On Maximizing Lateral Clearance of an Autonomous Vehicle in Urban Environments

Francesco Seccamonte, Juraj Kabzan and Emilio Frazzoli

Abstract— We consider the problem of maximizing distance to road agents for a self-driving car. To this extent, we employ a Model Predictive Control (MPC) approach for the steering tracking control of an Autonomous Vehicle (AV). Specifically, we first present a traditional MPC controller, which is then extended to encode the clearance maximization goal by manipulating its cost function and constraints. We provide insights on the additional information needed to achieve such goal, and how this modifies the structure of the original controller. Furthermore, a connection between commonly used safety metrics and clearance to road users is established. We implement the MPC controller using two off-the-shelf numerical solvers, assessing its computational feasibility. Finally, we show experimental results of the proposed approach on public roads in Boston and in Singapore.

I. INTRODUCTION

The execution of actions in a robotics system is often decomposed into a Planning and Decision-making layer (in the following, simply the planner), and a Controller module. The former is responsible for choosing a maneuver to be performed and for imposing restrictions in speed, time, and space on how this maneuver shall be executed. The latter is responsible for the optimal execution of the maneuver within the bounds supplied by the planner, possibly taking into account the physical limitations of the dynamical system as well. Such decomposition is common in Autonomous Vehicles (AVs), where the actions taken must obey specified rules and also represent a safe and socially acceptable behavior. The decision-making process is obviously based on the output of perception and prediction modules, which provide information about the current and predicted state of the world. Sampling-based motion planning algorithms [1] have become extremely popular in the AV domain as a solution for the decision-making problem, in particular the ones able to handle non-holonomic constraints, such as RRT [2] and its asymptotically optimal variant RRT* [3]. In an urban environment, such algorithms take on great importance, as they allow to encode safety rules, e.g., in a minimum-violation fashion [4], [5]. On the downside, navigating dense urban environments requires a thorough refinement of the trajectory to be executed.

Fig. 1 shows a narrow road traveled by our AV, with a construction zone on the right side. In such scenario it is extremely helpful to finely pick a path farther from the construction area. Indeed, being too close too it and interacting with pedestrians potentially moving outside of the area unexpectedly, or even construction material falling off

The authors are with nuTonomy, an Aptiv company.
{francesco, juraj, emilio}@nutonomy.com.

Fig. 1. The nuTonomy R&D platform navigating a road in Singapore. The AV "sees" a construction zone (right), and deviates from the reference to keep a higher distance (black arrows on the left).

might endanger safety of the AV's passengers and of other road users.

Refining the driven path at the sampling-based planning level results in a large number of samples needed, dramatically increasing its computational complexity. While speeding up motion planning algorithms is an active area of research [6], another possibility is lowering their complexity by deferring to lower layers (such as the controller) the refinement of trajectories for achieving additional goals. The second approach is the one we adopt for maximizing lateral clearance, as optimization over a continuous domain can be easily attained via gradient-based methods, which are widely used in control applications.

In a nutshell, such architecture[1] considers the combinatorial aspect of the decision-making process (e.g., to overtake a car or to stop for it) tackled at the higher decision-making level, while the local refinement of the trajectory is handled in the controller module. The goal of this paper is to present how the latter is performed; details on the planner employed, or on the perception and prediction modules used, are outside the scope of this paper.

---

[1]Please note that the functionality described is not necessarily representative of current and future products by nuTonomy, Aptiv and their partners. The scenarios presented are simplified for the purposes of exposition. The examples discussed are illustrative of the philosophy but not the precise specification we use. Moreover, we gloss over the extensive verification and validation processes that are performed before driving on public roads.

A. Related work

Model Predictive Control (MPC) is nowadays the standard for model-based optimal control. It can be used for both regulation and tracking problems, and it allows to explicitly account for constraints on the system. Applications in the Autonomous Driving domain range from motion planning for miniature race cars [7], to the investigation of different models suitable for an AV driving task on full size vehicles [8], and more recently an MPC-based planner which includes also a collision avoidance system [9]. This last system, experimentally demonstrated also on full sized vehicles, does not lead the AV to choose a safer path (that is, a path with larger clearance to obstacles): Avoiding the obstacle by $0.1\,\mathrm{m}$ or $10\,\mathrm{m}$ is considered equivalent.

Inducing the choice of a safer path can be achieved by modifying a traditional MPC architecture to specify additional objectives; Such approach is a fairly recent trend for autonomous systems. The authors in [10], [11] employ this framework to maximize targets visibility in drone cinematography; authors in [12] similarly encode the additional goal of keeping a point of interest in sight for a UAV. The MPC frameworks therein implemented serves as both planner and controller, given that the drone can simply move from point A to point B without any specific rule to obey.

In [13], the authors extend an MPC-based planner to account for visibility maximization while overtaking a static obstacle in an urban environment. Rule compliance is taken into account by means of a state machine which selects the maneuver that the MPC shall deliver. The paper includes simulation results showing the effectiveness of the approach in a scenario including only a single parked car. Extensions to more sophisticated scenarios can become very challenging; also, it is unclear how many agents such approach can handle simultaneously.

B. Contribution

Compared to the aforementioned approaches, our work differs in the way responsibilities are assigned to our AV's modules: The planning layer needs to ensure rule compliance, for example according to [14], and provide the controller a reference trajectory to be followed, as well as a rule compliant region in the phase space (in the following, the "tube"). Such region empowers the controller with some freedom, as it is both collision-free and rule compliant. The proposed controller is able to track the reference trajectory, and at the same time to retain the capability of changing the exact driven path, that is, *how* the AV performs the maneuver. It is worth mentioning that the proposed scheme is compatible with any planning algorithm, as no assumptions are made on its architecture (e.g., state machine or graph based). At the controller level, comfort and tracking objectives are then considered jointly with safety related goals, that is, maximizing clearance to other agents. This combination of goals requires a trade-off between tracking objective and clearance maximization goal, resulting in what we refer to as *biasing* [15].

Thanks to this additional functionality, the controller does not only follow the planner's trajectory, but is also directly aware of perceived objects. An illustration of this behavior is depicted in Fig. 2.


Fig. 2. Qualitative paths during an avoidance maneuver (left-hand traffic): planner's reference (blue) and controller's biased (red). The "tube" consists of the shaded green region.

Similarly to related approaches in the literature, an MPC control framework is chosen mostly for its ability to handle constraints, which are a key component to enable biasing functionality.

The contribution of this paper is threefold: Firstly, the MPC control framework is presented (Section III), together with its clearance maximization extension (Section IV). Insights on how to allow the MPC layer to consume objects information, as well as a connection between clearance to such objects and commonly used safety metrics, are given in Section V. Secondly, the resulting optimization problem is implemented using two different off-the-shelf numerical solvers, and its computational tractability is assessed (Section VI). Finally, extensive experimental results are shown using the nuTonomy R&D platform, demonstrating the validity of the proposed approach in real-world scenarios.

II. MATHEMATICAL NOTATION

We denote vectors with bold letters, and the column vector given by $[\mathbf{x}^T, \mathbf{y}^T]^T$ as $(\mathbf{x}, \mathbf{y})$. The time derivative of the function $f(t)$ is denoted as $\dot{f}(t)$. The sampling time of the system is denoted as T_s. We use $x_k = x(kT_s)$ to represent the k-th prediction step of variable x, where $k = \{0, 1, \dots, N\}$.

III. MPC FORMULATION

In this section the design of the lateral tracking controller is presented.

A. Vehicle's model

The first component of an MPC algorithm is a model of the dynamical system to be controlled. Since the operational domain considers an AV driving in urban scenarios, the maximum achievable speed is low (around $14\,\mathrm{m\,s^{-1}}$ depending on the regulations); similarly, acceleration is low as well. As highlighted in [8], in low speed and acceleration regimes the usage of a higher fidelity model does not provide any substantial benefit over a kinematic one. Hence, we choose to model the AV with a kinematic bicycle model [16] [17]. Denoting by x, y, θ the position and heading of the AV at the center of the rear axle in a global inertial frame, by v the speed at the rear axle, by δ its ground steering angle, we express the heading angle dynamics as a function of

speed and curvature at the rear axle κ, using the relation $\kappa(t) = \tan(\delta(t))/b$, with b denoting the wheelbase. Given the physical limits of the system, we can directly recover the ground steering angle $\delta(t)$ as $\delta(t) = \arctan(\kappa(t)b)$. Moreover, we model the actuation dynamics as a first-order dynamical system with time constant τ (with τ being a parameter identified using standard system identification techniques [18]), hence adding the additional variable κ_{des} denoting the desired curvature. This results in the following system of differential equations:

$$\dot{x}(t) = v(t)\cos(\theta(t)),\tag{1a}$$

$$\dot{y}(t) = v(t)\sin(\theta(t)),\tag{1b}$$

$$\dot{\theta}(t) = v(t)\kappa(t),\tag{1c}$$

$$\dot{\kappa}(t) = \frac{1}{\tau}(\kappa_{\text{des}}(t) - \kappa(t)),\tag{1d}$$

$$\dot{\kappa}_{\text{des}}(t) = u(t).\tag{1e}$$

In order to be employed in an MPC scheme, the continuous time differential equations (1) need to be discretized, and the Explicit Runge-Kutta numerical scheme is employed. Stacking together the state variables, we write the state vector $\mathbf{x} = (x, y, \theta, \kappa, \kappa_{\text{des}})$. Eq. (1e) shows the input u being the desired curvature rate. The model is depicted in Fig. 3.


Fig. 3. Kinematic bicycle model of a car.

Remark 1: The proposed scheme is a steering tracking controller, hence the speed $v(t)$ in Eq. 1 is an uncontrollable parameter and not an optimization variable.

B. Cost function

The proposed control architecture aims at addressing a tracking problem, hence the objective function is written in terms of tracking error. By denoting the reference trajectory vector at stage k as $\mathbf{z}_k = (\bar{x}_k, \bar{y}_k, \bar{\theta}_k, \bar{\kappa}_k)$, the cost to be minimized is the sum of the tracking error at each prediction step $k = 0, \ldots, N$, with N being the prediction horizon length (in time steps), that is

$$J_{\text{nom}} = \sum_{k=0}^{N-1} (q(\mathbf{x}_k, \mathbf{z}_k) + r(u_k)) + p(\mathbf{x}_N, \mathbf{z}_N).\tag{2}$$

The quantity $\mathbf{z}_k, k = \{0, \ldots, N\}$ is obtained by discretizing the continuous planner's reference trajectory, using the

predicted speed v_k the AV will drive at the corresponding time instants (which comes from a different module).

Similarly to what is performed in contouring control [7], we compute the lateral error at each prediction step k as

$$e_{\text{lat},k} = -(x_k - \bar{x}_k)\sin(\bar{\theta}_k) + (y_k - \bar{y}_k)\cos(\bar{\theta}_k),$$

and penalize it in the cost function.

Hence, by employing a quadratic cost function, the term $q(\mathbf{x}_k, \mathbf{z}_k)$ corresponds to

$$q_1 e_{\text{lat},k}^2 + q_2(\theta_k - \bar{\theta}_k)^2 + q_3(\kappa - \bar{\kappa}_k)^2,$$

with $q_1, q_2, q_3 \geq 0$. The term $p(\mathbf{x}_N, \mathbf{z}_N)$ is structured analogously (with different weights only), and the stage input cost is $r(u_k) = Ru_k^2$, $R > 0$.

C. Constraints

The following constraints are imposed on the system:

$$\underline{u} \leq u_k \leq \overline{u} \qquad \forall\, k = \{0, \ldots, N-1\},\tag{3a}$$

$$\underline{\kappa} \leq \kappa_k \leq \overline{\kappa} \qquad \forall\, k = \{0, \ldots, N\},\tag{3b}$$

$$\underline{e_{\text{lat},k}} \leq e_{\text{lat},k} \leq \overline{e_{\text{lat},k}} \qquad \forall\, k = \{0, \ldots, N\}.\tag{3c}$$

Constraints (3a) and (3b) bound curvature rate and curvature taking into account physical limits; (3c) constrains the (signed) lateral error to lie within time-varying bounds. In order to ensure feasibility, the inequalities (3c) are expressed as soft constraints, through the addition of a slack variable $\varepsilon_k, \varepsilon_k \geq 0$, which is also penalized in the cost function by means of linear and quadratic penalties [19].

IV. MAXIMIZING LATERAL CLEARANCE WITHIN MPC

In this section, we extend the previously derived MPC controller by manipulating its cost function and constraints.

A. Safety related auxiliary variable

Since we aim at maximizing lateral clearance in appropriate circumstances (e.g., in the proximity of agents whose motion might change unpredictably), as well as minimizing the tracking error, a natural approach is to extend the cost function by adding suitable terms. The simple and intuitive observation is the following: The higher distance a vehicle has with respect to other objects, the lower the likelihood of incurring accidents, hence the higher safety. More formally speaking, incrementing clearance increases the commonly used Time-To-React (TTR) metric quantity [20].

To this extent, we introduce the auxiliary optimization variable s representing achievable safety, that is, an empirical measure of the attainable safety to be correlated to the clearance given to objects. For ease of exposition, the case with one object on the right of the reference trajectory is considered first (illustration in Fig. 4); the case of one object on the left is analogous, and how to handle multiple road agents is described in Section V-B.

We denote by $d_{r,\text{ref}}$ the lateral distance from the reference to the object on its right. Such distance accounts for the car's footprint, and is non-negative. The actual distance from the AV to the agent is denoted as d_r, and is computed as

$d_r = d_{r,\text{ref}} + e_{\text{lat}}$. The lateral error e_{lat} is a signed quantity. For this reason, the case of the object being on the left side results in $d_l = d_{l,\text{ref}} - e_{\text{lat}}$. Fig. 4 depicts such quantities.


Fig. 4. Both longitudinal $d_{r,\text{lon}}$ and lateral $d_{r,\text{ref}}$ distances are computed from the discretized planner's reference (axes origin) to the object O and account for the car's and object's footprints. Biasing (dashed green rectangle) increases the actual lateral distance by e_{lat}.

The cost function (2) is extended via

$$J_{\text{bias}} = \sum_{k=0}^{N} (s_k - s_{\text{target}})^2 \ , \tag{4}$$

with s_{target} being a constant representing the target safety.

Safety and lateral distances are linked through the clearance-dependent safety constraint function $f_s(d)$. Hence, the following constraint is added:

$$s \leq f_s(d_r) + s_{r,\text{lon}} . \tag{5}$$

The safety term $s_{r,\text{lon}}$ is related to the longitudinal distance $d_{r,\text{lon}}$ via the function $f_{\text{lon}}(d_{\text{lon}})$. Such function is non-increasing in $(-\infty, 0)$ and non-decreasing in $[0, \infty)$, with $f_{\text{lon}}(0) = 0$. Since $d_{r,\text{lon}}$ is an uncontrollable real time parameter, the term $s_{r,\text{lon}}$ is computed outside of the MPC controller.

The function $f_s(d)$ is a function of the lateral clearance d and also non-decreasing; additional details are given below.

B. Shaping the clearance-dependent safety function

Depending on how the clearance-dependent safety function is shaped, the extent to which the car deviates from the reference will be different. We can use functions taking values in the range $[0, s_{\text{target}}]$: by relating safety to Time-To-React (TTR), a negative TTR (hence safety) relates to unavoidable collision states only. However, thanks to the Planner outputting collision-free tubes only, we can neglect the cases where TTR < 0. Moreover, safety cannot grow unbounded, since after a certain threshold a higher TTR does not provide any additional benefit. We employ sigmoid-like functions, e.g., $f_s(d) = \frac{s_{\text{target}}}{1+e^{-a(d-b)}}$, for both the lateral and longitudinal terms. Figure 5 depicts the value of the cost function (4) (for one time step only) associated to the scenario drawn in Fig. 4.

C. Control law

Putting together the dynamics equations (1) (in the form of equality constraints), the cost functions (2) and (4) and the constraints (3) and (5) (whose resulting feasible sets are compactly denoted as $\tilde{\mathcal{X}}$ and $\mathcal{U}$), we extend the state variables by adding the safety term, that is, $\tilde{\mathbf{x}} = (\mathbf{x}, s)$. Note that s has no dynamics, hence the equality constraints deriving from


Fig. 5. Value of the cost function J_{bias} for the example in Fig. 4. The green shaded area defines the tube, while the blue cube represents the obstacle.

(1) are unchanged. Finally, denoting the sequence of inputs $\{u_0, \ldots, u_{N-1}\}$ at time steps $\{0, \ldots, N-1\}$ as $\mathbf{u}$ we can write the full optimization problem:

$$\min_{\mathbf{u}} \quad J_{\text{nom}} + \alpha J_{\text{bias}} \tag{6a}$$

$$\text{s.t.} \quad \mathbf{x}_0 = \mathbf{x}(t) , \tag{6b}$$

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, u_k), \quad \forall k = \{0, \ldots, N-1\}, \tag{6c}$$

$$\tilde{\mathbf{x}}_k \in \tilde{\mathcal{X}}_k, \quad \forall k = \{0, \ldots, N\}, \tag{6d}$$

$$u_k \in \mathcal{U}, \quad \forall k = \{0, \ldots, N-1\} . \tag{6e}$$

The factor $\alpha \geq 0$ is a tuning parameter used to trade off between tracking performance and clearance maximization. At each time step, problem (6) is initialized with the measured state $\mathbf{x}(t)$ as initial condition (constraint (6b)). The new reference trajectory $\mathbf{z}_k$ as well as the distances $d_{l,\text{ref},k}, d_{r,\text{ref},k}, d_{l,\text{lon},k}, d_{r,\text{lon},k}, \forall k = \{0, \ldots, N\}$, are also updated. An optimal solution of (6) is $J^\star$, with associated optimizer $\mathbf{u}^\star = \{u_0^\star, \ldots, u_{N-1}^\star\}$. Problem (6) is solved in a receding horizon fashion [19], and only the first input $u_0^\star$ is applied to the system.

Remark 2: Given that constraints (3c) are soft, problem (6) is always feasible. This comes from the fact that constraints (3b) are related to physical limits of the vehicle (hence always feasible provided a feasible initial condition (6b)), and constraints (3a) and (5) limit optimization variables with no dynamics (thus they can vary freely and instantaneously). Such property is extremely important for the real-world deployment of the MPC controller.

V. CONSTRAINTS COMPUTATION

In this section we present details on the lateral error bounds (i.e., the "tube") and the bias-inducing terms. The tracked objects information (measured and predicted for the future k steps) are obtained from separate perception and prediction layers, which do not belong to the scope of this paper.

A. Examples of tube computation

The role of constraints (3c) is of paramount importance: They are needed to specify the maximum lateral error the AV can tolerate due to biasing. For this reason, such "tube" has to represent a collision-free region around the reference trajectory, where the car can freely move without

endangering safety or violating specifications. To this extent, all the available informations are used for tube derivation: offline (i.e., map-based) and online (i.e., perception-based). Moreover, in some situations a narrower tube is advisable: this is the case, for instance, of the AV navigating an intersection, and is due to the fact that the lane boundaries are often not well defined, and other cars do not follow their prescribed paths but tend to "cut". In such situations the tube is shrunk progressively as the AV approaches the maneuver requiring superior tracking performance to avoid abrupt swerves that would greatly impact comfort, or even endanger safety.


Fig. 6. Planner's trajectory (blue) and collision-free tube (green) accounting for a static (a) and a moving (b) obstacles, as well as road boundaries.

An example of the tube is depicted as the green shaded area in Fig. 6.

B. Handling multiple objects

In urban driving, it is not uncommon to have multiple road agents that the AV needs to account for. To this extent we introduce the notion of *most constraining object* on the left and on the right, o_l and o_r, respectively: that is, at each prediction time step k the object that constrains the safety variable s_k the most, for a fixed $e_{\text{lat}} = 0$. Note that the most constraining object at time step $(k+1)$ is not necessarily the same as at time step k, given its predicted motion.

Such objects are found via

$$
\begin{aligned}
o_{l,k} &= \underset{i \in \{1,...,L\}}{\arg\min} \left(f_s(d_{l,\text{ref},k}^i) + s_{l,\text{lon},k}^i \right), \\
o_{r,k} &= \underset{h \in \{1,...,M\}}{\arg\min} \left(f_s(d_{r,\text{ref},k}^h) + s_{r,\text{lon},k}^h \right).
\end{aligned}
\tag{7}
$$

Consequently, the distances used in the MPC controller are

$$
\begin{aligned}
d_{l,\text{ref},k} &= d_{l,\text{ref},k}^{o_{l,k}}, \\
d_{r,\text{ref},k} &= d_{l,\text{ref},k}^{o_{r,k}}, \\
d_{l,\text{lon},k} &= d_{l,\text{lon},k}^{o_{l,k}}, \\
d_{r,\text{lon},k} &= d_{r,\text{lon},k}^{o_{r,k}}.
\end{aligned}
$$

The proposed approach is therefore able to handle as many objects as the perception/prediction layers can process.

C. Diversifying the behavior for different classes of objects

In case the perception system provides classified objects (e.g., car, bicycle, pedestrian), we may want to diversify the behavior depending on the object the AV deals with. To this extent and since the number of classes is finite, extending the proposed framework is straightforward by adding as many constraints (5) as the number of classes and use different functions $f_s(d)$ (or the same function with different parameters) for different classes.

VI. RESULTS

We show experimental results obtained using the nuTonomy R&D platform (Renault Zoe). The experiments assume left-hand traffic (Singapore/UK regulations).

A. Comparison with MPC without biasing


(a) No biasing   (b) Biasing

Fig. 7. Comparison of the normal MPC and the extended version for a scenario with two pedestrians (blue boxes). The continuous blue line represents the driven path.

In the extended version, the parameters $a, b \geq 0$ in the clearance-dependent safety function (5) are tuned in simulation and verified through real car testing. In case no road agent is present, the values $d_{\text{lat}}, d_{\text{lon}}$ are set to ∞ and 0, respectively. This causes the cost (4) to equal 0. Similarly, if object i results in $f_s^i(d_{\text{ref}}) + s_{\text{lon}}^i > s_{\text{target}}$ (that is, is located very far), it is not accounted in Eq. (7). The parameter α in Eq. (6a) was tuned by simulating different scenarios, and refined through real-world testing.

The prediction horizon employed is $N = 60$, corresponding to a time prediction of $3\,\text{s}$. MPC controllers with longer prediction horizons (i.e., $70-80$ time steps) were also tested, without particular benefits in terms of performance.

Fig. 7 shows a real-world experiment consisting of our car driving with two pedestrians on the sides of the lane. In the first picture 7a, the simple MPC controller is used, and the car only follows the reference path. The closest lateral distance from its footprint to the pedestrian is about $1\,\text{m}$. In the second picture 7b, the extended MPC controller is employed. Additionally to following the reference trajectory, the car also maximizes lateral clearance to both pedestrians, and results in driving an S-shaped path (the blue line in the figure). The closest lateral distance from its footprint to the pedestrian is about $1.4\,\text{m}$, that is $40\,\%$ larger than in
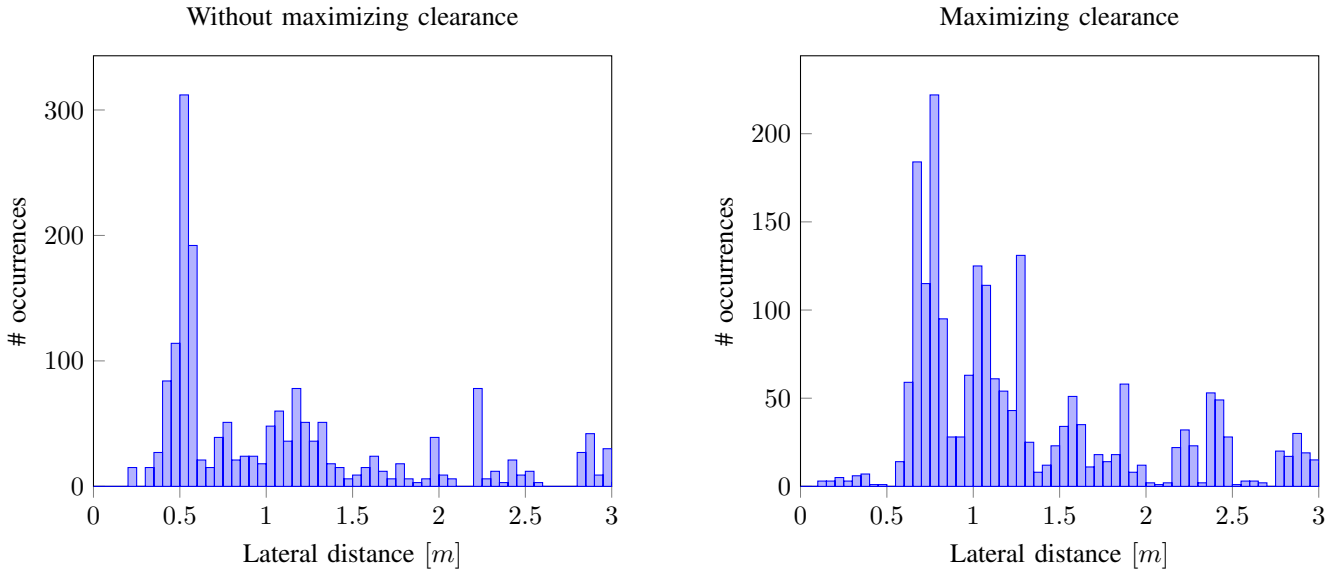
Fig. 8. Comparison of lateral clearance to other road agents during public road driving; the MPC with auxiliary goal (histogram on the right) consistently achieves higher lateral clearance.

the nominal case. In terms of input magnitude, there are no substantial differences, resulting in the same comfort level for passengers (same yaw rate and acceleration).

A video showing experimental results can be found at https://vimeo.com/348076680/be686d1397. For public road driving, the histograms in Fig. 8 show the lateral distance given to other road agents (pedestrians, bicyclists, cars) for both the MPC controllers. Such lateral distances are measured once the longitudinal displacement is equal to zero. It is evident that the extended MPC shows an increased lateral clearance for a large number of occurrences.

B. Numerical solvers used

In order to numerically solve problem (6), two off-the-shelf solvers tailored to MPC applications were tested, namely the ACADO [21] suite paired with QPOASES [22], and a proprietary Nonlinear Programming (NLP) solver. Both frameworks provide the possibility to automatically discretize the continuous time system dynamics in Eq. (1); the Explicit Runge-Kutta method of order 4 with a sampling time of $50\,\mathrm{ms}$, is employed. Even though both generate a high performing, tailored implementation of the MPC solver in C, they largely differ in terms of numerical algorithms employed. The analysis and comparison of such numerical algorithms are beyond the scope of this paper.

C. Solver comparison

Table VI-B shows the performance in terms of solvetime of the numerical solvers tested by comparing the traditional MPC tracking controller and its extended version with biasing functionality. Both solvers were compiled with all the optimizations turned on and tested in identical conditions, running on a $4.6\,\mathrm{GHz}$ i7 processor and sharing resources with other AV-relevant processes (perception, planning, etc.).

TABLE I

SOLVETIME OF THE TWO SOLVERS TESTED, WITH AND WITHOUT BIASING FUNCTIONALITY. HORIZON LENGTH $N = 60$.

	Average	Maximum
Proprietary NLP	2.1 ms	3.7 ms
Proprietary NLP (biasing)	2.3 ms	5.4 ms
ACADO+QPOASES	1.9 ms	3.9 ms
ACADO+QPOASES (biasing)	2.0 ms	5.0 ms

It is evident that both solvers are capable of finding a locally optimal solution to the problem within the time bounds provided (sampling time $T_s = 50\,\mathrm{ms}$), despite the extended MPC being a higher dimensional problem.

VII. CONCLUSION AND FUTURE WORK

This paper presents an approach enabling the specification of additional tasks in an MPC tracking controller of an AV. In particular, such task involves maximizing clearance to other road users, resulting in an increase of safety based on the TTR metric. Moreover, extensive experimental results are showed as well as comparisons with traditional MPC tracking controllers. The computational feasibility is assessed using two different numerical solvers.

Future development includes the implementation of the lane biasing functionality within a control framework commanding both speed and steering, as envisioned in [15]. The tractability of the resulting higher dimensional problem needs to be assessed, although the numerical results presented in Table VI-B are encouraging. Moreover, correlating the safety variable s with the maximum attainable speed might be an interesting extension, allowing to drive faster in case of larger clearance and consequently mimicking observed human behaviors [23] and achieving a more human-like

driving experience.

ACKNOWLEDGMENT

REFERENCES

[1] S. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.

[2] S. Lavalle. Rapidly-exploring random trees: A new tool for path planning. Technical Report 98-11, Computer Science Department, Iowa State University, 1998.

[3] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 30(7):846–894, June 2011.

[4] L. I. Reyes Castro, P. Chaudhari, J. Tumova, S. Karaman, E. Frazzoli, and D. Rus. Incremental sampling-based algorithm for minimum-violation motion planning. In *52nd IEEE Conference on Decision and Control (CDC)*, pages 3217–3224, Dec 2013.

[5] J. Tumova, G. C.Hall, S. Karaman, E. Frazzoli, and D. Rus. Least-violating control strategy synthesis with safety rules. In *Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control*, HSCC '13, pages 1–10. ACM, 2013.

[6] V. Varricchio and E. Frazzoli. Asymptotically optimal pruning for nonholonomic nearest-neighbor search. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 4459–4466, Dec 2018.

[7] A. Liniger, A. Domahidi, and M. Morari. Optimization-based autonomous racing of 1:43 scale rc cars. *Optimal Control Applications and Methods*, 36(5):628–647, 2015.

[8] J. Kong, M. Pfeiffer, G. Schildbach, and F. Borrelli. Kinematic and dynamic vehicle models for autonomous driving control design. In *IEEE Intelligent Vehicles Symposium (IV)*, pages 1094–1099, June 2015.

[9] B. Gutjahr, L. Gröll, and M. Werling. Lateral vehicle trajectory optimization using constrained linear time-varying mpc. *IEEE Transactions on Intelligent Transportation Systems*, 18(6):1586–1595, June 2017.

[10] T. Nägeli, J. Alonso-Mora, A. Domahidi, D. Rus, and O. Hilliges. Real-time motion planning for aerial videography with dynamic obstacle avoidance and viewpoint optimization. *IEEE Robotics and Automation Letters*, 2(3):1696–1703, July 2017.

[11] T. Nägeli, L. Meier, A. Domahidi, J. Alonso-Mora, and O. Hilliges. Real-time planning for automated multi-view drone cinematography. *ACM Trans. Graph.*, 36(4):132:1–132:10, July 2017.

[12] D. Falanga, P. Foehn, P. Lu, and D. Scaramuzza. PAMPC: Perception-aware model predictive control for quadrotors. In *IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, Oct 2018.

[13] H. Andersen et al. Trajectory optimization for autonomous overtaking with visibility maximization. In *IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–8, Oct 2017.

[14] A. Censi, K. Slutsky, T. Wongpiromsarn, D. Yershov, S. Pendleton, J. Fu, and E. Frazzoli. Liability, ethics, and culture-aware behavior specification using rulebooks. In *IEEE Conference on Robotics and Automation (ICRA)*, May 2019.

[15] F. Seccamonte, E. M. Wolff, E. Frazzoli, and J. Kabzan. Adjusting lateral clearance for a vehicle using a multi-dimensional envelope. *US Patent Application filed February 25th, 2019. Patent pending.*

[16] D. Schramm, M. Hiller, and R. Bardini. *Vehicle Dynamics: Modeling and Simulation*. Jan 2012.

[17] Jean-Paul P. Laumond. *Robot Motion Planning and Control*. Springer-Verlag, Berlin, Heidelberg, 1998.

[18] L. Ljung. *System Identification: Theory for the User*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1999.

[19] F. Borrelli, A. Bemporad, and M. Morari. *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press, New York, NY, USA, 1st edition, 2017.

[20] J. Hillenbrand, A. M. Spieker, and K. Kroschel. A multilevel collision mitigation approach - its situation assessment, decision making, and performance tradeoffs. *IEEE Transactions on Intelligent Transportation Systems*, 7(4):528–540, Dec 2006.

[21] B. Houska, H.J. Ferreau, and M. Diehl. ACADO Toolkit – An Open Source Framework for Automatic Control and Dynamic Optimization. *Optimal Control Applications and Methods*, 32(3):298–312, 2011.

[22] H.J. Ferreau, C. Kirches, A. Potschka, H.G. Bock, and M. Diehl. qpOASES: A parametric active-set algorithm for quadratic programming. *Mathematical Programming Computation*, 6(4):327–363, 2014.

[23] C. Llorca, A. Angel-Domenech, F. Agustin-Gomez, and A. Garcia. Motor vehicles overtaking cyclists on two-lane rural roads: Analysis on speed and lateral clearance. *Safety Science*, 92:302 – 310, 2017.