

Semiparametrical Gaussian Processes Learning of Forward Dynamical Models for Navigating in a Circular Maze

Romeres, Diego; Jha, Devesh K.; Dalla Libera, Alberto; Yezzunis, William S.; Nikovski, Daniel N.

TR2019-028 June 12, 2019

Abstract

This paper presents a problem of model learning for the purpose of learning how to navigate a ball to a goal state in a circular maze environment with two degrees of freedom. The motion of the ball in the maze environment is influenced by several non-linear effects such as dry friction and contacts, which are difficult to model physically. We propose a semiparametric model to estimate the motion dynamics of the ball based on Gaussian Process Regression equipped with basis functions obtained from physics first principles. The accuracy of this semiparametric model is shown not only in estimation but also in prediction at n -steps ahead and its compared with standard algorithms for model learning. The learned model is then used in a trajectory optimization algorithm to compute ball trajectories. We propose the system presented in the paper as a benchmark problem for reinforcement and robot learning, for its interesting and challenging dynamics and its relative ease of reproducibility.

IEEE International Conference on Robotics and Automation (ICRA)

© 2019 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Semiparametrical Gaussian Processes Learning of Forward Dynamical Models for Navigating in a Circular Maze

Diego Romeres¹, Devesh K. Jha¹, Alberto DallaLibera², Bill Yerazunis¹ and Daniel Nikovski¹

Abstract—This paper presents a problem of model learning for the purpose of learning how to navigate a ball to a goal state in a circular maze environment with two degrees of freedom. The motion of the ball in the maze environment is influenced by several non-linear effects such as dry friction and contacts, which are difficult to model physically. We propose a semiparametric model to estimate the motion dynamics of the ball based on Gaussian Process Regression equipped with basis functions obtained from physics first principles. The accuracy of this semiparametric model is shown not only in estimation but also in prediction at n -steps ahead and its compared with standard algorithms for model learning. The learned model is then used in a trajectory optimization algorithm to compute ball trajectories. We propose the system presented in the paper as a benchmark problem for reinforcement and robot learning, for its interesting and challenging dynamics and its relative ease of reproducibility.

I. INTRODUCTION

One challenge of learning in physical systems is that data are expensive to obtain, and system dynamics are affected by non-linear and discontinuous phenomena such as dry friction, contact, hysteresis, etc [1]. It is, in general, very difficult to accurately model and/or learn the effect of these non-linear phenomena. Our work is motivated by learning the dynamics of physical systems that exhibit (highly) non-linear and discontinuous dynamics for the purpose of precise optimizing control. Deriving accurate models from first principles of physics could be extremely challenging, and in its turn, learning a model solely from data could be prohibitively expensive and inaccurate in sparsely sampled regions of state-space. For this reason, semiparametric models which combine physical analytical models and data-driven techniques might be advantageous [2], [3], [4].

This paper focuses on learning a dynamical model for the purpose of controlling the motion of a ball in a circular maze environment (see Fig. 1). The motion of the ball in the maze environment is significantly influenced by viscous and dry friction, contacts with the walls of the maze, and delays in sensing and actuation. Furthermore, the geometry of the maze constrains the motion of the ball to be discontinuous in the vicinity of each of the openings in each of the rings (see Fig. 1). As a result, it exhibits rich, highly non-linear dynamics. In this paper, we present results for model learning of the ball dynamics in a circular maze environment (CME), and initial results for trajectory



Fig. 1: Circular Maze Environment (CME) with the ball

optimization using the learned models in a model-based reinforcement learning (MBRL) setting.

Related Work. RL has seen explosive growth in recent years with several groundbreaking results in the areas of video games [5], [6], robotics [7], etc. RL algorithms can be broadly classified into model-free and model-based [8], [9]. In the former class, policy is directly inferred using either a policy or a value-based method [10], [11]. In the latter, a model for the forward dynamics of the system is learned, and is then used to compute a policy to solve the desired task [12], [13], [14]. Model learning could result in better generalization across several different tasks. However, learning global dynamics for nonlinear systems is generally considered a challenging problem. The use of Gaussian processes for model learning followed by a suitable controller design has been studied in the literature. Some of the related works which are closely related to the approach presented in this paper can be found in [12], [15], [16], [17], [18], [19].

Apart from these, maze environments have been used to study RL algorithms [20], [21], [22]. The proposed CME presents a more complex geometry, which results in significantly more convoluted motion dynamics of the ball, as will be presented in the later sections. Furthermore, the action space in the circular maze system (CMS) is continuous, whereas the action spaces considered in the earlier work with mazes [20], [21], [22] were discrete. In [23], a deep RL-based end-to-end learning for the same CME is presented. The idea in this paper is to learn policies in simulation and transfer to the real system after minimum fine-tuning.

Contributions. Our contributions are the following. First, we present a semiparametric modeling technique [24], [25], [26] which combines the physical model of the system with data-driven Gaussian Process Regression [27] to learn the forward dynamics of the ball in a novel, real physical system.

¹Diego Romeres, Devesh K. Jha, Bill Yerazunis, Daniel Nikovski are with Mitsubishi Electric Research Laboratories, Cambridge, MA 02139. Email-{romeres, jha, yerazunis, nikovski}@merl.com

²Alberto D. Libera is with Dept. of Information Engineering, University of Padova, Padova, Italy Email-dallaliber@dei.unipd.it




Fig. 2: Circular Maze System (CMS) setup with the maze housed on the tip-tilt stage. Camera is mounted on top of the tip-tilt stage.

To the best of our knowledge, semi-parametric Gaussian process regression has only been used to learn inverse dynamics previously [28], [26], [24], [29]. Second, continuous-time forward dynamics are learned, and not discrete-time state transition functions. We estimate the ball acceleration signal, and then integrate it forward in time to predict the state; we do not directly estimate the successor state. Third, performance of the modeling choices are shown in long-horizon predictive capabilities (known to be a hard task for the compounding error effect) and compared with several other techniques to highlight its accuracy. Rollout accuracy is fundamental to obtain successful optimal control algorithms that use known dynamics. Fourth, we illustrate the control performance in the CME using the learned models and the Iterative Linear Quadratic Gaussian (iLQG) algorithm [30], [31]. We make simplifying assumptions to decompose the full goal-directed navigation into a sequence of sub-goals in each of the rings. We show that using the learned models, we can control the ball between arbitrary points in the maze environment, thus also navigating to the goal ring. In the remainder of the paper, CMS refers to the full system comprising of the CME and the tip-tilt platform.

II. EXPERIMENTAL CONFIGURATION

The electromechanical tip-tilt system is composed of the maze, a maze attachment plate, two gimbals, two RC servos, a 2D laser tilt sensor, a base plate and Arduino-based control processor (see Fig. 2). Camera input is provided by an Intel RealSense camera on an overhead mount (see Figure 2). All the communication from the camera to the image processing system, the laser sensor and the servo motors for measuring the inclination and driving the tip-tilt platform is done using the Robot Operating System (ROS) [32] at 30Hz.

The maze itself is a commercially available toy purchased from Amazon.com (see Fig. 1). A 3D printed plate with 2 positioning setscrews and 3 spring clips fits tightly inside the base of the maze and provides attachment points for the servo link rods and the gimbals. The gimbals constrain the platform to have pure rotational motion around its origin \mathbf{o}_m . The latter, represents the Cartesian coordinates of the rotational fulcrum identified by the gimbals in the inertial frame $(\mathbf{o}_b, \mathbf{x}, \mathbf{y}, \mathbf{z})$. W.l.o.g., we can assume that the origin




Fig. 3: Coordinate reference frames for the maze environment on the tip-tilt platform. Angles β and γ represent the orientation of the tip-tilt stage. The ball is shown in red.

of the inertial frame coincides with \mathbf{o}_m (see Fig. 3). The configuration of the platform could be fully described by the rotation matrix between frame $(\mathbf{o}_b, \mathbf{x}, \mathbf{y}, \mathbf{z})$ and the frame attached to the maze $(\mathbf{o}_m, \mathbf{x}', \mathbf{y}', \mathbf{z}')$; this can be described in a more compact fashion using the yaw, pitch and roll angles. The maze base is constrained by servo link rods, attached to the maze base through Traxxis ball joints; the lower end of each link rod is connected to a bellcrank arm of a hobbyist-grade HS-805BB RC servo motor with PWM-based control. The motors are positioned in perpendicular directions, so that the transformations between $(\mathbf{o}_b, \mathbf{x}, \mathbf{y}, \mathbf{z})$ and $(\mathbf{o}_m, \mathbf{x}', \mathbf{y}', \mathbf{z}')$ is described only by γ and β , the roll and pitch angles. The HS-805BB servos allow only position control. The control target of the motors, namely the input of our system, will be denoted by the vector $\mathbf{u} = (u^\beta, u^\gamma)$. The HS-805BB controller takes in the order of 200ms to move from an initial to a desired angle. Indirect measurements of γ and β are acquired using an inexpensive 5mW diode laser (Adafruit productid 1054), placed under the maze plate such that the direction of its optical axis is always equal to \mathbf{z}' . The laser projects downward onto the non-moving tilt position sensor, a $10\text{mm} \times 10\text{mm}$ position sensitive diode (a FirstSensor DL100-7-PCBA3 position sensitive photodiode). This sensor emits a pair of voltages corresponding to the x and y component of the impinging light. The map from the laser sensor readings to γ and β is learned by calibrating the laser readings with a precision (0.1°) clinometer.

The radius of the maze is 110mm and the diameter of the ball is 12.75mm . The width of each of the openings that connect rings is around 16mm . This suggests that a fairly accurate algorithm is needed in order to control the ball in the CME. The position of the ball is expressed in polar coordinates by r and θ , which represent the distance between \mathbf{o}_m and the center of the ball, and the rotation around \mathbf{z}' respectively. The position is measured using the camera and an off-the-shelf blob detection algorithm; angular velocity $\dot{\theta}$ is estimated using a Kalman filter. The entire CMS except for the camera is connected to a computer via an Arduino Mega 2560. The entire mechanical system with the exception of the bearings and servo motors were designed in-house with OpenSCAD, with most of the non-catalog parts printed in InPLA polymer on an Lulzbot TAZ 3 printer.

III. PROBLEM FORMULATION

In this section, we describe the task in the proposed CMS and the assumptions considered in order to decompose the original problem into sub-problems with lower complexity. Throughout the paper, we assume that the relationship between observations and state is deterministic and known. The system is thus fully defined by the combination of the state \mathbf{x}_k and the control inputs \mathbf{u}_k , and it evolves according to the dynamics $p(\mathbf{x}_{k+1}|\mathbf{x}_k, \mathbf{u}_k)$ which are composed of the ball dynamics in the maze and the tip-tilt platform dynamics. The goal is to learn an accurate model of the ball dynamics which can be used in a controller, $\pi(\mathbf{u}_k|\mathbf{x}_k)$, based on a trajectory optimization algorithm which allows the CMS to choose an action \mathbf{u}_k given the state observation \mathbf{x}_k to drive a ball from an initial condition to the target state. As a first simplification, we assume that the ball dynamics are independent of the radial dynamics in each of the individual rings – i.e., we quantize the radius of the ball position into the 5 rings of the maze, denominated from 1 to 5 starting from the most outer ring toward the most inner one. The radius of a ring and therefore the radius of the ball position is approximated with r_d , the mean radius of the ring in which the ball is moving. As discussed in Section II, the actuators for the tip-tilt platform allow only position control, and they experience significant delay (around 200 [ms]) in attaining the desired angle. Thus, we include the orientation of the tip-tilt stage as part of the state for our dynamical system obtaining a seven-dimensional state representation for the system, i.e., $\mathbf{x} = (r^d, \theta, \dot{\theta}, \beta, \dot{\beta}, \gamma, \dot{\gamma})$. It is noted that the radius r^d is a discrete variable, while the rest of the state variables are continuous. The input space for the system is the two dimensional vector \mathbf{u} defined in Section II.

To further simplify the problem, we decompose the problem into macro and micro goals. While the macro goal of the ball is to be able to reach the goal ring, this goal can be composed of a sequence of micro goals. For example, from any initial position, the ball should reach one of the openings and then move to the next ring. This sequence of movements should be repeated until the ball reaches the goal ring. This decomposition is motivated by discontinuity in ball dynamics when the ball moves to an adjacent ring through one of the openings. Such a decomposition of long horizon planning problems, with sparse discontinuities, are common in the robot learning literature [33], [34], [35]. While the full long-horizon control problem could be challenging to solve using a MBRL algorithm, the sub-problems for each of the micro-goals could be solved by model learning followed by trajectory optimization.

IV. PROPOSED LEARNING APPROACH

In this section, we describe the proposed solution to learn how to navigate a ball in the CME. First, we learn a Gaussian process-based model for movement in each ring. A full motion model for the entire maze is obtained by using an additional discrete model for the movement between rings. Second, the learned models are used to compute a controller using iLQG to learn a local policy to navigate in the maze.

A. Model Learning

We consider discrete-time systems:

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{e}_k \quad (1)$$

where $\mathbf{x}_k \in \mathbb{R}^7$ denotes the state, $\mathbf{u}_k \in \mathbb{R}^2$ the actions, and \mathbf{e}_k is assumed to be a zero mean white Gaussian noise with diagonal covariance that represents the uncertainty about the state at the discrete time instant $k \in [1, \dots, T]$. The transition function f maps the state-action pair to the successor state, which is assumed to evolve smoothly over time. Given the quantization of the radius introduced in Section III, we do not use the measurements of the radius of the ball's position. Even though this leads to information loss, preventing the possibility of modeling complex dynamics behaviours such as collisions, it reduces modeling complexity. The radial dynamics in the model will be incorporated in future works. The state \mathbf{x} for the CMS can be decomposed into two subsets: the ball state $\mathbf{x}^{ball} = [r^d, \theta, \dot{\theta}]$, and the tip-tilt platform state $\mathbf{x}^t = [\beta, \dot{\beta}, \gamma, \dot{\gamma}]$. This decomposition is convenient, because it is reasonable to assume that the tip-tilt platform dynamics are independent of the ball dynamics which is

$$p(\mathbf{x}_{k+1}^t | \mathbf{x}_k, \mathbf{u}_k) = p(\mathbf{x}_{k+1}^t | \mathbf{x}_k^t, \mathbf{u}_k) \quad (2)$$

In the proposed approach, we do not learn directly the static map (1), but we instead learn the evolution of $\dot{\mathbf{x}}_t$, which is a second order continuous-time system (subscript t denotes continuous time). This is different from earlier work presented in [12], [13] where the discrete-time map was learned. For our system, this reduces to the estimation of three ordinary differential equations for the acceleration functions of the positional states $[\theta, \beta, \gamma]$. This function can then be integrated forward in time to make predictions on future states. In our opinion, this choice has the advantage that $\dot{\mathbf{x}}_t$ is described by the underlying physical system, and this knowledge can be exploited more efficiently in learning dynamic models. Notice also that in discrete time the evolution of \mathbf{x}^{ball} given \mathbf{x}_k is not independent of \mathbf{u}_k , since we need to know the input \mathbf{u}_k until the next observation to make future predictions. In the continuous case, we have that $\dot{\mathbf{x}}_t^{ball}$ is independent from \mathbf{u}_t given \mathbf{x}_t , \mathbf{x}_t^{tt} and $\dot{\mathbf{x}}_t^{tt}$, which is

$$p(\dot{\mathbf{x}}_t^{ball} | \mathbf{x}_t, \mathbf{u}_t) = p(\dot{\mathbf{x}}_t^{ball} | \mathbf{x}_t, \dot{\mathbf{x}}_t^{tt})$$

The continuous time models are then discretized to compute the models used for control. In the following, we focus on learning acceleration models to be able to describe (1).

First, we model the dynamics of the tip-tilt platform. Because of (2), the tip-tilt stage dynamics (1) can be learned independently from the ball dynamics. The dynamics of β and γ have been shown to be modelled effectively from physical quantities with a linear second order continuous-time differential equation [36], [37]. Therefore, the acceleration functions $\ddot{\beta}_t = f^\beta(\dot{\beta}_t, \beta_t, \gamma_t, u_t^\beta)$ and $\ddot{\gamma}_t = f^\gamma(\dot{\gamma}_t, \gamma_t, \beta_t, u_t^\gamma)$ have been estimated using standard linear system identification techniques [38] using Prediction Error Methods, resulting in two linear ARX models for each angle β and γ .

Next, we focus on modelling the more complex ball dynamics. Additionally, we define a new state vector $\bar{\mathbf{x}}_t$ obtained by stacking together \mathbf{x}_t and $\dot{\mathbf{x}}_t^{tt}$

$$\ddot{\theta}_t = f^{ball}(\mathbf{x}_t, \dot{\mathbf{x}}_t^{tt}) + \mathbf{e}_t = f^{ball}(\bar{\mathbf{x}}_t) + \mathbf{e}_t \quad (3)$$

Most of the models introduced in this section are equipped with a set of unknown *hyperparameters* that we estimate using the maximization of the marginal likelihood. Since all the variables are referred to time t , from now on we will neglect the dependencies on t .

1) *Physics Inspired Model*: We are interested in determining an analytical expression of the forward dynamics (3), and using it for data-driven modeling. It is noted that this model only describes the motion of the ball in a circular ring, and doesn't consider the tip-tilt platform. As described in Section III, the radial dynamics of the ball is ignored, and the analytical model will be computed for a given radius r that is assumed to be constant. Using a standard Lagrangian approach [39], [40], [41], we obtain the following equations of motion (we skip a detailed derivation for brevity) :

$$\begin{aligned} \ddot{\theta} = & \ddot{\beta} \sin(\gamma) - \dot{\beta}^2 \sin(\theta) \cos(\gamma)^2 \cos(\theta) - 2\dot{\beta}\dot{\gamma} \cos(\gamma) \cos(\theta)^2 \\ & + 2\dot{\beta}\dot{\gamma} \cos(\gamma) + \frac{1}{2}\dot{\gamma}^2 \sin(2\theta)g \sin(\beta) \sin(\theta)/r \\ & - g \sin(\gamma) \cos(\beta) \cos(\theta)/r \end{aligned} \quad (4)$$

Equation (4) reflects the complex dynamics of ball movement due to the circular geometry of the maze when compared to relatively simple geometries considered earlier in [33], [34], [35]. We refer to this model as the Physical model, ‘‘P’’. While P in equation (4) may be inaccurate for the CME due to making simplifying assumptions, it does provide a well-defined structure to the model that can be hard to identify in data-driven models in an efficient fashion. This model structure is used and interpret each component of the right-hand side of equation (4) as basis functions of the feature space of a Bayesian linear model as defined in [27, Chp. 2]:

$$\begin{aligned} \phi^\top(\bar{\mathbf{x}}) = & [\dot{\beta} \sin(\gamma), -\dot{\beta}^2 \sin(\theta) \cos(\gamma)^2 \cos(\theta), 2\dot{\beta}\dot{\gamma} \cos(\gamma), \\ & - 2\dot{\beta}\dot{\gamma} \cos(\gamma) \cos(\theta)^2, 0.5\dot{\gamma}^2 \sin(2\theta) \sin(\beta) \sin(\theta)g/r, \\ & - \sin(\gamma) \cos(\beta) \cos(\theta)g/r] \end{aligned} \quad (5)$$

The expression of the forward dynamics (3) becomes

$$\ddot{\theta} = \phi^\top(\bar{\mathbf{x}})\mathbf{w} + \mathbf{e} \quad (6)$$

where $\mathbf{w} \sim \mathcal{N}(0, \Sigma_w)$ is a vector of weights modeled as a zero mean Gaussian prior with diagonal covariance matrix Σ_w . This model is denoted as Physics-Inspired model, ‘‘PI’’.

2) *Nonparametric Model*: According to the Gaussian Process Regression framework [27], the forward dynamics can be postulated in a data driven fashion as:

$$\ddot{\theta} = g(\bar{\mathbf{x}}) + \mathbf{e} \quad (7)$$

where $g(\bar{\mathbf{x}})$ is a Gaussian process with zero mean and covariance function $k(\bar{\mathbf{x}}, \bar{\mathbf{x}}') := \mathbb{E}[g(\bar{\mathbf{x}})g(\bar{\mathbf{x}}')]$ defined through a radial basis function (RBF) kernel. This data-driven model class is known to have high flexibility in modelling functions and also high prediction performance, because it extrapolates the dynamics directly from the data without relying on any physical assumptions. However, this method can perform poorly in sparsely sampled regions of the state-space. This model is denoted as the Nonparametric model, ‘‘NP’’.

3) *Semiparametric Model*: Semiparametric models attempt to combine the benefits of both the physical and data-driven model classes, which are good global behavior and accuracy, respectively. There can be several ways of combining the two methods and choosing the a priori distribution, as shown in [24], but one of the most effective has been shown to be:

$$\ddot{\theta} = \phi^\top(\bar{\mathbf{x}})\mathbf{w} + g(\bar{\mathbf{x}}) + \mathbf{e} \quad (8)$$

where $\phi^\top(\bar{\mathbf{x}})$ are the basis function suggested by the physics and defined as (5), \mathbf{w} is a vector of parameters with a zero mean Gaussian prior and a diagonal covariance matrix Σ_w and $g(\bar{\mathbf{x}})$ is a zero mean Gaussian process with RBF covariance function. The covariance of the whole process is $k(\bar{\mathbf{x}}, \bar{\mathbf{x}}') := \phi^\top(\bar{\mathbf{x}})\Sigma_w\phi(\bar{\mathbf{x}}) + \mathbb{E}[g(\bar{\mathbf{x}})g(\bar{\mathbf{x}}')]$. This class of models is also known as semi-parametrical models, and it has been shown to be a successful technique in model learning for robotic applications in the recent literature [24], [25], [26]. This model is denoted as Semiparametric model, ‘‘SP’’.

B. Radial Movement

The model of the continuous movement of the ball in each of the rings was described in Section IV-A. Next, we model the discrete movement of the ball between rings. To learn this movement, we first learn the spatial clusters where there is a transition in r^d in the collected trajectories of the ball. Thus, we obtain two clusters in the vicinity of each opening, based on the sign of Δr^d . The radial transition between rings is then approximated by the vector joining the centroids of the two clusters on the two sides of an opening, and the control action \mathbf{u} is approximated as the maximum signal whose projection on the horizontal plane is parallel to the direction of the radial transition vector.

C. Trajectory Optimization using iLQG

In this section, we briefly describe the iLQG algorithm which is used for trajectory optimization. The discrete time dynamics $\mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k)$ and the cost function are used to compute local linear models and a quadratic cost function for the system along a trajectory. These linear models are then used to compute optimal control inputs and local gain matrices by iteratively solving the associated LQG problem. We skip the mathematical details for space limitations. For more details, interested readers are referred to [31].

V. EXPERIMENTAL RESULTS

In this section, we will provide experimental results for the models presented in Section IV-A, and we will show their efficacy in the real system, CMS, by learning a policy as described in Section IV-C.

Data collection consisted of around 50 minutes of operation on CMS, when control actions were applied at 30Hz, and measurements from the laser and camera were recorded. During data collection, control actions were generated as a sum of 50 sine and cosine waves with random sampled frequency between $[0, 1.5]$ Hz and shift phases in $[0, 2\pi]$. The data set consists of the control actions matrix $\mathbf{U} \in$

Ring	N_{tr}	c_{DS}	nRMSE Training					N_{test}	nRMSE Test				
			RF	P	PI	NP	SP		RF	P	PI	NP	SP
1	3874	5	0.18	2.07	0.49	0.31	0.28	19368	0.49	1.97	0.48	0.38	0.35
2	4365	4	0.22	1.56	0.63	0.31	0.29	17459	0.59	1.57	0.63	0.45	0.43
3	4606	2	0.19	1.32	0.66	0.19	0.16	9213	0.53	1.35	0.65	0.42	0.40
4	2814	1	0.15	1.10	0.79	0.02	0.01	2815	0.67	1.08	0.81	0.71	0.65

TABLE I: nRMSE performance in training and in test data in all the rings for all the presented models and a Random Forest model.

$\mathbb{R}^{N \times 2}$ and the state matrix $\mathbf{X} \in \mathbb{R}^{N \times 7}$ with $N = 90702$. Finally, the output vector Y , given by the values of $\ddot{\theta}$, is computed using an acausal filter for our supervised model learning algorithm. Note that $\ddot{\theta}$ are not needed online while controlling CMS and an acausal filter is important in order to reduce the delay and improve the accuracy of the signal to be learned. Finally, the data collected are divided by rings, and split into two to obtain 4 training and testing data sets (ring 5 is the target, so it is not considered in the learning algorithm). The time of use of the CMS needed by our learning algorithm is less than 25 minutes. To limit the computational complexity of the algorithms, the 4 training data sets are uniformly downsampled with coefficient c_{DS} , in order to obtain training sets of size N_{tr} less than 5000 points; the size of the test set is denoted by N_{test} . During data collection, the ball naturally tends to stay more in the outermost ring, with the effect that the number of data per ring decreases going from ring 1 to ring 4 (see Table I). Each of the models described in Section IV-A and denoted by P, PI, NP, and SP, has been learned and tested with the data set for each ring and in the real CMS. The quality of the learned models is evaluated in terms of nRMSE in training and test data for each ring and each model in Table I. Model P represents only an approximation of $\ddot{\theta}$ and, as expected, its prediction performance is significantly worse than the other methods. Similarly, PI performs poorly, however, it outperforms P, confirming the benefit of using the nonlinear behaviour of the state suggested by the physics as basis functions. The best methods are NP and SP, and we can notice that SP always outperforms NP. This shows the higher robustness of the semiparametric model due to its use of physical basis functions. Moreover, as the ring becomes smaller, the prediction performance deteriorates for all the models, and the gap in the performances of NP and SP between training and test increases when the amount of data available decreases, namely in ring 3 and 4. This agrees with theoretical expectations, and is likely due to the lower variability present in the data set. Finally, for comparison purposes, we added the performance of Random Forest regression (RF) [42], a standard machine learning method based on a different paradigm than GPR. Table I highlights that NP and SP exhibit better performance than RF. In Figure 4, we compare the learning curves of the two best performing models (SP and NP) as the number of training data increases for ring 1. SP outperforms NP in




Fig. 4: Learning curves of SP and NP in training and test as the number of training data increases for ring 1.

terms of nRMSE at each evaluated N_{tr} , and achieves the same performance with half the data (note SP performance at $N_{tr} = 2152$ and NP at $N_{tr} = 3874$).

After the evaluation of the prediction of acceleration, we are now interested on evaluating the n-steps-ahead prediction performance of the actual state variables θ and $\dot{\theta}$. This hard modelling task is fundamental for trajectory optimization control algorithms. Results are shown over a 40-step horizon, which is commensurate with the time needed for the ball to reach an opening, as we will see later in the experiments on the CME. From the test datasets, we randomly chose trajectories of length 40 steps that terminate in front of an opening (gate). The models are initialized with the initial conditions of these trajectories, and rolled out with the same control actions. In Figure 5, the results of all the models in ring 1 are shown, evaluating the absolute errors in position $|\theta_k - \hat{\theta}_k|$ and velocity $|\dot{\theta}_k - \hat{\dot{\theta}}_k|$ at each prediction step $[1, \dots, 40]$, and reporting the mean over the 226 trajectories. This error index has been chosen for better interpretability of the control performance.




Fig. 5: Rollout performance. The average absolute error in θ and $\dot{\theta}$ at k -steps ahead are shown for all the models, on the left and right-hand side, respectively.

r	PI [rad]		NP [rad]		NPd [rad]		SP [rad]	
	n=20	n=40	n=20	n=40	n=20	n=40	n=20	n=40
1	0.40	1.32	0.11	0.26	1.36	3.01	0.11	0.25
2	0.21	0.57	0.15	0.33	0.87	2.09	0.15	0.33
3	0.24	0.64	0.19	0.34	0.99	2.26	0.18	0.32
4	1.81	1.21	0.46	1.76	0.65	2.45	1.74	1.18

TABLE II: Rollout performance for all the rings evaluated as the average absolute error in θ at 20 and 40 steps ahead.

In this comparison, we considered also the estimation of the discrete static map (1) using nonparametric GPR described in Section IV-A.2, trained on the same data sets as all the other models and using the same kernel function of




Fig. 6: Comparison of the iLQG controller with PD in the outer-most ring of the maze over 10 runs. The steady error for both controllers is comparable to the size of the ball.

NP. This model is denoted as Nonparametric discrete, “NPd”.

All models diverge rapidly from the ground truth, except for SP and NP, which achieve competitive performance. Model SP consistently outperforms all the others with an average error in position of 14 degrees after 40 steps. The confidence interval at 95% for model SP and NP is also shown, from which we can see that there are trajectories where contacts and bounces make the dynamics hard to predict. Finally, in Table II, the average of the absolute error $|\hat{\theta}_k - \theta_k|$ of the same trajectories is computed at 20 and 40 steps ahead for all the models (except P which performs poorly) in all the rings. Consistently, model SP and NP outperform the others in all of the rings. The significantly higher performance of NP w.r.t. NPd shows the benefit of formulating the learning problem as a dynamical system, as in (3), and not as a static map, as in (1). The performance in ring 4 deteriorates significantly, even for model SP, because as explained before, the data collected in this ring are much less and of lower quality in comparison to the other rings. This problem will become clear while controlling the ball in CME. Next, we present results for trajectory optimization using iterative LQG [30], [31]. To demonstrate correctness of the models that were learned in the last section, we first show several motions to a target point in the outermost ring of the maze using the learned model. We use the cost function proposed by [31], which penalizes any deviation from the target state, as well as saturation of the control signal. The state cost is the smooth-abs function $\ell(x) = \sqrt{x^2 + \alpha^2} - \alpha$, where x represents state error and α is a parameter to control the smoothness around origin. The control cost is the function $\ell(u) = \nu^2(\cosh(u/\nu) - 1)$ which penalizes input saturation, where u is the control input and by varying ν , we can limit the control signal to a particular volume of control space. We compare the results with a well-tuned PD controller. The error plots for 10 such trajectories, for a movement of $\pi/4$ radians in the outer-most ring, are

shown in Figure 6. We can see that the iLQG controller is behaving in a nearly time-optimal fashion using the model learned earlier. We repeat the same tests for all the rings for a movement of $\pi/4$ radians. With iLQG, we achieve a settling time of 1.5s and 2.73s, and steady state error of 0.013 and 0.05 radians, for rings 2 and 3, respectively. With PD, we achieved 2.0s and 2.70s with error of 0.066 and 0.012 radians in rings 2 and 3, respectively. In the fourth ring, our model isn’t very accurate, and, without feedback tracking, we were not able to regulate the position of the ball for very small movements. To solve the full problem of navigating to the inner-most goal state of the maze, we implement the trajectories computed for each of the rings with trajectory tracking using a PD controller. The radial transition was implemented when the ball was in front of an opening using the control signal described in Section IV-B. This resulted in a reasonably good performance. However, as the trajectories are local, there were some failure runs due to contacts with the walls. We tried 10 runs for the full trajectory, where 7 runs were successful with an average time of 7.71s, best time of 6.2s, and worst time of 10.13s (see attached video for successful and unsuccessful runs). While the navigation problem could be solved by a PD controller and a simple planner, the CME offers the possibility to deal with rich complicated dynamics and interesting planning problems. This gives the opportunity to test the state-of-the-art RL algorithms while having a base line to compare with.

VI. CONCLUSION AND DISCUSSION

In this paper, the problem of learning how to control a ball in a CME is presented. Our approach offers an initial solution to navigate the ball to the goal-state of the CME. The goal of reaching the final target state is decomposed into a sequence of micro-goals which consist of navigating in each ring, each of which is solved by a trajectory optimization algorithm using the learned models. Accurate semiparametric GP models are learned to describe the acceleration function of the dynamical system. While the current controller can navigate the ball to the goal state in the CME reasonably well, it is not robust to changes in initial conditions and other disturbances due to contacts. This is left as future work. Moreover, we are investigating on how to apply GP online learning techniques [43], [44], [45], [28], [46], [47], [48] to improve the model while iteratively improving the policy and comparing our approach with a method composing the prior physics knowledge with a neural network framework [49]. Apart from this, the proposed CMS could be used to study several other interesting problems. For example, how to leverage the models learned for a single ball in the environment to multiple balls or solving the navigation problems composing reusable skills [50], [51], [52].

The CMS is suitable to study several problems in RL and robot learning presenting several challenges such as nonlinear and fast dynamics, contacts and planning problem, thus it could possibly be used as a benchmark system for its ease of reproducibility as it is starting to be more and more required by the community [53], [54], [55].

REFERENCES

- [1] G. Gilardi and I. Sharf, "Literature survey of contact dynamics modelling," *Mechanism and machine theory*, vol. 37, no. 10, pp. 1213–1239, 2002.
- [2] A. Kloss, S. Schaal, and J. Bohg, "Combining learned and analytical models for predicting action effects," *arXiv preprint arXiv:1710.04102*, 2017.
- [3] N. Fazeli, S. Zapolsky, E. Drumwright, and A. Rodriguez, "Learning data-efficient rigid-body contact models: Case study of planar impact," in *Conference on Robot Learning*, pp. 388–397, 2017.
- [4] A. Ajay, J. Wu, N. Fazeli, M. Bauza, L. P. Kaelbling, J. B. Tenenbaum, and A. Rodriguez, "Augmenting physical simulators with stochastic neural networks: Case study of planar pushing and bouncing," *arXiv preprint arXiv:1808.03246*, 2018.
- [5] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, *et al.*, "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, p. 354, 2017.
- [6] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, p. 529, 2015.
- [7] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [8] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*, vol. 1. MIT press Cambridge, 1998.
- [9] J. Kober, J. A. Bagnell, and J. Peters, "Reinforcement learning in robotics: A survey," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013.
- [10] M. P. Deisenroth, G. Neumann, J. Peters, *et al.*, "A survey on policy search for robotics," *Foundations and Trends® in Robotics*, vol. 2, no. 1–2, pp. 1–142, 2013.
- [11] A. Antos, C. Szepesvári, and R. Munos, "Fitted Q-iteration in continuous action-space MDPs," in *Advances in neural information processing systems*, pp. 9–16, 2008.
- [12] M. P. Deisenroth, D. Fox, and C. E. Rasmussen, "Gaussian processes for data-efficient learning in robotics and control," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 2, pp. 408–423, 2015.
- [13] M. Deisenroth and C. E. Rasmussen, "PILCO: A model-based and data-efficient approach to policy search," in *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pp. 465–472, 2011.
- [14] S. Levine and V. Koltun, "Guided policy search," in *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pp. 1–9, 2013.
- [15] J. Boedecker, J. T. Springenberg, J. Wülfing, and M. Riedmiller, "Approximate real-time optimal control based on sparse gaussian process models," in *Adaptive Dynamic Programming and Reinforcement Learning (ADPRL), 2014 IEEE Symposium on*, pp. 1–8, IEEE, 2014.
- [16] S. Kamthe and M. P. Deisenroth, "Data-efficient reinforcement learning with probabilistic model predictive control," *arXiv preprint arXiv:1706.06491*, 2017.
- [17] G. Lee, S. S. Srinivasa, and M. T. Mason, "GP-ILQG: Data-driven robust optimal control for uncertain nonlinear dynamical systems," *arXiv preprint arXiv:1705.05344*, 2017.
- [18] Y. Pan and E. Theodorou, "Probabilistic differential dynamic programming," in *Advances in Neural Information Processing Systems 27 (Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, eds.)*, pp. 1907–1915, Curran Associates, Inc., 2014.
- [19] D. Meger, J. C. G. Higuera, A. Xu, P. Gigue, and G. Dudek, "Learning legged swimming gaits from experience," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2332–2338, May 2015.
- [20] M. Stolle and C. Atkeson, "Policies based on trajectory libraries," in *Proceedings of the International Conference on Robotics and Automation*, January 2006.
- [21] M. Zucker and J. A. Bagnell, "Reinforcement planning: RL for optimal planners," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pp. 1850–1855, IEEE, 2012.
- [22] D. C. Bentivegna and C. G. Atkeson, "Learning from observation using primitives," in *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, vol. 2, pp. 1988–1993, IEEE, 2001.
- [23] J. van Baar, A. Sullivan, R. Cordorel, D. Jha, D. Romeres, and D. Nikovski, "Sim-to-Real Transfer Learning using Robustified Controllers in Robotic Tasks involving Complex Dynamics," *ArXiv e-prints*, Sept. 2018.
- [24] D. Romeres, M. Zorzi, R. Camoriano, and A. Chiuso, "Online semi-parametric learning for inverse dynamics modeling," in *Conference on Decision and Control (CDC), 2016 IEEE 55th*, pp. 2945–2950, IEEE, 2016.
- [25] T. Wu and J. Movellan, "Semi-parametric Gaussian process for robot system identification," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 725–731, 2012.
- [26] D. Nguyen-Tuong and J. Peters, "Using model knowledge for learning inverse dynamics," in *IEEE International Conference on Robotics and Automation*, 2010.
- [27] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [28] D. Romeres, M. Zorzi, R. Camoriano, S. Traversaro, and A. Chiuso, "Derivative-free online learning of inverse dynamics models," *ArXiv e-prints*, Sept. 2018.
- [29] D. Nguyen-Tuong and J. Peters, "Model learning for robot control: a survey," *Cognitive Processing*, vol. 12, no. 4, pp. 319–340, 2011.
- [30] E. Todorov and W. Li, "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems," in *American Control Conference, 2005. Proceedings of the 2005*, pp. 300–306, IEEE, 2005.
- [31] Y. Tassa, T. Erez, and E. Todorov, "Synthesis and stabilization of complex behaviors through online trajectory optimization," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pp. 4906–4913, IEEE, 2012.
- [32] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.
- [33] O. Kroemer, H. Van Hoof, G. Neumann, and J. Peters, "Learning to predict phases of manipulation tasks as hidden states," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pp. 4009–4014, IEEE, 2014.
- [34] R. Lioutikov, G. Neumann, G. Maeda, and J. Peters, "Probabilistic segmentation applied to an assembly task," in *Humanoid Robots (Humanoids), 2015 IEEE-RAS 15th International Conference on*, pp. 533–540, IEEE, 2015.
- [35] O. Kroemer, C. Daniel, G. Neumann, H. Van Hoof, and J. Peters, "Towards learning hierarchical skills for multi-phase manipulation tasks," in *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pp. 1503–1510, IEEE, 2015.
- [36] T. Anjali and S. S. Mathew, "Implementation of optimal control for ball and beam system," in *Emerging Technological Trends (ICETT), International Conference on*, pp. 1–5, IEEE, 2016.
- [37] D. U. Campos-Delgado, "Optimal tracking controller for ball and beam system," 2004.
- [38] L. Ljung, *System Identification - Theory for the User*. Upper Saddle River, N.J.: Prentice-Hall, 2nd ed., 1999.
- [39] B. Siciliano, L. Sciacivico, L. Villani, and G. Oriolo, *Robotics: modelling, planning and control*. Springer Science & Business Media, 2010.
- [40] F. Bullo and A. D. Lewis, *Geometric control of mechanical systems: modeling, analysis, and design for simple mechanical control systems*, vol. 49. Springer Science & Business Media, 2004.
- [41] J. Hollerbach, W. Khalil, and M. Gautier, "Model identification," in *Springer Handbook of Robotics*, pp. 321–344, Springer, 2008.
- [42] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [43] N. Lawrence, M. Seeger, and R. Herbrich, "Fast sparse Gaussian process methods: The informative vector machine," in *Proceedings of the 15th International Conference on Neural Information Processing Systems, NIPS'02*, (Cambridge, MA, USA), pp. 625–632, MIT Press, 2002.
- [44] E. Snelson and Z. Ghahramani, "Sparse gaussian processes using pseudo-inputs," in *Advances in Neural Information Processing Systems*, pp. 1257–1264, MIT press, 2006.
- [45] A. Ranganathan, M. H. Yang, and J. Ho, "Online sparse Gaussian process regression and its applications," *IEEE Transactions on Image Processing*, vol. 20, pp. 391–404, Feb 2011.
- [46] D. Romeres, G. Prando, G. Pilonetto, and A. Chiuso, "On-line bayesian system identification," in *Control Conference (ECC), 2016 European*, pp. 1359–1364, IEEE, 2016.

- [47] G. Prando, D. Romeres, and A. Chiuso, "Online identification of time-varying systems: A bayesian approach," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 3775–3780, Dec 2016.
- [48] A. Gijsberts and G. Metta, "Incremental learning of robot dynamics using random features," in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 951–956, 2011.
- [49] A. Ajay, J. Wu, N. Fazeli, M. Bauza, L. P. Kaelbling, J. B. Tenenbaum, and A. Rodriguez, "Augmenting Physical Simulators with Stochastic Neural Networks: Case Study of Planar Pushing and Bouncing," *ArXiv e-prints*, Aug. 2018.
- [50] P. Bacon, J. Harb, and D. Precup, "The option-critic architecture," *CoRR*, vol. abs/1609.05140, 2016.
- [51] J. Andreas, D. Klein, and S. Levine, "Modular multitask reinforcement learning with policy sketches," *CoRR*, vol. abs/1611.01796, 2016.
- [52] H. Sahni, S. Kumar, F. Tejani, and C. L. Isbell, "Learning to compose skills," *CoRR*, vol. abs/1711.11289, 2017.
- [53] A. R. Mahmood, D. Korenkevych, B. J. Komer, and J. Bergstra, "Setting up a reinforcement learning task with a real-world robot," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4635–4640, IEEE, 2018.
- [54] P. Henderson, R. Islam, P. Bachman, J. Pineau, D. Precup, and D. Meger, "Deep reinforcement learning that matters," *CoRR*, vol. abs/1709.06560, 2017.
- [55] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *CoRR*, vol. abs/1606.01540, 2016.