Experience Selection Using Dynamics Similarity for Efficient Multi-Source Transfer Learning Between Robots

Michael J. Sorocky^{*}, Siqi Zhou^{*}, and Angela P. Schoellig

Abstract-In the robotics literature, different knowledge transfer approaches have been proposed to leverage the experience from a source task or robot-real or virtual-to accelerate the learning process on a new task or robot. A commonly made but infrequently examined assumption is that incorporating experience from a source task or robot will be beneficial. In practice, inappropriate knowledge transfer can result in negative transfer or unsafe behaviour. In this work, inspired by a system gap metric from robust control theory, the ν -gap, we present a data-efficient algorithm for estimating the similarity between pairs of robot systems. In a multi-source inter-robot transfer learning setup, we show that this similarity metric allows us to predict relative transfer performance and thus informatively select experiences from a source robot before knowledge transfer. We demonstrate our approach with quadrotor experiments, where we transfer an inverse dynamics model from a real or virtual source quadrotor to enhance the tracking performance of a target quadrotor on arbitrary hand-drawn trajectories. We show that selecting experiences based on the proposed similarity metric effectively facilitates the learning of the target quadrotor, improving performance by 62% compared to a poorly selected experience.

I. INTRODUCTION

Knowledge transfer or transfer learning has been used in robotics to accelerate the learning process for acquiring new skills, or to further improve a robots' performance (e.g., [1], [2]). At its core, transfer learning seeks to leverage existing experiences to initialize learning in a new domain that would otherwise be learned from scratch, thereby reducing the time and cost for training on physical or virtual robots. Examples of transfer learning in robotics include both (*i*) inter-task transfer learning, in which the experience learned for one task is used for learning new tasks on the same robot [2]–[4], and (*ii*) inter-robot transfer learning, where the experience from one robot is used to accelerate or improve the learning on a second robot [5]–[7].

While transfer learning has been demonstrated to improve robot performance, an implicit assumption that is less often examined is the similarity between the source and the target task or domain. As discussed in different transfer learning work [8]–[11], transferring inappropriately selected experiences can deteriorate learning and lead to worse performance compared to not using any source data this phenomenon is termed *negative transfer*. Approaches to avoid negative transfer include filtering out irrelevant source information [10], [11], or using data from multiple sources



Fig. 1. An illustration of the ν -gap inspired dynamics similarity characterization between five target (red) and source (blue) robot pairs—points closer together on the sphere (open dots) are seen as more similar. In this work, we consider a multi-source inter-robot transfer learning setup, where we use the proposed similarity characterization to informatively select a source experience from a repository that best facilitates the learning process on the target robot. A video of this work can be found here: http://tiny.cc/similarity

to increase the chance of positive transfer [12]. Although there exist approaches to encourage positive transfer and suppress negative transfer [8], negative transfer is usually detected only *after* the transferred experience is applied to the target. This after-the-fact negative transfer detection can be inefficient and unsafe for robotic applications, particularly when physical robots are involved in the training process.

As noted in [8], [10], the effectiveness of transfer learning typically depends on the similarity between the source and the target domain. Characterizing the similarity between the source and the target task or robot is thus crucial for understanding when transfer learning is safe and effective. In this paper, we consider a multi-source inter-robot transfer learning problem. Inspired by a notion of system gap from robust control theory, the ν -gap, we propose a data-efficient algorithm for characterizing similarity between two robots with unknown dynamics (Fig. 1). We show that the proposed similarity characterization, computed with simple experiments on periodic trajectories, allows us to predict relative transfer performance and thus informatively select the source robot that best facilities learning on the target robot. Our contributions in this paper are as follows:

1) Inspired by the ν -gap metric from robust control theory [13], we propose a data-efficient approach for characterizing the similarity between two robot sys-

^{*}Equal contribution. The authors are with the Dynamic Systems Lab (http://www.dynsyslab.org), Institute for Aerospace Studies, University of Toronto, Canada. The authors are also affiliated with the Vector Institute for Artificial Intelligence, Toronto. Emails: msorocky@robotics.utias.utoronto.ca, siqi.zhou@robotics.utias.utoronto.ca

tems whose exact dynamics are not necessarily known.

- Based on this notion of similarity, we introduce an optimization framework that allows us to systematically evaluate experience from different source robots before transferring the experience to the target robot.
- 3) We show in quadrotor experiments that an online learned inverse dynamics model with knowledge transferred from the most similar source quadrotor allows the target quadrotor to better track arbitrary handdrawn trajectories in single attempts, as compared to transfer from less similar robots.

II. RELATED WORK

Transfer learning has been considered in a wide range of robot learning problems, which include, but are not limited to, robot model learning [5], [14], [15], policy learning [1], [2], [16], and objective function learning [17], [18]. Whether it is inter-task transfer [2], [17], inter-robot transfer [5], [14]–[16], or transferring experience from simulation to reality [18], [19], the common goal in transfer learning is to leverage existing experiences to accelerate or improve the learning process of the target task or robot.

While in an inter-task transfer framework there exist discussions on characterizing task similarity for multi-source experience selection via a maximum likelihood model selection approach [17], or based on the difference of task descriptor representations [2], discussions on a similarity characterization of robot dynamics for inter-robot transfer problems are rare. As similarity and transfer performance are positively correlated concepts [8], [10], in inter-robot transfer problems, especially those involving robot model learning, characterizing the *dynamics similarity* of the robots is essential for predicting the effectiveness of transferring an experience. In this work, we consider a multi-source transfer learning setup similar to [2], [17], but we focus on an interrobot model transfer problem, where our goal is to select the best source robot to facilitate the learning of a target robot.

A concept that can be used for quantifying the distance between two dynamical systems is the notion of system gap from robust control theory [13]. In robust control, the incentive of characterizing system similarity is to provide theoretical guarantees in the presence of model uncertainties. For instance, for a classical feedback control architecture, it has been shown that an uncertain system can be stabilized by a controller designed based on a model that is sufficiently close in the sense of system gap [13], [20]. In addition to providing stability guarantees, gap metrics have also been applied to multi-linear model control, where a set of linear system models are used to approximate the dynamics of a nonlinear system. The control input to the nonlinear system is comprised of adding control inputs computed based on each linear system, each weighted with the gap metric between the linear system and a linearization of the nonlinear system at the current operating point [21], [22].

Although gap metrics have been shown as an appropriate measure of dynamics similarity between systems, the computation of the gap metrics requires either detailed system models [13], [23] or sufficiently large amounts of data [24], which can be challenging or non-economical to obtain for practical robot systems. In this work, inspired by the notion of system gap [13], we propose a data-efficient optimization algorithm for quantifying the similarity between two robots whose dynamics are unknown. We further leverage this similarity characterization in an inter-robot transfer learning problem, and show that the proposed similarity characterization allows us to effectively select the best source robot for improving the performance of the target robot.

III. PROBLEM FORMULATION

We consider a multi-source inter-robot transfer learning problem, where a target robot improves its performance through learning and has access to experience from a set of N source robots, which can be leveraged to facilitate its own learning. Our goal is to develop an algorithm that allows us to quantify the similarity between each source and the target robot such that we can efficiently select experiences from a source robot that best facilitates the learning on the target robot. We aim to show that selecting an experience from a source robot similar to the target robot in the sense of the proposed metric will result in better transfer performance.

To analyze the proposed similarity characterization, we consider an inverse dynamics learning framework shown in Fig. 2. Through leveraging the inverse dynamics from a source robot, we aim to enhance the impromptu tracking performance of the target robot on arbitrary hand-drawn trajectories using *minimal online data* from the target robot.

We denote the input and output of the inverse dynamics module by ξ and γ , and assume the inverse dynamics of the target robot are represented by a Gaussian Process (GP),

$$F_T(\xi) \sim \mathcal{GP}\left(\mu(\xi), \kappa(\xi, \xi')\right),$$
 (1)

where $\mu(\cdot)$ and $\kappa(\cdot, \cdot)$ are the prior mean and covariance function of the GP, respectively. To facilitate the training of the GP inverse dynamics model of the target robot, we leverage the inverse dynamics model of a previously trained source robot as the prior mean function. The posterior mean $f_T(\xi)$ and variance $\sigma_T^2(\xi)$ of the GP at a new input ξ_* are

$$f_{T}(\xi_{*}) = f_{S_{n}}(\xi_{*}) + \mathbf{k}^{T}(\xi_{*}) \left(\mathbf{K} + \sigma_{n}^{2}I\right)^{-1} (\boldsymbol{\gamma} - \mathbf{f}_{S_{n}})$$

$$\sigma_{T}^{2}(\xi_{*}) = \kappa(\xi_{*}, \xi_{*}) - \mathbf{k}^{T}(\xi_{*}) \left(\mathbf{K} + \sigma_{n}^{2}I\right)^{-1} \mathbf{k}(\xi_{*}),$$
 (2)

where $f_{S_n}(\cdot)$ is a deterministic inverse dynamics model of source robot $n, \gamma = [\gamma_1, \ldots, \gamma_D]^T$ and $\mathbf{f}_{S_n} = [f_{S_n}(\xi_1), \ldots, f_{S_n}(\xi_D)]^T$ with $\{\xi_d, \gamma_d\}_{d=1}^{d=D}$ being the inputoutput data collected online from the target robot, $\mathbf{k}(\xi_*) := [\kappa(\xi_*, \xi_1), \ldots, \kappa(\xi_*, \xi_D)]^T$, $\mathbf{K} \in \mathbb{R}^{D \times D}$ is the covariance matrix with entries $[\mathbf{K}]_{ij} = \kappa(\xi_i, \xi_j), \sigma_n^2$ is the noise variance, and $I \in \mathbb{R}^{D \times D}$ is the identity matrix.

Following the discussion in [25], for the inverse dynamics learning approach to be safely applied, we assume that the target and the source robot systems are minimum-phase (i.e., are stable and have stable inverse dynamics). This assumption is not restrictive and is satisfied by practical robot systems such as quadrotors [25] and robot manipulators [26].



Fig. 2. Block diagram of an inverse learning approach for enhancing the tracking performance of black-box systems on arbitrary trajectories [25], [27]. The inverse dynamics module (blue box) represents the inverse of the target robot baseline system (grey box); it is pre-cascaded to the target robot baseline system to adjusts the reference signal y_r to reduce the trajectory tracking error of the baseline system (i.e., $e = y_d - y_a$, where y_d denotes the desired output and y_a denotes the actual output of the target robot system). The inverse dynamics module is implemented as a GP and the inverse dynamics model of a source system is used as the prior mean function to facilitate online learning of the target inverse dynamics (Eqn. (2)).

IV. EXPERIENCE SELECTION THROUGH A CHARACTERIZATION OF ROBOT DYNAMICS SIMILARITY

In this section, we outline an algorithm for prior mean selection based on a characterization of similarity between robot dynamics. To facilitate our discussion, in Sec. IV-A, we first present a brief background on the ν -gap metric from robust control theory. Inspired by the ν -gap metric, we formulate an optimization problem based on a characterization of system similarity for experience selection in Sec. IV-B. In Sec. IV-C, we propose a Bayesian Optimization (BO)-based algorithm with a novel acquisition function to solve this optimization problem for robots with unknown dynamics.

A. Background on the ν -Gap Metric

One notion of system similarity is the ν -gap metric, which provides a measure of distance between two linear dynamical systems [28]. In this subsection, we present the essential concepts to facilitate our discussion. An in-depth analysis of the metric can be found in [13], [28], [29].

For linear dynamical systems, the response to a periodic input u is a periodic output y with the same frequency but with different magnitude and phase. Given a periodic input with frequency ω to a dynamical system, we define a complex number $z(\omega) = P(j\omega) = A(\omega)e^{j\theta(\omega)}$ to characterize the input-output response at the input frequency ω , where P(s) is the transfer function of the system, $A(\omega)$ and $\theta(\omega)$ are the amplitude gain and phase shift in the input-output response, and j is the imaginary unit.

In [13], it is shown that for two minimum phase systems P_1 , P_2 , under the condition that $||P_1(-j\omega)P_2(j\omega)||_{\infty} < 1$, the ν -gap metric can be computed by

$$\delta_{\nu}(z_1, z_2) := \sup_{\omega \in \mathbb{R}} \psi(z_1(\omega), z_2(\omega)), \tag{3}$$

where z_1 , z_2 characterize the input-output response of P_1 , P_2 for an input frequency ω , and

$$\psi(z_1(\omega), z_2(\omega)) = \frac{|z_1(\omega) - z_2(\omega)|}{\sqrt{1 + |z_1(\omega)|^2}\sqrt{1 + |z_2(\omega)|^2}}.$$
 (4)

The function ψ in Eqn. 4 can be given a geometric interpretation as the Euclidean distance between the projection of two points $z_1(\omega), z_2(\omega) \in \mathbb{C}$ onto the Riemann sphere, a sphere of radius 0.5 tangent to the origin of the complex plane. The solid red and blue dots in Fig. 1 are projected onto the sphere by connecting a line between each point and the point (0, 0, 1) (the *North Pole*) on the sphere — the intersection of these lines and the sphere (the open red/blue dots in Fig. 1) are their respective projections. The function ψ gives the Euclidean distance between the projected points (open red and blue dots in Fig. 1) — this distance is termed the *chordal distance*. The ν -gap metric represents the largest chordal distance over all frequencies ω . It can be shown that δ_{ν} satisfies $0 \le \delta_{\nu}(z_1, z_2) \le 1$ [13]. We use the ν -gap metric as a measure of similarity, interpreting values of δ_{ν} closer to 0 meaning the two systems are more similar.

B. Experience Selection as an Optimization Problem

Given a set of N source robots S, we aim to select an experience from the source robot with the most similar dynamics to the target to improve the performance of the target robot. We denote the *n*th source robot by R_{S_n} , and the target robot by R_T . The gain A_R and phase shift θ_R of a robot's response to an input with frequency ω are characterized by

$$z_R(\omega) = A_R(\omega) e^{j\theta_R(\omega)},$$
(5)

where R denotes the target robot R_T or a source robot R_{S_n} .

Inspired by the ν -gap metric, we propose the following characterization of dynamics similarity between a source and target robot over a set of operating frequencies of interest $\mathcal{W} := [\omega_{\min}, \omega_{\max}]$:

$$\psi_n^* := \max_{\omega \in \mathcal{W}} \psi_n(\omega), \tag{6}$$

where using Eqn. (4) we define

$$\psi_n(\omega) := \psi(z_T(\omega), z_{S_n}(\omega)). \tag{7}$$

Analogous to the ν -gap metric, it holds that $0 \le \psi_n^* \le 1$, and we interpret values of ψ_n^* close to 0 indicating the two systems are similar. In contrast to $\delta_{\nu}(\cdot, \cdot)$ in Eqn. (3), we consider a finite frequency range of interest \mathcal{W} .

With this notion of dynamics similarity, we formulate experience selection as the following optimization problem:

$$R_S^* := \underset{n}{\operatorname{arg\,min}} \ \psi_n^*, \tag{8}$$

where ψ_n^* is the characterization of dynamics similarity in Eqns. (6) and (7). Intuitively, the optimization problem in Eqn. (8) selects the source system R_S^* whose maximum chordal distance (i.e., the worst-case distance) from the target system over the range W is the minimal. Given the inverse dynamics learning framework in Fig. 2, by considering the worst-case chordal distance between the target and the source robot systems, we select a source inverse dynamics model that facilitates high-performance tracking for arbitrary trajectories over the range W.

Note that in robust control, system similarity metrics are often computed based on detailed models of the systems. In the next subsection, we outline a data-efficient algorithm to estimate dynamics similarity and select an experience based on Eqn. (8) for robots with possibly unknown dynamics.

C. Experience Selection for Source and Target Robot Systems with Unknown Dynamics

We first outline an approach to estimate the similarity metric in Eqn. (6) using data from a source robot and a target robot, then extend the solution to the multi-source experience selection problem in Eqn. (8). To solve the experience selection problem, we first note that while the objective function $\psi_n(\omega)$ is unknown, we can sample from it at selected frequencies ω_{sample} . In particular, given a frequency ω_{sample} , we can send a periodic input to the source and the target robots to estimate the amplitude gain $A(\omega_{\text{sample}})$ and phase shift $\theta(\omega_{\text{sample}})$ from input-output data. This information allows us to calculate $z_T(\omega_{\text{sample}})$ and $z_{S_n}(\omega_{\text{sample}})$ defined in Eqn. (5), and $\psi_n(\omega_{\text{sample}})$ defined in Eqn. (7).

For solving the problem in Eqn. (6), we propose a Bayesian Optimization (BO)-based algorithm to find the maximum of ψ_n . In particular, we model ψ_n by a GP and find its maximum by iteratively sampling the objective function based on an acquisition function. We denote the GP posterior mean and variance by $\hat{\psi}_n(\omega)$ and $\hat{\sigma}_n^2(\omega)$. Given M previous samples $\{\psi_n(\omega_m)\}_{m=1}^M$, we define $\psi_{n,\max} := \max_k \psi_n(\omega_k)$. A common BO acquisition function is the expected improvement (EI), which is given for $\omega \in \mathcal{W}$ by

$$\operatorname{EI}_{n}(\omega) = (\hat{\psi}_{n}(\omega) - \psi_{n,\max} - \xi)\Phi(Z) + \hat{\sigma}_{n}(\omega)\phi(Z) \quad (9)$$

if $\hat{\sigma}_n(\omega) \neq 0$, and 0 otherwise, where $Z := \frac{\hat{\psi}_n(\omega) - \psi_{n,\max} - \xi}{\hat{\sigma}_n(\omega)}$, Φ and ϕ are the cumulative and probability density functions of the standard normal distribution, and ξ is an exploration parameter [30].

The maximum of ψ_n is found through an iterative process: we sample ψ_n at $\omega_{\text{sample}} := \arg \max_{\omega \in \mathcal{W}} EI(\omega)$, we refit the GP with the new sample, and the process is repeated. Upon convergence, the estimate of similarity $\hat{\psi}_n^*$ between the target robot and the *n*th source robot is given by the maximum of the posterior mean of its GP approximation:

$$\hat{\psi}_n^* := \max_{\omega \in \mathcal{W}} \hat{\psi}_n(\omega), \tag{10}$$

with a smaller ψ_n^* representing a more similar source robot.

For solving the experience selection problem in Eqn. (8), we use a GP to model the unknown chordal distance function ψ_n for each source and target robot pair. To account for the outer-loop minimization problem in Eqn. (8), we modify the acquisition function as follows:

$$\alpha(\omega) := \max_{R_{S_n} \in \mathcal{S}} \operatorname{EI}_n(\omega) / \hat{\psi}_n^*, \tag{11}$$

where $\operatorname{EI}_n(\omega)$ is the expected improvement of the point ω for the *n*-th GP, and $\hat{\psi}_n^*$ is the similarity estimate for R_{S_n} . Intuitively, the acquisition function in Eqn. (11) encourages the selection of a sample ω_{sample} at the worst-case chordal distance for more similar robots.

Based on this BO framework, we outline our proposed algorithm to determine system similarity for experience selection in Alg. 1. The BO algorithm is initialized with a uniformly sampled ω_{sample} . We sample from each GP by running periodic trajectories on each source and target system until convergence. Given the estimates of similarity

Algorithm 1 Source/Target Robot Similarity Estimation

- **Input:** One target robot, N source robots and a frequency operating range W
- **Output:** A similarity estimate of each source robot to the target, $\hat{\psi}_n^*$, and the source robot R_S^* most similar to the target over \mathcal{W}
- **Initialize:** Empty datasets: $\mathcal{D}_{S_n} \leftarrow \emptyset \ (n = 1, \ \dots, \ N)$
- 1: Compute initial sample location: $\omega_{\text{sample}} \sim \mathcal{U}(\omega_{\min}, \omega_{\max})$
- 2: while not converged do 3: Estimate $z_T(\omega_{\text{sample}})$ 4: for n = 1, ..., N do 5: Estimate $z_{S_n}(\omega_{\text{sample}})$ 6: Compute $\psi_n(\omega_{\text{sample}})$ (Eqn. (7))
- 7: $\mathcal{D}_{S_n} \leftarrow \mathcal{D}_{S_n} \cup \{(\omega_{\text{sample}}, \psi_n(\omega_{\text{sample}}))\}$
- 8: Fit *n*th GP with new data \mathcal{D}_{S_n}
- 9: end for
- 10: Compute $\omega_{\text{sample}} = \arg \max_{\omega \in \mathcal{W}} \alpha(\omega)$

11: end while
12:
$$\hat{\psi}_n^* := \max_{\omega} \hat{\psi}_n(\omega)$$
 (Eqn. (10)), and $R_s^* = \arg \min_n \hat{\psi}_n^*$

between each source robot and the target robot, the most similar robot R_S^* has minimum $\hat{\psi}_n^*$: $R_S^* = \arg \min_n \hat{\psi}_n^*$, solving the experience selection problem posed in Eqn. (8).

V. SIMULATION

We first illustrate the proposed algorithm (Alg. 1) for characterizing dynamics similarity between robots in simulation.

A. Simulation Setup

We consider a target robot and a set of three source robots whose dynamics are represented by

$$\mathbf{x}(k+1) = \begin{bmatrix} 1 & T_s - \beta \\ 0 & 1 - \beta \end{bmatrix} \mathbf{x}(k) + \alpha \begin{bmatrix} \frac{1}{2}T_s^2 \\ T_s \end{bmatrix} u(k)$$

$$y_a(k) = \begin{bmatrix} 1 & 0 \end{bmatrix} \mathbf{x}(k),$$
(12)

where $k \in \mathbb{Z}_{\geq 0}$ is the discrete-time index, $\mathbf{x} \in \mathbb{R}^2$ is the state, $y_a \in \mathbb{R}$ is the output, $u \in \mathbb{R}$ is the command sent to the robot, $T_s = 0.015$ sec is the sampling time, and (α, β) are two system parameters that differentiate the source and the target robot dynamics. In this simulation, the parameters for the target robot are (0.85, 0.003); for the three source robots, the parameters are (1.0, 0.003), (0.97, 0.004), and (0.9, 0.001), respectively. As shown in Fig. 2, we aim to improve a baseline tracking controller using online learning. For illustration purposes, we use a simple PD baseline tracking controller defined as follows: $u(k) = k_p e(k) + k_d \dot{e}(k)$, where $k_p = 5$ and $k_d = 4.5$ are the baseline controller parameters and $e(k) = y_r(k) - y_a(k)$. Note that, when applying the proposed algorithm, we assume that the dynamics of the source and the target robot baseline systems are unknown: we only use the input and output data of the source and the target robots from frequency tests.

Motivated by the inverse dynamics learning application (Fig. 2), we are interested in characterizing the similarity between the source and the target robot baseline systems (i.e., the dynamics from y_r to y_a). In this study, we apply Alg. 1 to estimate the similarity of the three target/source robot pairs. For the simulation, we set $\mathcal{W} = [0, 10]$ rad/sec,



Fig. 3. Illustration of the proposed dynamics similarity characterization outlined in Alg. 1. The top and bottom panels respectively show the similarity estimation in an earlier and a later iteration of the algorithm. The true chordal distance $\psi_n(\omega)$ between a source and a target pair and its GP approximation $\hat{\psi}_n(\omega)$ are represented by the dashed line and solid line, respectively. The maximum of $\psi_n(\omega)$ for each source and target robot pair is indicated by '+', and the sampled training points are indicated by '•'.

and the algorithm is initialized with a randomly generated input frequency ω_{sample} . In each iteration, we send a periodic reference trajectory y_T with frequency ω_{sample} to the source and target robots. Based on the input-output response, we estimate z_T and z_{S_n} to calculate $\psi_n(\omega_{\text{sample}})$ in Eqn. (7). For each source robot, iteratively augmented data \mathcal{D}_{S_n} is used to fit a GP model and estimate the similarity $\hat{\psi}_n^*$ of the source and target robot pair. We use GPs with a zero prior mean function and the Matérn 3/2 kernel; the kernel hyperparameters are optimized in each iteration to maximize the marginal likelihood of the GP model. As formulated in Eqn. (8), the 'best' source robot is the one that minimizes the worst-case chordal distance, $\max_{\omega \in \mathcal{W}} \psi_n(\omega)$.

B. Simulation Results

Figure 3 shows the true chordal distance ψ_n (dashed line), and its GP approximation ψ_n (solid line) for each source/target pair in an earlier and a later BO iteration. As can be seen from Fig. 3, the proposed algorithm encourages sampling around the maxima of the chordal distance curves for each source/target pair. This allows us to effectively estimate the worst-case distance between the source and the target robot and quantify their similarity. The similarity estimation converges to the maxima of the true function ψ_n (indicated by '+') in seven iterations. In contrast to traditional approaches that require either an fully identified model or sufficient excitation of the robots, we see that the proposed algorithm requires minimal data from a few simple frequency tests. The converged similarity estimates $\hat{\psi}_n^*$ are 0.069, 0.017, and 0.087, respectively, for the three source/target robot pairs, respectively, with $R_S^* = R_{S_2}$. In the next section, we demonstrate the similarity characterization for experience selection in quadrotor tracking experiments.



Fig. 4. Similarity estimates between the target quadrotor and the five real/virtual source quadrotors: $\hat{\Psi}_n^* \sim \mathcal{N}(\hat{\psi}_n^*, \hat{\sigma}_n^{*2})$, where $\hat{\psi}_n^*$ is the similarity estimate defined in Eqn. (10) and $\hat{\sigma}_n^{*2}$ is the corresponding variance. The thin lines correspond to estimates at the second iteration of Alg. 1, while the thick lines correspond to their converged distributions in the ninth (final) iteration. Initially, sources 1 and 2 are equally similar to the target, while estimates for sources 3–5 are very uncertain. At the final iteration, the estimates are more certain that sources 1 and 2 are more similar to the target robot, in contrast to sources 3–5 which are less similar. The first source robot is most similar to the target (i.e., $R_s^* = R_{S_1}$).

VI. QUADROTOR EXPERIMENTS

In this section, we demonstrate the experience selection algorithm from Sec. IV-C with quadrotor experiments.

A. Experiment Setup

The goal of the experiments is to improve the tracking performance for a target quadrotor R_T Bitcraze Crazyflie 2.0 (mass 36g) over its feasible operating frequencies $\mathcal{W} =$ [0.1, 5.4] rad/sec. We consider five source quadrotors shown in Fig. 1: (R_{S_1}) Bitcraze Crazyflie 2.0 (mass 31g), (R_{S_2}) Crazyflie Simulator, (R_{S_3}) Parrot ARDrone 2.0, (R_{S_4}) ARDrone Simulator, and (R_{S_5}) Parrot Bebop 2. They have different baseline tracking controller designs and physical properties. The mass of source quadrotors ranges between 0.86–14.5x that of the target, and the rotor-to-rotor distance of the source quadrotors ranges between 1-3.9x that of the target. As in our earlier work [25], [27], each source quadrotor has a DNN inverse dynamics model previously trained on a rich dataset consisting of 2100 data points. The DNNs of the source quadrotors are fully-connected networks with four hidden layers of 128 ReLU neurons. To facilitate online learning of the target inverse dynamics model (Fig. 2), we are interested in characterizing the similarity between the source and the target robot baseline systems to select the source quadrotor with the most similar closed-loop dynamics.

B. Quadrotor Experience Selection

In this subsection, we present the results of applying Alg. 1 to the inter-quadrotor transfer problem. To estimate $z_T(\omega_{\text{sample}})$ and $z_{S_n}(\omega_{\text{sample}})$, we run a sinusoidal trajectory of frequency ω_{sample} along the x-axis on each target and source system, using their respective *baseline* tracking controllers (i.e., we do *not* use any inverse modules in this step). We assume that the response along the y-axis of each quadrotor is similar to their response along the x-axis. We estimate the magnitude $A(\omega_{\text{sample}})$ and phase $\theta(\omega_{\text{sample}})$ for each quadrotor from its input-output response data.

We ran Alg. 1 for nine iterations and observed that the estimated similarity $\hat{\psi}_n^*$ between the source and the target robot systems converged to approximately constant values. Fig. 4 shows the posterior distributions $\hat{\Psi}_n^* \sim \mathcal{N}(\hat{\psi}_n^*, \hat{\sigma}_n^{*2})$ of the estimated similarity in the second and the final iteration



Fig. 5. Five hand drawings for generating desired trajectories (y_d) in the quadrotor experiments. Trajectories are scaled based on three maximum speeds to test the performance for different aggressiveness of maneuvers.



Fig. 6. A summary of error reduction as compared to the baseline controller using the inverse dynamics model from each source system at different operating speeds. Using an inverse dynamics model from quadrotors deemed more similar by the proposed metric (Fig. 4) yields better tracking performance, while using the inverse from a less similar quadrotor (source 5) can even result in *worse* performance as compared to the baseline controller.

of the algorithm, where $\hat{\psi}_n^*$ is the similarity estimate defined in Eqn. (10), and $\hat{\sigma}_n^{*^2}$ is the corresponding variance. We observe the progression of the similarity estimates from the second to the final iteration of Alg. 1. As seen in Fig. 4, sources 1 and 2 are more similar to the target robot and sources 3–5 are less similar. The most similar robot is R_{S_1} , and is expected to best facilitate the target quadrotor inverse learning. As we verify in the next set of experiments, the proposed experience selection algorithm allows us to infer the *relative* performance of transferring experiences from different source robots *before* testing on the target robot.

C. Inverse Dynamics Learning and Tracking Experiments

In this subsection, we verify the efficacy of the experience selection algorithm in the framework of Fig. 2. Specifically, we validate that using an inverse dynamics module from a more similar robot yields improved performance compared to using a module from a less similar robot. We modify the reference y_r to the target baseline system using a GP inverse dynamics model trained on small amounts of online data from the target robot. The GP has a prior mean function given by a previously trained source DNN inverse dynamics model, and a Matérn 3/2 kernel function. The kernel hyperparameters are initialized for all source robots as $\sigma_n^2 = 0.002$ (measurement noise), $\sigma_f^2 = 0.1$ (prior variance), and $\ell = 0.5$ (lengthscale) [31]. We implement the online module using the Python GPy library, and optimize hyperparameters online by maximizing the marginal likelihood of the GP.

We select the input and output of the inverse dynamics module based on [25], and fit the GP with the latest 50 data points at each timestep. We evaluate tracking performance on 15 hand-drawn trajectories in the x-y plane by scaling the drawings in Fig. 5 based on three maximum speeds

TABLE I: Summary of Transfer Performance Over 15 Trajectories

	Source 1	Source 2	Source 3	Source 4	Source 5
$(\hat{\psi}_n^*)$	(0.07)	(0.19)	(0.44)	(0.48)	(0.52)
e _{rms} [m]	0.09 67.2%	0.10 66.2%	0.12 57.0%	0.14 52.1%	0.25
-70 Keu.	07.270	00.2%	57.9%	52.170	12.3%

* *e*_{rms} denotes the average RMS tracking error of the target robot with different source experience, and '% Red.' is the average tracking error reduction as compared to the baseline performance.

 $\{0.3, 0.6, 0.9\}$ m/s. We use the root-mean-square (RMS) tracking error as the performance metric.

Figure 6 summarizes the performance of the target quadrotor with inverse dynamics modules transferred from the five different source quadrotors (Fig. 2) on the fifteen test trajectories. We note a general correlation between the proposed similarity metric and performance: using the inverse dynamics module from more similar robots results in improved performance. Table I summarizes the performance of the target robot over the 15 test trajectories when the experience from different source quadrotors is used, again demonstrating that source experience selected by the proposed similarity metric leads to improved transfer performance. Using the experience from the most similar source quadrotor (R_{S_1}) yields 62% better performance than using the experience from an otherwise poorly selected source quadrotor $(R_{S_{\pi}})$. These results confirm that the proposed experience selection algorithm based on dynamics similarity is able to predict relative transfer performance with multiple source robots.

VII. DISCUSSION

In this work, we showed that the proposed experience selection algorithm can inform the *relative* transfer performance with different source robots. However, in practice, it would be beneficial to have a threshold on robot similarity that will guarantee positive transfer. One may show that the ν -gap metric is linked to a lower bound on the tracking error for linear systems. A valuable extension is providing theoretical guarantees on the transfer performance of the target robot. Moreover, we note that the ν -gap is developed for linear dynamical systems. While we showed the proposed similarity characterization effectively improves quadrotor tracking performance, we would like to explore experience selection approaches for more challenging nonlinear robot systems.

VIII. CONCLUSION

Inspired by the ν -gap metric from robust control theory, we introduced an experience selection algorithm using a characterization of dynamics similarity between robots that can be determined through simple experiments on periodic trajectories. Using this characterization of similarity, we proposed a BO-based algorithm for experience selection when the dynamics of the source and target robots are unknown. We showed in experiments that transferring an inverse dynamics module selected based on the proposed algorithm yields 62% better performance than an otherwise inappropriately selected experience, illustrating the importance of dynamics similarity in knowledge transfer.

REFERENCES

- C. Devin, A. Gupta, T. Darrell, P. Abbeel, and S. Levine, "Learning modular neural network policies for multi-task and multi-robot transfer," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 2169–2176.
- [2] S. Stark, J. Peters, and E. Rueckert, "Experience reuse with probabilistic movement primitives," *arXiv preprint arXiv:1908.03936*, 2019.
- [3] J. Fu, S. Levine, and P. Abbeel, "One-shot learning of manipulation skills with online dynamics adaptation and neural network priors," in *Proc. of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 4019–4026.
- [4] M. Hamer, M. Waibel, and R. D'Andrea, "Knowledge transfer for high-performance quadrocopter maneuvers," in *Proc. of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2013, pp. 1714–1719.
- [5] B. Bocsi, L. Csató, and J. Peters, "Alignment-based transfer learning for robot models," in *Proc. of the IEEE International Joint Conference* on Neural Networks (IJCNN), 2013, pp. 1–7.
- [6] H. B. Ammar, E. Eaton, P. Ruvolo, and M. E. Taylor, "Unsupervised cross-domain transfer in policy gradient reinforcement learning via manifold alignment," in *Proc. of the AAAI Conference on Artificial Intelligence*, 2015.
- [7] K. Pereida, M. K. Helwa, and A. P. Schoellig, "Data-efficient multirobot, multitask transfer learning for trajectory tracking," *IEEE Robotics and Automation Letters (RA-L)*, vol. 3(2), pp. 1260–1267, 2018.
- [8] L. Torrey and J. Shavlik, "Transfer learning," in *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques.* IGI Global, 2010, pp. 242–264.
- [9] M. E. Taylor and P. Stone, "Transfer learning for reinforcement learning domains: A survey," *Journal of Machine Learning Research* (*JMLR*), vol. 10(1), pp. 1633–1685, 2009.
- [10] M. T. Rosenstein, Z. Marx, L. P. Kaelbling, and T. G. Dietterich, "To transfer or not to transfer," in *Proc. of Neural Information Processing Systems (NIPS) Workshop on Transfer Learning*, 2005.
- [11] Z. Wang, Z. Dai, B. Póczos, and J. Carbonell, "Characterizing and avoiding negative transfer," in *Proc. of the IEEE Conference on Computer Vision and Patter Recognition (CVPR)*, 2019, pp. 11293– 11302.
- [12] Y. Yao and G. Doretto, "Boosting for transfer learning with multiple sources," in *Proc. of the IEEE Conference on Computer Vision and Patter Recognition (CVPR)*, 2010, pp. 1855–1862.
- [13] K. Zhou and J. C. Doyle, *Essentials of Robust Control*. Prentice Hall Upper Saddle River, 1998.
- [14] M. K. Helwa and A. P. Schoellig, "Multi-robot transfer learning: A dynamical system perspective," in *Proc. of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 4702– 4708.
- [15] N. Makondo, B. Rosman, and O. Hasegawa, "Knowledge transfer for learning robot models via local procrustes analysis," in *Proc.* of the IEEE-RAS International Conference on Humanoid Robots (Humanoids), 2015, pp. 1075–1082.

- [16] A. Gupta, C. Devin, Y. Liu, P. Abbeel, and S. Levine, "Learning invariant feature spaces to transfer skills with reinforcement learning," in *Proc. of the International Conference on Learning Representations* (ICLR), 2017.
- [17] R. Pautrat, K. Chatzilygeroudis, and J.-B. Mouret, "Bayesian optimization with automatic prior selection for data-efficient direct policy search," in *Proc. of the IEEE International Conference on Robotics* and Automation (ICRA), 2018, pp. 7571–7578.
- [18] A. Marco, F. Berkenkamp, P. Hennig, A. P. Schoellig, A. Krause, S. Schaal, and S. Trimpe, "Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with bayesian optimization," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 1557–1563.
- [19] M. Cutler and J. P. How, "Efficient reinforcement learning for robots using informative simulated priors," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 2605– 2612.
- [20] G. Zames and A. El-Sakkary, "Uncertainty in unstable systems: The gap metric," *IFAC Proceedings Volumes*, vol. 14(2), pp. 149–152, 1981.
- [21] E. Arslan, M. C. Çamurdan, A. Palazoglu, and Y. Arkun, "Multi-model control of nonlinear systems using closed-loop gap metric," in *Proc. of the IEEE American Control Conference (ACC)*, 2004, pp. 2374–2378.
 [22] J. Du and T. A. Johansen, "A gap metric based weighting method for
- [22] J. Du and T. A. Johansen, "A gap metric based weighting method for multimodel predictive control of mimo nonlinear systems," *Journal of Process Control*, vol. 24(9), pp. 1346–1357, 2014.
- [23] T. T. Georgiou, "On the computation of the gap metric," in *Proc.* of the IEEE Conference on Decision and Control (CDC), 1988, pp. 1360–1361.
- [24] T. Koenings, M. Krueger, H. Luo, and S. X. Ding, "A data-driven computation method for the gap metric and the optimal stability margin," *IEEE Transactions on Automatic Control*, vol. 63(3), pp. 805–810, 2017.
- [25] S. Zhou, M. K. Helwa, and A. P. Schoellig, "Design of deep neural networks as add-on blocks for improving impromptu trajectory tracking," in *Proc. of the IEEE Conference on Decision and Control (CDC)*, 2017, pp. 5201–5207.
- [26] M. W. Spong, S. Hutchinson, and M. Vidyasagar, *Robot Modeling and Control.* Wiley, NJ, 2006.
- [27] Q. Li, J. Qian, Z. Zhu, X. Bao, M. K. Helwa, and A. P. Schoellig, "Deep neural networks for improved, impromptu trajectory tracking of quadrotors," in *Proc. of the IEEE International Conference on Robotics* and Automation (ICRA), 2017, pp. 5183–5189.
- [28] G. Vinnicombe, "Frequency domain uncertainty and the graph topology," *IEEE Transactions on Automatic Control*, vol. 38(9), pp. 1371– 1383, 1993.
- [29] —, Uncertainty and Feedback: H [infinity] Loop-Shaping and the [nu]-Gap Metric. World Scientific, 2001.
- [30] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, "Taking the human out of the loop: A review of Bayesian optimization," *Proc. of the IEEE*, vol. 104(1), pp. 148–175, 2015.
- [31] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.