_d$. We consider a linear time-invariant double integrator system to describe the system dynamics. Alternatively, a time-varying linearization of the system dynamics along the reference trajectory can also be used to model the system dynamics without loss of generality. Both discrete and continuous time linear quadratic regulator/tracker can be used to follow the desired trajectory. The discrete time formulation, however, gives numerically stable results for a wide range of values of $\mathbf{R}$. The control law $\mathbf{u}_t^*$ that minimizes the cost function in Eq. (5) under finite horizon subject to the linear dynamics in discrete time is given as,

$$\mathbf{u}_t^* = \mathbf{K}_t (\hat{\mu}_t - \xi_t) + \mathbf{u}_t^{\text{FF}}, \quad (6)$$

where $\mathbf{K}_t = [\mathbf{K}_t^P, \mathbf{K}_t^V]$ are the full stiffness and damping matrices for the feedback term, and $\mathbf{u}_t^{\text{FF}}$ is the feedforward term (see App. 7.2 for computing the gains). Fig. 2




Fig. 3: Task-parameterized formulation of HSMM: four demonstrations on left are observed from two coordinate systems that define the start and end position of the demonstration (starting in purple position and ending in green position for each demonstration). The generative model is learned in the respective coordinate systems. The model parameters in respective coordinate systems are adapted to the new unseen object positions by computing the products of linearly transformed Gaussian mixture components. The resulting HSMM is combined with LQT for smooth retrieval of manipulation tasks.

shows the results of applying discrete LQT on the desired step-wise sequence of states sampled from an HSMM encoding of the Z-shaped demonstrations. Note that the gains can be precomputed before simulating the system if the reference trajectory does not change during the reproduction of the task. The resulting trajectory ξ_t^* smoothly tracks the step-wise reference trajectory $\hat{\mu}_t$ and the gains K_t^p, K_t^v locally stabilize ξ_t along ξ_t^* in accordance with the precision required during the task.

3 Invariant Task-Parameterized HSMMs

Conventional approaches to encode task variations such as change in the pose of the object is to augment the state of the environment with the policy parameters [21]. Such an encoding, however, does not capture the geometric structure of the problem. Our approach exploits the problem structure by introducing the task parameters in the form of coordinate systems that observe the demonstrations from multiple perspectives. Task-parameterization enables the model parameters to adapt in accordance with the external task parameters that describe the environmental situation, instead of hard coding the solution for each new situation or handling it in an *ad hoc* manner [31]. When a different situation occurs (pose of the object changes), changes in the task parameters/reference frames are used to modulate the model parameters in order to adapt the robot movement to the new situation.

3.1 Model Learning

We represent the task parameters with F coordinate systems, defined by $\{A_j, b_j\}_{j=1}^F$, where A_j denotes the orientation of the frame as a rotation matrix and b_j represents the origin of the frame. We assume that the coordinate frames are specified by the user, based on prior knowledge about the carried out task. Typically, coordinate frames will be attached to objects, tools or locations that could be relevant in the execution of a task. Each datapoint ξ_t is observed from the viewpoint of F different experts/frames,

with $\boldsymbol{\xi}_t^{(j)} = \mathbf{A}_j^{-1}(\boldsymbol{\xi}_t - \mathbf{b}_j)$ denoting the datapoint observed with respect to frame j . The parameters of the task-parameterized HSMM are defined by

$$\theta = \left\{ \left\{ \boldsymbol{\mu}_i^{(j)}, \boldsymbol{\Sigma}_i^{(j)} \right\}_{j=1}^F, \left\{ a_{i,m} \right\}_{m=1}^K, \boldsymbol{\mu}_i^S, \boldsymbol{\Sigma}_i^S \right\}_{i=1}^K,$$

where $\boldsymbol{\mu}_i^{(j)}$ and $\boldsymbol{\Sigma}_i^{(j)}$ define the mean and the covariance matrix of i -th mixture component in frame j . Parameter updates of the task-parameterized HSMM algorithm remain the same as HSMM, except the computation of the mean and the covariance matrix is repeated for each coordinate system separately. The emission distribution of the i -th state is represented by the product of the probabilities of the datapoint to belong to the i -th Gaussian in the corresponding j -th coordinate system. The forward variable of HMM in the task-parameterized formulation is described as

$$\alpha_{t,i}^{\text{TP-HMM}} = \left(\sum_{j=1}^K \alpha_{t-1,j}^{\text{HMM}} a_{j,i} \right) \prod_{j=1}^F \mathcal{N}(\boldsymbol{\xi}_t^{(j)} | \boldsymbol{\mu}_i^{(j)}, \boldsymbol{\Sigma}_i^{(j)}). \quad (7)$$

Similarly, the backward variable $\beta_{t,i}^{\text{TP-HMM}}$, the smoothed node marginal $\gamma_{t,i}^{\text{TP-HMM}}$, and the smoothed edge marginal $\zeta_{t,i,j}^{\text{TP-HMM}}$ can be computed by replacing the emission distribution $\mathcal{N}(\boldsymbol{\xi}_t | \boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$ with the product of probabilities of the datapoint in each frame $\prod_{j=1}^F \mathcal{N}(\boldsymbol{\xi}_t^{(j)} | \boldsymbol{\mu}_i^{(j)}, \boldsymbol{\Sigma}_i^{(j)})$. The duration model $\mathcal{N}(s | \boldsymbol{\mu}_i^S, \boldsymbol{\Sigma}_i^S)$ is used as a replacement of the self-transition probabilities $a_{i,i}$. The hidden state sequence over all demonstrations is used to define the duration model parameters $\{\boldsymbol{\mu}_i^S, \boldsymbol{\Sigma}_i^S\}$ as the mean and the standard deviation of staying s consecutive time steps in the i -th state.

3.2 Model Adaptation in New Situations

In order to combine the output from coordinate frames of reference for an unseen situation represented by the frames $\{\tilde{\mathbf{A}}_j, \tilde{\mathbf{b}}_j\}_{j=1}^F$, we linearly transform the Gaussians back to the global coordinates with $\{\tilde{\mathbf{A}}_j, \tilde{\mathbf{b}}_j\}_{j=1}^F$, and retrieve the new model parameters $\{\tilde{\boldsymbol{\mu}}_i, \tilde{\boldsymbol{\Sigma}}_i\}$ for the i -th mixture component by computing the products of the linearly transformed Gaussians (see Fig. 3)

$$\mathcal{N}(\tilde{\boldsymbol{\mu}}_i, \tilde{\boldsymbol{\Sigma}}_i) \propto \prod_{j=1}^F \mathcal{N}(\tilde{\mathbf{A}}_j \boldsymbol{\mu}_i^{(j)} + \tilde{\mathbf{b}}_j, \tilde{\mathbf{A}}_j \boldsymbol{\Sigma}_i^{(j)} \tilde{\mathbf{A}}_j^\top). \quad (8)$$

Evaluating the products of Gaussians represents the observation distribution of HSMM, whose output sequence is decoded and combined with LQT for smooth motion generation as shown in the previous section.

$$\tilde{\boldsymbol{\Sigma}}_i = \left(\sum_{j=1}^F \left(\tilde{\mathbf{A}}_j \boldsymbol{\Sigma}_i^{(j)} \tilde{\mathbf{A}}_j^\top \right)^{-1} \right)^{-1}, \quad \tilde{\boldsymbol{\mu}}_i = \tilde{\boldsymbol{\Sigma}}_i \sum_{j=1}^F \left(\tilde{\mathbf{A}}_j \boldsymbol{\Sigma}_i^{(j)} \tilde{\mathbf{A}}_j^\top \right)^{-1} \left(\tilde{\mathbf{A}}_j \boldsymbol{\mu}_i^{(j)} + \tilde{\mathbf{b}}_j \right). \quad (9)$$




Fig. 4: Parameters representation of a diagonal, full and mixture of factor analyzers decomposition of covariance matrix. Filled blocks represent non-zero entries.

4 Latent Space Representations

Dimensionality reduction has long been recognized as a fundamental problem in unsupervised learning. Model-based generative models such as HSMMs tend to suffer from the *curse of dimensionality* when few datapoints are available. We use statistical subspace clustering methods that reduce the number of parameters to be robustly estimated to address this problem. A simple way to reduce the number of parameters would be to constrain the covariance structure to a diagonal or spherical/isotropic matrix, and restrict the number of parameters at the cost of treating each dimension separately. Such decoupling, however, cannot encode the important motor control principles of coordination, synergies and action-perception couplings [32].

Consequently, we seek out a latent feature space in the high-dimensional data to reduce the number of model parameters that can be robustly estimated. We consider three formulations to this end: 1) low-rank decomposition of the covariance matrix using *Mixture of Factor Analyzers* (MFA) approach [16], 2) partial tying of the covariance matrices of the mixture model with the same set of basis vectors, albeit different scale with semi-tied covariance matrices [8,26], and 3) Bayesian non-parametric sequence clustering under small variance asymptotics [14,24,27]. All the decompositions can readily be combined with invariant task-parameterized HSMM and LQT for encapsulating reactive autonomous behaviour as shown in the previous section.

4.1 Mixture of Factor Analyzers

The basic idea of MFA is to perform subspace clustering by assuming the covariance structure for each component of the form,

$$\Sigma_i = \Lambda_i \Lambda_i^\top + \Psi_i, \quad (10)$$

where $\Lambda_i \in \mathbb{R}^{D \times d}$ is the *factor loadings matrix* with $d < D$ for parsimonious representation of the data, and Ψ_i is the diagonal noise matrix (see Fig. 4 for MFA representation in comparison to a diagonal and a full covariance matrix). Note that the mixture of probabilistic principal component analysis (MPPCA) model is a special case of MFA with the distribution of the errors assumed to be isotropic with $\Psi_i = I\sigma_i^2$ [29]. The MFA model assumes that ξ_t is generated using a linear transformation of d -dimensional vector of latent (unobserved) factors f_t ,

$$\xi_t = \Lambda_i f_t + \mu_i + \epsilon, \quad (11)$$

where $\mu_i \in \mathbb{R}^D$ is the mean vector of the i -th factor analyzer, $f_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ is a normally distributed factor, and $\epsilon \sim \mathcal{N}(\mathbf{0}, \Psi_i)$ is a zero-mean Gaussian noise with diagonal covariance Ψ_i . The diagonal assumption implies that the observed variables are independent given the factors. Note that the subspace of each cluster is not spanned by orthogonal vectors, whereas it is a necessary condition in models based on eigendecomposition such as PCA. Each covariance matrix of the mixture component has its own subspace spanned by the basis vectors of Σ_i . As the number of components increases to encode more complex skills, an increasing large number of potentially redundant parameters are used to fit the data. Consequently, there is a need to share the basis vectors across the mixture components as shown below by semi-tying the covariance matrices of the mixture model.

4.2 Semi-Tied Mixture Model

When the covariance matrices of the mixture model share the same set of parameters for the latent feature space, we call the model a *semi-tied* mixture model [26]. The main idea behind semi-tied mixture models is to decompose the covariance matrix Σ_i into two terms: a common latent feature matrix $\mathbf{H} \in \mathbb{R}^{D \times D}$ and a component-specific diagonal matrix $\Sigma_i^{(\text{diag})} \in \mathbb{R}^{D \times D}$, i.e.,

$$\Sigma_i = \mathbf{H} \Sigma_i^{(\text{diag})} \mathbf{H}^\top. \quad (12)$$

The latent feature matrix encodes the locally important synergistic directions represented by D non-orthogonal basis vectors that are shared across all the mixture components, while the diagonal matrix selects the appropriate subspace of each mixture component as convex combination of a subset of the basis vectors of $\mathbf{H}$. Note that the eigen decomposition of $\Sigma_i = \mathbf{U}_i \Sigma_i^{(\text{diag})} \mathbf{U}_i^\top$ contains D basis vectors of Σ_i in $\mathbf{U}_i$. In comparison, semi-tied mixture model gives D globally representative basis vectors that are shared across all the mixture components. The parameters $\mathbf{H}$ and $\Sigma_i^{(\text{diag})}$ are updated in closed form with EM updates of HSMM [8].

The underlying hypothesis in semi-tying the model parameters is that similar coordination patterns occur at different phases in a manipulation task. By exploiting the spatial and temporal correlation in the demonstrations, we reduce the number of parameters to be estimated while locking the most important synergies to cope with perturbations. This allows the reuse of the discovered synergies in different parts of the task having similar coordination patterns. In contrast, the MFA decomposition of each covariance matrix separately cannot exploit the temporal synergies, and has more flexibility in locally encoding the data.

4.3 Bayesian Non-Parametrics under Small Variance Asymptotics

Specifying the number of latent states in a mixture model is often difficult. Model selection methods such as cross-validation or Bayesian Information Criterion (BIC) are typically used to determine the number of states. Bayesian non-parametric approaches comprising of Hierarchical Dirichlet Processes (HDPs) provide a principled model selection procedure by Bayesian inference in an HMM with infinite number of states [28].




Fig. 5: Bayesian non-parametric clustering of Z-shaped streaming data under small variance asymptotics with: (left) online DP-GMM, (right) online DP-MPPCA. Note that the number of clusters and the subspace dimension of each cluster is adapted in a non-parametric manner.

These approaches provide flexibility in model selection, however, their widespread use is limited by the computational overhead of existing sampling-based and variational techniques for inference. We take a **small variance asymptotics** approximation of the Bayesian non-parametric model that collapses the posterior to a simple deterministic model, while retaining the non-parametric characteristics of the algorithm.

Small variance asymptotic (SVA) analysis implies that the covariance matrix Σ_i of all the Gaussians is set to the isotropic noise σ^2 , i.e., $\Sigma_i \approx \lim_{\sigma^2 \rightarrow 0} \sigma^2 \mathbf{I}$ in the likelihood function and the prior distribution [14,3]. The analysis yields simple deterministic models, while retaining the non-parametric nature. For example, SVA analysis of the Bayesian non-parametric GMM leads to the DP-means algorithm [14]. Similarly, SVA analysis of the Bayesian non-parametric HMM under Hierarchical Dirichlet Process (HDP) yields the segmental k -means problem [24].

Restricting the covariance matrix to an isotropic/spherical noise, however, fails to encode the coordination patterns in the demonstrations. Consequently, we model the covariance matrix in its intrinsic affine subspace of dimension d_i with projection matrix $\Lambda_i^{d_i} \in \mathbb{R}^{D \times d_i}$, such that $d_i < D$ and $\Sigma_i = \lim_{\sigma^2 \rightarrow 0} \Lambda_i^{d_i} \Lambda_i^{d_i \top} + \sigma^2 \mathbf{I}$ (akin to DP-MPPCA model). Under this assumption, we apply the small variance asymptotic limit on the remaining $(D - d_i)$ dimensions to encode the most important coordination patterns while being parsimonious in the number of parameters (see Fig. 5). Performing small-variance asymptotics of the joint likelihood of HDP-HMM yields the maximum a posteriori estimates of the parameters by iteratively minimizing the loss function*

$$\begin{aligned} \mathcal{L}(\mathbf{z}, \mathbf{d}, \boldsymbol{\mu}, \mathbf{U}, \mathbf{a}) = & \sum_{t=1}^T \text{dist}(\boldsymbol{\xi}_t, \boldsymbol{\mu}_{z_t}, \mathbf{U}_{z_t}^{d_i})^2 + \lambda(K - 1) \\ & + \lambda_1 \sum_{i=1}^K d_i - \lambda_2 \sum_{t=1}^{T-1} \log(a_{z_t, z_{t+1}}) + \lambda_3 \sum_{i=1}^K (\tau_i - 1), \end{aligned}$$

where $\text{dist}(\boldsymbol{\xi}_t, \boldsymbol{\mu}_{z_t}, \mathbf{U}_{z_t}^{d_i})^2$ represents the distance of the datapoint $\boldsymbol{\xi}_t$ to the subspace of cluster z_t defined by mean $\boldsymbol{\mu}_{z_t}$ and unit eigenvectors of the covariance matrix $\mathbf{U}_{z_t}^{d_i}$ (see App. 7.3). The algorithm optimizes the number of clusters and the subspace dimen-

*Setting $d_i = 0$ by choosing $\lambda_1 \gg 0$ gives the loss function formulation with isotropic Gaussian under small variance asymptotics [24].




Fig. 6: (left) Baxter robot picks the glass plate with a suction lever and places it on the cross after avoiding an obstacle of varying height, (centre-left) reproduction for previously unseen object and obstacle position, (centre-right) left-right HMM encoding of the task with duration model shown next to each state ($s^{\max} = 100$), (right) rescaled forward variable evolution of the forward variable over time.

sion of each cluster while minimizing the distance of the datapoints to the respective subspaces of each cluster. The λ_2 term favours the transitions to states with higher transition probability (states which have been visited more often before), λ_3 penalizes for transition to unvisited states with τ_i denoting the number of distinct transitions out of state i , while λ and λ_1 are the penalty terms for increasing the number of states and the subspace dimension of each output state distribution.

The analysis is used here for scalable online sequence clustering that is non-parametric in the number of clusters and the subspace dimension of each cluster. The resulting algorithm groups the data in its low dimensional subspace with non-parametric mixture of probabilistic principal component analyzers based on Dirichlet process, and captures the state transition and state duration information in a HDP-HSMM. The cluster assignment and the parameter updates at each iteration minimize the loss function, thereby, increasing the model fitness while penalizing for new transitions, new dimensions and/or new clusters. An interested reader can find more details of the algorithm in [27].

5 Experiments, Results and Discussion

We now show how our proposed work enables a Baxter robot to learn a pick-and-place task from a few human demonstrations. The objective of the task is to place the object in a desired target position by picking it from different initial poses of the object, while adapting the movement to avoid the obstacle. The setup of pick-and-place task with obstacle avoidance is shown in Fig. 6. The Baxter robot is required to grasp the glass plate with a suction lever placed in an initial configuration as marked on the setup. The obstacle can be vertically displaced to one of the 8 target configurations. We describe the task with two frames, one frame for the object initial configuration with $\{A_1, b_1\}$ and other frame for the obstacle $\{A_2, b_2\}$ with $A_2 = I$ and b_2 to specify the centre of the obstacle. We collect 8 kinesthetic demonstrations with different initial configurations of the object and the obstacle successively displaced upwards as marked with the visual tags in the figure. Alternate demonstrations are used for the training set, while the rest are used for the test set. Each observation comprises of the end-effector Cartesian position,




Fig. 7: Task-Parameterized HSMM performance on pick-and-place with obstacle avoidance task: (top) training set reproductions, (bottom) testing set reproductions.

quaternion orientation, gripper status (open/closed), linear velocity, quaternion derivative, and gripper status derivative with $D = 16$, $P = 2$, and a total of 200 datapoints per demonstration. We represent the frame $\{A_1, b_1\}$ as

$$A_1^{(n)} = \begin{bmatrix} R_1^{(n)} & \mathbf{0} & \mathbf{0} & \mathbf{0} & 0 \\ \mathbf{0} & \mathcal{E}_1^{(n)} & \mathbf{0} & \mathbf{0} & 0 \\ \mathbf{0} & \mathbf{0} & R_1^{(n)} & \mathbf{0} & 0 \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathcal{E}_1^{(n)} & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}, b_1^{(n)} = \begin{bmatrix} p_1^{(n)} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ 1 \end{bmatrix}, \quad (13)$$

where $p_1^{(n)} \in \mathbb{R}^3$, $R_1^{(n)} \in \mathbb{R}^{3 \times 3}$, $\mathcal{E}_1^{(n)} \in \mathbb{R}^{4 \times 4}$ denote the Cartesian position, the rotation matrix and the quaternion matrix in the n -th demonstration respectively. Note that we do not consider time as an explicit variable as the duration model in HSMM encapsulates the timing information locally.

Performance setting in our experiments is as follows: $\{\pi_i, \mu_i, \Sigma_i\}_{i=1}^K$ are initialized using k-means clustering algorithm, $R = 9I$, where I is the identity matrix, learning converges when the difference of log-likelihood between successive demonstrations is less than 1×10^{-4} . Results of regenerating the movements with 7 mixture components are shown in Fig. 7. For a given initial configuration of the object, the model parameters are adapted by evaluating the product of Gaussians for a new frame configuration. The reference trajectory is then computed from the initial position of the robot arm using the forward variable of HSMM and tracked using LQT. The robot arm moves from its initial configuration to align itself with the first frame $\{A_1, b_1\}$ to grasp the object, and follows it with the movement to avoid the obstacle and subsequently, align with the second frame $\{A_2, b_2\}$ before placing the object and returning to a neutral position. The model exploits variability in the observed demonstrations to statistically encode different phases of the task such as reach, grasp, move, place, return. The imposed




Fig. 8: Latent space representations of invariant task-parameterized HSMM for a randomly chosen demonstration from the test set. Black dotted lines show human demonstration, while grey line shows the reproduction from the model (see supplementary materials for details).

Table 1: Performance analysis of invariant hidden Markov models with training MSE, testing MSE, number of parameters for pick-and-place task. MSE (in meters) is computed between the demonstrated trajectories and the generated trajectories (lower is better). Latent space formulations give comparable task performance with much fewer parameters.

Model	Training MSE	Testing MSE	Number of Parameters
pick-and-place via obstacle avoidance ($K = 7, F = 2, D = 16$)			
HSMM	0.0026 ± 0.0009	0.014 ± 0.0085	2198
Semi-Tied HSMM	0.0033 ± 0.0016	0.0131 ± 0.0077	1030
MFA HSMM ($d_k = 1$)	0.0037 ± 0.0011	0.0109 ± 0.0068	742
MFA HSMM ($d_k = 4$)	0.0025 ± 0.0007	0.0119 ± 0.0077	1414
MFA HSMM ($d_k = 7$)	0.0023 ± 0.0009	0.0123 ± 0.0084	2086
SVA HDP HSMM ($K = 8, \bar{d}_k = 3.94$)	0.0073 ± 0.0024	0.0149 ± 0.0072	1352

structure with task-parameters and HSMM allows us to acquire a new task in a few human demonstrations, and generalize effectively in picking and placing the object.

Table 1 evaluates the performance of the invariant task-parameterized HSMM with latent space representations. We observe significant reduction in the model parameters, while achieving better generalization on the unseen situations compared to the task-parameterized HSMM with full covariance matrices (see Fig. 8 for comparison across models). It is seen that the MFA decomposition gives the best performance on test set with much fewer parameters.

6 Conclusions

Learning from demonstrations is a promising approach to teach manipulation skills to robots. In contrast to deep learning approaches that require extensive training data, generative mixture models are useful for learning from a few examples that are not explicitly labelled. The formulations are inspired by the need to make generative mixture models easy to use for robot learning in a variety of applications, while requiring considerably less learning time.

We have presented formulations for learning invariant task representations with hidden semi-Markov models for recognition, prediction, and reproduction of manipulation tasks; along with learning in latent space representations for robust parameter estimation of mixture models with high-dimensional data. By sampling the sequence of states from the model and following them with a linear quadratic tracking controller, we are able to autonomously perform manipulation tasks in a smooth manner. This has enabled a Baxter robot to tackle a pick-and-place via obstacle avoidance problem from previously unseen configurations of the environment. A relevant direction of future work is to not rely on specifying the task parameters manually, but to infer generalized task representations from the videos of the demonstrations in learning the invariant segments. Moreover, learning the task model from a small set of labelled demonstrations in a semi-supervised manner is an important aspect in extracting meaningful segments from demonstrations.

Acknowledgements: This work was, in large part, carried out at Idiap Research Institute and Ecole Polytechnique Federale de Lausanne (EPFL) Switzerland. This work was in part supported by the DexROV project through the EC Horizon 2020 program (Grant 635491), and the NSF National Robotics Initiative Award 1734633 on Scalable Collaborative Human-Robot Learning (SCHooL). The information, data, comments, and views detailed herein may not necessarily reflect the endorsements of the sponsors.

References

1. Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robot. Auton. Syst.*, 57(5):469–483, May 2009.
2. Francesco Borrelli, Alberto Bemporad, and Manfred Morari. *Predictive control for linear and hybrid systems*. Cambridge University Press, 2011.
3. Tamara Broderick, Brian Kulis, and Michael I. Jordan. Mad-bayes: Map-based asymptotic derivations from bayes. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pages 226–234, 2013.
4. S. Calinon. A tutorial on task-parameterized movement learning and retrieval. *Intelligent Service Robotics*, 9(1):1–29, 2016.
5. Yan Duan, Marcin Andrychowicz, Bradley C. Stadie, Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. One-shot imitation learning. *CoRR*, abs/1703.07326, 2017.
6. Nadia Figueroa and Aude Billard. Transform-invariant non-parametric clustering of covariance matrices and its application to unsupervised joint segmentation and action discovery. *CoRR*, abs/1710.10060, 2017.
7. Roy Fox, Sanjay Krishnan, Ion Stoica, and Ken Goldberg. Multi-level discovery of deep options. *CoRR*, abs/1703.08294, 2017.
8. Mark J. F. Gales. Semi-tied covariance matrices for hidden markov models. *IEEE Transactions on Speech and Audio Processing*, 7(3):272–281, 1999.
9. Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *CoRR*, abs/1606.03476, 2016.
10. A. Ijspeert, J. Nakanishi, P. Pastor, H. Hoffmann, and S. Schaal. Dynamical movement primitives: Learning attractor models for motor behaviors. *Neural Computation*, (25):328–373, 2013.
11. K. Khokar, R. Alqasemi, S. Sarkar, K. Reed, and R. Dubey. A novel telerobotic method for human-in-the-loop assisted grasping based on intention recognition. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4762–4769, 2014.

12. S. Krishnan, R. Fox, I. Stoica, and K. Goldberg. DDCO: Discovery of Deep Continuous Options for Robot Learning from Demonstrations. *CoRR*, 2017.
13. D. Kulic, W. Takano, and Y. Nakamura. Incremental learning, clustering and hierarchy formation of whole body motion patterns using adaptive hidden markov chains. *Intl Journal of Robotics Research*, 27(7):761–784, 2008.
14. Brian Kulis and Michael I. Jordan. Revisiting k-means: New algorithms via bayesian non-parametrics. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, pages 513–520, New York, NY, USA, 2012. ACM.
15. D. Lee and C. Ott. Incremental motion primitive learning by physical coaching using impedance control. In *Proc. IEEE/RSJ Intl Conf. on Intelligent Robots and Systems (IROS)*, pages 4133–4140, Taipei, Taiwan, October 2010.
16. G. J. McLachlan, D. Peel, and R. W. Bean. Modelling high-dimensional data by mixtures of factor analyzers. *Computational Statistics and Data Analysis*, 41(3-4):379–388, 2003.
17. Jose Medina R. and Aude Billard. Learning Stable Task Sequences from Demonstration with Linear Parameter Varying Systems and Hidden Markov Models. In *Conference on Robot Learning (CoRL)*, 2017.
18. Chrystopher L. Nehaniv and Kerstin Dautenhahn, editors. *Imitation and social learning in robots, humans, and animals: behavioural, social and communicative dimensions*. Cambridge University Press, 2004.
19. Scott Niekum, Sarah Osentoski, George Konidaris, and Andrew G Barto. Learning and generalization of complex tasks from unstructured demonstrations. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5239–5246, 2012.
20. Takayuki Osa, Joni Pajarinen, Gerhard Neumann, Andrew Bagnell, Pieter Abbeel, and Jan Peters. *An Algorithmic Perspective on Imitation Learning*. Now Publishers Inc., Hanover, MA, USA, 2018.
21. Alexandros Paraschos, Christian Daniel, Jan R Peters, and Gerhard Neumann. Probabilistic movement primitives. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 2616–2624. Curran Associates, Inc., 2013.
22. L. R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE*, 77:2:257–285, 1989.
23. M. Racca, J. Pajarinen, A. Montebelli, and V. Kyrki. Learning in-contact control strategies from demonstration. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 688–695, 2016.
24. Anirban Roychowdhury, Ke Jiang, and Brian Kulis. Small-variance asymptotics for hidden markov models. In *Advances in Neural Information Processing Systems 26*, pages 2103–2111. Curran Associates, Inc., 2013.
25. A. K. Tanwani. *Generative Models for Learning Robot Manipulation Skills from Humans*. PhD thesis, Ecole Polytechnique Federale de Lausanne, Switzerland, 2018.
26. A. K. Tanwani and S. Calinon. Learning robot manipulation tasks with task-parameterized semitized hidden semi-markov model. *IEEE Robotics and Automation Letters*, 1(1):235–242, 2016.
27. A. K. Tanwani and S. Calinon. Small variance asymptotics for non-parametric online robot learning. *CoRR*, abs/1610.02468, 2016.
28. Yee Whye Teh, Michael I. Jordan, Matthew J. Beal, and David M. Blei. Hierarchical dirichlet processes. *Journal of the American Statistical Association*, 101(476):1566–1581, 2006.
29. M. E. Tipping and C. M. Bishop. Mixtures of probabilistic principal component analyzers. *Neural Computation*, 11(2):443–482, 1999.
30. Yining Wang and Jun Zhu. DP-space: Bayesian nonparametric subspace clustering with small-variance asymptotics. In *Proc. of the 32nd International Conference on Machine Learning, ICML*, pages 862–870, 2015.

31. A. D. Wilson and A. F. Bobick. Parametric hidden Markov models for gesture recognition. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 21(9):884–900, 1999.
32. D. M. Wolpert, J. Diedrichsen, and J. R. Flanagan. Principles of sensorimotor learning. *Nature Reviews*, 12:739–751, 2011.
33. Danfei Xu, Suraj Nair, Yuke Zhu, Julian Gao, Animesh Garg, Li Fei-Fei, and Silvio Savarese. Neural task programming: Learning to generalize across hierarchical tasks. *CoRR*, abs/1710.01813, 2017.
34. C. Yang, J. Luo, C. Liu, M. Li, and S. Dai. Haptics electromyography perception and learning enhanced intelligence for teleoperated robot. *IEEE Transactions on Automation Science and Engineering*, pages 1–10, 2018.
35. S.-Z. Yu. Hidden semi-Markov models. *Artificial Intelligence*, 174:215–243, 2010.

7 Appendix

7.1 EM updates of HMM

The intermediary variables, namely **forward variable** $\alpha_{t,i}^{\text{HMM}}$, **backward variable** $\beta_{t,i}^{\text{HMM}}$, **smoothed node marginal** $\gamma_{t,i}^{\text{HMM}}$, and **smoothed edge marginal** $\zeta_{t,i,j}^{\text{HMM}}$ are mathematically represented as:

$$\begin{aligned}\alpha_{t,i}^{\text{HMM}} &= \left(\sum_{j=1}^K \alpha_{t-1,j}^{\text{HMM}} a_{j,i} \right) \mathcal{N}(\xi_t | \mu_i, \Sigma_i), & \beta_{t,i}^{\text{HMM}} &= \sum_{j=1}^K a_{i,j} \mathcal{N}(\xi_{t+1} | \mu_j, \Sigma_j) \beta_{t+1,j}^{\text{HMM}}, \\ \gamma_{t,i}^{\text{HMM}} &= \frac{\alpha_{t,i}^{\text{HMM}} \beta_{t,i}^{\text{HMM}}}{\sum_{k=1}^K \alpha_{t,k}^{\text{HMM}} \beta_{t,k}^{\text{HMM}}}, & \zeta_{t,i,j}^{\text{HMM}} &= \frac{\alpha_{t,i}^{\text{HMM}} a_{i,j} \mathcal{N}(\xi_{t+1} | \mu_j, \Sigma_j) \beta_{t+1,j}^{\text{HMM}}}{\sum_{k=1}^K \sum_{l=1}^K \alpha_{t,k}^{\text{HMM}} a_{k,l} \mathcal{N}(\xi_{t+1} | \mu_l, \Sigma_l) \beta_{t+1,l}^{\text{HMM}}}.\end{aligned}\quad (14)$$

The expected complete log-likelihood of HMMs for a set of M demonstrations, $\mathcal{Q}(\theta, \theta^{\text{old}}) = \mathbb{E} \left\{ \sum_{m=1}^M \sum_{t=1}^T \log \mathcal{P}(\xi_{m,t}, z_t | \theta) \mid \xi, \theta^{\text{old}} \right\}$, is maximized in an EM manner with

$$\begin{aligned}\mathcal{Q}(\theta, \theta^{\text{old}}) &= \sum_{i=1}^K \sum_{m=1}^M \gamma_{m,1,i}^{\text{HMM}} \log \Pi_i + \sum_{i=1}^K \sum_{j=1}^K \sum_{m=1}^M \sum_{t=1}^T \zeta_{m,t,i,j}^{\text{HMM}} \log a_{i,j} + \\ &\quad \sum_{m=1}^M \sum_{t=1}^T \sum_{i=1}^K \mathcal{P}(z_t = i | \xi_{m,t}, \theta^{\text{old}}) \log \mathcal{N}(\xi_{m,t} | \mu_i, \Sigma_i).\end{aligned}\quad (15)$$

$$E\text{-step: } \gamma_{m,t,i}^{\text{HMM}} = \frac{\alpha_{t,i}^{\text{HMM}} \beta_{t,i}^{\text{HMM}}}{\sum_{k=1}^K \alpha_{t,k}^{\text{HMM}} \beta_{t,k}^{\text{HMM}}},$$

$$\begin{aligned}M\text{-step: } \Pi_i &\leftarrow \frac{\sum_{m=1}^M \gamma_{m,1,i}^{\text{HMM}}}{M}, & a_{i,j} &\leftarrow \frac{\sum_{m=1}^M \sum_{t=1}^{T_m-1} \zeta_{m,t,i,j}^{\text{HMM}}}{\sum_{m=1}^M \sum_{t=1}^{T_m-1} \gamma_{m,t,i}^{\text{HMM}}}, \\ \mu_i &\leftarrow \frac{\sum_{m=1}^M \sum_{t=1}^{T_m} \gamma_{m,t,i}^{\text{HMM}} \xi_{m,t}}{\sum_{m=1}^M \sum_{t=1}^{T_m} \gamma_{m,t,i}^{\text{HMM}}}, & \Sigma_i &\leftarrow \frac{\sum_{m=1}^M \sum_{t=1}^{T_m} \gamma_{m,t,i}^{\text{HMM}} (\xi_{m,t} - \mu_i)(\xi_{m,t} - \mu_i)^\top}{\sum_{m=1}^M \sum_{t=1}^{T_m} \gamma_{m,t,i}^{\text{HMM}}}.\end{aligned}$$

Note that numerical underflow issues occur with a naive implementation of the above algorithm. In practice, a simple approach to avoid this issue is to rely on scaling factors during the computation of the forward and backward variables, which get canceled out when normalizing the posterior [22].

7.2 Linear Quadratic Tracking

The discrete-time dynamical system for the double integrator is defined as,

$$\overbrace{\begin{bmatrix} \xi_{t+1} \\ x_{t+1} \\ x_{t+2} \end{bmatrix}} = \overbrace{\begin{bmatrix} A_d \\ I & \Delta t \\ 0 & I \end{bmatrix}} \overbrace{\begin{bmatrix} \xi_t \\ x_t \\ x_{t+1} \end{bmatrix}} + \overbrace{\begin{bmatrix} B_d \\ I \frac{1}{2} \Delta t^2 \\ I \Delta t \end{bmatrix}} u_t. \quad (16)$$

The control law $\mathbf{u}_t^*$ that minimizes the cost function in Eq. (5) under **finite horizon** subject to the linear dynamics in discrete time is given as,

$$\begin{aligned} \mathbf{u}_t^* &= -(\mathbf{R} + \mathbf{B}_d^\top \mathbf{P}_t \mathbf{B}_d)^{-1} \mathbf{B}_d^\top \mathbf{P}_t \mathbf{A}_d (\boldsymbol{\xi}_t - \hat{\boldsymbol{\mu}}_t) - (\mathbf{R} + \mathbf{B}_d^\top \mathbf{P}_t \mathbf{B}_d)^{-1} \mathbf{B}_d^\top (\mathbf{P}_t (\mathbf{A}_d \hat{\boldsymbol{\mu}}_t - \hat{\boldsymbol{\mu}}_t) + \mathbf{d}_t), \\ &= \mathbf{K}_t^\mathcal{P} (\hat{\boldsymbol{\mu}}_t^x - \mathbf{x}_t) + \mathbf{K}_t^\mathcal{V} (\hat{\boldsymbol{\mu}}_t^{\dot{x}} - \dot{\mathbf{x}}_t) - (\mathbf{R} + \mathbf{B}_d^\top \mathbf{P}_t \mathbf{B}_d)^{-1} \mathbf{B}_d^\top (\mathbf{P}_t (\mathbf{A}_d \hat{\boldsymbol{\mu}}_t - \hat{\boldsymbol{\mu}}_t) + \mathbf{d}_t), \end{aligned} \quad (17)$$

where $[\mathbf{K}_t^\mathcal{P}, \mathbf{K}_t^\mathcal{V}] = -(\mathbf{R} + \mathbf{B}_d^\top \mathbf{P}_t \mathbf{B}_d)^{-1} \mathbf{B}_d^\top \mathbf{P}_t \mathbf{A}_d$ are the full stiffness and damping matrices for the feedback term, and $(\mathbf{R} + \mathbf{B}_d^\top \mathbf{P}_t \mathbf{B}_d)^{-1} \mathbf{B}_d^\top (\mathbf{P}_t (\mathbf{A}_d \hat{\boldsymbol{\mu}}_t - \hat{\boldsymbol{\mu}}_t) + \mathbf{d}_t)$ is the feedforward term. $\mathbf{P}_t$ and $\mathbf{d}_t$ are respectively obtained by solving the Riccati differential equation and linear differential equation backwards in discrete time from terminal conditions $\mathbf{P}_{T_p} = \mathbf{Q}_{T_p}$ and $\mathbf{d}_{T_p} = \mathbf{0}$,

$$\mathbf{P}_{t-1} = \mathbf{Q}_t - \mathbf{A}_d^\top (\mathbf{P}_t \mathbf{B}_d (\mathbf{R} + \mathbf{B}_d^\top \mathbf{P}_t \mathbf{B}_d)^{-1} \mathbf{B}_d^\top \mathbf{P}_t - \mathbf{P}_t) \mathbf{A}_d, \quad (18)$$

$$\mathbf{d}_{t-1} = (\mathbf{A}_d^\top - \mathbf{A}_d^\top \mathbf{P}_t \mathbf{B}_d (\mathbf{R} + \mathbf{B}_d^\top \mathbf{P}_t \mathbf{B}_d)^{-1} \mathbf{B}_d^\top) (\mathbf{P}_t (\mathbf{A}_d \hat{\boldsymbol{\mu}}_t - \hat{\boldsymbol{\mu}}_{t+1}) + \mathbf{d}_t) \quad (19)$$

For the **infinite horizon** case with $T \rightarrow \infty$ and the desired pose $\hat{\boldsymbol{\mu}}_t = \hat{\boldsymbol{\mu}}_{t_0}$, the control law in (17) remains the same except the feedforward term is set to zero and $\mathbf{P}_{t-1} = \mathbf{P}_t = \mathbf{P}$ is the steady-state solution obtained by eigen value decomposition of the discrete algebraic Riccati equation (DARE) in (18) [2]. To solve DARE, we define the symplectic matrix,

$$\mathbf{H}_b = \begin{bmatrix} \mathbf{A}_d + \mathbf{B}_d \mathbf{R}^{-1} \mathbf{B}_d^\top (\mathbf{A}_d^{-1})^\top \mathbf{Q} & \mathbf{B}_d \mathbf{R}^{-1} \mathbf{B}_d^\top (\mathbf{A}_d^{-1})^\top \\ -(\mathbf{A}_d^{-1})^\top \mathbf{Q} & (\mathbf{A}_d^{-1})^\top \end{bmatrix}. \quad (20)$$

The eigenvectors of $\mathbf{H}_b$ corresponding to eigenvalues lying inside the unit circle are used to solve DARE. Let $[\mathbf{V}_1^\top \mathbf{V}_{21}^\top]^\top$ denote the corresponding subspace of $\mathbf{H}_b$, then the solution of DARE is, $\mathbf{P} = \mathbf{V}_{21} \mathbf{V}_1^{-1}$ and the control law takes the form,

$$\mathbf{u}_t^* = -(\mathbf{R} + \mathbf{B}_d^\top \mathbf{P} \mathbf{B}_d)^{-1} \mathbf{B}_d^\top \mathbf{P} \mathbf{A}_d (\boldsymbol{\xi}_t - \hat{\boldsymbol{\mu}}_t). \quad (21)$$

Both discrete and continuous time linear quadratic regulator/tracker can be used to follow the desired pose/trajectory. The discrete time formulation, however, gives numerically stable results for a wide range of values of $\mathbf{R}$.

7.3 Distance to Cluster Subspace vs Distance to Cluster Mean

The distance of a datapoint $\boldsymbol{\xi}_t$ to an existing cluster with mean $\boldsymbol{\mu}_i$ is represented as: $\|\boldsymbol{\xi}_t - \boldsymbol{\mu}_i\|_2^2$. In contrast, we define the distance of a datapoint from the subspace of a cluster, $\text{dist}(\boldsymbol{\xi}_t, \boldsymbol{\mu}_i, \mathbf{U}_i^{d_i})^2$, as the difference between the mean-centered datapoint and the mean-centered datapoint projected upon the subspace $\mathbf{U}_i^{d_i} \in \mathbb{R}^{D \times d_i}$ spanned by the d_i unit eigenvectors of the covariance matrix, i.e.,

$$\text{dist}(\boldsymbol{\xi}_t, \boldsymbol{\mu}_i, \mathbf{U}_i^{d_i}) = \left\| (\boldsymbol{\xi}_t - \boldsymbol{\mu}_i) - \rho_i \mathbf{U}_i^{d_i} \mathbf{U}_i^{d_i \top} (\boldsymbol{\xi}_t - \boldsymbol{\mu}_i) \right\|_2, \quad (22)$$

where

$$\rho_i = \exp \left(-\frac{\|\boldsymbol{\xi}_t - \boldsymbol{\mu}_i\|_2^2}{b_m} \right)$$

weighs the projected mean-centered datapoint according to the distance of the datapoint from the cluster center ($0 < \rho_i \leq 1$). Its effect is controlled by the bandwidth parameter b_m . If b_m is large, then the far away clusters have a greater influence; otherwise nearby clusters are favored. Note that ρ_i assigns more weight to the projected mean-centered datapoint for the nearby clusters than the distant clusters to limit the size of the cluster/subspace. Our subspace distance formulation is different from [30] as we weigh the subspace of the nearby clusters more than the distant clusters. This allows us to avoid clustering all the datapoints in the same subspace (near or far) together.

7.4 Pick-and-Place with Obstacle Avoidance Results




Fig. 9: Latent space representations of invariant task-parameterized HSMM for a randomly chosen demonstration from the test set. Black dotted lines show human demonstration, while grey line shows the reproduction from the model.