

Iterative Semi-parametric Dynamics Model Learning For Autonomous Racing

Ignat Georgiev[‡], Christoforos Chatzikomis[†], Timo Völkl^{†*}, Joshua Smith[‡], Michael Mistry[‡]

[†]Arrival

[‡]School of Informatics, The University of Edinburgh

Abstract: Accurately modeling robot dynamics is crucial to safe and efficient motion control. In this paper, we develop and apply an iterative learning semi-parametric model, with a neural network, to the task of autonomous racing with a Model Predictive Controller (MPC). We present a novel non-linear semi-parametric dynamics model where we represent the known dynamics with a parametric model, and a neural network captures the unknown dynamics. We show that our model can learn more accurately than a purely parametric model and generalize better than a purely non-parametric model, making it ideal for real-world applications where collecting data from the full state space is not feasible. We present a system where the model is bootstrapped on pre-recorded data and then updated iteratively at run time. Then we apply our iterative learning approach to the simulated problem of autonomous racing and show that it can safely adapt to modified dynamics online and even achieve better performance than models trained on data from manual driving.

Keywords: Robots, Dynamics, Model Learning, Autonomous Driving

1 Introduction

Robots are one of the most fascinating human-made machines. They raise hopes of exploring extraterrestrial planets, replacing human workers in dangerous environments, assisting us with laborious tasks, and even saving human lives [1]. One of the first instances of potential for massive adaption of this technology has been autonomous cars, which are promising to alleviate traffic congestion, reduce road incidents, and lower pollution [2]. A challenging aspect of driverless vehicle development is controlling them accurately, safely, and efficiently. One popular approach to solve this problem is using model-based control, where the system dynamics are modeled and then used by some controller to achieve the desired task. In recent years, the most promising direction for these systems has been Model Predictive Control (MPC), which forward simulate the dynamics to predict the future states [3]. However, for this approach to work, a highly accurate and computationally efficient dynamics model is necessary. Furthermore, robotic systems deployed in the real world exhibit dynamical changes over time, and as such, these models must also adapt [4].

Traditionally for such systems, *parametric models* are utilized based on known physics to which parameters are fitted from data. These types of models are explainable, efficient, and generalize well to large state spaces but usually do not capture all dynamical properties [5]. On the other hand, *non-parametric models* that can use any function approximation technique to express the dynamics of the system. These data-driven models have the capability of capturing dynamics much more accurately and without requirements for extensive dynamics knowledge beforehand. However, these types of models are also not intuitively explainable and generalize poorly to unknown state spaces. A particular variety of non-parametric models that have seen success and popularity recently has been neural networks with several layers of depth, allowing them to capture highly non-linear dynamics. Finally, one of the more recent proposed categories of dynamics models has been semi-parametric models, which are a combination of the two models explained above [6, 7, 8].

*Author's affiliation has changed since this research was done. New affiliation with Munich Electrification

The model type and structure are important, but equally as crucial is learning their parameters. Traditionally this is done by collecting large datasets of real-world interactions, processing the data, learning the parameters from it, and then deploying the robotics system with these parameters fixed. This approach has successfully been used for decades. However, it is inherently limited when deployed outside a laboratory environment where the robot dynamics are continually changing from both wear and tear of the robot itself and external factors such as winds and surface friction. To address this problem, the dynamics model must be adapted online, which is a non-trivial problem due to the non-stationary learning target where any updates to the parametric part of the model modify the learning target of the non-parametric part [8]. Furthermore, online learning is an even bigger issue when applied to neural networks due to the problem of catastrophic forgetting, which is their tendency to forget the old data when learning new ones [9, 10, 11].

The novel contribution of this paper is applying an iterative learning non-linear semi-parametric model, with a neural network, to the task of autonomous racing with a Model Predictive Controller (MPC). We propose a process for bootstrapping the model from previously recorded data, deploying it, and then iteratively updating the model online. We experimentally show that the model successfully adapts to long-term changes in dynamics, achieves better performance than a purely parametric model, and generalizes better than a purely non-parametric model. The main contributions of this paper are:

1. A novel efficient semi-parametric model that utilizes a neural network that can be used as part of an MPC.
2. An iterative learning system that adapts the model to modified dynamics while taking into account the previously encountered state space.

2 Related Work

Semi-parametric models are one of the more recently proposed models for dynamics learning and are a combination of parametric and non-parametric models. These models vary in form, but the general idea is that the parametric part captures the known (usually linear) structure of the dynamics, and the non-parametric part of the model captures the unmodeled dynamics (usually highly non-linear). This approach to modeling dynamics combines the generalization and efficiency capabilities of parametric models with the accuracy of non-parametric models. It has been successfully applied to various robotics platforms using different non-parametric functions, including Least Squares [7], Gaussian Processes Regression (GPR) [6] and Gaussian Mixture Models (GMM) [8].

Iterative learning (also known as online learning) for robotic systems has also seen an increase in interest recently due to its high desirability for real system deployment [7, 8, 12]. Model adaption of purely parametric models is straightforward [13], but doing so for semi-parametric models imposes the challenge of non-stationary target distributions. None the less, there exist examples of successful applications of this as a standalone model [7] and to Composite Adaptive Control [8]. However, no approaches are using an iterative learning semi-parametric with a neural network applied to an MPC. The most notable similar work on this is applying a purely non-parametric neural network model to an iterative learning system using an MPC [12], which serves as the basis to this work.

Iterative learning for neural networks has also received significant attention in recent years due to the rise in the popularity of deep learning. Many directions have been explored here, including controlling how parameters are updated [11, 14] and meta-learning, where there is another machine learning algorithm, is responsible for training the main model [15, 16]. However, the more mature family of techniques that explicitly tackle the issue of catastrophic forgetting is that of rehearsal method. The idea is that old data is stored and used to train the model in addition to the new data [9]. However, maintaining all old data is infeasible, which has given rise to pseudo-rehearsal methods that do not explicitly store data but instead model it as data distribution and can generate random samples from it [17].

3 Problem Formulation

A robot operating in an environment with a model-based controller utilizes a dynamics model. The state \mathbf{x}_t and executed controls \mathbf{u}_t at the current time influence the state derivative $\dot{\mathbf{x}}_t$. Assuming

the model operates in discrete time with a timestep Δt , then the state transition function is given by:

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \dot{\mathbf{x}}_t \Delta t = \mathbf{x}_t + \mathbf{F}(\mathbf{x}_t, \mathbf{u}_t; \theta) \Delta t \quad (1)$$

Where \mathbf{F} is the full system dynamics model with some parameters θ . To learn the dynamics, our model must predict the state derivative $\dot{\mathbf{x}}$. To obtain the discrete version of that and thus our prediction target Y we reformulate Equation 1 into:

$$Y_t = \dot{\mathbf{x}}_t = \mathbf{F}(X; \theta) = \mathbf{F}(\mathbf{x}_t, \mathbf{u}_t; \theta) = \frac{\mathbf{x}_{t+1} - \mathbf{x}_t}{\Delta t} \quad (2)$$

where $X = (\mathbf{x}_t; \mathbf{u}_t)$ are the inputs into the model. For semi-parametric models, \mathbf{F} can be expanded into specific parametric and non-parametric parts. Without making any assumptions about the inputs to each, we rewrite Equation 1 for a semi-parametric model:

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \left(\underbrace{f(X_t^f; \theta^f)}_{\text{parametric part}} + \underbrace{g(X_t^g; \theta^g)}_{\text{non-parametric part}} \right) \Delta t \quad (3)$$

Note that in this general notation, both have different inputs and parameters. Assuming that the X_t^g is a function of X_t^f , then the issue of non-stationary learning targets arises. To update the full model iteratively, both parts of the model need to be updated independently. Since neural networks are updated via batches, one can first update the parametric model and then fully train the neural network to account for all changes. However, this is not feasible for a control system, which must adapt to such changes in real-time. The closest to that is to update the neural network with only a few mini-batches, but that is not guaranteed to account for the changes in the parametric model. Thus this results in the non-parametric model trying to learn a constantly changing target distribution, which is challenging [8].

When updating the neural network parameters we also have to tackle the issue of catastrophic forgetting where if only newly collected data $X \sim \mathcal{P}_L(\mathcal{X})$ from the *Local Operating Distribution* $\mathcal{P}_L(\mathcal{X})$ is used for learning, will result in the neural network forgetting prior dynamics. To tackle this with rehearsals [9], we assume that another *System Identification Distribution* $\mathcal{P}_{ID}(\mathcal{X})$ is available which captures data from all vehicle operations similar to [12].

4 Model

In this section, we describe the employed dynamics model. We chose to utilize a sequential semi-parametric modeling approach where the parametric part of the model first makes a state derivative prediction $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$ which is then fed as an input into the non-parametric part of the model $g(\cdot)$ resulting in the final semi-parametric model represented as:

$$\dot{\mathbf{x}} = \underbrace{f(\mathbf{x}, \mathbf{u})}_{\text{parametric part}} + \underbrace{g(\dot{\mathbf{x}}, \mathbf{u})}_{\text{non-parametric part}} \quad (4)$$

For our model, we define the state \mathbf{x} consisting of longitudinal and lateral position, yaw angle, longitudinal velocity, lateral velocity, and yaw rate; the control inputs \mathbf{u} are defined as acceleration and steering commands:

$$\mathbf{x} = (x \quad y \quad \psi \quad \dot{x} \quad \dot{y} \quad \dot{\psi}) \quad \mathbf{u} = (u_a \quad u_\delta) \quad (5)$$

4.1 Parametric part

Due to the inherent symmetry of cars, it is common to simplify 4-wheeled motors to a 2-wheeled one where the front and rear set of tires are lumped into only one at the front and one at the rear. This is commonly referred to as the *bicycle model* [18]. The model can be summarized with the

following state transition:

$$\dot{\mathbf{x}} = \begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{\psi} \\ \ddot{x} \\ \ddot{y} \\ \dot{\psi} \end{pmatrix} = \begin{pmatrix} \sin(\psi)\dot{x} - \sin(\psi)\dot{y} \\ \sin(\psi)\dot{x} + \cos(\psi)\dot{y} \\ \dot{\psi} \\ \dot{\psi} + u_a \\ -\dot{\psi}\dot{x} + \frac{2}{m}(F_{yf} \cos u_\delta + F_{yr}) \\ \frac{2}{I_z}(l_f F_{yf} - l_r F_{yr}) \end{pmatrix} \quad (6)$$

where F_{yf} and F_{yr} are the lateral tire forces at the front and rear respectively, and m , I_{zz} , l_f , l_r are the mass, yaw inertia, distance from Center of Gravity (C.G) to front axle and distance from C.G. to rear axle parameters respectively.

Except for aerodynamic drag and gravity, tire forces are the only external force acting on the vehicle. The tire forces have highly non-linear behavior at the limits of friction; however, it is also on that limit that tires exhibit maximum force, which is required for aggressive driving. Thus, correct tire modeling is essential for which we use the *tire brush model*. In particular, the variant proposed by Fromm [19] which defines forces as:

$$F_y = \begin{cases} -C_s \tan \alpha + \frac{C_s^2}{3\mu F_z} |\tan \alpha| \tan \alpha - \frac{C_s^3}{27\mu^2 F_z^2} \tan^3 \alpha & |\alpha| \leq \alpha_{crit} \\ -\mu F_z \operatorname{sgn}(\alpha) & |\alpha| > \alpha_{crit} \end{cases} \quad (7)$$

where C_s is the cornering stiffness, μ is the friction coefficient between the road surface and the tire, α is the slip angle of the front, and rear tires is given by:

$$\alpha_f = u_\delta - \arctan \frac{\dot{y} + \dot{\psi} l_f}{\dot{x}} \quad \alpha_r = -\arctan \frac{\dot{y} - \dot{\psi} l_r}{\dot{x}} \quad \alpha_{crit} = \arctan \frac{3\mu F_z}{C_s} \quad (8)$$

α_{crit} is the smallest magnitude slip angle at which the entire contact patch is at the friction limit. F_z is the vertical force for the tire which for our simplified model can be calculated for the front and rear axles with:

$$F_{zf} = \frac{l_r}{2(l_f + l_r)} mg \quad F_{zr} = \frac{l_f}{2(l_f + l_r)} mg \quad (9)$$

where m is the mass of the vehicle, and g is the gravitational force. Combining this tire model expressed by its lateral forces F_{yf} and F_{yr} with Equation 6 results in our final parametric model.

4.2 Non-parametric part

Intending to enhance the performance of the parametric model, we chose to use a neural network with 2 hidden layers as it can accurately capture the highly non-linear dynamics missed by the parametric model and do so with excellent computational efficiency, as shown in [12]. For the purposes of this project, the non-parametric part predicts only the velocities and yaw rate. In Figure 1, it was experimentally found that a network with hidden layer sizes of (20, 20) and *tanh* activation functions performs sufficiently well while maintaining a small set of parameters. For the other hyper-parameter, we also experimentally identified that the following parameters produce the best results: batch size of 100, epoch size of 1000, Adam optimizer [20] with a learning rate of 10^{-3} , L2 regularization with a weight decay coefficient of 10^{-3} and data normalization. The reasoning for these choices can be found in Appendix A. We found it particularly useful to use significant L2 regularization when training the full semi-parametric model. The reasoning for this is that the neural network is only learning residuals and does not need to make significant corrections to the prior prediction of the parametric part of the model.

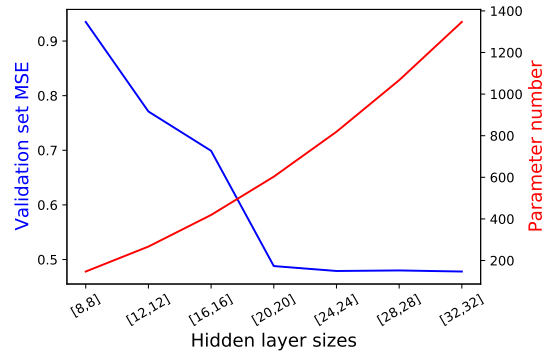


Figure 1: Accuracy comparison between networks of different hidden layer sizes.

5 Iterative learning

In designing this system, we had two goals in mind (1) the model must adapt to changes swiftly but also (2) not wholly ignore prior dynamics. The justification for this is that the system must adapt to changes quickly but also be resilient to short-term disturbances such as wind gusts. We first assume that the model is bootstrapped with data generated from the $\mathcal{P}_{ID}(\mathcal{X})$ distribution, where the parametric part of the model is fit first, and the neural network learns its residuals. Then we assume that the dynamics remain similar throughout the operation of the robot and that the parametric model parameters will not change significantly; as such, we chose only iteratively to update the neural network. Although this is a strong assumption, it guarantees that we do not encounter the problem of non-stationary targets (Section 3) while still learning an accurate dynamics model, albeit potentially slower.

To address the issue of catastrophic forgetting in the neural network (Section 3), we chose to use pseudo-rehearsal methods where the model is trained simultaneously on the old data (drawn from $\mathcal{P}_{ID}(\mathcal{X})$) and the new data (drawn from $\mathcal{P}_L(\mathcal{X})$). For generating inputs $X_{ID} \sim \mathcal{P}_{ID}(\mathcal{X})$, a **Gaussian Mixture Model (GMM)** [21] is used, and the neural network generates the corresponding targets Y_{ID} . The GMM is fit using expectation-minimization (EM) [22]. To guarantee consistency between the two data distributions, we utilize a **constrained gradient update rule** introduced in [12]. The similarity between the two data distributions is enforced via the cosine of the angle between the update directions of the gradients computed from both datasets:

$$\theta_{i+1}^g = \theta_i^g - \gamma(\alpha G_L + G_{ID}) \quad (10)$$

$$\alpha = \max_{\alpha \in [0,1]} \text{s.t. } \langle \alpha G_L + G_{ID}, G_{ID} \rangle \geq 0 \quad (11)$$

where γ is the learning rate, G_L and G_{ID} are the gradients of the model parameters computed from the two data distributions:

$$G_L = \nabla_{\theta_i^g} \|Y_L - g(X_L^g, \theta_i^g)\|^2 \quad (12)$$

$$G_{ID} = \nabla_{\theta_i^g} \|Y_{ID} - g(X_{ID}^g, \theta_i^g)\|^2 \quad (13)$$

The new data is stored in the **Local Operating Set Buffer** of fixed length and is used as a full batch along with a batch of equal size generated by the GMM to train the neural network. Training initiates online when the **buffer** is full of new data. We empirically found that it is best to use the same optimizer and hyper-parameters as used during the bootstrapping. After training has finished, we use the new data to **update the GMM via an incremental EM training approach** [23]. This way, the memory footprint of the data is kept low while ensuring the model never forgets old data. A visual guide of the system is presented in Figure 2.

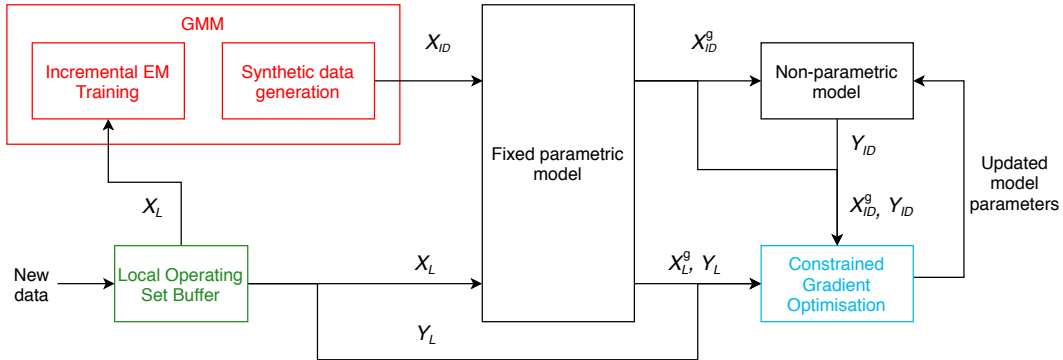


Figure 2: Iterative learning system for the non-parametric model.

6 Experimental Validation

6.1 Experimental Setup

We test our approach to the task of learning dynamics as part of an MPC in an autonomous racing scenario where the goal is to drive as fast as possible without crashing or leaving the race track

boundaries. This is a particularly difficult problem as aggressive driving attempts to exploit the dynamical limits of the vehicle, which are highly non-linear, rarely observed, and challenging to model.

The target platform for our experiments is a full-scale electric vehicle fitted with all the necessary sensors and compute for the task. Dynamically the vehicle is capable of achieving $11,8 \text{ m/s}^2$ lateral acceleration, $4,9 \text{ m/s}^2$ longitudinal acceleration, $8,8 \text{ m/s}^2$ longitudinal deceleration (braking) and top speeds of 200 kph (56 m/s). Sadly due to external factors, we were unable to execute real-world experiments, and as such, this section is only limited to simulated ones. A proprietary simulator is used with an extremely accurate vehicle dynamics model. Bootstrapping and validation data is collected from this simulation by manual driving to maintain real-world feasibility. The data is then filtered, smoothed, and prepared for model learning.

For the remainder of this section, we have conducted three types of experiments. (1) *semi-parametric model experiments* which aim to validate the benefits of semi-parametric models. (2) *offline iterative learning experiments* where the system is bootstrapped on a dataset and then trained iteratively on another dataset collected with modified dynamics. We evaluate the initial and adapted models on the latter dataset. (3) *online iterative learning experiments* where the system is bootstrapped on a dataset and then deployed along with a controller in the simulation with modified dynamics. For this, we use the open-source Model Predictive Path Integral (MPPI) controller, which is a sampling-based, gradient-free method that computes the optimal trajectory as a cost-weighted average over randomly sampled trajectories [24].

6.2 Semi-parametric Model Experiments

In this subsection, we validate the generalization and efficiency of using semi-parametric models. Firstly, we split the dataset collected into a train, validation, and test sets such that each set contains data from a unique part of the state space. This is done by sorting all the data by longitudinal velocity data and splitting it into the sets as proportions of the full spectrum of velocities. This is a plausible real-world scenario as it is usually not feasible and expensive to collect large datasets exploring the entire system state space. Further explanation and visualizations can be found in Appendix B. Next, we evaluate the generalization capabilities of the semi-parametric model by comparing it to only its parametric part and only its non-parametric parts. The results of this training procedure plotted against the neural network epoch number can be seen in Figure 3.

The purely parametric model exhibits the worst training set loss. However, taking a look at the validation loss performance, the purely parametric model shows superior generalization capabilities when compared to the strictly non-parametric model, which overfits the training state space. The semi-parametric model achieves similar accuracy to the purely non-parametric model on the training set and is slightly more efficient at learning while also showing excellent generalization capabilities on the validation set. It is worth noting that here for the purely non-parametric model, we use a larger neural network with 2 hidden layers of 32 neurons each as otherwise, the model is not able to capture the dynamics adequately.

6.3 Offline Iterative Learning Experiments

In this subsection, we use the bootstrapping and validation datasets described earlier to evaluate the proposed iterative learning offline without the added stochasticity of introducing a controller. We first use the bootstrapping dataset to train the model and then let it adapt to the changed

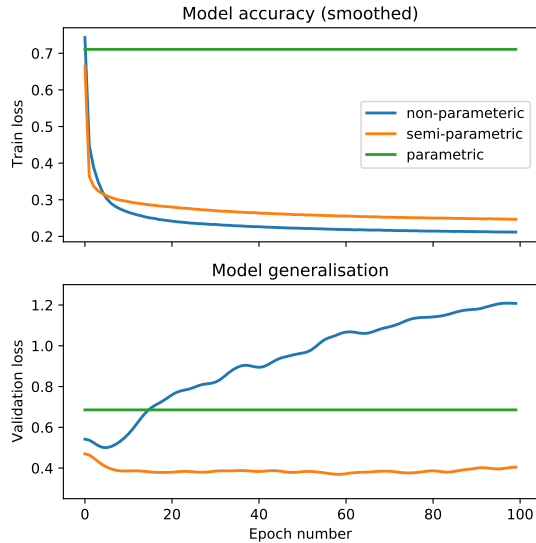


Figure 3: Comparison between the learning and generalization capabilities of the different model types. The purely parametric model shows a constant error since it is not trained in batches.

dynamics in the validation set where the mass of the vehicle changed from 1350 kg to 1430 kg and the friction coefficient μ from 1,0 to 0,8. This is a plausible real-world scenario found during empirical testing where a passenger can increase the overall mass of the vehicle, and a different road surface can reduce tire friction. After the model has adapted through the whole dataset, we evaluate its prediction accuracy on the bootstrapping data, on the validation data, and finally on a third dataset, which we will call test data consisting again of modified dynamics. Empirically we found that a local operating set of 500 data entries (10 seconds of driving) worked well combined with a mini-batch size of 100 and 3 epochs of training every time the local operating set is full. To highlight the generalization capabilities of our proposed approach, we also show the results of using pure Stochastic Gradient Decent (SGD) for learning with the same hyper-parameters as our proposed approach. The final results are shown in Table 1, where we can see that both of the iterative learning models adapt successfully to the modified dynamics of the validation dataset; however, the performance of the simple SGD approach deteriorates when evaluated on the test data. On the other hand, our proposed iterative learning approach produces excellent results on the validation data while also generalizing to the test data.

Table 1: Offline iterative learning experimental results.

	Bootstrap data MSE	Validation data MSE	Test data MSE
Bootstrapped	0.504	1.037	1.050
SGD	2.148	0.121	0.442
Iterative learning	1.423	0.181	0.170

6.4 Online Iterative Learning Experiments

In this subsection, we deploy the semi-parametric model and the iterative learning system as part of the MPPI controller. First, we use the trained semi-parametric model from Section 6.2 without iterative learning; we call this the *normal* scenario. Then we run another experiment with a simulation with the same dynamics modifications introduced in the previous subsection; we call this the *modified* experiment. We chose the control parameters to be safe ones to ensure that all models complete the course successfully. These two experiments serve as the lower and upper bound of performance we should expect. Finally, we run our iterative learning system bootstrapped with data of unmodified dynamics. Figure 4 shows the results of 10 consecutive laps of driving averaged over five different experiments for each scenario to account for the stochastic nature of the controller and the dynamics learning system. Table 2 shows one of the iterative learning experiments in more detail where the MSE is calculated on the dataset of modified dynamics collected over 10 laps. To showcase the model learning, we have saved snapshots of the model at the beginning of an experiment (0 s), at 60 s, and 120 s of learning and ran these as fixed models for one lap. Then we plot the resulting GG diagrams in Figure 5, which indicate how well the controller harnesses the full vehicle dynamics.

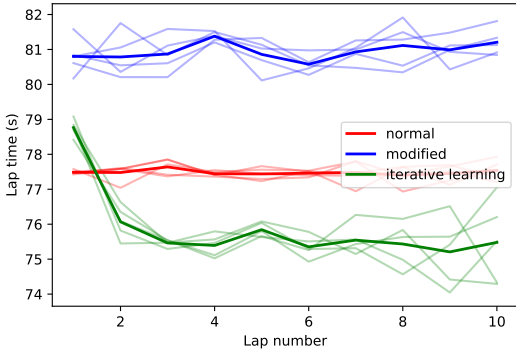


Figure 4: Consecutive lap times of iterative learning system compared to the fixed models. Each model is ran for 5 different experiments, and their average is shown with a solid line.

Lap	Lap Time	Model MSE
0	-	1.050
1	78.69	0.210
2	75.93	0.194
3	75.52	0.193
4	75.53	0.192
5	74.94	0.194
6	75.55	0.192
7	75.55	0.192
8	75.44	0.192
9	75.03	0.193
10	75.24	0.193

Table 2: Lap time performance of the iterative learning system and model MSE over one of the 10-lap experiments from the Figure on the left. Lap 0 stands for the bootstrapped model.

From these results, we can see that iterative learning is working successfully and adapts to the changing dynamics*. Furthermore, it consistently outperforms the *normal* experiments, which were previously considered the upper bound of performance. This could be because the model fits more specifically to the current racetrack or because the controlled explores the state space better than the human driver. However, there is also a negative effect - the iterative learning model becomes progressively more unstable, which is shown via the variance in the lap times of the iterative learning system. We do not explicitly consider comparisons to simple SGD model updates as those usually result in incorrectly learned dynamics and the controller being unable to finish a single lap.

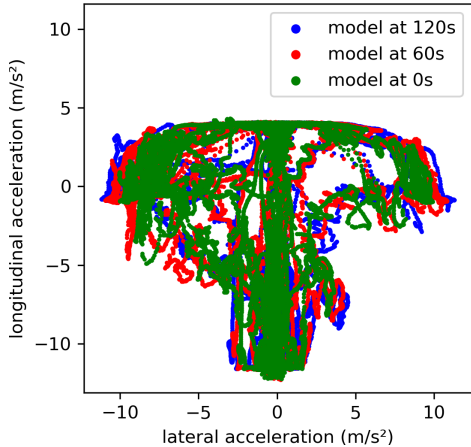


Figure 5: GG plot of accelerations achieved by the iteratively learned model at different stages. Larger lateral acceleration magnitudes indicate better and more aggressive performance.

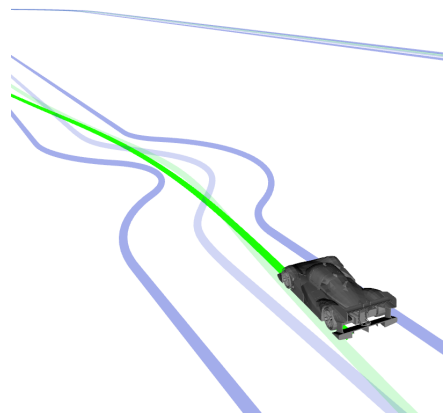


Figure 6: Data visualization from the simulation. The solid green line is the optimized trajectory by MPPI, and the transparent green line is the reference path.

7 Conclusion

In this paper, we present a novel approach to modeling robot dynamics by using a non-linear semi-parametric model with a neural network, resulting in an accurate and computationally efficient model. We developed a system for iteratively updating the non-parametric part of the model using pseudo-rehearsals and apply it to the case of Model Predictive Control (MPC) in autonomous racing. We show that our model achieves excellent accuracy in comparison to purely parametric models while also retaining superb generalization capabilities in contrast to a strictly non-parametric model. To test our iterative learning system, we use several datasets to evaluate our approach and show that our model overcomes the issue of catastrophic forgetting common for neural networks. Furthermore, it adapts to changing dynamics with batch learning and also generalizes to unseen state spaces. Finally, we deploy our system along with an MPC in a simulation and show that our system can safely adapt changing dynamics at operation time and improve overall controller performance.

Due to external circumstances, we were unable to apply our approach to a real system. However, we would like to note that our iterative learning system allows for straightforward simulation-to-real-world transfer where the model can be bootstrapped in simulation and then left to adapt in the real-world iteratively. Even though our simulation shows positive results, dynamics modeling with neural network has two short-comings - instability and batch-learning. The first issue is caused by data overfitting and exploding gradients, which, although can be mitigated with regularization techniques, is still an underlying concern when deploying systems with long life-spans. The second issue of batch-learning is inherent to neural networks. It renders them incapable of learning short-term changes to the system dynamics, which is desirable but of uncertain feasibility. Additionally, in our approach, we do not update the parametric part of the model iteratively and leave its error to be absorbed into the non-parametric part of the model. Updating the full semi-parametric model is desirable but difficult [8].

*A supplementary video is available online - <https://youtu.be/aLWXa95AEKE>.

Acknowledgments

The MPC used throughout this paper is based on the open-source implementation of the Model Predictive Path Integral (MPPI) controller by G. Williams et al. [25]

References

- [1] R. Siegwart, I. R. Nourbakhsh, D. Scaramuzza, and R. C. Arkin. *Introduction to autonomous mobile robots*. MIT press, 2011.
- [2] D. J. Fagnant and K. Kockelman. Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. *Transportation Research Part A: Policy and Practice*, 77: 167–181, 2015.
- [3] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli. A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on intelligent vehicles*, 1 (1):33–55, 2016.
- [4] B. Siciliano and O. Khatib. *Springer handbook of robotics*, pages 190–192. Springer, 2016.
- [5] B. Siciliano and O. Khatib. *Springer handbook of robotics*, pages 113–136. Springer, 2016.
- [6] D. Nguyen-Tuong and J. Peters. Using model knowledge for learning inverse dynamics. In *2010 IEEE international conference on robotics and automation*, pages 2677–2682. IEEE, 2010.
- [7] D. Romeres, M. Zorzi, R. Camoriano, and A. Chiuso. Online semi-parametric learning for inverse dynamics modeling. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 2945–2950. IEEE, 2016.
- [8] J. Smith and M. Mistry. Online simultaneous semi-parametric dynamics model learning. *IEEE Robotics and Automation Letters*, 5(2):2039–2046, 2020.
- [9] A. Robins. Catastrophic forgetting in neural networks: the role of rehearsal mechanisms. In *Proceedings 1993 The First New Zealand International Two-Stream Conference on Artificial Neural Networks and Expert Systems*, pages 65–68. IEEE, 1993.
- [10] I. J. Goodfellow, M. Mirza, D. Xiao, A. Courville, and Y. Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.
- [11] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- [12] G. Williams, B. Goldfain, J. M. Rehg, and E. A. Theodorou. Locally weighted regression pseudo-rehearsal for online learning of vehicle dynamics. *arXiv preprint arXiv:1905.05162*, 2019.
- [13] R. Johansson. Parametric models of linear multivariable systems for adaptive control. *IEEE transactions on automatic control*, 32(4):303–313, 1987.
- [14] S. Wan and L. E. Banta. Parameter incremental learning algorithm for neural networks. *IEEE transactions on neural networks*, 17(6):1424–1438, 2006.
- [15] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*, 2017.
- [16] R. Vilalta and Y. Drissi. A perspective view and survey of meta-learning. *Artificial intelligence review*, 18(2):77–95, 2002.
- [17] A. Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2): 123–146, 1995.
- [18] M. Althoff. Commonroad: Vehicle models. *Technische niversität München, Garching*, pages 1–25, 2017.

- [19] R. Hadekel. The mechanical characteristics of pneumatic tires. *Clearinghouse Fed Sci & Tech Info*, 1952.
- [20] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [21] D. A. Reynolds. Gaussian mixture models. *Encyclopedia of biometrics*, 741, 2009.
- [22] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977.
- [23] Y. Zhang, L. Chen, and X. Ran. Online incremental em training of gmm and its application to speech processing applications. In *IEEE 10th INTERNATIONAL CONFERENCE ON SIGNAL PROCESSING PROCEEDINGS*, pages 1309–1312. IEEE, 2010.
- [24] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou. Information-theoretic model predictive control: Theory and applications to autonomous driving. *IEEE Transactions on Robotics*, 34(6):1603–1622, 2018.
- [25] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou. Aggressive driving with model predictive path integral control. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1433–1440. IEEE, 2016.

Appendices

A Semi-parametric model hyper-parameter search

Due to that and computational restrictions, we chose to explore hidden layer sizes of (8,8), (12,12), (16,16), (20,20), (24,24), (28,28) and (32,32) and learning rates of 0.0025, 0.001 and 0.00075. The other hyper-parameters are shown below, and the results are shown in Table 3.

- batch size of 100 and epoch size of 1000
- tanh activation function
- Adam optimizer
- 10^{-5} weight decay coefficient
- data normalization only for the non-parametric part
- 70%, 20%, 10% train, validation and test dataset split

Table 3: Semi-parametric training experiments with different hidden layer sizes and learning rates. Average lap times are estimated over five consecutive laps. A failure is designated if the car span or went off track; in that case, lap times are omitted from the table. The best hyper-parameters for each model size is shown in bold. The model which has achieved the best lap times is italicized.

Hidden layer size	Learning rate	Validation MSE	Average lap time	Lap time std.
8,8	0.0025	0.981	-	-
8,8	0.001	0.935	-	-
8,8	0.00075	1.044	-	-
12,12	0.0025	0.801	-	-
12,12	0.001	0.771	126.45	0.744
12,12	0.00075	0.794	-	-
16,16	0.0025	0.733	-	-
16,16	0.001	0.699	124.82	0.24
16,16	0.00075	0.7368	125.06	0.13
20,20	0.0025	0.502	126.2583	0.8365
20,20	0.001	0.491	120.65	0.35
20,20	0.00075	0.496	121.98	0.484
24,24	0.0025	0.508	126.98	0.836
24,24	0.001	0.479	121.12	0.348
24,24	0.00075	0.491	121.27	0.461
28,28	0.0025	0.511	-	-
28,28	0.001	0.480	121.64	0.488
28,28	0.00075	0.485	120.88	0.374
32,32	0.0025	0.525	-	-
32,32	0.001	0.479	125.6762	0.6672
32,32	0.00075	0.478	121.93	0.248

We can see that the bolded model with hidden layer sizes (20, 20) achieves sufficiently good performance in both Validation MSE and average lap times while maintaining computational efficiency. However, we are still not finished with identifying hyper-parameters. The semi-parametric model differs from the non-parametric model as in that it learns not the full model, but only the residuals leftover from the parametric part. As such, one would expect the weights of such a model to be smaller. Thus, here we consider one more important hyper-parameter - the *weight decay coefficient*, which is responsible for the regularization of the neural network. In the case of this project, we use L2 regularization due to its prior availability and excellent results. We put this to the test with the previously identified best semi-parametric model and show the results in Table 4. From there, we can identify two good values of the parameter: $\lambda = 10^{-4}$ produces the best result, but $\lambda = 10^{-3}$ provides more stable lap times as training error plateaus. As stability is crucial for the vehicle model, we chose to utilize the latter.

Table 4: Comparison between the model trained with different weight decay coefficients. Average lap times are estimated over five consecutive laps. A failure is designated if the car span or went off track; in that case, lap times are omitted from the table.

Weight decay coeff. (λ)	Validation MSE	Average lap time	Lap time std.
10^{-1}	1.721	-	-
10^{-2}	0.951	128.73	2.19
10^{-3}	0.504	121.12	0.12
10^{-4}	0.488	120.65	0.35
10^{-5}	0.491	120.94	0.47

B Data split for different state spaces

To test the generalization and accuracy benefits of the semi-parametric model in Section 6.2, we had to modify the typical machine learning validation process where data is randomized into train, validation and test sets. Instead, we split our data based on sorted longitudinal velocities, which nicely separates the whole state space such that the validation set includes data that is not present in the test set. Thus, this approach allows us to validate the generalization capabilities of dynamics models and is reflective of real-world applications where obtaining data from the full state space is not feasible. An example of this is shown in Figure 7.

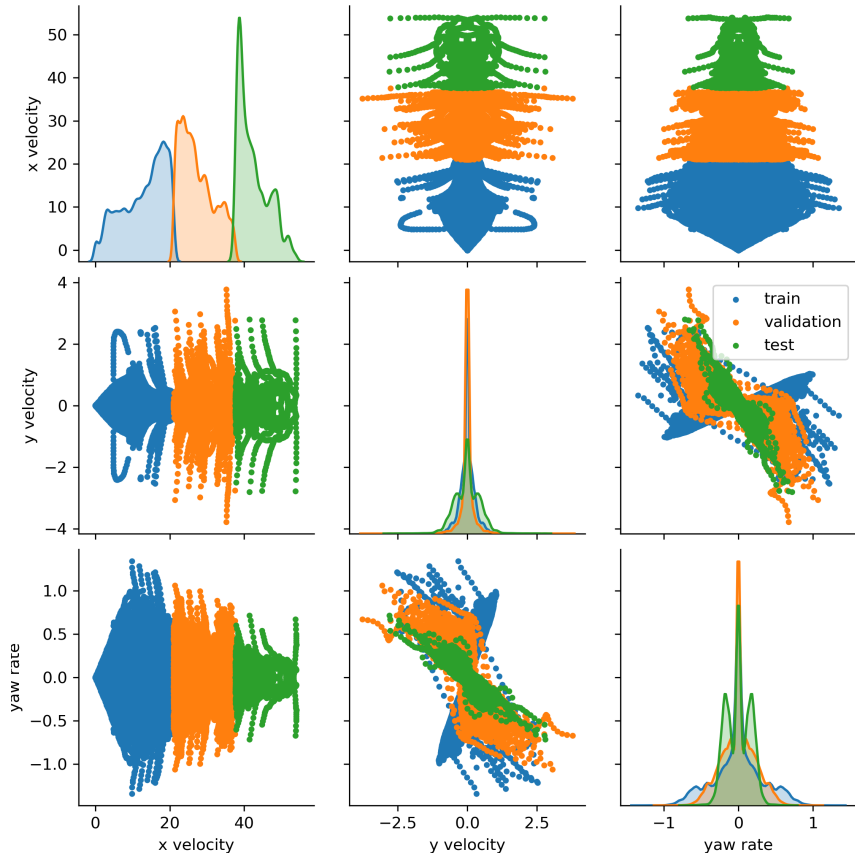


Figure 7: Example of how the dataset is split into train, validation, and test sets such that they occupy a different part of the overall state space. In this particular example, the data is divided according to the x velocity field, which can be seen in the top left plot. In this plot, the train, validation, and test datasets have 60%, 35%, 5% proportions. It is worth noting that these proportions are not accurately represented in the diagonal kernel density plots.