
MPC-based Imitation Learning for Safe and Human-like Autonomous Driving

Flavia Sofia Acerbo^{1 2 3} Jan Swevers^{2 4} Tinne Tuytelaars³ Tong Duy Son¹

Abstract

To ensure user acceptance of autonomous vehicles (AVs), control systems are being developed to mimic human drivers from demonstrations of desired driving behaviors. Imitation learning (IL) algorithms serve this purpose, but struggle to provide safety guarantees on the resulting closed-loop system trajectories. On the other hand, Model Predictive Control (MPC) can handle non-linear systems with safety constraints, but realizing human-like driving with it requires extensive domain knowledge. This work suggests the use of a seamless combination of the two techniques to learn safe AV controllers from demonstrations of desired driving behaviors, by using MPC as a differentiable control layer within a hierarchical IL policy. With this strategy, IL is performed in closed-loop and end-to-end, through parameters in the MPC cost, model or constraints. Experimental results of this methodology are analyzed for the design of a lane keeping control system, learned via behavioral cloning from observations (BCO), given human demonstrations on a fixed-base driving simulator.

1. Introduction

Currently, the design of autonomous vehicles (AVs) is primarily focused on safety. However, to ensure a broad use by the general public, AVs should also be comfortable and enjoyable. Comfort in AVs is a complex and multifaceted entity, concerning attributes such as motion sickness, perceived safety, understandability, trust and many others (Elbanhawi et al., 2015). Hence, comfort optimization covers different areas of AVs research, such as personalized

¹Siemens Digital Industries Software, Leuven, Belgium
²KU Leuven, Mechanical Engineering Department, Leuven, Belgium
³KU Leuven, Electrical Engineering Department, Leuven, Belgium
⁴DMMS-M core-lab Flanders Make, Leuven, Belgium. Correspondence to: Flavia Sofia Acerbo <flavia.acerbo@siemens.com>.

driving controllers, novel cabin and seat design, active integrated chassis control or motion sickness minimization. In this work, we consider *human-like driving*, a solution to improve AVs comfort which entails an implicit personalization of AVs behavior inspired by human driving behaviors (Bellem et al., 2016; Kalabic et al., 2019). This can be realized in a model-based or in a data-driven approach; intensive research has been done on human driving models, with sensory, cognitive and muscular level of detail (Nash & Cole, 2019). Though accurate, these models are usually very task-specific and their identification for different tasks is expensive (Kolekar et al., 2020). Therefore, only a relatively small part of research on human-like AVs is implementing human driving models in the controllers, e.g. (Lazcano et al., 2021). For higher level of autonomy, research is primarily focusing on data-driven approaches, such as imitation learning (IL), which can obtain policies mimicking driving demonstrations experienced or labeled by the user as adequate. In this domain, various solutions from both academia and industry (Hawke et al., 2020) show great potential on learning policies end-to-end, i.e. from sensory data to control actions or states.

Imitation learning solutions for human-like driving usually involve deep learning (DL) in their architecture, and this limits the verifiability of the safety and stability properties of the closed-loop system. To provide guarantees, many works have proposed a hierarchical framework, where the DL components are combined with safety filters, frequently model-based controllers. In (Pan et al., 2017; Chen et al., 2019; Pulver et al., 2021; Acerbo et al., 2020) such DL components represent trajectories fed to tracking controllers, but learning is done without directly considering their effect, hence disrupting the end-to-end learning. In the majority of these works, model predictive control (MPC) is used, due to its ability to explicitly handle hard constraints. Moreover, MPC is a generic controller, which can be designed to have parameters in its cost function, model or constraints. In this regard, there have been recent advances on the use of Model Predictive Control (MPC), as a differentiable policy in the context of IL (Amos et al., 2018) and reinforcement learning (RL) (Gros & Zanon, 2020).

This work is a preliminary study to extend existing hierarchical structures of IL for AVs to differentiable MPC, which allows for end-to-end learning while guaranteeing safety

and stability of the controlled system. Specifically, MPC is used as a reactive control layer and its parameters are provided by other DL indirect controllers, which map from sensory data to the parameters. The paper is structured as follows: first, we provide an introduction on related works concerning IL for AVs and previous studies on MPC with IL and RL. Then, we present the generic formulation of the proposed approach, namely MPC-IL, which leverages the MPC differentiability and the hierarchical decomposition. After that, we detail the approach under the behavioral cloning from observations algorithm (MPC-BCO) and finally we provide preliminary experimental results of MPC-BCO applied to the design of a human-like lane keeping controller from human demonstrations.

2. Background and Related Work

2.1. Imitation Learning for Autonomous Driving

Let us consider a fully observable Markov Decision Process characterized by a set of continuous states $s \in \mathcal{S} \subset \mathbb{R}^{n_s}$, which can be controlled with a set of continuous actions $a \in \mathcal{A} \subset \mathbb{R}^{n_a}$ and with a state transition probability distribution $s_{t+1} \sim P(\cdot|s_t, a_t)$. An agent can control the system with a policy $\pi(a_t|s_t)$. We denote the *state occupancy measure* $\rho_\pi(s)$ induced by a policy π as the density of occurrence of states, while following that policy, over an infinite time horizon, discounted by the factor γ : $\rho_\pi^s = \sum_{t=0}^{\infty} \gamma^t P(s|\pi)$. Given an expert policy π_E to be imitated, IL tries to find a policy π minimizing the occupancy measure distance formulated as:

$$\min_{\pi} \mathbb{E}_{s \sim \rho_{\pi}^s} [\mathcal{L}(\pi_E, \pi)] = \mathcal{L}(\rho_{\pi_E}^s, \rho_{\pi}^s). \quad (1)$$

As demonstrated by (Ghasemipour et al., 2020), all IL algorithms can be read under this formulation. However, in the context of AVs, behavioral cloning (BC) is the algorithm with the most promising history of success (Bansal et al., 2018; Scheel et al., 2021). Standard BC provides a straightforward solution to the IL problem, by learning the policy in a supervised way through direct maximum likelihood, i.e. by minimizing the distance between the action distributions under the expert state occupancy measure, as: $\min_{\pi} \mathbb{E}_{s \sim \rho_{\pi_E}^s} [\mathcal{L}(\pi_E, \pi)]$. This suffers from the covariate shift problem, since the policy π generates a different occupancy measure when run in closed-loop, and then becomes unpredictable in its own induced distribution of states. Among different solutions to mitigate this problem, it has been shown by (Ghasemipour et al., 2020), that learning the marginal state distribution of the expert in addition to the action distribution, can heavily influence the results. Hence, behavioral cloning from observations (BCO) (Torabi et al., 2018) aims at minimizing the Kullback–Leibler (KL) divergence on the state occupancy measures as: $\min_{\pi} \text{KL}(\rho_{\pi_E}^s || \rho_{\pi}^s)$. Nevertheless, BCO is more

expensive than BC, since it requires to run simulations in the learning loop and needs solutions to estimate the policy gradient from the states.

2.2. Model Predictive Control and Reinforcement Learning

Model predictive control (MPC) is becoming a popular technique in AV research and industry. It is an optimal controller with receding horizon denoted in its most general form as:

$$\begin{aligned} a = \arg \min_{x, u} \quad & \sum_{k=0}^{N-1} l(\theta, x_k, u_k) + l_N(\theta, x_k, u_k) \quad (2) \\ \text{s. t.} \quad & x_0 = \hat{x} \\ & x_{k+1} = f(x_k, u_k, \theta) \\ & h(x_k, u_k, \theta) \leq 0 \\ & a = u_0. \end{aligned}$$

At each time step, the current state of the system is estimated as \hat{x} . Then, MPC minimizes a cost function across a prediction horizon N , where the evolution of the model is simulated with f and where its states and controls should not violate the constraints specified by h . Usually, when talking about learning-based MPC, we refer to system identification of the underlying model f , given data from the controlled system. However, it is possible to parametrize also its cost function and constraints. In this regard, MPC was recently proposed as a solution for safety and sample efficiency issues in IL and RL. (Gros & Zanon, 2020) showed that using a highly parametric MPC as a differentiable function approximator for RL results in safe and stable RL policies and allows MPC to optimize its closed-loop performance, even when its underlying model f is wrong and system identification is not performed. (Amos et al., 2018) have benchmarked the use of differentiable MPC for classical IL/RL problems showing the advantages in terms of sample efficiency and learning flexibility with respect to generic IL and system identification. However, these studies did not consider learning MPC parameters in the form of goals, references or other generic complex behavior to be imitated and therefore considered MPC as the entire policy, without combining it with other learning components.

3. MPC-based Imitation Learning

This section presents the formulation and generic algorithmic framework of our MPC-IL method. Here, MPC is used as the differentiable final control layer of a hierarchical policy. This means that MPC parameters are given by preceding indirect controllers, which are learnt with end-to-end imitation learning. Moreover, a specific IL algorithm, namely behavioral cloning from observations (BCO), is detailed as an example of the MPC-IL approach.

3.1. Problem Formulation

As defined in Section 2.1 we consider IL as the minimization of the occupancy measure distance, according to a certain metric \mathcal{L} . The system under study is controlled by the human with an (unknown) stochastic policy π_H as $a_t \sim \pi_H(\cdot|s_t)$, while the mimicking controller controls it with a deterministic policy parametrized by θ as $a_t = \pi_\theta(s_t)$. We now formulate this problem in the context where the controller part of the policy is represented by a MPC, optimizing its control action a according to the state s of the system and the parameters θ that are part of the objective function, model and constraints:

$$\begin{aligned} \min_{\theta} \quad & \mathbb{E}_{s \sim \rho_{\pi_\theta}^s} [\mathcal{L}(\pi_H, \pi_\theta)] = \mathcal{L}(\rho_{\pi_H}^s, \rho_{\pi_\theta}^s) \\ \text{with} \quad & \pi_\theta \in \arg \min_{x, u} \sum_{k=0}^{N-1} l(\theta, x_k, u_k) + l_N(\theta, x_k, u_k) \\ \text{s. t.} \quad & x_0 = s \\ & x_{k+1} = f(x_k, u_k, \theta) \\ & h(x_k, u_k, \theta) \leq 0 \\ & a = u_0. \end{aligned} \quad (3)$$

Equation (3) describes a *bilevel optimization problem*, where the upper-level problem is the imitation of policy π with respect to the human policy π_H and the lower-level problem in the MPC optimization, given a state s . Among gradient-based solution techniques it is common to replace the lower level optimization problems with their Karush-Kuhn-Tucker (KKT) conditions as constraints (Hatz et al., 2012). However, directly solving a constrained optimization problem is not trivial in the context of imitation and reinforcement learning. Moreover, it requires second order information (i.e. Hessian matrices) and that does not scale well when neural networks and large amount of data are involved. Conversely, the adopted solution is based on first-order information only (Amos et al., 2018). The core idea is to compute the gradient of the parametrized solution of the MPC with respect to the optimization variables of the upper-level imitation problem, and then solve it as an unconstrained one with gradient descent steps of size α as:

$$\begin{aligned} \theta &\leftarrow \theta - \alpha \left[\frac{\partial \mathcal{L}}{\partial \theta} \right]^T \\ \left[\frac{\partial \mathcal{L}}{\partial \theta} \right]^T &= \left[\frac{\partial \pi_\theta}{\partial \theta} \right]^T \left[\frac{\partial \mathcal{L}}{\partial \rho_{\pi_\theta}^s} \frac{\partial \rho_{\pi_\theta}^s}{\partial \pi_\theta} \right]^T. \end{aligned} \quad (4)$$

The computation of $\left[\frac{\partial \mathcal{L}}{\partial \theta} \right]^T$ requires:

- $\left[\frac{\partial \mathcal{L}}{\partial \rho_{\pi_\theta}^s} \frac{\partial \rho_{\pi_\theta}^s}{\partial \pi_\theta} \right]^T$, which depends on the adopted IL algorithm, defining the distance \mathcal{L} and how to approximate the occupancy measure ρ .

- $\left[\frac{\partial \pi_\theta}{\partial \theta} \right]^T$, i.e. the derivative of the MPC problem at its (locally) optimal solution. In the next section we describe how to efficiently compute it via automatic differentiation.

3.2. Backpropagation through the MPC Problem

We now detail how to compute the gradient of a policy π_θ with a MPC control layer. Let us introduce the Lagrangian associated to the MPC problem of Equation (3):

$$\begin{aligned} \mathcal{L}(x, u, \mu, \lambda, \theta) &= \sum_{k=0}^{N-1} l(\theta, x_k, u_k) + l_N(\theta, x_k, u_k) \\ &+ \mu^T h(x, u, \theta) + \lambda_0^T (x_0 - s) \\ &+ \sum_{k=0}^{N-1} \lambda_{k+1}^T [f(x_k, u_k, \theta) - x_{k+1}], \end{aligned} \quad (5)$$

where μ, λ are the multipliers associated to the dual solution. For simplicity of notation, we denote as $z = [x, u, \lambda, \mu]$ the vector of the primal and dual variables associated with the optimal solution, each of them with N components e.g., x_0, x_1, \dots, x_N . At a local optimum z^* and if the problem satisfies the linear independence constraint qualification (LICQ) and the second-order sufficient conditions (SOCS) (Gros & Zanone, 2020), then the following Karush–Kuhn–Tucker (KKT) conditions hold:

$$\nabla_{x, u} \mathcal{L}(z^*, \theta) = 0 \quad (6)$$

$$\begin{aligned} x_0^* &= s, \\ x_{k+1}^* &= f(x_k^*, u_k^*, \theta) \text{ for } k = 0, \dots, N-1, \end{aligned} \quad (7)$$

$$\begin{aligned} h(z^*, \theta) &\leq 0 \\ \mu^* &\geq 0 \end{aligned} \quad (8)$$

$$\mu^{*T} h(z^*, \theta) = 0. \quad (9)$$

If the active set of the inequalities constraints at z^* is known, then we can write the previous equations as an implicit function $F(z^*, \theta) \leftrightarrow z^* = \pi_\theta$ as:

$$F(z^*, \theta) = \begin{bmatrix} \nabla_{x, u} \mathcal{L}(z^*, \theta) \\ I_h h(z^*, \theta) - I_h^0 \mu^* \\ \mathbf{f}(z^*, \theta) \end{bmatrix} = 0, \quad (10)$$

where I_h and I_h^0 are diagonal matrices denoting active and inactive inequality constraints, respectively. \mathbf{f} represents the matrix form of the N equality constraints from Equation (7). According to the implicit function theorem, we can write the Jacobian of π_θ as:

$$\frac{\partial \pi_\theta}{\partial \theta} = - \frac{\partial F^{-1}}{\partial z} \frac{\partial F}{\partial \theta}. \quad (11)$$

The full $\frac{\partial \pi_\theta}{\partial \theta}$ may be very expensive to compute. However, we are interested in the Jacobian-times-vector product

$\left[\frac{\partial \pi_\theta}{\partial \theta}\right]^T \bar{z}$. This is also known as adjoint sensitivity and can be computed efficiently within automatic differentiation (AD) engines at the cost of a linear system solve and a reverse mode sweep (Andersson & Rawlings, 2018) as: $\frac{\partial \pi_\theta}{\partial \theta} \bar{z} = -\frac{\partial F}{\partial \theta}^T \frac{\partial F}{\partial z}^{-T} \bar{z}$.

Remark 3.1. In order to compute the gradient we must ensure that $\frac{\partial F}{\partial z}$ is invertible. This means that, alongside the conditions of LICQ and SOCS, it is required that the optimal solution z^* where we compute the gradient is at least locally unique. This is not guaranteed in the case of nonlinear MPC and therefore we should impose some sort of regularization on the primal variables.

3.3. Hierarchical Decomposition

In the available literature the parameters θ refer to physical properties in the model or to the weight matrices in a quadratic regulator cost, and are therefore static (Amos et al., 2018). Nevertheless, to learn more complex and strategic behaviors, as it is the case for end-to-end AV control from demonstrations, we suggest the need to combine MPC control with other differentiable layers, as graphically represented in Figure 1. Let us consider some context information about the system at time t in the form of a latent variable χ_t , inferrable from the system current state as $\chi_t = Q(s_t)$; for example, χ_t may be a high-dimensional input (e.g. camera images), from which it would be convenient to learn end-to-end. For this purpose, we suggest to learn a function $\theta = g(\chi_t)$ that maps the latent variable χ_t to the MPC parameters θ . For example, θ can represent the goal state of the system or a trajectory to follow, which changes depending on the current situation at time t ; We denote this approach as *hierarchical decomposition*; the MPC is responsible for the reactive control behavior, i.e. the mapping $s_t \rightarrow a_t$, while the preceding layers learn subtasks that indirectly influence the closed-loop behavior of the system. Such layers are said to be *indirect controllers*. Then, the policy is written as: $a_t = \pi_{\theta(\chi_t)}(s_t)$.

3.4. MPC-based Behavioral Cloning from Observations (MPC-BCO)

It is now detailed how we use the MPC-IL method to learn a deterministic hierarchical policy π_θ with a MPC control layer using closed-loop policy learning via behavioral cloning from observations (BCO) and a differentiable simulator.

Algorithm We consider rollout state trajectories of length T , induced by the policy π_θ : $\tau_{\pi_\theta} = s_0, s_1, \dots, s_T$, and state trajectories sampled from the human demonstrations: $\tau_{\pi_H} = s_0^*, s_1^*, \dots, s_T^*$. We frame the problem as the minimization of the L2 state error $L(s_t, s_t^*) = \|s_t - s_t^*\|_2^2$ between the human and policy on sequences of induced

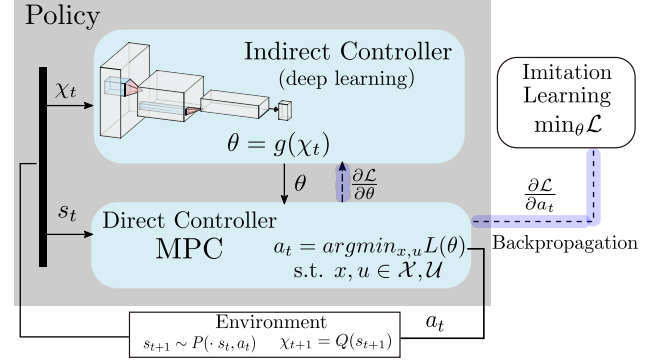


Figure 1. Hierarchical decomposition of the policy with a MPC control layer: the MPC provides the mapping $s_t \rightarrow a_t$, based on a constrained optimization. The indirect controller outputs the MPC parameters θ based on the latent variable χ_t . The dashed line show the flow of the gradient through the whole policy, possible thanks to the MPC differentiability.

states, when starting from the same initial state s_0^* . These sequences are acquired through a simulation model F , providing the state transitions $s_{t+1} = F(s_t, a_t)$. Let us assume that the first derivative of F around a certain s_t, a_t can be computed (this assumption is clarified in the last paragraph of this section). Therefore, we suggest computing $\nabla_\theta L(s_t)$ with backpropagation through time (BPTT) as:

$$\begin{aligned} \nabla_\theta L|_{s=s_t} &= \nabla_s L|_{s=s_t} \left(\nabla_a F \nabla_\theta a|_{a=a_{t-1}} + \right. \\ &\quad \left. + \nabla_a F \nabla_s a \nabla_\theta s|_{a=a_{t-1}, s=s_{t-1}} + \right. \\ &\quad \left. + \nabla_s F \nabla_\theta s|_{s=s_{t-1}} \right). \end{aligned} \quad (12)$$

Hence, with BPTT future state errors are accounted by the earlier actions taken by the policy.

We define $J(\theta) = \mathbb{E}_{s_t \sim \rho_{s_t}^\theta} \mathbb{E}_{s_t^* \sim \rho_{s_t^*}^H} \left[\sum_{t=0}^T L(s_t, s_t^*) \right]$ the sum of the L2 pose errors along a trajectory. Then, to update the policy, we compute the derivatives of J over (s_t, a_t, s_{t+1}) transitions with BPTT, leveraging on the differentiability of the policy and of the simulator. From (Heess et al., 2015), the gradient is computed by the following recursive formulas:

$$\begin{aligned} \frac{\partial J_t}{\partial s_t} &= \frac{\partial L}{\partial s_t} + \frac{\partial J_{t+1}}{\partial s_{t+1}} \left(\frac{\partial F}{\partial s_t} + \frac{\partial F}{\partial a_t} \frac{\partial \pi_\theta}{\partial s_t} \Big|_{\chi_t} \right) \\ \frac{\partial J_t}{\partial \theta} &= \frac{\partial J_{t+1}}{\partial s_{t+1}} \frac{\partial F}{\partial a_t} \frac{\partial \pi_\theta}{\partial \theta} \Big|_{\chi_t} + \frac{\partial J_{t+1}}{\partial \theta}. \end{aligned} \quad (13)$$

The final gradient $J(\theta)$ is computed by applying the equations from (13) recursively, starting from $t = T$ up to $t = t_s$. We suggest to not backpropagate until $t_s = 0$ to avoid bias; since the initial state s_0^* belongs to the human state occupancy metric, there may be a transitory part before the policy reaches a steady state corresponding to its actual occupancy

metric. The complete algorithm is described in Algorithm 1.

Remark 3.2. Note that in Equation (13) the loss and its gradients are computed only with respect to the state s and *not with respect to the latent variable* χ , which we provide to the indirect controllers of the policy. This is because we assume that χ is a nondifferentiable variable, coming from a black-box part of the simulation. As an example, we can consider χ as the result of an object detection. Since the policy actions cannot influence the object detection itself, it is useless to compute $\frac{\partial J_t}{\partial \chi_t}$. On the other hand, we may want to keep a specific distance to the specific detected object, which would be influenced by the policy actions and that would be included in the state s_t .

Algorithm 1 MPC-BCO

Input: Human sampled trajectories $D(\tau_{\pi_H}) = \tau_1^*, \tau_2^*, \dots, \tau_m^* \sim \rho_{\pi_H}$ where $\tau^* = s_0^*, s_1^*, \dots, s_T^*$
Initialize π_θ with θ_0
for $\tau^* \sim D(\tau_{\pi_H})$ **do**
 $s_0 = s_0^*, \chi_0 = Q(s_0^*)$
 for $t = 0$ **to** $T - 1$ **do**
 {Policy Simulation Rollout}
 $a_t = \pi_\theta(\chi_t)(s_t)$
 $s_{t+1} = F(s_t, a_t)$
 $\chi_{t+1} = Q(s_{t+1})$
 end for
 $\frac{\partial J_{T+1}}{\partial s_{T+1}} = \frac{\partial J_{T+1}}{\partial \theta} = 0$
 for $t = T$ **down to** t_s **do**
 {BPTT}
 $\frac{\partial J_t}{\partial s_t} = \frac{\partial L}{\partial s_t} + \frac{\partial J_{t+1}}{\partial s_{t+1}} \left(\frac{\partial F}{\partial s_t} + \frac{\partial F}{\partial a_t} \frac{\partial \pi_\theta}{\partial s_t} \right)$
 $\frac{\partial J_t}{\partial \theta} = \frac{\partial J_{t+1}}{\partial s_{t+1}} \frac{\partial F}{\partial a_t} \frac{\partial \pi_\theta}{\partial \theta} + \frac{\partial J_{t+1}}{\partial \theta}$
 end for
 $\theta \leftarrow \theta - \alpha \left[\frac{\partial J_{t_s}}{\partial \theta} \right]^T$ {Gradient Update}
end for

Backpropagation through the simulation model The MPC-BCO algorithm hinges on the first-order differentiability of the simulation model $s_{t+1} = F(s_t, a_t)$. Here, we describe two possibilities that enable this: differentiating through the simulator itself and differentiating through the MPC model. Recently there have been interesting developments towards the implementation of differentiable simulators for autonomous driving, i.e. extensions of existing numerical solvers with the functionality to compute gradients with respect to their inputs. These simulators should consist of both differentiable traffic scenario and ego vehicle dynamics. As to the former, examples can be found in (Suo et al., 2021) and (Scheel et al., 2021). As to the latter, high-fidelity vehicle dynamics models can be used without their explicit knowledge, as the only thing that’s required is their evaluation for a given set of states and their

respective Jacobians, which both are provided for example by the Functional Mock-up Interface (FMI) standard version 2.0 descriptions (Blockwitz et al., 2012).

If a differentiable simulator is not available, it is possible to exploit again the MPC structure. Indeed, given that the MPC relies on a good enough model, it is possible to simulate forward and backward by mapping $s, a \rightarrow x, u \Rightarrow x_{k+1} = f(x_k, u_k, \theta) \rightarrow s_{t+1} = f(s_t, u_t, \theta)$ and therefore compute a good approximation of the system gradients through f . However, as it can be seen from the previous formulation, the model can depend on the parameters θ which are not guaranteed to converge to the actual system parameter during learning, since we are optimizing only a closed-loop metric and not performing system identification. The possibility to combine system identification with imitation learning, similar to what is presented by (Martinsen et al., 2020), is beyond the scope of this paper and will be object of future research.

4. MPC-BCO for Lane Keeping Control

This section presents experimental results of the algorithm described in Section 3.4. The considered application is the design of a lane keeping control system, whose characteristics are learned from human demonstrations, collected on a fixed-base driving simulator. The MPC is designed as a path follower in Frenet coordinates and is considered as the control layer of a policy learned with Behavioral Cloning from Observations (BCO), in closed-loop. The following results are presented: learning of static parameters in the MPC stage cost and learning of an indirect controller for the MPC terminal cost. The latter shows capability in learning complex behavior and in adapting its output to the MPC dynamics to obtain the desired closed-loop trajectories.

4.1. Human Curve Driving

Although lane keeping may seem like a trivial automated control function to realize, it is quite hard to understand and realize it in a human-like way. Extensive research is available on human steering behaviors in curves. (Godthelp et al., 1984) showed that during curve driving a measure related to a visual feature, namely the time to lane crossing (TLC), is always kept above a minimum threshold, characterizing different driving styles. The TLC is defined as: $TLC = \frac{D}{V}$, where V is the longitudinal speed of the vehicle and the distance to lane crossing D is defined as the distance where the vehicle would exit the lane if no actions were taken from the current state. To influence the TLC the driver can modify its speed V and/or its distance to lane crossing D , which is determined by the lateral centerline deviation d and heading error θ , defined in Figure 2. As shown by (Barendswaard et al., 2019), the desired values of these variables change according to different phases of curve driving




Figure 2. Bicycle Model in Frenet coordinates.

(estimation, anticipation, entry and exit) and the transitory behaviour before curve entry is often overlooked by steering controllers. These considerations motivate the use of the MPC-BCO algorithm to learn complex behaviors in lane keeping controllers from human demonstrations.

4.2. Human Demonstrations

Human demonstrations of lane keeping control are collected using a fixed-base driving simulator. Firstly, the track and the visual feedback presented to the driver are designed with Simcenter Prescan. The track is 1200m long and made of 7 curved roads with different lengths and curvatures. Secondly, the ego vehicle is modelled via Simcenter Amesim as a 15DOF high-fidelity model of a Ford Focus. Finally, the two softwares are directly interfaced and can be run simultaneously through a C++, Python or Simulink interface, to create the whole simulation environment. During the simulation, the driver controls the ego vehicle via a Logitech G27 steering wheel with force feedback, while the longitudinal speed is kept constant at 50km/h. Therefore, the human can control its desired TLC only by the lateral distance and heading error of the vehicle with respect to the road. The following two main datasets are collected: 1. **D1**: 10 laps around the track, with a lane width $w = 4.5m$ and asking the driver to just keep the lane, and 2. **D2**: 10 laps around the track, with a lane width $w = 8m$ and asking the driver to keep the lane while travelling on the inner radius of the curves.

4.3. Vehicle Modelling

The MPC requires a model of the system to be controlled. In our scenario, this model includes the vehicle dynamics and kinematics and the road itself. To do this, we use Frenet coordinates to describe the vehicle kinematic with respect to the centerline of the road. The coordinate system is defined with the following variables, shown in Figure 2: 1. the arc length s , representing the travelled distance along the road, 2. the centerline deviation d , representing the lateral signed

position on the road with respect to the centerline, according to the $\mathbf{n} - \mathbf{t}$ frame and 3. the heading error θ , representing the difference between the yaw angle of the vehicle and the heading of the road. The kinematics of the vehicle evolve according to the changing curvature of the road $\kappa(s)$, that we assume as known. Moreover, we assume a constant longitudinal velocity $v_x = 50km/h$. The state equations for the kinematics are (Qian et al., 2016):

$$\begin{aligned} \dot{s} &= \frac{v_x \cos \theta - v_y \sin \theta}{1 - \kappa(s)d} \\ \dot{d} &= v_x \sin \theta - v_y \cos \theta \\ \dot{\theta} &= \dot{\psi} - \kappa(s) \frac{v_x \cos \theta - v_y \sin \theta}{1 - \kappa(s)d} \end{aligned} \quad (14)$$

As to the dynamics, we employ a linearized bicycle model, with lateral velocity v_y and yaw rate $\dot{\psi}$ as states (Milliken et al., 1995). The steering angle at the wheel is denoted as δ .

4.4. MPC Formulation

Based on the model chosen in the previous section, we can define the states and controls of the MPC as: $x = (v_y, \dot{\psi}, s, d, \theta, \delta)$, $u = \dot{\delta}$. The control input has been chosen to be the steering rate $\dot{\delta}$ as this is known to improve the smoothness of the control action and allows us to directly impose cost and constraints on its value. The MPC is formulated as in the following equations:

$$\min_{x,u} \sum_{k=0}^{N-1} l(\theta, x_k, u_k) + l_N(\theta, x_k, u_k) \quad (15)$$

$$\text{s. t. } x_{k+1} = f(x_k, u_k) \quad (16)$$

$$-w/2 \leq d \leq w/2 \quad (17)$$

where: eq. (16) is the Runge-Kutta discretization (sample time of 0.1s) of the vehicle model presented in Section 4.3; eq. (17) represents the safety constraints related to the lane boundaries, whose width is w ; eq. (15) is the objective function whose parameters are learned with BCO and whose structure is defined for each experiment in Section 4.5. The other states and controls are also box-constrained according to physical limits. The control horizon is chosen such that the controller has a lookahead distance of 30m, therefore $N = 20$. The MPC is implemented using Rockit (Gillis et al., 2020), an interface to CasADi (Andersson et al., 2018), aimed at optimal control problems definition. The problem is then translated into a PyTorch module, by extending its automatic differentiation engine to the one of CasADi. In this way, it is possible to seamlessly use a MPC layer within the PyTorch framework, though there is still the need of moving the tensors onto the CPU. A more efficient implementation of the MPC problem supporting GPU operations may be object of future research.

4.5. Results

This section shows results of MPC-BCO on the demonstration data D1 and D2 described in Section 4.2. Both datasets are made of 10 laps, denoted as Lap $\{i\}$. The laps are divided into different trajectories of length T and MPC-BCO is performed with stochastic gradient descent (with Adam algorithm) on batches of 10 trajectories each. The imitation loss L is chosen as the L2 error between the observed and demonstrated lateral centerline deviation, i.e. d and d^* respectively; hence, $L = \|d - d^*\|_2^2$. The rollout trajectories are of length $T = 10s$ and $t_s = 5s$ is used for BPTT.

Firstly, we learn static parameters in the cost function using the D1 dataset. The designed cost function of the MPC for D1 is $l_{D1} = \sum_{k=0}^{N-1} W_d(d_k - \bar{d})^2 + W_\theta \theta_k^2 + W_\delta \delta_k^2$, so only a stage cost and no terminal cost involved. The learnable parameters are $(W_d, W_\theta, W_\delta, \bar{d})$. To ensure that the cost function is plausible, we project the parameters into a different space. Specifically, the weights $(W_d, W_\theta, W_\delta)$ are mapped onto a softplus space $\log(1 + e^z)$, which guarantees non-negativity, and the offset \bar{d} is mapped onto a $\tanh(\bar{d})$ space, multiplied by the limits of the lane boundaries $w/2$. The initial parameters are $(1, 1, 1, 0)$. The numerical results in Table 1 show a consistent behavior of the learned parameters when learned of different laps; the weight on the lateral centerline deviation is decreased, while the ones on the control amplitude and the heading error are increased. The parameter that influences the most the closed-loop behavior is \bar{d} , determining the average preferred lateral centerline offset on the road. Visually, \bar{d} influences the driver perspective on the road during the constant curvature traits and therefore also the distance to lane crossing D . In Figure 3, we show a complete lap of the track, performed by the human on Lap5 (Demonstration), by the MPC with initial parameters (Before Imitation) and by the MPC with parameters learned on Lap5-6. In the lap performed by the latter, d is shifted towards negative values, i.e. to the right side of the lane, as the demonstrations show this driving preference. In Table 2, we show some relevant validation metrics compared with a BCO baseline, similarly to (Scheel et al., 2021), where π is chosen as a feedforward policy network with two hidden layers of 32 units, as in (Fu et al., 2017). The metrics reflect the main attributes of good human-like autonomous driving policy; first, the imitation performance, computed through L , then the safety performance, computed by the number of constraints violations and finally the comfort performance, computed as the integral of the squared lateral acceleration. Their values are computed on and averaged over 5 validation sub-trajectories taken randomly from all the laps. MPC-BCO outperforms BCO on all these attributes; by using the same amount of demonstrations, MPC-BCO can exploit its model and cost to learn with more accuracy than a typical network. Additionally, thanks to the explicit constraints, MPC-BCO is able to avoid any violation of the

lane boundaries, while BCO still violates them quite often. Finally, MPC-BCO improves the comfort of the trajectories; thanks to its cost on the control amplitude, it provides a smoother behavior, with a comfort metric value much closer to the human one, with respect to pure learning-based BCO.

Laps (Sub-datasets)	W_d	W_θ	W_δ	\bar{d}
Lap1-2	0.887	1.105	1.125	-0.191
Lap3-4	0.839	1.181	1.181	-0.381
Lap5-6	0.875	1.127	1.121	-0.412
Lap6-8	0.829	1.209	1.196	-0.444

Table 1. Learning static parameters for l_{D1} : learned values for different subsets of 2 laps each.

Metric	BCO	MPC-BCO	Human
Imitation	3.0566m	0.5832m	-
Safety	90	0	0
Comfort	758.10	67.58	59.49

Table 2. Validation results of MPC-BCO compared to BCO. The imitation performance is computed as $L = \|d - d^*\|_2^2$, safety by the number of times $|d| > w/2$ and comfort by $\int a_y^2$.

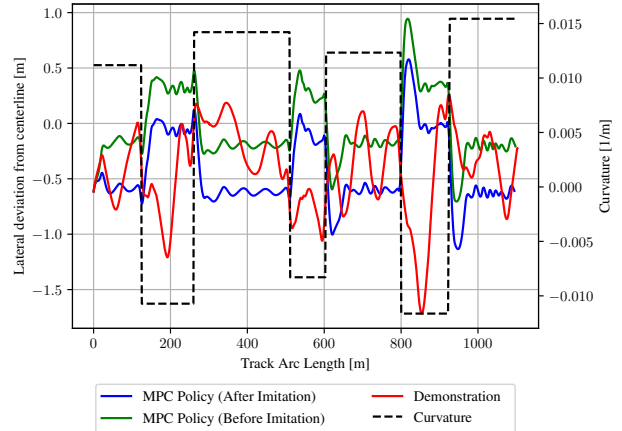


Figure 3. Learning parameters for l_{D1} : demonstrated and learned d on Lap5.

Secondly, we use demonstrations from D2, which provide a more complex behavior to imitate, depending on the road curvature ahead κ . Following the findings from the previous experiment, we learn only \bar{d} , and set the weights W_θ, W_δ equal to 1. Moreover, instead of providing a planned reference for all discretization steps, we consider

\bar{d} in the terminal cost as the lateral deviation goal to reach 2s ahead. Hence, we consider the following cost function: $l_{D2} = \sum_{k=0}^{N-1} \theta_k^2 + \delta_k^2 + (d_N - \bar{d}(\kappa))^2$. The indirect controller $\bar{d}(\kappa)$ is chosen to be represented as a piecewise linear function, approximated as a multilayer perceptron (MLP) with two layers of 50 hidden neurons, with a rectifying linear unit (ReLU) in between. Its inputs are chosen to be seven road curvature values, starting from the current one up to 30m ahead, equally spaced of 5m between each other. As starting point, $\bar{d}(\kappa)$ is learned in open-loop. The laps are divided into trajectories of 2s each and supervised learning (SL) is done, according to a mean squared error, to match $\bar{d}(\kappa)$ to the final d in the 2s trajectory, given the κ inferred from the initial state of the trajectory. Then, as in the previous experiments, laps are again divided into 10s trajectories and MPC-BCO is run to refine the MLP parameters optimizing the closed-loop performance. In Figure 5, we show the convergence results of this learning, starting from the imitation loss that is obtained when using $\bar{d}(\kappa)$ learned with SL in the closed-loop system. From that point, MPC-BCO reduces the imitation loss by 50%, by taking into consideration the closed-loop effects. In Figure 4 we show a complete lap of the track, performed by the human on Lap5 (Demonstration), by the MPC with $\bar{d}(\kappa)$ learned with SL and by the MPC with $\bar{d}(\kappa)$ learned with MPC-BCO. It is interesting to point out that the more the learned reference violates the safety boundary, the more the imitation loss on the closed-loop trajectories improve, though staying safe. This hints that, for a better closed-loop balance between safety and performance, the indirect controller should not be constrained on safety. This also highlights the difference to other approaches that learn a safe planning module followed by an idealized tracking controller, whereas MPC-IL leverages on the end-to-end differentiability of the policy to directly optimize the closed-loop behavior through MPC parameters.

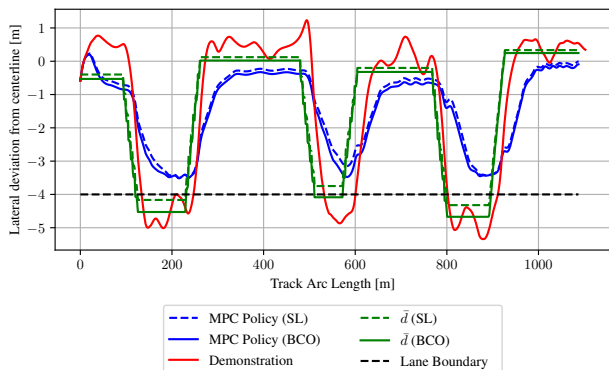


Figure 4. Learning an indirect controller for l_{D2} : demonstrated and learned d (both from SL and MPC-BCO) on Lap5.

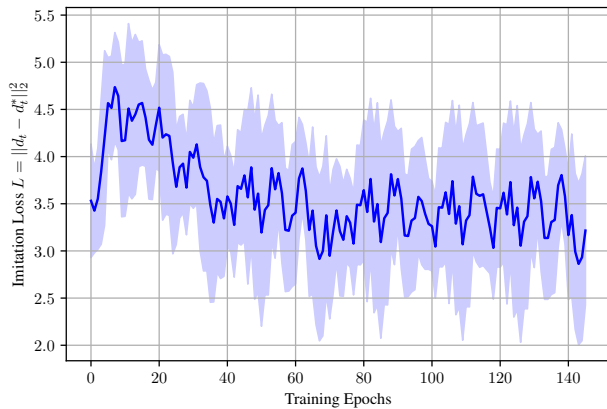


Figure 5. Convergence result of MPC-BCO. The initial loss is the imitation loss obtained when using $\bar{d}(\text{SL})$ into the closed-loop system.

5. Conclusions and Future Works

In this paper we have explored the use of MPC as a controller within an autonomous driving policy, learned via imitation learning. Firstly, we have outlined the methodology of MPC-IL consisting of a policy with a MPC reactive controller, preceded by indirect imitation learning-based controllers that based on the current latent information, output the optimal parameters in the MPC cost function, model or constraints. A novelty of this approach lies in the combination of both learning-based indirect controllers and the end-to-end differentiability of the framework, possible by leveraging on implicit differentiation. In this way, the safety of the MPC control and the flexibility of end-to-end learning can be combined in closed-loop imitation learning algorithms for autonomous driving. Preliminary experimental results have been shown for a lane keeping controller. This is an introductory work and it has its limitations; for example, the experimental results were obtained considering a simple, non-interactive and artificially constructed scenario where a predefined structure of the MPC cost function and the behavioral cloning loss were possible. Moreover, we see potential in using camera images as input to the indirect controllers, in order to compare its performance and safety with respect to popular end-to-end controllers for AVs, directly mapping images to control actions.

Acknowledgements

This work was supported by the Flemish Agency for Innovation and Entrepreneurship (VLAIO) in the context of MIMIC (huMan IMItation for autonomous driving Comfort) Baekeland Mandaat [Project nr. HBC.2020.2263].

References

- Acerbo, F. S., der Auweraer, H. V., and Son, T. D. Safe and computational efficient imitation learning for autonomous vehicle driving. In *2020 American Control Conference (ACC)*. IEEE, jul 2020. doi: 10.23919/acc45564.2020.9147256.
- Amos, B., Jimenez, I., Sacks, J., Boots, B., and Kolter, J. Z. Differentiable mpc for end-to-end planning and control. *Advances in neural information processing systems*, 31, 2018.
- Andersson, J. A. and Rawlings, J. B. Sensitivity analysis for nonlinear programming in CasADi. *IFAC-PapersOnLine*, 51(20):331–336, 2018. doi: 10.1016/j.ifacol.2018.11.055.
- Andersson, J. A. E., Gillis, J., Horn, G., Rawlings, J. B., and Diehl, M. CasADi: a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, jul 2018. doi: 10.1007/s12532-018-0139-4.
- Bansal, M., Krizhevsky, A., and Ogale, A. Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst. *arXiv preprint arXiv:1812.03079*, 2018.
- Barendswaard, S., van Breugel, L., Schelfaut, B., Sluijter, J., Zuiker, L., Pool, D. M., Boer, E. R., and Abbink, D. A. Effect of velocity and curve radius on driver steering behaviour before curve entry. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*. IEEE, oct 2019. doi: 10.1109/smc.2019.8914263.
- Bellem, H., Schöenberg, T., Krems, J. F., and Schrauf, M. Objective metrics of comfort: Developing a driving style for highly automated vehicles. *Transportation Research Part F: Traffic Psychology and Behaviour*, 41:45–54, aug 2016. doi: 10.1016/j.trf.2016.05.005.
- Blockwitz, T., Otter, M., Akesson, J., Arnold, M., Clauss, C., Elmqvist, H., Friedrich, M., Junghanns, A., Mauss, J., Neumerkel, D., Olsson, H., and Viel, A. Functional mockup interface 2.0: The standard for tool independent exchange of simulation models. In *Linköping Electronic Conference Proceedings*. Linköping University Electronic Press, nov 2012. doi: 10.3384/ecp12076173.
- Chen, J., Yuan, B., and Tomizuka, M. Deep imitation learning for autonomous driving in generic urban scenarios with enhanced safety. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, nov 2019. doi: 10.1109/iros40897.2019.8968225.
- Elbanhawi, M., Simic, M., and Jazar, R. In the passenger seat: Investigating ride comfort measures in autonomous cars. *IEEE Intelligent Transportation Systems Magazine*, 7(3):4–17, 2015. doi: 10.1109/imits.2015.2405571.
- Fu, J., Luo, K., and Levine, S. Learning robust rewards with adversarial inverse reinforcement learning. *arXiv preprint arXiv:1710.11248*, 2017.
- Ghasemipour, S. K. S., Zemel, R., and Gu, S. A divergence minimization perspective on imitation learning methods. In Kaelbling, L. P., Kragic, D., and Sugiura, K. (eds.), *Proceedings of the Conference on Robot Learning*, volume 100 of *Proceedings of Machine Learning Research*, pp. 1259–1277. PMLR, 30 Oct–01 Nov 2020.
- Gillis, J., Vandewal, B., Pipeleers, G., and Swevers, J. Effortless modeling of optimal control problems with rokit. In *39th Benelux Meeting on Systems and Control, Date: 2020/03/10-2020/03/12, Location: Elspeet, The Netherlands*, 2020.
- Godthelp, H., Milgram, P., and Blaauw, G. J. The development of a time-related measure to describe driving strategy. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 26(3):257–268, jun 1984. doi: 10.1177/001872088402600302.
- Gros, S. and Zanon, M. Data-driven economic NMPC using reinforcement learning. *IEEE Transactions on Automatic Control*, 65(2):636–648, feb 2020. doi: 10.1109/tac.2019.2913768.
- Hatz, K., Schlöder, J. P., and Bock, H. G. Estimating parameters in optimal control problems. *SIAM Journal on Scientific Computing*, 34(3):A1707–A1728, jan 2012. doi: 10.1137/110823390.
- Hawke, J., Shen, R., Gurau, C., Sharma, S., Reda, D., Nikolov, N., Mazur, P., Mickelthwaite, S., Griffiths, N., Shah, A., and Kndall, A. Urban driving with conditional imitation learning. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, may 2020. doi: 10.1109/icra40945.2020.9197408.
- Heess, N., Wayne, G., Silver, D., Lillicrap, T., Erez, T., and Tassa, Y. Learning continuous control policies by stochastic value gradients. *Advances in neural information processing systems*, 28, 2015.
- Kalabic, U., Chakrabarty, A., Quirynen, R., and Cairano, S. D. Learning autonomous vehicle passengers’ preferred driving styles using g-g plots and haptic feedback. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, oct 2019. doi: 10.1109/itsc.2019.8917328.
- Kolekar, S., de Winter, J., and Abbink, D. Human-like driving behaviour emerges from a risk-based driver model. *Nature Communications*, 11(1), sep 2020. doi: 10.1038/s41467-020-18353-4.

- Lazcano, A. M., Niu, T., Akutain, X. C., Cole, D., and Shyrokau, B. MPC-based haptic shared steering system: A driver modeling approach for symbiotic driving. *IEEE/ASME Transactions on Mechatronics*, 26(3):1201–1211, jun 2021. doi: 10.1109/tmech.2021.3063902.
- Martinsen, A. B., Lekkas, A. M., and Gros, S. Combining system identification with reinforcement learning-based mpc. *IFAC-PapersOnLine*, 53(2):8130–8135, 2020.
- Milliken, W. F., Milliken, D. L., et al. *Race car vehicle dynamics*, volume 400. Society of Automotive Engineers Warrendale, PA, 1995.
- Nash, C. J. and Cole, D. J. Identification and validation of a driver steering control model incorporating human sensory dynamics. *Vehicle System Dynamics*, 58(4):495–517, mar 2019. doi: 10.1080/00423114.2019.1589536.
- Pan, Y., Cheng, C.-A., Saigol, K., Lee, K., Yan, X., Theodorou, E., and Boots, B. Agile autonomous driving using end-to-end deep imitation learning. *arXiv preprint arXiv:1709.07174*, 2017.
- Pulver, H., Eiras, F., Carozza, L., Hawasly, M., Albrecht, S. V., and Ramamoorthy, S. PILOT: Efficient planning by imitation learning and optimisation for safe autonomous driving. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, sep 2021. doi: 10.1109/iros51168.2021.9636862.
- Qian, X., de La Fortelle, A., and Moutarde, F. A hierarchical model predictive control framework for on-road formation control of autonomous vehicles. In *2016 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, jun 2016. doi: 10.1109/ivs.2016.7535413.
- Scheel, O., Bergamini, L., Wolczyk, M., Osinski, B., and Ondruska, P. Urban driver: Learning to drive from real-world demonstrations using policy gradients. In *Conference on Robot Learning (CoRL)*, 2021.
- Suo, S., Regalado, S., Casas, S., and Urtasun, R. Trafficsim: Learning to simulate realistic multi-agent behaviors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10400–10409, 2021.
- Torabi, F., Warnell, G., and Stone, P. Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*, 2018.