

# PlanNetX: Learning an Efficient Neural Network Planner from MPC for Longitudinal Control

**Jasper Hoffmann**

**Diego Fernandez Clausen**

**Julien Brosseit**

*Neurorobotics Lab, University of Freiburg, Germany.*

HOFFMAJA@INFORMATIK.UNI-FREIBURG.DE

FERNANDI@INFORMATIK.UNI-FREIBURG.DE

BROSSEIT@INFORMATIK.UNI-FREIBURG.DE

**Julian Bernhard**

**Klemens Esterle**

**Moritz Werling**

**Michael Karg**

*BMWGroup, Unterschleissheim, Germany.*

JULIAN.BERNHARD@BMW.DE

KLEMENS.ESTERLE@BMW.DE

MORITZ.WERLING@BMW.DE

MICHAEL.KARG@BMW.DE

**Joschka Bödecker**

*Neurorobotics Lab, University of Freiburg, Germany.*

JBOEDECK@INFORMATIK.UNI-FREIBURG.DE

**Editors:** A. Abate, M. Cannon, K. Margellos, A. Papachristodoulou

## Abstract

Model predictive control is a powerful, optimization-based approach for controlling dynamical systems. However, the computational complexity of online optimization can be problematic on embedded devices. Especially, when we need to guarantee a fixed control frequency. Thus, previous work proposed to reduce the computational burden using imitation learning approximating the MPC policy by a neural network. In this work, we instead learn the whole planned trajectory of the MPC. We introduce a combination of a novel neural network architecture PlanNetX and a simple loss function based on the state trajectory that leverages the parameterized optimal control structure of the MPC. We validate our approach in the context of autonomous driving by learning a longitudinal planner and benchmarking it extensively in the CommonRoad simulator using synthetic scenarios and scenarios derived from real data. Our experimental results show that we can learn the open-loop MPC trajectory with high accuracy while improving the closed-loop performance of the learned control policy over other baselines like behavior cloning.

**Keywords:** imitation learning, model predictive control, autonomous driving

## 1. Introduction

Motion planning for autonomous driving is challenging due to the high uncertainty of the surrounding traffic requiring constant re-planning to adjust the plan to newly obtained information. Model predictive control (MPC) is an optimization-based approach that relies on rapidly resolving an optimal control problem (OCP) numerically, planning a trajectory at each step. However, online optimization can be a limiting computational factor if a control needs to be provided in a fixed time interval and especially if a lot of different candidate trajectories need to be planned in parallel.

To reduce the computational cost of MPC, a lot of effort was put into computing the MPC solutions offline and storing them in an appropriate representation to allow fast online evaluation during deployment. One example of this is *explicit MPC*. In the case of linear MPC, the MPC policy,

defined as the first control of the MPC solution, is a piece-wise linear function that can be stored as a look-up table [Bemporad et al. \(1999\)](#). However, for non-linear dynamics or constraints, the exact representation of an MPC policy is challenging and often done only approximately [Johansen and Grancharova \(2003\)](#). Another prominent example is imitation learning (IL). Here, we try to approximate the MPC policy by using function approximators like neural networks. Previous work focused on improving the safety aspects [Mamedov et al. \(2022\)](#) of the learned policy or replacing the loss function of standard behavior cloning (BC) [Carius et al. \(2020\)](#); [Ghezzi et al. \(2023\)](#), which tries to minimize the squared distance to the first control of the MPC solution, using objectives that are aware of the underlying OCP structure of the MPC.

In this work, instead of just learning the first control we want to learn the whole planned trajectory of the MPC solution. This is useful in the field of autonomous driving, where it is common to separate the tasks of motion planning from the task of low-level vehicle control to simplify the control problem. Additionally, we find that learning the whole planned trajectory allows us to define a simple loss formulation that uses the underlying OCP structure during training being potentially aware of the network’s approximation errors. Furthermore, learning the trajectory could be used to warm start the MPC solver for significantly faster convergence and better safety guarantees. The goal of this work is to general improve the planning time in comparison to numerically solving the OCP. Thus, we deploy common neural network compression techniques like pruning and quantization to reduce the inference time.

The contributions of this work are threefold: **(1)** We propose an imitation learning framework called `PlanNetX` specifically for learning the planned trajectory of an MPC planner, consisting of a simple loss function and a specialized neural network architecture. **(2)** We validate the approach by learning a longitudinal planner for autonomous driving and extensively testing it on a handcrafted CommonRoad, [Althoff et al. \(2017\)](#), benchmark set consisting of real and synthetic driving data. **(3)** We analyze whether common neural network compression techniques like pruning and quantization allow us to reduce the inference time even further.

## 1.1. Related Work

Recent autonomous driving systems incorporate more and more learning-based control [Kiran et al. \(2022\)](#), by learning control policies via IL or reinforcement learning (RL) either from simulations or human driving data. For autonomous driving, RL was used to warm-start the MPC solver in order to mitigate the problem of the solver getting stuck in local minima [Hart et al. \(2019\)](#).

In the context of IL from an MPC policy, various aspects were considered: (1) *Safety and Stability*: [Mamedov et al. \(2022\)](#) use a safety filter formulation similar to [Wabersich and Zeilinger \(2021\)](#) to guarantee the safety of the final control policy. Using the concept of control-barrier functions, [Cosner et al. \(2022\)](#) show that under certain conditions the safety guarantees of a robust MPC expert can be transferred to the learned policy. (2) *Verification*: [Schwan et al. \(2023\)](#) provide an upper bound on the approximation error of the learned policy to the original MPC policy by solving a mixed-integer formulation, whereas in [Hertneck et al. \(2018\)](#); [Karg et al. \(2021\)](#) a probabilistic bound is given. (3) *Improved Loss Function*: Further research tries to replace the surrogate loss of standard behavior cloning, which is often the squared distance, with an objective that takes into account the underlying OCP problem: In [Carius et al. \(2020\)](#); [Reske et al. \(2021\)](#) a loss based on the control Hamiltonian is introduced, whereas in [Ghezzi et al. \(2023\)](#) a Q-loss is introduced by calculating the value of a fixed first control with an MPC formulation.

IL from trajectory optimization was considered by [Mordatch and Todorov \(2014\)](#); [Levine et al. \(2016\)](#) where an alternating direction method of multipliers (ADMM) method jointly optimizes a task and policy reconstruction cost. Similarly, in [Levine and Koltun \(2013a,b\)](#) an iterative linear quadratic regulator (iLQR) [Li and Todorov \(2004\)](#) is used for guiding an RL policy to low costs regions. A related learning formulation to the one presented in this work is [Drgoňa et al. \(2021, 2024\)](#) where the learning objective is derived from an OCP by differentiating through the dynamics to solve a tracking formulation. Our approach differs in that we first solve the OCP via numerical optimization and then use the solution as a tracking signal.

In general, apart from [Drgoňa et al. \(2021, 2024\)](#); [Vaupel et al. \(2020\)](#) most approaches focus on learning the control policy but not the planned trajectory.

## 2. Preliminaries

In this section, we give a quick introduction to MPC and IL.

### 2.1. Model Predictive Control

In the following, we introduce the notion of a parameterized discrete-time OCP. At each time step, an MPC planner observes a current state  $\bar{x}_0 \in \mathbb{R}^{n_x}$  to generate an optimal control and state trajectory  $U \in \mathbb{R}^{n_u \times N}$  and  $X \in \mathbb{R}^{n_x \times N+1}$  by solving the OCP

$$\min_{X, U} T(x_N) + \sum_{k=0}^{N-1} L_k(x_k, u_k) \quad (1a)$$

$$\text{s.t. } x_{k+1} = f(x_k, u_k), \quad k = 0, \dots, N-1, \quad (1b)$$

$$h(x_k, u_k, p_k) \leq 0, \quad k = 0, \dots, N-1, \quad (1c)$$

$$x_0 = \bar{x}_0, \quad (1d)$$

with a planning horizon length  $N$ , stage costs  $L_k$ , terminal costs  $T$ , and known, potentially simplified, system dynamics  $f$ . Furthermore, in (1c) we assume that we can parameterize the inequality constraints  $h$  at each planning step with additional parameters  $P \in \mathbb{R}^{n_p \times N}$ . For the rest of the work, we denote  $X^* \in \mathbb{R}^{n_x \times N+1}$  and  $U^* \in \mathbb{R}^{n_u \times N}$  as the optimal solution found by the MPC when solving the above OCP.

### 2.2. Imitation Learning

With IL we try to approximate the behavior of an expert by using a function approximator like a neural network. In the context of this work, the MPC planner is the expert. We differentiate between learning the *control policy* or control law  $\bar{x}_0 \mapsto u_0^*$  of the MPC, which is already enough to control the system, and learning the *planning policy* of the MPC planner, where we want to approximate the mapping  $\bar{x}_0 \mapsto (X^*, U^*)$ . Note, that the second formulation contains the first one as a sub-problem. In more mathematical terms, we want to approximate either the MPC control or plan by the following parameterized policies

$$\begin{aligned} \pi_{\theta}^{\text{policy}} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_p \times N} &\rightarrow \mathbb{R}^{n_u}, & \pi_{\theta}^{\text{plan}} : \mathbb{R}^{n_x} \times \mathbb{R}^{n_p \times N} &\rightarrow \mathbb{R}^{n_x \times (N+1)} \times \mathbb{R}^{n_u \times N}, \\ \bar{x}_0, P &\mapsto \hat{u}_0^{\theta} & \bar{x}_0, P &\mapsto \hat{X}^{\theta} \times \hat{U}^{\theta} \end{aligned} \quad (2)$$

where  $\theta$  denotes the parameter vector of the policy which could be for example the weights of a neural network. For the planning policy  $\pi_\theta^{\text{plan}}$ , we assume that  $\hat{x}_0^\theta = \bar{x}_0$ . The goal of both learning formulations is to minimize the following two objectives:

$$\mathcal{L}^{\text{policy}}(\theta) = \mathbb{E}_{\bar{x}_0, P \sim \mathcal{D}} \left[ \ell \left( \hat{u}_0^\theta, u_0^* \right) \right], \quad \mathcal{L}^{\text{plan}}(\theta) = \mathbb{E}_{\bar{x}_0, P \sim \mathcal{D}} \left[ \ell \left( \hat{U}^\theta, \hat{X}^\theta, U^*, X^* \right) \right], \quad (3)$$

where  $\mathcal{D}$  is a given state and parameter distribution over  $\mathbb{R}^{n_x} \times \mathbb{R}^{n_p \times (N+1)}$  whereas  $\ell$  is an arbitrary pointwise loss functions.

**Behavior Cloning:** For learning just the control policy, the simplest formulation is BC which does not use any information about the underlying OCP. Instead, here we minimize the distance of our predicted control  $\hat{u}_0^\theta$  to the optimal control  $u_0^*$  quadratically

$$\mathcal{L}^{\text{BC}}(\theta) := \mathbb{E}_{\bar{x}_0, P \sim \mathcal{D}} \left[ \|\hat{u}_0^\theta - u_0^*\|_2^2 \right]. \quad (4)$$

Recent work tries to replace the quadratic distance with a distance measure that considers the underlying OCP structure as in [Ghezzi et al. \(2023\)](#); [Carius et al. \(2020\)](#); [Reske et al. \(2021\)](#).

### 3. Learning an MPC-based Planner

In the following, we want to discuss how we can use the knowledge of the experts underlying OCP structure to learn the planned trajectory. We first discuss different loss formulations and then continue with the PLANNETX architecture tailored to learning the MPC planning policy  $\pi^{\text{plan}}$  of (2).

#### 3.1. Loss Formulation

Given the initial state  $\bar{x}_0$  and the dynamics model  $f$ , predicting the controls  $\hat{U}^\theta$  is already enough to also generate the according state trajectory  $\hat{X}^\theta$ . One idea would be to minimize the distance  $\hat{U}^\theta$  between  $U^*$  as done in [Vaupel et al. \(2020\)](#). For this, we define the *control trajectory loss*

$$\mathcal{L}^u(\theta) := \mathbb{E}_{\bar{x}_0, P \sim \mathcal{D}} \left[ \frac{1}{N} \sum_{k=0}^{N-1} \gamma^k \left\| \hat{u}_k^\theta - u_k^* \right\|_W^2 \right], \quad (5)$$

where  $\gamma$  is a discount factor, with  $0 < \gamma \leq 1$  allowing to put more weights on early controls in the loss, and  $W$  is a positive-definite weight matrix, defining a norm via  $\|x\|_W^2 = x^T W x$ . However, the loss function defined in (5) does not consider any approximation errors of the neural network. Thus, if there is a mismatch early in the predicted trajectory, there is no incentive for the network to correct it in the later parts of the predicted trajectory.

Instead of minimizing the distance between  $\hat{U}^\theta$  and  $U^*$  in control space, we instead minimize the distance in the state space, thus between the predicted state trajectory  $\hat{X}^\theta$  and  $X^*$ . We define the *state trajectory loss* with

$$\mathcal{L}^x(\theta) := \mathbb{E}_{\bar{x}_0, P \sim \mathcal{D}} \left[ \frac{1}{N} \sum_{k=1}^N \gamma^k \left\| \hat{x}_k^\theta - x_k^* \right\|_W^2 \right]. \quad (6)$$

If there are specific constraints and costs regarding the control trajectory, one could also potentially combine the two loss formulations (5) and (6). Note, that using the above loss formulations we are implicitly able to learn the underlying expert cost function and the respective constraints.




Figure 1: The two proposed architectures `PlanNetX` and `PlanNetXEnc`. The `PlanNetXEnc` includes an additional Transformer encoder layer that contains potentially the information of the parameters of all different time steps.

### 3.2. PlanNetX

The motivation for the `PlanNetX` architecture comes from taking a look at the original OCP formulation of (1) that we solve for MPC inference. The idea is to replace the free decision variables  $X, U$  with predictions of a neural network policy  $\pi$ , by starting a roll-out from the initial state  $\bar{x}_0$  and iteratively applying our learned policy  $\pi_\theta$  and the dynamics  $f$ . We then apply the state trajectory loss on the obtained roll-out resulting in the following formulation:

$$\min_{\theta} \mathbb{E}_{\bar{x}_0, P \sim \mathcal{D}} \left[ \frac{1}{N} \sum_{k=1}^N \gamma^k \left\| \hat{x}_k^\theta - x_k^* \right\|_W^2 \right] \quad (7a)$$

$$\hat{x}_{k+1} = f(\hat{x}_k, \hat{u}_k), \hat{u}_k = \pi_\theta(\hat{x}_k, z_k), z_k = \psi_k^\theta(P, T), \hat{x}_0 = \bar{x}_0. \quad (7b)$$

The policy  $\pi^\theta$  takes an information vector  $z_k$  containing the necessary information of the inequality constraint parameter  $P$ . The information vector is derived from an additional constraint network  $\psi_k^\theta$  that bundles all the information of the inequality constraint parameter  $P$ . We additionally introduce a time vector  $T = (0, t_d, 2t_d, \dots, (N-1)t_d)^T \in \mathbb{R}^N$  to give the policy the information, for which time step it predicts a control, where  $t_d$  denotes the discretization time of the MPC planner. Note, that optimizing (7) requires backpropagating through the dynamics function  $f$ .

We derive two neural network architectures, see figure 1. The `PlanNetX` architecture does not use an additional parameter network  $\psi_\theta$ . Instead, we directly give the parameter  $p_k$  and time  $t_k$  of stage  $k$  as an input. Note, that in some scenarios this is not enough to predict the control  $\hat{u}_k$  correctly, and the information of the full parameter matrix  $P$  consisting out of  $p_1, \dots, p_{N-1}$  is required. Thus, the second architecture `PlanNetXEnc` adds a Transformer encoder as described in Vaswani et al. (2017) to compress the whole parameter matrix  $P$  and the time vector  $T$  into a latent vector  $Z \in \mathbb{R}^{n_z \times N}$ , providing the policy  $\pi_\theta$  with the required information. For the Transformer each parameter and time tuple  $p_k$  can be considered as one token, which is first embedded with a linear layer like in Dosovitskiy et al. (2021) and then processed by the self-attention layer. The



Figure 2: Exemplary cut-in scenario of the CommonRoad benchmark created from recordings of the HighD dataset. The green car is the ego car, the red car is performing a cut-in in front of the ego car and the yellow region is the goal position the ego car needs to reach. The surrounding traffic is depicted in blue.

time information is added by a learned positional encoding, as often used in Transformer models [Devlin et al. \(2019\)](#). We further normalize the input features by min-max normalization where the minimum and maximum values are derived from the values that would be obtained when solving the OCP numerically.

## 4. Experiments

In the following, we present the empirical results of this work. After introducing the longitudinal MPC Planner, we continue with our benchmarking setup, and then present and discuss the results in terms of open-loop planning performance, closed-loop performance, and worst-case inference time.

### 4.1. Longitudinal Planner

The following MPC formulation is adopted from [Gutjahr et al. \(2017\)](#) and [Pek and Althoff \(2021\)](#). We define the longitudinal state of the ego car with  $x = [s, v, a, j]^T \in \mathbb{R}^4$ , where  $s$  is the longitudinal position of the front of the car,  $v$  is the velocity,  $a$  is the acceleration, and  $j$  is the jerk. The input of the system is the snap given by  $u(t) = \ddot{a}(t) \in \mathbb{R}$  which is the derivative of the jerk resulting in the linear time-invariant system:

$$\frac{d^4}{dt^4} s(t) = u(t). \quad (8)$$

Integrating the dynamics with a time discretization  $t_d$  and assuming piecewise constant inputs we get the discrete-time dynamics system

$$x_{k+1} = f(x_k, u_k) = A_d x_k + B_d u_k. \quad (9)$$

The cost function  $L$  is a mixture of quadratic costs penalizing high accelerations, jerk, and controls with weights  $w_a, w_j, w_u > 0$  for comfort and encourages the progress made in the longitudinal direction with a linear weight  $w_s > 0$ , resulting in the following loss

$$L_k(x, u) = \gamma^k (w_a a^2 + w_j j^2 + w_u u^2 - w_s s). \quad (10)$$

An additional discount factor  $\gamma$ , can be used to favor cost minimization of short-term over long-term costs.

To guarantee the feasibility of the planned trajectory, we assume the following box constraints for the velocity  $v_{\min} \leq v_k \leq v_{\max}$ , acceleration  $a_{\min} \leq a_k \leq a_{\max}$  and jerk  $j_{\min} \leq j_k \leq j_{\max}$  for  $k = 0, \dots, N$ . We further added a terminal acceleration box constraint,  $0 \leq a_N \leq 0$ , such that the

final acceleration is zero. We describe the rear position of the *lead vehicle* that is in front of us by  $\bar{x}_0^{\text{lead}} = [s^{\text{lead}}, v^{\text{lead}}, a^{\text{lead}}]^T$  and use a simple forward prediction  $\hat{X}^{\text{lead}} \in \mathbb{R}^{n_{x,\text{lead}}}$  that assumes that the lead car keeps its acceleration for  $t_{\text{acc}}$  seconds and then has an acceleration of zero. With the forward prediction of the lead vehicle  $\hat{X}^{\text{lead}}$ , a minimum distance  $d_{\text{min}}$  and a reaction time  $t_{\text{brake}}$ , the safety distance to the leading vehicle is assured by the constraint

$$h_k^{\text{dist}}(x_k, \hat{x}_k^{\text{lead}}) := \max \left( \frac{v_k^2 - (\hat{v}_k^{\text{lead}})^2}{2|a_{\text{max}}|} + v_k t_{\text{brake}}, d_{\text{min}} \right) - (\hat{s}_k^{\text{lead}} - s_k). \quad (11)$$

Additionally, to respect speed limits we use the following constraint

$$h_k^v(x_k, v^{\text{max}}) := v_k - v^{\text{max}}(s_k), \quad \text{with } v^{\text{max}}(s_k) := \begin{cases} v_{\text{max}, 1} & \text{if } s_k < s_{\text{change}} \\ v_{\text{max}, 2} & \text{else} \end{cases} \quad (12)$$

where  $v^{\text{max}}$  is a position-based function encoding the current speed limit, with  $v_{\text{max}, 1}$  being the current speed limit,  $v_{\text{max}, 2}$  the next speed limit, and  $s_{\text{change}}$  the position where the speed limit changes. Note, that in the context of (1a) the parameter  $p_k$  is given by the forward prediction  $\hat{x}_k^{\text{lead}}$  and the constant speed limit parameters  $v_{\text{max}, 1}$ ,  $v_{\text{max}, 2}$ ,  $s_{\text{change}}$ . Additionally, we introduced a quadratic slack variable  $\zeta_{\text{dist}}$  to loosen the distance constraint of (11) with a weight  $w_{\zeta, \text{dist}} > 0$  and a slack variable  $\zeta_{a_N}$  for the terminal acceleration constraint with a weight  $w_{\zeta, a_N} > 0$ .

## 4.2. CommonRoad Benchmark

To test the closed-loop performance, we benchmarked our learned models with the composable benchmark simulator CommonRoad Althoff et al. (2017) on real and synthetic scenarios. The real scenarios were created from the HighD Krajewski et al. (2018) dataset, a collection of vehicle trajectories recorded on German highways, and then transformed into the CommonRoad format. As the real scenarios do not cover the whole space of potential driving maneuvers we additionally created synthetic scenarios: (1) Breaking, involving uniformly sampled lead cars with decreasing acceleration, (2) Synthetic speed limit changes, which include uniformly sampled lead cars with either decreasing or increasing speed limit changes and (3) Synthetic cut-ins, where cut-ins in front of the ego car are generated with velocities uniformly sampled. The trajectories of the lead car were generated via the intelligent driver model (IDM) Treiber et al. (2000). The average scenario duration is around 6.5 seconds and we tested on 1540 scenarios. In figure 3 we can further see that the predicted open-loop trajectory takes into account constraints such as the terminal acceleration constraint  $a_N = 0$ .

## 4.3. Experimental Setup

**Dataset** For solving the MPC formulation we used the NLP solver *acados* Verschueren et al. (2021) with *HPiPM* Frison and Diehl (2020) using a horizon length of  $N = 30$ . For creating training data, we uniformly sampled the initial state  $\bar{x}_0$  from the box constraints of the longitudinal planner and filtered out scenarios where a crash was inevitable. With a probability of 1/3, we sampled a speed limit change and with a probability of 1/3 we sampled that the lead car made a cut-in, leading to a different forward prediction. We created a training dataset  $\mathcal{D}$  with  $10^5$  samples, a validation dataset  $\mathcal{D}_{\text{valid}}$  with 33333, samples and a test set  $\mathcal{D}_{\text{test}}$  with 33333 samples.

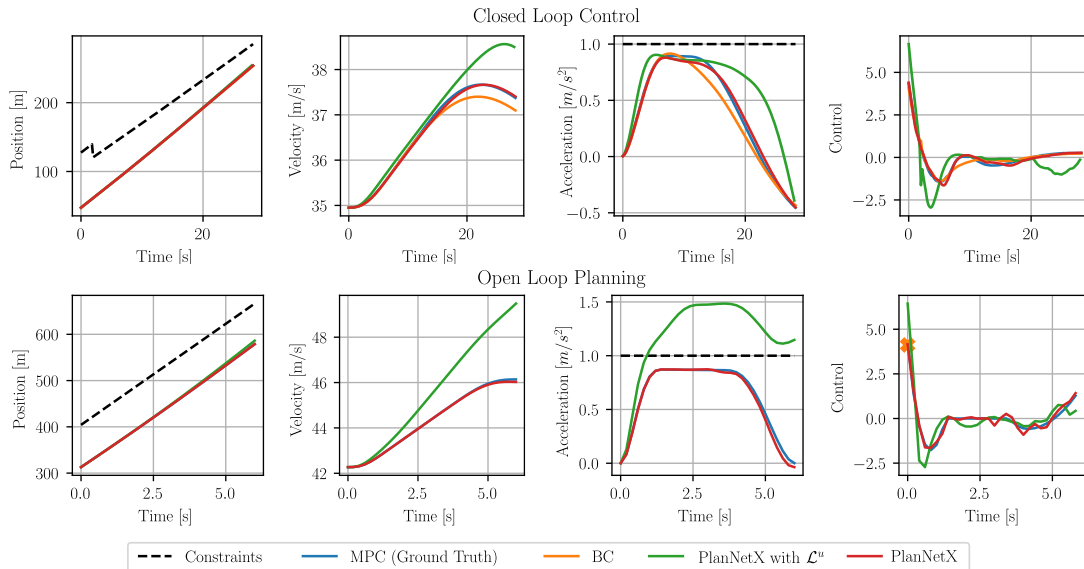


Figure 3: The closed-loop control plots show a trajectory generated in the CommonRoad simulator. The open-loop plan shows a predicted trajectory with sampled inputs from the test set  $\mathcal{D}_{\text{test}}$ . For the open-loop behavior of BC, we only report the first control, marked with an orange cross. The position constraints are either the forward prediction (open-loop) or the driven trajectory of the leading vehicle (closed-loop). In the closed-loop control example, the dashed line jumps because of a cut-in, thus a different car becomes the leading vehicle.

**Metrics** We use the following metrics averaged over three training seeds: 1) *Trajectory MSE*: Is the mean squared error between the predicted trajectory and the ground truth  $\hat{X}^\theta$  and  $X^{\text{MPC}}$  on the test set  $\mathcal{D}_{\text{test}}$ . 2) *Policy MSE*: Is the mean squared error between  $\hat{u}_0^\theta$  and  $u_0^{\text{MPC}}$  on the test set  $\mathcal{D}_{\text{test}}$ . 3) *Planning Time*: We randomly sampled 1000 inputs and estimate their inference times. As there is some additive CPU noise e.g. from background processes, we repeated the inference benchmark 20 times for each input, taking the minimum. Across all samples, we took the 95% quantile to estimate the worst-case runtime. 4) *Policy Time*: Similar to the planning time estimation, we estimated the inference time, but only for the first control. 5) *Avg.  $\|s\|_2$ , Avg.  $\|v\|_2$  and Avg.  $\|a\|_2$* : Are the average  $\ell^2$ -distances between the driven trajectories of the closed-loop MPC policy and the learned policy, in terms of the position  $s$ , velocity  $v$ , and acceleration  $a$ .

**Hyperparameters** For each neural network architecture and loss combination, we did a hyperparameter search using the Optuna library Akiba et al. (2019). The search space was defined by  $[1 \times 10^{-5}, 1 \times 10^{-2}]$  for the learning rate,  $\{1, 2, 3\}$  for the number of hidden layers in  $\pi_\theta$ ,  $\{32, 64, 128, 256, 512\}$  for the width of the hidden layers and  $\{8, 16, 32, 64\}$  for the batch size. For the policy part of the PlanNetX architectures and the baselines, we found that a width of 512 and a depth of 3 worked the best, while using ReLU activation functions. Thus, this was also the model



that we used for the pruning and quantization experiments. For the `PlanNetXEnc` experiments we used three self-attention Transformer layers with embedding dimension 32 and 4 attention heads. For the hyperparameter search, we used 100 training epochs, whereas for the full training, we used 300 epochs. In all experiments, we used the Adam optimizer [Kingma and Ba \(2015\)](#). For the weight matrix  $W$  in (5) and (6), we always used the identity matrix, as this gave us the best performance. For the discount factor  $\gamma$  we always used a value of 0.98. For measuring the inference time, we used an AMD Ryzen 9 5950X @ 3.40 GHz CPU.

**Neural Network Compression** To reduce inference time further, we considered pruning and quantization: Pruning removes the weights of the neural network that contribute little to the prediction. We used iterative  $\ell^1$ -structured pruning to prune the columns of the linear layers, [Li et al. \(2017\)](#), which removes the columns with the smallest  $\ell^1$ -norm of one linear layer. In our experiments, we set the number of parameters to be pruned to 90%. Quantization is a technique that involves reducing the precision of the parameters of the network, typically from 32-bit floating-point (FP32) to 16-bit floating point (FP16) or 8-bit integer (INT8) [van Baalen et al. \(2023\)](#). We used uniform quantization with INT8 precision, a common method with low overhead where the mapping consists of scaling, rounding and adding an offset. To estimate the clipping ranges and scaling factors, we employ post-training *static quantization* using the validation dataset  $\mathcal{D}_{\text{valid}}$  to collect statistics on the data before online inference [Gholami et al. \(2021\)](#).

#### 4.4. Results

**Open-loop performance** For the open-loop planning performance, we measured the *Policy MSE* and *Trajectory MSE* on the test dataset  $\mathcal{D}_{\text{test}}$ . From table 1, we can see that the proposed methods `PlanNetX` and `PlanNetXEnc` can approximate the MPC plan policy very well especially when compared to the `PlanNetX` network with the  $\mathcal{L}^u$  loss. Interestingly, the *Policy MSE* of the `PlanNetX` methods that use the state-trajectory

| Method                                     | Policy MSE | Trajectory MSE |
|--|------------|----------------|
| BC   | 0.66       | -              |
| <code>PlanNetX</code>                      | 107.68     | 0.01           |
| <code>PlanNetXEnc</code>                   | 96.72      | 0.02           |
| <code>PlanNetX</code> with $\mathcal{L}^u$ | 9.37       | 0.84           |

Table 1: Open-loop performance

loss is significantly higher than of BC, indicating that sometimes the `PlanNetX` deviates from the MPC plan to find its control sequence that leads to better tracking in state space. We noticed that this was especially true when the MPC planner showed bang-bang behavior for the control sequence.

**Closed-loop performance** When applying the approach `PlanNetX` in simulation for closed-loop control, we find that the driven trajectories are closer to the trajectories of the MPC policy than when just using the BC baseline, see figure 3 and table 1. The `PlanNetXEnc` achieves a similar performance then `PlanNetX`, whereas the control trajectory loss from (5) leads to significantly worse results, indicating that learning the full control trajectory  $U^*$  does not necessarily help for learning the first control  $u_0^*$ . For the inference time of the control policy, we achieve a significant speed-up over the MPC solver. Even though, for the planning time, the speed up is not as significant, we were able to reduce the planning time using structured  $\ell^1$ -pruning, without losing significantly in closed-loop performance. The quantization with INT8 precision resulted in faster inference, but at the cost of a severe degradation in accuracy. We suspect that during the roll-out of the `PlanNetX` architecture the rounding errors accumulate, leading to sub-optimal training results. We therefore further tested quantization by using FP16, which maintains approximately the same accuracy with less memory consumption. However, on a CPU we did not achieve any speedup.

| Method                        | Avg. $\ \Delta s\ _2$ [m] | Avg. $\ \Delta v\ _2$ [m/s] | Avg. $\ \Delta a\ _2$ [m/s <sup>2</sup> ] | Policy Time [ms] | Planning Time [ms] |
|-------------------------------|---------------------------|-----------------------------|---|------------------|--------------------|
| MPC                           | -                         | -                           | -   | 6.44             | 6.44               |
| MPC warm start                | -                         | -                           | -   | 5.08             | 5.08               |
| BC                            | 0.65                      | 0.17                        | 0.05                                      | 0.06             | -                  |
| PlanNetX                      | 0.24                      | 0.05                        | 0.04                                      | 0.10             | 2.62               |
| PlanNetXEnc                   | 0.29                      | 0.05                        | 0.03                                      | 0.31             | 2.79               |
| PlanNetX with $\mathcal{L}^u$ | 2.00                      | 0.56                        | 0.13                                      | 0.10             | 2.61               |
| PlanNetX quant.               | 1.19                      | 0.39                        | 0.19                                      | 0.07             | 1.94               |
| PlanNetX $\ell^1$ -str        | 0.27                      | 0.11                        | 0.06                                      | 0.06             | 1.83               |

Table 2: The closed-loop performance for the proposed PlanNetX and PlanNetXEnc approaches and some further ablations, where we tested the control trajectory loss  $\mathcal{L}^u$ , a pruned and INT8 quantized network as well as warm starting the MPC.

**When to use the encoder?** We further tested the capabilities of the encoder architecture to adapt to varying forward predictions of the lead car  $\hat{X}^{\text{lead}}$ . For this, we introduced a jump in the forward prediction of the lead car at a randomized time point, c.f. figure 4. In such a scenario, the PlanNetX can not infer the information of when the jump occurs from only the stage-wise parameter  $p_k$ , thus taking an average guess to predict  $\hat{x}_k$  and  $\hat{u}_k$ . As we can see in figure 4, by additionally using the encoder, the PlanNetXEnc is able to fit the MPC plan, as  $z_k$  encodes the combined information of the parameters of all stages. We defer a more detailed analysis of the generalization capabilities of the encoder for future research.




Figure 4: Trajectory predictions for varying forward pred.  $\hat{X}^{\text{lead}}$ .

## 5. Conclusion

We found that our proposed PlanNetX framework can imitate the trajectories of the longitudinal MPC planner to high accuracy while achieving a better closed-loop control in simulation than BC and having a significantly lower inference time than the MPC planner. Using compression techniques like pruning and quantization further reduced the inference time, however only with pruning we were able to obtain comparable performance to the uncompressed network. Further, directions could be to extend the method to MPC with a non-linear model where stability issues due to vanishing or exploding gradients might arise. Here, reducing the roll-out horizon, or using an MPC-based value function Ghezzi et al. (2023) might further stabilize the training. Other interesting directions would be to investigate the generalization capabilities of the PlanNetXEnc architecture for more MPC formulations with a potentially more complex parameterization of the MPC or to compare to other state-of-the-art IL methods like Diffusion models Janner et al. (2022).

## Acknowledgments

The authors thank Rudolf Reiter for providing helpful feedback throughout this work. This project has been supported by the European Union’s Horizon 2020 research and innovation programme, Marie Skłodowska-Curie grant agreement 953348, ELO-X, and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) with grant number 428605208.

## References

- Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A Next-generation Hyperparameter Optimization Framework. In Ankur Teredesai, Vipin Kumar, Ying Li, Rómer Rosales, Evimaria Terzi, and George Karypis, editors, *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4-8, 2019*, pages 2623–2631. ACM, 2019.
- Matthias Althoff, Markus Koschi, and Stefanie Manzingler. CommonRoad: Composable benchmarks for motion planning on roads. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 719–726, June 2017.
- A. Bemporad, F. Borrelli, and M. Morari. The explicit solution of constrained LP-Based receding horizon control. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, Sydney, Australia, 1999.
- Jan Carius, Farbod Farshidian, and Marco Hutter. MPC-Net: A First Principles Guided Policy Search. *IEEE Robotics and Automation Letters*, 5(2):2897–2904, April 2020.
- Ryan K. Cosner, Yisong Yue, and Aaron D. Ames. End-to-End Imitation Learning with Safety Guarantees using Control Barrier Functions. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 5316–5322, Cancun, Mexico, December 2022. IEEE.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *9th International Conference for Learning Representations*. International Conference for Learning Representations, 2021.
- Ján Drgoňa, Aaron Tuor, Elliott Skomski, Soumya Vasisht, and Draguna Vrabie. Deep Learning Explicit Differentiable Predictive Control Laws for Buildings. *IFAC-PapersOnLine*, 54(6):14–19, January 2021.

- Ján Drgoňa, Aaron Tuor, and Draguna Vrabie. Learning Constrained Parametric Differentiable Predictive Control Policies With Guarantees. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, pages 1–12, 2024.
- G. Frison and M. Diehl. HPIPM: A high-performance quadratic programming framework for model predictive control. In *Proceedings of the IFAC World Congress*, Berlin, Germany, July 2020.
- Andrea Ghezzi, Jasper Hoffman, Jonathan Frey, Joschka Boedecker, and Moritz Diehl. Imitation Learning from Nonlinear MPC via the Exact Q-Loss and its Gauss-Newton Approximation. In *Proceedings of the IEEE Conference on Decision and Control (CDC)*, 2023.
- Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W. Mahoney, and Kurt Keutzer. A Survey of Quantization Methods for Efficient Neural Network Inference. In *Low-Power Computer Vision*, pages 291–326. Chapman and Hall/CRC, June 2021.
- Benjamin Gutjahr, Lutz Gröll, and Moritz Werling. Lateral Vehicle Trajectory Optimization Using Constrained Linear Time-Varying MPC. *IEEE Transactions on Intelligent Transportation Systems*, 18(6):1586–1595, June 2017.
- Patrick Hart, Leonard Rychly, and Alois Knoll. Lane-Merging Using Policy-based Reinforcement Learning and Post-Optimization. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 3176–3181, Auckland, New Zealand, October 2019. IEEE.
- Michael Hertneck, Johannes Kohler, Sebastian Trimpe, and Frank Allgower. Learning an Approximate Model Predictive Controller With Guarantees. *IEEE Control Systems Letters*, 2(3):543–548, July 2018.
- Michael Janner, Yilun Du, Joshua Tenenbaum, and Sergey Levine. Planning with Diffusion for Flexible Behavior Synthesis. In *International Conference on Machine Learning*, pages 9902–9915. PMLR, 2022.
- T.A. Johansen and A. Grancharova. Approximate explicit constrained linear model predictive control via orthogonal search tree. *IEEE Trans. Automatic Control*, 48:810–815, 2003.
- Benjamin Karg, Teodoro Alamo, and Sergio Lucia. Probabilistic performance validation of deep learning-based robust NMPC controllers. *International Journal of Robust and Nonlinear Control*, 31(18):8855–8876, 2021.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *3rd International Conference for Learning Representations*, volume 3rd. International Conference for Learning Representations, 2015.
- B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Yogamani, and Patrick Pérez. Deep Reinforcement Learning for Autonomous Driving: A Survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):4909–4926, June 2022.
- Robert Krajewski, Julian Bock, Laurent Kloeker, and Lutz Eckstein. The highD Dataset: A Drone Dataset of Naturalistic Vehicle Trajectories on German Highways for Validation of Highly Automated Driving Systems. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 2118–2125, 2018.

- Sergey Levine and Vladlen Koltun. Guided Policy Search. In *Proceedings of the 30th International Conference on Machine Learning*, pages 1–9. PMLR, May 2013a.
- Sergey Levine and Vladlen Koltun. Variational Policy Search via Trajectory Optimization. In *Advances in Neural Information Processing Systems*, volume 26, 2013b.
- Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research*, 17(39):1–40, 2016.
- Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning Filters for Efficient ConvNets. In *5th International Conference on Learning Representations*. International Conference on Learning Representations, 2017.
- Weiwei Li and Emanuel Todorov. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *First International Conference on Informatics in Control, Automation and Robotics*, volume 2, pages 222–229. SciTePress, 2004.
- Shamil Mamedov, Rudolf Reiter, Moritz Diehl, and Jan Swevers. Safe Imitation Learning of Non-linear Model Predictive Control for Flexible Robots. *CoRR*, abs/2212.02941, 2022.
- Igor Mordatch and Emo Todorov. Combining the benefits of function approximation and trajectory optimization. In *Robotics: Science and Systems X*. Robotics: Science and Systems Foundation, July 2014.
- Christian Pek and Matthias Althoff. Fail-Safe Motion Planning for Online Verification of Autonomous Vehicles Using Convex Optimization. *IEEE Transactions on Robotics*, 37(3):798–814, June 2021.
- Alexander Reske, Jan Carius, Yuntao Ma, Farbod Farshidian, and Marco Hutter. Imitation Learning from MPC for Quadrupedal Multi-Gait Control. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5014–5020, May 2021.
- Roland Schwan, Colin N. Jones, and Daniel Kuhn. Stability Verification of Neural Network Controllers Using Mixed-Integer Programming. *IEEE Transactions on Automatic Control*, pages 1–16, 2023.
- Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested Traffic States in Empirical Observations and Microscopic Simulations. *Physical Review E*, 62(2):1805–1824, August 2000.
- Mart van Baalen, Andrey Kuzmin, Suparna S. Nair, Yuwei Ren, Eric Mahurin, Chirag Patel, Sundar Subramanian, Sanghyuk Lee, Markus Nagel, Joseph Soriaga, and Tijmen Blankevoort. FP8 versus INT8 for efficient deep learning inference. *CoRR*, abs/2303.17951, 2023.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Yannic Vaupel, Nils C. Hamacher, Adrian Caspari, Adel Mhamdi, Ioannis G. Kevrekidis, and Alexander Mitsos. Accelerating nonlinear model predictive control through machine learning. *Journal of Process Control*, 92:261–270, August 2020.

Robin Verschueren, Gianluca Frison, Dimitris Kouzoupis, Jonathan Frey, Niels van Duijkeren, Andrea Zanelli, Branimir Novoselnik, Thivaharan Albin, Rien Quirynen, and Moritz Diehl. Acados – a modular open-source framework for fast embedded optimal control. *Mathematical Programming Computation*, October 2021.

Kim Peter Wabersich and Melanie N. Zeilinger. A predictive safety filter for learning-based control of constrained nonlinear dynamical systems. *Automatica*, 129, July 2021.