Deriving Spatial Policies for Overtaking Maneuvers with Autonomous Vehicles

Jayanth Bhargav
Electrical & Systems Engineering
University of Pennsylvania
Philadelphia, USA
jayanthb@seas.upenn.edu

Johannes Betz
Electrical & Systems Engineering
University of Pennsylvania
Philadelphia, USA
joebetz@seas.upenn.edu

Hongrui Zehng
Electrical & Systems Engineering
University of Pennsylvania
Philadelphia, USA
hongruiz@seas.upenn.edu

Rahul Mangharam
Electrical & Systems Engineering
University of Pennsylvania
Philadelphia, USA
rahulm@seas.upenn.edu

2022 14th International Conference on COMmunication Systems & NETworkS (COMSNETS) | 978-1-6654-2104-1/22/$31.00 ©2022 IEEE | DOI: 10.1109/COMSNETS53615.2022.9668548

***Abstract*—Planning an accurate and safe trajectory is a crucial element in autonomous driving. To execute complex driving maneuvers like overtaking, motion planning requires an enhanced decision-making algorithm that decides the when, where and how of the overtaking maneuver. This paper proposes an algorithm that increases the likelihood of a safe overtaking maneuver by learning spatial information. Here, spatial information refers to the track portion/curve and the position of the ego vehicle with reference to that. The technique is applied to an autonomous racing setup where vehicles have to detect and operate at the limits of dynamic handling. To learn the spatial information, offline experiments of a 2-player race are conducted to generate probability distributions of overtaking maneuvers conditioned on speed and relative-position of the ego vehicle with respect to the opponent. Furthermore, a Switched Model Predictive Contouring Controller (SMPCC) is proposed for incorporating the policy learning algorithm into the path planning and control setup. Extensive simulations show that the proposed algorithm is able to achieve an increased number of overtakes at different track portions on known and unknown race tracks.**

***Index Terms*—autonomous systems, automobiles, intelligent vehicles, optimal control, path planning**

I. Introduction

A. Autonomous Racing

Autonomous racing has become popular over the recent years and competitions like Roborace [1] or the Indy Autonomous Challenge as well as small-scale competitions like F1Tenth [2] provide platforms for evaluating autonomous driving algorithms and software. The overall goal of all these competitions is that researchers and engineers can develop algorithms that operate vehicles at the edge: high speeds, high accelerations, high computation power, adversarial environments. The algorithms that were developed in the field of autonomous racing so far are mostly focusing on single vehicle only that try to achieve a human-like lap time. The field of high dynamic overtaking maneuver with dynamic opponents are less displayed. In addition, achieving a human-like behavior (e.g. like a Formula 1 race driver) that makes the decision about an overtaking maneuver and executes a secure and reliable maneuver at high speeds is still an unsolved problem.

B. Contributions

In this paper, an approach to learn spatial information for overtaking maneuvers in autonomous vehicles is presented. This work has three primary contributions:

1) Design of Experiments (DoE) for offline policy learning.
2) An application of autonomous driving to learn effective overtaking maneuvers for autonomous race cars. Discretization of selected f1 tracks into a category of turns/curves and simulations of 2-player race to derive overtaking policies for different track portions.
3) A Switched Model Predictive Contouring Controller (SMPCC) setup based on [3], which combines a receding horizon control algorithm and specific driving behaviours.

II. Related Work

Dixit et al. [4] provide a state of the art review of trajectory planning and control for autonomous overtaking maneuvers. The authors state finally in their review, that two important aspects of trajectory planning for high-speed overtaking need to be addressed: (i) inclusion of vehicle dynamics and environmental constraints and (ii) accurate knowledge of the environment and surrounding obstacles.

Although the state of the art displays a plethora of algorithms for path and behavioral planning of autonomous vehicles, explicit algorithm development for autonomous race cars is relatively lesser. As part of the Roborace competition [5], [6] [7] presented a planning and control system for real life autonomous racing cars. Both approaches focused on a holistic software architecture that is capable of dynamic overtaking. Nevertheless none of them realized a head to head race with the vehicles. As a part of the same competition, [8] presented a nonlinear model predictive control (NMPC) for racing. The overtaking strategy was implemented as a term in the objective

function. The NMPC has the freedom to choose the side for an overtake and was mainly relying on the obstacles velocity to perform the overtaking maneuver. In [9] a simple Q-Learning algorithm is applied to learn the behavior of an virtual opponent car to apply an effective overtaking strategy on either the straight or right before tight bend.
In [10] a method to plan overtaking maneuvers in autonomous racing based on gaussian processes is presented. This machine learning method is able to learn the behavior of the opponent vehicle. Based on the outputs of this process, a stochastic MPC plans optimistic trajectories that lead to a controlled overtaking maneuver of the lead vehicle.
In multi vehicle racing, [11] presented a non-cooperative game theory approach where autonomous racing, formulated as racing decisions is a non-cooperative nonzero-sum game. Liniger et al. [11] displayed that different games can be modelled that achieve successfully different racing behaviors and generate interesting racing situations e.g. blocking and overtaking. Notomista et al. [12] considered a two-player racing game where the ego vehicle is based on a Sensitivity-ENhanced NAsh equilibrium seeking (SENNA) method, which uses an iterated best response algorithm in order to optimize for a trajectory in a two-car racing game. Jung et al. [13] present a game-theoretic MPC approach for head-to-head autonomous racing that consists of a (1) game-based opponents' trajectory predictor, (2) high-level race strategy planner, and (3) MPC-based low-level controller. Based on the results of the prediction, the high-level race strategy planner plans several behaviors to respond to various race circumstances.

The state of the art displays that the autonomous racing community is focusing on integrating effective learning techniques and strategies into dynamic path and behavioral planning to make the car faster, more reliable and more interactive [14] [15] [16]. Improvements in planning/control for overtaking maneuvers have not yet been explored extensively and learning from spatial information (track portions and position of the vehicle on the track) has not been examined before.

III. Design of Experiments

We propose an offline experiment setup which will create emphasis on specific track portions and examine the overtaking maneuvers. With these offline experiments, it is possible to create track-based policies that can be used in a high level decision maker, behavior or motion planner.

A. Track Portions

Based on our racetrack application, in the first step we define the track portions that we will examine. We will use four high level definitions of track portions that are the most common kinds of curves/turns found on racetracks: *Straight*, *Sweeper Curve*, *Hairpin Curve* and *Chicane*

For example, consider the racetrack in Budapast, Hungary. In figure 1 we display 11 different track portions on the track that are marked with labels 1 to 11.

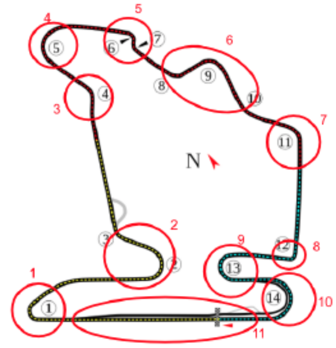

Fig. 1. Example racetrack for offline experiments: Budapest Circuit, Hungary

Track portions are defined by $\mathbf{T} = \{\tau | 1 \leq \tau \leq 11, \tau \in \mathbb{N}\}$ and are categorized in the four track segment types:

- *Straight*: 11
- *Sweeper Curve*: 3, 7, 8
- *Hairpin Curve*: 1, 2, 4, 9, 10
- *Chicane*: 5, 6

B. Sampling based trajectory rollouts

To examine these defined track portions we setup an offline simulation that varies different parameters visualized in figure 2.

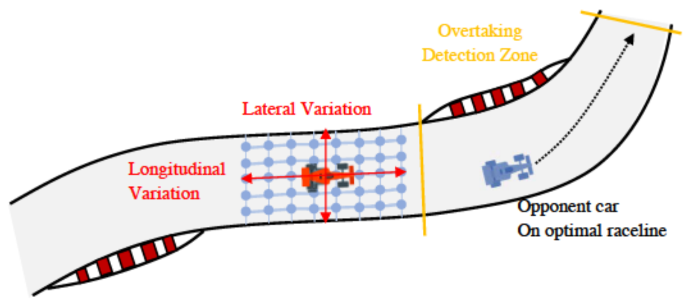

Fig. 2. Ego vehicle (red car) starting behind the opponent vehicle (blue car) on the track. Apart from the velocity, lateral and longitudinal positions of the ego vehicle are varied as shown in the figure. The opponent vehicle follows a pre-computed race line and is non-interactive. The overtaking maneuver is examined in the overtaking detection zone (marked in yellow).

The opponent vehicle follows a curvature optimal pre-computed race line based on [17] and is non-interactive. For every track portion $\tau \in \mathbf{T}$, a uniformly sampled set of positions $\mathcal{P} : \mathcal{XXY} \subset \mathbf{R}^2$ are chosen as the starting position for the ego. The obstacle vehicle speed is varied as $v_{obs} = v_{baseline} * (1 + s)$ where $s \in \{-0.2, 0, +0.2\}$, $v_{baseline}$ being the speed of the obstacle from the pre-computed optimal race line.

The agents are initialised with these positions and set to start the simulation. The ego vehicle synthesises dynamic trajectories based on the MPCC planner with obstacle avoidance. A fully observable model is used for the ego vehicle i.e. the ego vehicle will have the information of the track portion τ which it is driving in and the current state of the obstacle $X_{obs} = [x_{obs}, y_{obs}, \phi_{obs}]$

In this setup, we conduct simulations based on the following parameter variations:

- *Lateral offset:* The position of the ego vehicle is varied lateral across the track with an offset from the centerline.
- *Longitudinal offset:* The position of the ego vehicle is varied longitudinal along the centerline of the track.
- *Opponent speed change:* The opponent speed is varied with $\pm 20\%$ from baseline

With a high expectation, the ego vehicle will succeed in an overtaking maneuver when the obstacle speed is 20% lower than its baseline. This verifies the fact that speed advantage always helps in overtaking (e.g. DRS zones in F1). The next set of parameters that influence the overtaking maneuver is the position. In convoluted race tracks, we can display that starting off at a specific position gives us a higher chance of an overtaking maneuver. For each track portion, we define four regions of interest: $\mathcal{R}_1$, $\mathcal{R}_2$, $\mathcal{R}_3$ and $\mathcal{R}_4$. Starting positions of the ego vehicle are uniformly sampled in all the four regions to generate experimental data.

IV. Planning and Control Setup

Continuous time system dynamics is used to develop a constrained optimal controller to steer the vehicle in the track. The optimal planner plans the path for a horizon of N steps ahead, steers the vehicle with the first step, and again repeats the process for the specified amount of time. This is a modified form of the Model Predictive Controller (MPC).

A. Model Predictive Contouring Control

The MPCC problem defined in [3] is re-formulated into a finite-continuous time optimal control problem as follows:

$$\min \int_0^T \begin{bmatrix}\epsilon_c^{lin}(t) & \epsilon_l^{lin}(t)\end{bmatrix} \begin{bmatrix} Qc & 0 \\ 0 & Ql \end{bmatrix} \begin{bmatrix}\epsilon_c^{lin}(t) \\ \epsilon_l^{lin}(t)\end{bmatrix}$$
$$-Q_\theta \dot{\theta}(t) + u^T(t) R u(t) dt$$

$$\begin{aligned} \text{s.t.} \qquad \dot{x} &= f(x, u, \Phi) \\ b_{lower} \preceq x(t) &\preceq b_{\text{upper}} \\ l_{\text{lower}} \preceq u(t) &\preceq l_{\text{upper}} \\ h(x, \Phi) &\leq 0 \end{aligned}$$

given the system dynamics f and the arclength parametrization of the contour (the track) Φ. A single-track bicycle model is used. Here $x(t)$ denotes the system state, $u(t)$ the inputs to the system, b the box constraints on the state, l the box constraints on the input and h captures the track boundary constraints. The state of the system is augmented with the advancing parameter θ and the virtual input $\dot{\theta}$ is appended to the inputs from the original system dynamics.
Q_c, Q_l, Q_θ and R are the cost-function parameters of the MPC controller.

The track boundary constraint is realized as a convex disk constraint.

$$h(x, \Phi) = \left(x - x_t^{lin}(\theta)\right)^2 + \left(y - y_t^{lin}(\theta)\right)^2 - r_\Phi(\hat{\theta})^2$$

Here $r_\Phi(\hat{\theta})$ is the half-width of the track at the last predicted arc length.

The contouring error ϵ_c^{lin} and lag error ϵ_l^{lin} described in [3] are modified by linearizing them around the previous solution θ.

The MPCC is optimizing to move the position of a virtual point $\theta(t)$ along the track to achieve as much progress as possible while steering the model of the vehicle to keep contouring and lag errors as small as possible.
The center-line of the track is given in way-points (X-and Y-position). To implement MPCC an arc-length parametrization Φ is required. This is realized by interpolating the way-points using cubic splines with a cyclic boundary condition, and creating a dense lookup table with the track location and the linearization parameters.

B. Switched Model Predictive Contouring Control (SMPCC)

To achieve more control over the path planning of the ego vehicle, the proposed SMPCC setup is displayed in figure 3.

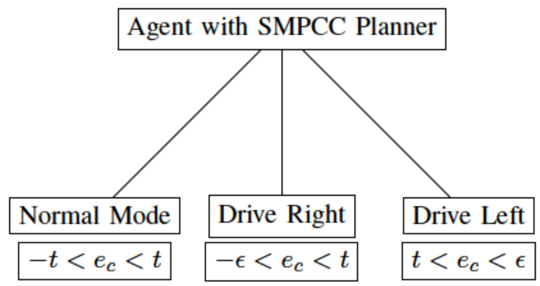

Fig. 3. Modes of SMPCC

In this, the agent switches between different modes defined by different solver formulations. The constraints for the modes as shown in fig. 3 are added to the problem formulated in section IV-A. These constraints are tuned with the slack variable (ϵ) to ensure that the planner does not get stuck into an in-feasibility loop leading to a crash. The ego can therefore overtake on both left and right side of the opponent vehicle. The MPCC control problem is solved by efficient interior point solvers in FORCES [18].

V. Results and Discussion

A. Offline Spatial Policy Learning

In the following section, we present the results from the offline experiments. Algorithm 1 elucidates the offline experiment based policy learning developed in this paper.
$\mathcal{X}$, $\mathcal{Y}$ are the set of x and y coordinate offsets (expressed as percentage of track width) and $\mathcal{S} = \{-0.2, 0, +0.2\}$ is the speed offset expressed as a percentage change from the baseline obstacle speed.

The obstacle update model is $g(\cdot)$, which is a pre-computed curvature optimal race line of the race track under consideration. Define Ψ : {Silverstone circuit (England), Hungaroring circuit (Budapest), Catalunya circuit (Spain) and Nürburgring circuit (Germany)}, the set of race circuits on which the race is conducted for learning the policies. In total we simulate 576 experiments based on the 16 lateral, 12 longitudinal and 3 velocity variations for each track portion in each of the four racetracks. The algorithm 1 populates the policy map Π_k with the track regions for each of the curves, having highest probability of overtakes for all $k \in \Psi$.

Algorithm 1: Offline Spatial Policy Learning

```
Function MPCC Planner(X_obs):
    solve MPCC problem defined in Section IV
    return u*;
Function Check Overtake(X_obs, X_ego):
    project X_obs, X_ego as s_1, s_2 on track
    return bool(s_1 > s_2);
for k ∈ Ψ do
    initialize: Π_k = {}
    for τ ∈ T do
        initialize: p = {}, overtakes = {}, total = {}
        for x, y, s ∈ 𝒳X𝒴X𝒮 do
            initialize: X_ego, X_obs
            for t = 0 to T_sim do
                u* = MPCC Planner (X_obs)
                steer ego: X_ego^+ = f(X_ego, u*)
                update obstacle pos: X_obs^+ = g(X_obs)
                X_obs, X_ego = X_obs^+, X_ego^+
                identify track region i ∈ {0, 1, 2, 3}
                if Check Overtake(X_obs, X_ego) then
                    overtakes[R_i] ++;
                total[R_i] ++ ;
            end
        end
        compute p[R_i] = overtakes[R_i]/total[R_i], ∀i
        Π_k[τ] = argmax(p)
    end
end
```

Figure 4 describes an example of the four track regions and the overtaking probabilities that are evaluated based on algorithm 1. The overtaking probabilities for two of the racetracks with their respective track portions are displayed in tables I and II.

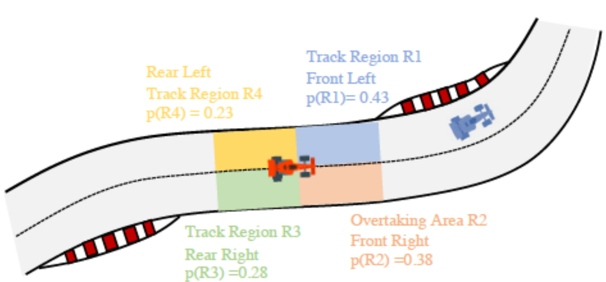

Fig. 4. Predefined track regions of interest for overtaking maneuver at a specific turn with overtaking success probabilities.

Results from the experiments conducted on the four race-tracks are summarized in figures 5, 6 and 7. They display the overtaking probability distribution for each overtaking zone

TABLE I
OVERTAKING PROBABILITIES FOR ALL TRACK PORTIONS - RACETRACK 1 (SILVERSTONE, ENGLAND)

Track Portion (τ)	Track Portion Type	$p(\mathcal{R}_1)$	$p(\mathcal{R}_2)$	$p(\mathcal{R}_3)$	$p(\mathcal{R}_4)$
1	Sweeper	**0.99**	0.93	0.52	0.59
2	Hairpin	**0.63**	0.52	0.31	0.33
3	Hairpin	**0.41**	0.38	0.25	0.21
4	Sweeper	0.65	**0.67**	0.57	0.59
5	Chicane	**0.25**	0.21	0.14	0.21
6	Straight	0.99	**1.0**	0.95	0.99
7	Sweeper	0.47	**0.52**	0.33	0.31
8	Hairpin	**0.40**	0.36	0.37	0.38

TABLE II
OVERTAKING PROBABILITIES FOR ALL TRACK PORTIONS - RACETRACK 2 (BUDAPEST, HUNGARY)

Track Portion (τ)	Track Portion Type	$p(\mathcal{R}_1)$	$p(\mathcal{R}_2)$	$p(\mathcal{R}_3)$	$p(\mathcal{R}_4)$
1	Hairpin	**0.52**	0.49	0.34	0.29
2	Hairpin	0.31	**0.43**	0.27	0.21
3	Sweeper	**0.71**	0.62	0.54	0.59
4	Hairpin	**0.67**	0.53	0.39	0.41
5	Chicane	0.41	**0.43**	0.22	0.25
6	Chicane	**0.27**	0.21	0.19	0.14
7	Sweeper	**0.67**	0.64	0.35	0.37
8	Sweeper	**0.59**	0.48	0.40	0.36
9	Hairpin	0.44	**0.58**	0.35	0.29
10	Hairpin	**0.65**	0.60	0.44	0.38
11	Straight	0.97	**0.99**	0.95	0.94

($\mathcal{R}_1$ -$\mathcal{R}_4$) at chicanes, sweeper curves and hairpins respectively across all the race-tracks considered for policy learning. From statistical analysis on the offline experiments for four different racetracks and their 39 track portions, we have the following observations:

- An overtaking maneuver will be more successful if we are closer to the opponent vehicle. This can be seen in both the raw numbers from table I and II as well in the higher median and maximum in the boxplots for $\mathcal{R}_1$ & $\mathcal{R}_2$.
- On the *straight*, it does not really matter which side we are on the track when trying to overtake, we just need stay closer to the opponent.
- The *sweeper curve* generally has a high overtaking probability due to the high speeds of the car. We only get an advantage here if we are close enough to the car and therefore we need to be in region $\mathcal{R}_1$ or $\mathcal{R}_2$.
- We can see that in each *hairpin* we have the highest overtaking probability in either in $\mathcal{R}_1$ or $\mathcal{R}_2$ depending on the nature of the hairpin turn: right or left. This is due to the fact that being on the inside of the curve near a hairpin, the car is able to achieve a better trajectory through the hairpin.
- A *chicane* generally has the lowest overtaking probability due to the fact that it is a complex region for the car to maneuver and with lesser space for overtaking. Since the chicane is a highly convoluted turn, curvature direction does not indicate any better start regions.

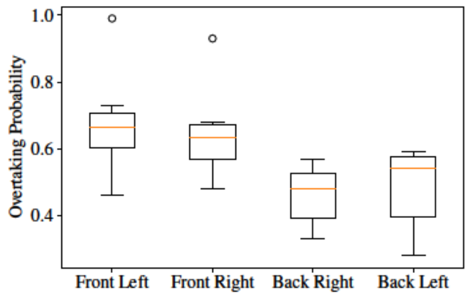

Fig. 5. Overtaking probability distribution for *sweeper curves.*

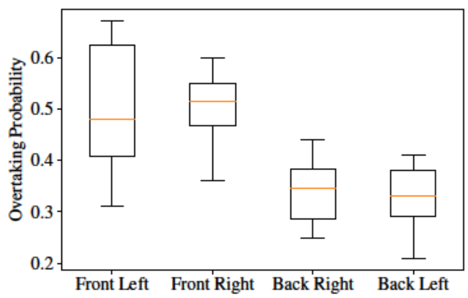

Fig. 6. Overtaking probability distribution for *hairpins.*

B. Online Policy Execution

Algorithm 1 furnishes a policy map for all the different curves/track portions from different race tracks considered during the policy learning phase. We will now compute the best policy $\Pi(s)$ as a function of the track portion s and integrate it into the SMPCC setup on the ego vehicle. We then compare the number of overtakes with and without the spatial policy based MPCC controller to verify the effectiveness of our algorithm.

Algorithm 2: Online Evaluation with SMPCC

Function SMPCC (X_{obs}, *mode*) **:**
    if *mode* = *'normal'* **then**
        $u^* =$ solve MPCC problem in Section IV
    else
        modify MPCC to integrate policy (see fig. 3)
        $u^* =$ solve modified MPCC problem
    end
    return u^*;
initialize X_{ego}, X_{obs} mode = 'normal'
for $t = 0$ *to* T_{sim} **do**
    $u^* =$ SMPCC (X_{obs},mode)
    steer the ego: $X_{ego}^{+} = f(X_{ego}, u^*)$
    update obstacle position: $X_{obs}^{+} = g(X_{obs})$
    $X_{obs}, X_{ego} = X_{obs}^{+}, X_{ego}^{+}$
    identify track portion τ where ego is present
    policy lookup: mode = $\Pi[\tau]$
end

The results display that the offline policy learning approach has been successful by showing an increased number of overtaking maneuvers at all racetracks and track portions. We can observe that overtaking on the straights is usually easy (even without policy). Since sweeper curves are usually wide track portions that allow high speeds and do not involve complicated maneuvers, both with and without switching policy, we achieve higher overtaking maneuvers. Although we see that the switching policy leads to more overtaking maneuvers because having the right position for the overtaking maneuver is crucial here, too. We see the highest impact of our algorithm at hairpins and chicanes. This is mainly due to the fact that overtaking at these curves is usually complicated and needs a good strategy beforehand. We see that our algorithm can nearly double the amount of overtaking maneuvers in the chicane (track portion 10 of Catalunya) which substantiates the fact that having the right starting position for an overtaking maneuver is indispensable.

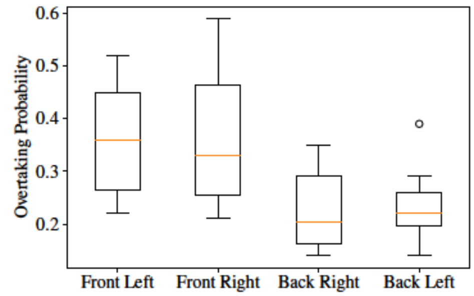

Fig. 7. Overtaking probability distribution for *chicanes.*

TABLE III
RESULTS: NUMBER OF OVERTAKES WITH AND WITHOUT POLICY ON RACETRACK 1 (SILVERSTONE, ENGLAND)

Track Portion (τ)	Track Portion Type	Number of Overtakes Policy OFF	Number of Overtakes Policy ON
1	Sweeper	436	452
2	Hairpin	256	337
3	Hairpin	308	426
4	Sweeper	342	357
5	Chicane	117	302
6	Straight	565	566
7	Sweeper	237	283
8	Hairpin	218	394

TABLE IV
RESULTS: NUMBER OF OVERTAKES WITH AND WITHOUT POLICY ON RACETRACK 2 (BUDAPEST, HUNGARY)

Track Portion (τ)	Track Portion Type	Number of Overtakes Policy OFF	Number of Overtakes Policy ON
1	Hairpin	286	375
2	Hairpin	297	404
3	Sweeper	410	446
4	Hairpin	337	398
5	Chicane	270	361
6	Chicane	180	329
7	Sweeper	372	438
8	Sweeper	239	297
9	Hairpin	288	326
10	Hairpin	301	374
11	Straight	558	563

C. Evaluation on an Unknown Track

As an ultimate test, we now apply this online policy algorithm to the agent racing in an unknown racetrack (Sakhir Circuit, Bahrain). The results from Sakhir Circuit show an increase in the number of overtakes with the policy at all track portions as displayed in table V.

TABLE V
RESULTS: NUMBER OF OVERTAKES WITH AND WITHOUT POLICY ON UNKNOWN RACETRACK 5 (SAKHIR, BAHRAIN)

Track Portion (τ)	Track Portion Type	Turn Direction (Left/ Right)	Policy Region	Number of Overtakes Policy OFF	Number of Overtakes Policy ON
1	Chicane	Right	$\mathcal{R}_2$	254	298
2	Sweeper	Right	$\mathcal{R}_1$	278	367
3	Chicane	Left	$\mathcal{R}_1$	214	391
4	Hairpin	Right	$\mathcal{R}_1$	319	421
5	Hairpin	Left	$\mathcal{R}_2$	224	327
6	Straight	Left	$\mathcal{R}_2$	551	560
7	Sweeper	Right	$\mathcal{R}_1$	348	418
8	Sweeper	Right	$\mathcal{R}_1$	297	368
9	Straight	Left	$\mathcal{R}_2$	559	549
10	Sweeper	Right	$\mathcal{R}_1$	311	396
11	Straight	Right	$\mathcal{R}_1$	552	561

Additional offline policy evaluations have shown that a generalization from turn directions and overtaking zones is only partially useful. The generalised policies did not lead to successful overtakes always and were not feasible in some cases. A possible refinement could be the focus on using a parametric curvature of the turn rather than the high level definition of left or right.

VI. CONCLUSION AND FUTURE WORK

In this paper, an algorithm for spatial policy learning from offline experiments is proposed to learn effective overtaking strategies based on position advantage. Extensive simulations on real world racetrack layouts show that the proposed algorithm is able to learn regions of high probabilities on a racetrack for successful and safe overtaking maneuvers. The (SMPCC) setup, that has the driving policies integrated into the motion planning and control stack of the vehicle resulted in an increase in the number of overtakes. Specifically, the policy based algorithm was found to be highly effective for convoluted track portions like chicanes, where a positional advantage plays a major role in a successful overtaking maneuver. In summary, with the setup defined in this paper, one can create more realistic and better overtaking maneuvers for autonomous vehicles. This brute-force technique of learning spatial information serves as a fundamental result and ground truth for future work. Extensions to this work will include learning-based algorithms based on reinforcement learning techniques to identify the overtaking probability based on the curvature information of upcoming turns and can therefore applied to behavioral planners for passenger autonomous vehicles. One can consider non-reactive, reactive and aggressive opponents which are defensive and sophisticated to overtake and therefore deriving a holistic strategy for overtaking on e.g. highways. With this setup, one can learn and integrate complex human-like overtaking maneuvers for autonomous vehicles in a safe and reliable manner.

REFERENCES

[1] J. Betz, A. Wischnewski, A. Heilmeier, F. Nobis, T. Stahl, L. Hermansdorfer, B. Lohmann, and M. Lienkamp, "What can we learn from autonomous level-5 motorsport?" in *Proceedings*. Springer Fachmedien Wiesbaden, Sep. 2018, pp. 123–146.
[2] M. OKelly, H. Zheng, D. Karthik, and R. Mangharam, "F1tenth: An open-source evaluation environment for continuous control and reinforcement learning," in *Proceedings of the NeurIPS 2019 Competition and Demonstration Track*, ser. Proceedings of Machine Learning Research, vol. 123. PMLR, 2020, pp. 77–89.
[3] A. Liniger, A. Domahidi, and M. Morari, "Optimization-based autonomous racing of 1: 43 scale rc cars," *Optimal Control Applications and Methods*, vol. 36, no. 5, pp. 628–647, 2015.
[4] S. Dixit, S. Fallah, U. Montanaro, M. Dianati, A. Stevens, F. Mccullough, and A. Mouzakitis, "Trajectory planning and tracking for autonomous overtaking: State-of-the-art and future prospects," *Annual Reviews in Control*, vol. 45, pp. 76–86, 2018.
[5] J. Betz, A. Wischnewski, A. Heilmeier, F. Nobis, L. Hermansdorfer, T. Stahl, T. Herrmann, and M. Lienkamp, "A software architecture for the dynamic path planning of an autonomous racecar at the limits of handling," in *2019 IEEE International Conference on Connected Vehicles and Expo (ICCVE)*. IEEE, Nov. 2019.
[6] T. Stahl, A. Wischnewski, J. Betz, and M. Lienkamp, "Multilayer graph-based trajectory planning for race vehicles in dynamic scenarios," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, Oct. 2019.
[7] D. Caporale, L. Venturini, A. Fagiolini, L. Pallottino, A. Settimi, A. Biondo, F. Amerotti, F. Massa, S. D. Caro, and A. Corti, "A planning and control system for self-driving racing vehicles," in *2018 IEEE 4th International Forum on Research and Technology for Society and Industry (RTSI)*. IEEE, Sep. 2018.
[8] A. Buyval, A. Gabdulin, R. Mustafin, and I. Shimchik, "Deriving overtaking strategy from nonlinear model predictive control for a race car," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Sep. 2017.
[9] D. Loiacono, A. Prete, P. L. Lanzi, and L. Cardamone, "Learning to overtake in TORCS using simple reinforcement learning," in *IEEE Congress on Evolutionary Computation*. IEEE, Jul. 2010.
[10] T. Brüdigam, A. Capone, S. Hirche, D. Wollherr, and M. Leibold, "Gaussian process-based stochastic model predictive control for overtaking in autonomous racing," 2021.
[11] A. Liniger and J. Lygeros, "A noncooperative game approach to autonomous racing," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 3, pp. 884–897, May 2020.
[12] G. Notomista, M. Wang, M. Schwager, and M. Egerstedt, "Enhancing game-theoretic autonomous car racing using control barrier functions," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2020.
[13] C. Jung, S. Lee, H. Seong, A. Finazzi, and D. H. Shim, "Game-theoretic model predictive control with data-driven identification of vehicle model for head-to-head autonomous racing," 2021.
[14] U. Rosolia and F. Borrelli, "Learning how to autonomously race a car: A predictive control approach," *IEEE Transactions on Control Systems Technology*, vol. 28, no. 6, pp. 2713–2719, Nov. 2020.
[15] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, "Learning-based model predictive control for autonomous racing," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3363–3370, Oct. 2019.
[16] N. R. Kapania and J. C. Gerdes, "Learning at the racetrack: Data-driven methods to improve racing performance over multiple laps," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 8, pp. 8232–8242, Aug. 2020.
[17] A. Heilmeier, A. Wischnewski, L. Hermansdorfer, J. Betz, M. Lienkamp, and B. Lohmann, "Minimum curvature trajectory planning and control for an autonomous race car," *Vehicle System Dynamics*, vol. 58, no. 10, pp. 1497–1527, Jun. 2019.
[18] A. Zanelli, A. Domahidi, J. Jerez, and M. Morari, "Forces nlp: an efficient implementation of interior-point methods for multistage nonlinear nonconvex programs," *International Journal of Control*, vol. 93, no. 1, pp. 13–29, 2020.