

Following Cars With Elastic Bands

Fritz Ulbrich¹, Stephan Sundermann¹, Tobias Langner¹, Daniel Goehring¹, and Raúl Rojas¹

Abstract—We propose a trajectory planning approach for autonomous vehicles in highly dynamic traffic scenarios, using elastic bands to follow the observed trajectories of other vehicles. The focus of this paper is on the initialization of the elastic band. The proposed method does not rely on a map. We tested our method using recorded urban traffic data. The results show that the presented approach is valid and the proposed initialization process is clearly superior to naive initialization.

I. INTRODUCTION

Traditionally, path and trajectory planning algorithms for autonomous cars rely heavily on high precision maps. Even though map providers have been coming up with high precision maps recently, creating these high precision maps is still a tedious task and those maps are not available everywhere. The given map may be outdated, temporarily invalid due to accidents, construction sites or road work. Weather conditions may influence the drivability of roads or the code of conduct regarding e.g. speed limits or using specific lanes. Also, a source for global localization may be temporarily unavailable, making it impossible to correlate the vehicle's position with the map.

The approach presented in this paper aims at solving the problem of (temporarily) not having a valid map. It is based on a simple but effective strategy: following other cars and imitating their behavior. This can be seen as some kind of swarm behavior, such as ants following the pheromone tracks of their fellow species. Especially in the near future, where there will be a mixed traffic of (mostly) driver-operated and autonomous vehicles. Autonomous cars should make use of the exceptional abilities of humans to perceive uncommon environments and adopt to unforeseen events. But also in an all autonomous environment, there may be other cars with a valid map or a superior sensor setup.

The concept of elastic bands for trajectory optimization is perfectly suited for the task of following other vehicles' trajectories. They enable planning in continuous space and are still efficient compared to searching a huge discrete state space. On the other hand, they can only find a local optimum, which can be mitigated by providing a good estimate for initialization. This publication focuses on this process of initializing the elastic band in more complicated scenarios. Figure 1 gives an overview of the proposed process of generating a trajectory.

Section II gives a brief overview of approaches for path planning, especially ones using elastic bands. Section III describes briefly the concept of timed elastic band and




Fig. 1: Overview of the proposed trajectory planning. (a) A target vehicle (dark blue) is chosen. The ego vehicle (gray) is supposed to follow the target's trajectory. (b) An initial elastic band is generated based on the target's trajectory. (c) The trajectory is generated by optimizing the initial band. The optimization is taking into account dynamic obstacles and attracts the band to all observed vehicles' trajectories.

introduces algorithms for initializing the band based on the trajectories of other vehicles. Experimental results are presented in IV, followed by a conclusion in V.

II. RELATED WORK

Many planning algorithms for road vehicles require maps, as e.g. in [2]. For lane changes and other maneuvers they often use state machines or other hand-crafted, rule based solutions. Often, those hand-crafted rules are too conservative to work well in certain real-world scenarios, because of high traffic dynamics, small distances between vehicles, or because of occluded areas. One example for a planning approach which relies on maps and which is based on the A* planner has been described in [16]. An extension to A* planners, the ARA* planner has been described in [7].

An example for A* planners with applications to free form navigation is given by [8]. In the work of [5] the authors present a solution to overcome the separation between route based and trajectory planning while introducing a concept called pseudo-homology. A key focus of their work has been put on the question of how to yield, follow or overtake other agents. In [14] an idea was presented how to create driveable paths by observing other objects using lidar sensors. Another approach which takes advantage of potential fields and MPC to plan around obstacles has been presented in [10]. Other popular approaches to plan in dynamic environments are spatiotemporal state lattices [17].

The trajectory planning approach described in this paper uses the concept of stable timed elastic bands with loose

¹DCMLR, Freie Universität Berlin, Germany, {fritz.ulbrich | daniel.goehring | stephan.sundermann | tobias.langner | raul.rojas}@fu-berlin.de

ends (STEBLE) proposed in [15]. It largely builds on the timed elastic band (TEB) framework introduced by Rösman et. al. [11] and objective functions implemented in the ROS package `teb_local_planner` [1]. TEB extends the original elastic band approach, introduced by Quinlan et. al. [9], by augmenting the elastic band with a sequence of time intervals. This enables formulating constraints related to temporal aspects. In [12] Rösman et. al. propose an extension to avoid local minima by planning multiple topologically distinctive trajectories simultaneously and [13] combines timed elastic bands with MPC. Algorithms which involve timed elastic bands have been applied recently for collision avoidance in Keller et al. [6] as well as for urban autonomous driving [4].

III. STABLE TIMED ELASTIC BANDS WITH LOOSE ENDS

The approach described in this paper relies on stable timed elastic bands with loose ends (STEBLE) to generate locally optimal trajectories. The concept is presented and described in detail in [15]. The main idea behind STEBLE is, that the time between the poses is constant. As a result the band length itself is fixed regarding the temporal aspect. A fixed time interval is defined between a fixed number of spatial configurations. This results in a much smaller search space for (locally) optimal solutions. Another major aspect is, that the last configuration of the band (i.e. the goal) is not set to a fixed position and orientation during the optimization process. A spatially fixed goal configuration would imply the unnecessary constraint of being at a specific time at a specific position. Instead, to follow other cars it is sufficient to be somewhere on their trajectory at a given time. STEBLE uses an objective function that attracts the band to other vehicles' trajectories to achieve such a behaviour.

The elastic band itself is described as a sequence $Q = \{X_i\}_{i=0}^n$ of spatial robot configurations in a three-dimensional workspace with $\mathbf{x}_i = (x_i, y_i, \beta_i)$ (two coordinates for the position on a plane plus one angle for the orientation). The number of configurations n is set initially and does not change during the optimization. Section III-C describes how the initial number and values of configurations are determined in detail. The first configuration - i.e. the start - is always fixed during the optimization process. That is because the planned trajectory must always start at the ego vehicle's current position and orientation.

Furthermore, a time interval ΔT is chosen, denoting the time needed by the ego vehicle to transit from one configuration \mathbf{x}_i to another configuration \mathbf{x}_{i+1} . ΔT is fixed for one elastic band and has the same value for all transitions. Therefore the total duration - i.e. the length regarding the temporal aspect - of the resulting trajectory is determined at initialization.

The resulting timed elastic band $B = (Q, \Delta T)$ is then optimized in terms of spatial configurations with regard to a weighted multi-objective function $f(B)$, which is the weighted sum of components f_k described in the following. This optimization problem is represented as a hyper-graph (explained in detail in [11]) and solved using the `g2o` framework [3].

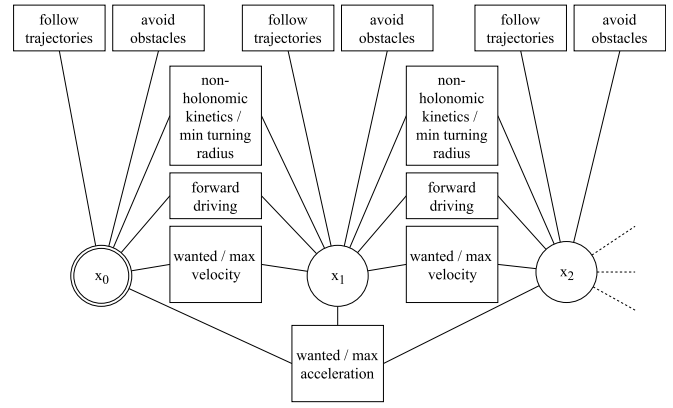


Fig. 2: Representation of the elastic band optimization problem as hyper-graph. Configurations of the band are represented as nodes (circles). The components of the objective function are represented as unary, binary or ternary edges (squares). The first configuration is fixed during the optimization process.

A. Objective Functions

As stated above, the final timed elastic band is optimized in terms of spatial configurations. At the core of this optimization process is the weighted multi-objective function (1). It calculates the weighted sum of several objective functions f_k with corresponding weights γ_k .

$$f(B) = \sum_k \gamma_k f_k(B) \quad (1)$$

All objective functions used for the experiments below are listed in Table I.

Each function f_k represents an objective or a constraint that is associated with one or more consecutive configurations \mathbf{x}_i . The constraints are formulated as objectives in terms of a “piecewise continuous, differentiable cost function that penalizes the violation of a constraint” [11].

Although in the approach used in this paper only the spatial configurations are optimized, it is still possible to formulate constraints regarding dynamic aspects, such as velocity, acceleration or moving obstacles, due to the time interval ΔT associated with the transition between two configurations.

Several components f_k used in this paper were proposed by Rösman et. al. in [11], i.e., restrictions for maximum translational velocity, maximum translational and rotational acceleration and non-holonomic kinematics. Also, existing functions from the implementation in [1] were used to penalize backward-driving, falling below the minimum turning radius and maintaining a specific velocity at the start and goal configuration. For the STEBLE approach presented in [15] more objective functions regarding the centripetal acceleration, dynamic obstacles and following other vehicles' trajectories were added. The objective functions mentioned above are described in detail in the respective papers.

For the experiments described in this paper the objective function regarding dynamic obstacles was modified, so that for the prediction of dynamic obstacles the rotational velocity

around the vertical axis (yaw rate) is taken into account (in addition to the linear velocity).

Furthermore, objective functions regarding the wanted velocity and acceleration were added. These are the same as the corresponding functions for maximum acceleration, but with lower thresholds and weights, as they represent constraints for comfortable driving and not the vehicle limits.

B. Selecting The Target Vehicle

A fundamental decision when following other vehicles is which one to follow. In the scope of this paper the general idea is to choose the vehicle whose state resembles most the current state of ego vehicle. Three criteria were chosen, i.e. distance, orientation, velocity. Also, the time we were already following a vehicle in previous iterations is considered. This discourages frequent changes of the target vehicle. Instead we evaluate the current state of observed vehicles and the previous state s_i on their respective trajectories closest to the current position of the ego vehicle. For this a tracking over time was implemented.

The input data for the trajectory planning module includes position, orientation, linear and angular velocity and acceleration of the ego vehicle as well as objects surrounding the ego vehicle. We only select objects classified as cars as the target vehicle. Each object o_{id} is identified over consecutive iterations of the planning module by a unique id.

Based on their id, the state $s_i = (t_i, \mathbf{y}_i, v_i)$ consisting of the timestamp t_i , the configuration $\mathbf{y}_i = (x_i, y_i, \beta_i)$ (two-dimensional position and orientation, resembling the configurations \mathbf{x}_i on the elastic band), and the velocity v_i is tracked over time and stored for each object o_{id} observed. While the velocity v_i can be calculated from the configurations and time stamps, it is stored for convenience. For each o_{id} a sequence of states $S_{id} = \{s_i\}_{i=0}^n$ is maintained. During the selection process, for each obstacle the state s_i closest to the current position of the ego vehicle is evaluated. For each of those closest states the following evaluation variables are calculated:

- dist** The euclidean distance from the current ego vehicle position. The distance is normalized in the range $[0, 1]$ in a way that the state closest to the ego vehicle is 1 and the one farthest away is 0.
- orient** The angular difference to the ego vehicle orientation. Normalized in the range $[0, 1]$, where matching orientations are 1 and opposite orientations are 0.
- vel** The correspondence with the current ego vehicle velocity. Similar to distance, velocity is normalized in the range $[0, 1]$ in a way that the (linear) velocity closest to the ego vehicle is 1 and the one with the largest absolute difference is 0.
- trust** The trust variable measures the time a vehicle was followed in previous iterations. For each second an object with a specific id is followed, the trust value is incremented by 0.1. For each second the object is not followed, the trust value is decremented by 0.1. The value is clamped in the range $[0, 1]$.

For each object o_{id} a weighted sum Ω_{id} is calculated based on the evaluation variables:

$$\Omega_{id} = \omega_0 * dist_{id} + \omega_1 * orient_{id} + \omega_2 * vel_{id} + \omega_3 * trust_{id} \quad (2)$$

All observed objects o_{id} are then stored in a priority list, sorted by their according Ω_{id} . The object with the highest priority, i.e. the highest value Ω_{id} , is then chosen as target vehicle o_{target} for the process of initializing the band, described below. If for some reason no trajectory can be generated for o_{target} , the next object in the priority list is chosen as target.

C. Initializing the Band

One of the most important steps, when using elastic bands for optimizing trajectories, is the initialization of the band, i.e. providing an initial band $B_{init} = (Q_{init}, \Delta T)$, with Q_{init} consisting of a sequence of initial configurations \mathbf{x}_i . This is mostly due to the property of elastic bands, to find only a local optimum. Thus, the initialization determines some properties of the optimized trajectory, e.g. if a specific obstacle is passed on the left or right. Furthermore, a initialization with a good estimate regarding the dynamic constraints for velocity and acceleration, reduces the number of iterations necessary to converge to the local optima. This is important, since the number of iterations is limited in the approach evaluated in this paper (compare Section III-D).

Since the general idea of the presented approach is to follow other vehicles' trajectories, an obvious choice for an initial sequence of configurations is the trajectory of the vehicle chosen as target vehicle (compare Section III-B). This trajectory has a high probability of being close to the current ego vehicle configuration \mathbf{x}_0 - especially if we were following that vehicle for some time - but there is no guarantee on the distance. As the elastic band is required to start at exactly at \mathbf{x}_0 , a procedure to generate an initial band, smoothly transitioning from \mathbf{x}_0 to the target vehicles' trajectory - in the following called target trajectory S_{target} - is described. An overview of this process is shown in Figure 3.

Filtering the Target Trajectory: The target trajectory S_{target} of the target vehicle o_{target} (i.e. a sequence of states s_i with respective configurations y_i , time stamps t_i , and velocities v_i) can be easily obtained, as it is maintained for the process of selecting the target vehicle (compare Section III-B). The problem is to identify, which states to use. States which are not in the direction of the orientation of \mathbf{x}_0 should be omitted, but only if they are within a certain distance. If they are far away, it is desirable for the ego vehicle to turn and reach the target vehicles' trajectory as soon as possible. On the other hand, the vehicle should not be required to turn away from the target trajectory or drive in loops, but instead prefer a direct transition, i.e. a simple or S-shaped curve in the direction of the target (compare Figure 5). Thus, to determine which states of S_{target} are feasible, the following algorithm is used:

Objective Function	Weight	Threshold	Description
non-holonomic kinetics	1000	n/a	Ensures that all configurations on the band follow the principle of non-holonomic kinetics. Each pair of adjacent configurations \mathbf{x}_i and \mathbf{x}_{i+1} is required to be on a common arc segment of constant curvature. These functions penalize the deviation of the position and orientation from such an arc segment, as well as the radius of the arc segment being below the specified threshold.
min turning radius	1000	4 m	
forward driving	1000	n/a	Enforces a “only forward driving” policy for the whole band. For each pair of consecutive configurations \mathbf{x}_i and \mathbf{x}_{i+1} , it is required that the second pose is “in front” of the first. If the dot product of the vector pointing from the location of \mathbf{x}_i to \mathbf{x}_{i+1} and the orientation vector of \mathbf{x}_i is negative, the return value is the squared dot product. Otherwise it is zero.
avoid obstacles	200	2 m	Repels the band from obstacles closer than the specified threshold. The function is associated with each configuration \mathbf{x}_i . Takes size of the obstacles (and ego vehicle) into account by representing them as line segments with associated radius (compare [15]). For each obstacle closer to the ego vehicle than the threshold distance, the amount of the euclidean distance falling below the threshold is squared and summed up as return value. For dynamic obstacles the location of the obstacle is predicted assuming constant velocity. The specific interval for which the obstacle has to be predicted for each configuration \mathbf{x}_i can be calculated by multiplying the corresponding index i with ΔT .
max velocity	100	determined on initialization	Keeps the maximum velocity on the band below the specified threshold. The function is associated with each pair of consecutive configurations \mathbf{x}_i and \mathbf{x}_{i+1} . The return value is the amount of the velocity exceeding the threshold, calculated from the two configurations and ΔT . The threshold in the experiments below was set to 110 % of the maximum velocity value of the initial elastic band (compare Section III-C).
max centripetal acceleration	100	2 m/s ²	Keeps the respective maximum acceleration on the band below the specified threshold. The respective functions are associated with each triple of consecutive configurations \mathbf{x}_i , \mathbf{x}_{i+1} and \mathbf{x}_{i+2} . The return value is the amount of the respective acceleration exceeding the threshold, calculated from the three spatial configurations and ΔT . Additionally, the first and last pair of configurations on the whole band are associated with a similar function based on set start and goal velocity.
max rotational acceleration	100	2 m/s ²	
max longitudinal acceleration	50	2 m/s ²	
follow trajectories	40	n/a	Attracts the band to the trajectories of other vehicles. This objective is associated with each configuration \mathbf{x}_i . The trajectories are represented by a set of line segments, stored in an R-tree for faster lookup. The return value is the euclidean distance between \mathbf{x}_i and the closest line segment.
wanted velocity	10	determined on initialization	Similar to the “max velocity” constraint, but penalizes deviation from the threshold in both directions (lower and higher). The value of the threshold is determined dynamically when the band is initialized, as it is used to influence the follow distance.
wanted centripetal acceleration	2	0 m/s ²	Same as the corresponding “max acceleration” functions, but with different thresholds and weights.
wanted rotational acceleration	2	0 m/s ²	
wanted longitudinal acceleration	1	0 m/s ²	

TABLE I: OBJECTIVE FUNCTIONS

- Find state \mathbf{s}_i which is closest to the current ego vehicle position and prune all states with smaller index (i.e. older) from the trajectory.
- For each remaining state $\mathbf{s}_i = (t_i, \mathbf{y}_i, v_i)$:
 - Get the maximum of the ego vehicle velocity and the velocity in state \mathbf{s}_i : $v_{max} = \max(v_{ego}, v_i)$
 - Calculate the radius r of a minimum turning circle with centripetal acceleration a_{max} with respect to v_{max} : $r = \frac{v_{max}^2}{a_{max}}$.
 - Construct a minimum turning circle C_x with radius r (from previous step) and centre point c (compare Figure 4). The position of c is chosen, so that the vector $\vec{x}\hat{c}$ from the position of \mathbf{x}_0 to c has a length of r (i.e. \mathbf{x}_0 is on the perimeter of the circle) and is perpendicular to the orientation vector \vec{o}_x of \mathbf{x}_0 . This results in two possible options for c : c_l and c_r , either orthogonal left or right of \mathbf{x}_0 with regard to the orientation vector \vec{o}_x . To determine which of c_l and c_r to use, another vector is needed: $\vec{x}\hat{y}$, pointing from the position of \mathbf{x}_0 the

position of \mathbf{y}_i . With $\vec{o}_x = (a_0, a_1)$ and $\vec{x}\hat{y} = (b_0, b_1)$, the sign of $d_x = a_0 \cdot b_1 - a_1 \cdot b_0$ determines if y (the point in front of \mathbf{y}_i) is to the left or right of \mathbf{x}_0 with regard to \vec{o}_x .

If $d > 0$, y is left of \mathbf{x}_0 and $c = c_l$ is chosen accordingly, otherwise $c = c_r$. The circle C_x

- Construct a circle C_y , similar to C_x , but with \mathbf{y}_i on the circle perimeter and d_y determining which of center c_l and c_r to use (orthogonal left or right of \mathbf{y}_i in distance r with respect to orientation \vec{o}_y), based on the vector $\vec{y}\hat{x} = -\vec{x}\hat{y}$, pointing from the position of \mathbf{y}_i to the position of \mathbf{x}_0 (compare Figure 4).
- If the two circles C_x and C_y do intersect, the configuration \mathbf{y}_i is not feasible, i.e. it cannot be reached with centripetal acceleration below the given a_{max} at the given velocity, without turning away from the target trajectory at some point. If \mathbf{y}_i is not feasible, prune it from the trajectory and proceed with the next state \mathbf{s}_{i+1} . If \mathbf{y}_i is feasible, the loop can be left and all remaining states used for the initial elastic band.




Fig. 3: Overview of the initialization of the elastic band. (a) depicts the target (blue) and the according sequence of tracked states. The ego vehicle is gray. (b) All states before the one closest to the ego vehicle and within 180° of the ego vehicles orientation are discarded. (c) For each remaining state two minimum turning circles are constructed. One is attached to the ego vehicle's position (gray) and one to the respective state (blue). All states before the first one, where the circles do not intersect, are discarded. (d) An additional state is predicted close to the ego vehicle and a cubic spline is constructed, using the remaining states as support points. (e) The resulting spline is sampled to generate the configurations of the initial band.

Some examples for the resulting circles are depicted in Figure 5. The result of the process of selecting feasible states from the target trajectory S_{target} is a sequence $S_{feasible}$, which consists of all states s_i with index larger than some state s_{cf} . s_{cf} is the first state of S_{target} (starting from the oldest one) which is feasible according to the algorithm stated above and has a larger index than the state closest to the ego vehicle.

If the resulting trajectory $S_{feasible}$ has less than two configurations, it is rejected, as at least two configurations are needed for the subsequent steps. In this case, the next object o_{id} in the priority list is chosen as o_{target} (compare Section III-B) and the process of filtering the target trajectory is repeated.

Sampling With Constant Time Intervals: As the elastic band $B = (Q, \Delta T)$ is defined by the constant time interval ΔT between adjacent configurations, the initial elastic band B_{init} should reflect this property. The filtered trajectory $S_{feasible}$ generated in the previous step does not necessarily guarantee this, as the sample rate of the states of the observed objects may be lower due to constraints of the sensors or post-processing steps. Additionally, the distance between the first configuration x_0 to the first feasible configuration y_0 of $S_{feasible}$ is variable, i.e. it is highly probable that the according (estimated) time interval ΔT_0 does not equal




Fig. 4: In the process of filtering the target trajectory S_{target} , two minimum turning circles C_x and C_y are constructed and placed to the left or right of the respective configuration, to evaluate if the configuration y_i can be reached from the configuration x_0 without a maneuver requiring to recede from the target configuration.

ΔT . For those reasons, a cubic spline (satisfying continuity requirements up to the second derivative) is constructed for each of the two spatial dimensions over time, using the respective values from x_0 and $y_i \in S_{feasible}$ as support points, as described in the following.

In a first step, the configuration x_1 is estimated by predicting x_0 by ΔT based on the current linear and angular velocity of the ego vehicle. This additional configuration provides some more stability to the constructed spline, as the uncertainty in the measurements of the configurations y_i may introduce some oscillation.

The next step is to estimate the time interval ΔT_1 between x_1 and y_0 . It is derived from the estimated travel-distance s and the average velocity v_a between those two configurations. While v_a can be calculated directly from the current velocity of the ego vehicle and the respective value stored in $S_{feasible}$, s is estimated by the arc length of a circular segment connecting x_1 and y_0 . The circular segment is defined by the chord length c , which is the euclidean distance between the positions of the two configurations, and the angle difference $\Delta \alpha$ between the respective orientations. s , i.e. the arc length of the circular segment, can be calculated as follows:

$$s = \left| \frac{\Delta \alpha \cdot c}{(2 * \sin(\frac{\Delta \alpha}{2}))} \right| \quad (3)$$

With this, $\Delta T_1 = s * v_a$ can be calculated. According to ΔT_1 (and $\Delta T_0 = \Delta T$ between x_0 and x_1) the time stamps t_i corresponding to the respective states $s_i \in S_{feasible}$ are updated, so that they are relative to the current time, i.e. the start time of the elastic band.

The sequence of pairs of configurations and time stamps $\{(x_0, 0), (x_1, \Delta T), (y_0, t_0), \dots, (y_n, t_n)\}$ is then used to construct two cubic splines S_x and S_y (one for each spatial dimension). The support points consist of the coordinate in the respective dimension and the according time stamp. As stated above, the splines are continuous in the first and second derivative by construction.

B_{init} is then obtained by sampling a number of initial

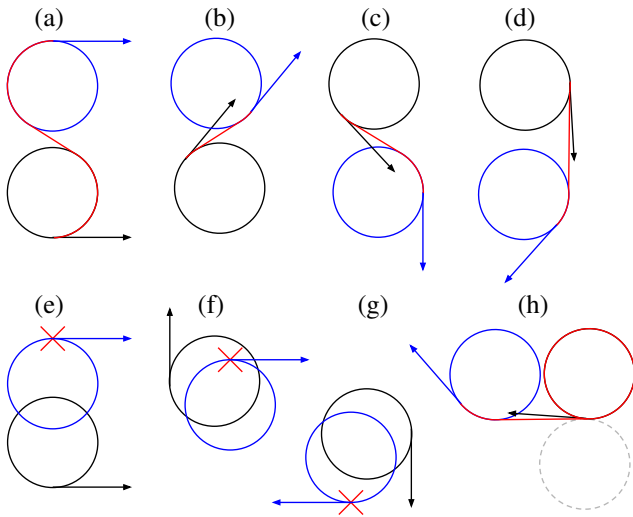


Fig. 5: Examples for the evaluation of feasible target configurations. The orientation vector \vec{o}_x and the minimum turning circle C_x are shown in black, the target configurations' \vec{o}_y and C_y in blue, respectively. The assumed trajectory transitioning between the configurations is red.

If the two circles do not intersect, a direct transition is possible, i.e. without receding from the target ((a)-(d)). If the circles intersect ((e),(f)), it would be required to initially move away from the target, so the respective configurations are omitted from the filtered target trajectory. (g) depicts a rare case, where the circles intersect, but a direct transition is possible. This is accepted in favor of a more simple algorithm.

On the other hand, (h) depicts a case, where the circles do not intersect, but receding from the target would be required (if following the perimeter of C_x). This is not a problem, since for all those cases a circle C'_x (dotted gray) on the opposite side of \vec{o}_x can be constructed, which also does not intersect with C_y . When following the perimeter of C'_x , a direct transition is possible.

configurations $\mathbf{x}_i = (x_i, y_i, \beta_i)$ from the splines S_x and S_y with a step size of ΔT . The values for x_i and y_i are the values of the splines at the respective time. The value for β_i can be calculated from the first derivatives of the splines, as those can be interpreted as the components of the respective orientation vector. The number of configurations sampled can be limited (compare Section III-D) to reduce the computational complexity of the optimization.

D. Parameters, Thresholds and Weights

There are two major parameters influencing the computational complexity of the optimization problem and thus the run time of the algorithm: The number of configurations n and the number of iterations for optimizing the graph (compare [3]). The number of configurations is fixed during optimization and is determined at initialization by the value for ΔT and the maximum overall duration of the band. In the context of this paper ΔT was set to 0.2 s. This is the maximum value, which still guarantees that no object can “slip” between two consecutive configurations at maximum velocity in an urban environment (taking into account a threshold of 2 m for minimum obstacle separation). The resulting separation of two configurations at 60 km/h is ~ 3.3 m. The maximal duration T_{max} of the elastic band was set to 5 s which still enables foresight regarding evasive maneuvers. This results in a maximum of $n = 25$ configurations per elastic band. Due to limitation of computational

power, the number of iterations for the optimization process was set to 50, although a higher value would be preferable.

All thresholds and weights used for the objective functions are given in table I. The weights are reflecting the priority of the objectives. Kinematic constraints have the highest priority, since it makes no sense to generate trajectories, the ego vehicle cannot follow. Obviously, avoiding obstacles also has a high priority, but in favor of a drivable trajectory it can be accepted to violate the safety margin. The maximum centripetal and rotational acceleration are weighted higher than the longitudinal acceleration, i.e. braking is preferred for obstacle avoidance. Objectives regarding the wanted acceleration and velocity have lower priority, since they can be regarded as comfort features. The weight for following a trajectory can be adjusted to control how strictly other vehicles' trajectories are followed.

The thresholds for the minimal turning radius and maximum acceleration are directly related to the vehicle dynamic and acceptable comfort for passengers, respectively. The threshold for obstacle avoidance was set to 2 m, reasoning on this can be found in [15]. Thresholds for wanted acceleration were set to 0 m/s², as we want to prevent any acceleration, when not necessary.

Two of the thresholds are not set to static values, but are determined during the initialization process: the maximum velocity v_{max} and the wanted velocity v_{wanted} . They are set after the initial band B_{init} is constructed. In most cases the maximum velocity is regulated by traffic rules - in urban scenarios it is usually much lower than the maximum velocity the ego vehicle can reach. In the presented approach we do not want to make any assumptions on the traffic rules, but instead adapt to the behaviour of other drivers. Therefore the threshold for maximum velocity is set relative to the velocities on the initial band B_{init} , which incorporates the velocity of the target vehicle. In the experiments below, v_{max} was set to 110% of the maximum velocity on B_{init} .

The other threshold determined during the initialization - the wanted velocity v_{wanted} - is set based on the current velocity v_{target} of the target O_{target} and the distance d between the ego vehicle and O_{target} . It directly influences the velocity on the optimized band, so it is used to catch up to O_{target} if d is larger than a specified follow-distance d_{follow} , and fall back if d is smaller, respectively. In the experiments below, v_{wanted} is calculated according to the following equations:

$$\begin{aligned} d_{follow} &= \max(5 \text{ m}, v_{ego} \cdot 1 \text{ s}) \\ v_{offset} &= 0.1 \text{ s}^{-1} \cdot (d - d_{follow}) \\ v_{wanted} &= \min(v_{max}, v_{target} + v_{offset}) \end{aligned} \quad (4)$$

d_{follow} is set to the amount of the current velocity of the ego vehicle (in m/s). This corresponds to following the target with a temporal distance of 1 s. Also, d_{follow} cannot fall below a specified minimum follow-distance of 5 m. The value of v_{wanted} is limited from above, so that it cannot exceed v_{max} .

IV. EXPERIMENTAL RESULTS

Experiments on the general validity of the approach described in this publication were conducted in a simulation environment combined with live data. While the information on observed objects is taken from recorded data of the autonomous vehicle “MadeInGermany”, the trajectory of the ego vehicle was simulated.

The examined data set consists of multiple transits of the roundabout “Großer Stern” in the urban center of Berlin, Germany. It has a total length of ca. 39 minutes. An example of the observed objects is shown in Figure 6.




Fig. 6: A typical scene of the object data recorded at “Großer Stern”, Berlin. The ego vehicle is white. Objects classified as car are shown in orange, their trajectories red. Other objects are colored in yellow.

The trajectory of the ego vehicle was simulated assuming that the car follows perfectly the planned trajectory. The pose of the simulated ego vehicle was reset to the real vehicle’s pose, if the respective positions deviate more than 20 m. This was done so that the simulated ego vehicle is always close to the center of observation.

In the whole data set an average of 6.3 other vehicles were observed. An average of 2.8 objects allowed the generation of an initial trajectory, i.e. was a valid target. In 80.1% of the iterations of the trajectory planning module at least one target was available. The absence of valid targets occurred mostly when the real car was doing a U-turn to reenter the roundabout (what no other drivers were doing) and when standing in the front row at a traffic light.

Table II shows the results for six transits of the roundabout, where the simulated trajectory was not resetted (i.e. the simulated vehicle was always within a range of 20 m of the real vehicle) and there was always at least one valid target to follow available. For measuring the quality of the trajectories, three values are most important: the longitudinal acceleration, the lateral acceleration, and the distance to obstacles.

The results show that the values for the simulated vehicle and the human driven real vehicle are quite comparable. On the other hand, the thresholds specified in the respective objective functions (2 m/s^2 maximum acceleration and 2 m distance to obstacles) are exceeded in many cases. This is expected to some degree, as the objective functions do not represent hard constraints, but only values above the thresholds are penalized. Nonetheless, there are some very high values for the longitudinal acceleration in scenario 1 and 6. This correlates with very low minimal distances to obstacles in those scenarios. The reason is that the elastic

band gets stuck in a local minimum between two observed vehicles (compare Figure 7).

This is an inherent problem of using fixed time intervals between the configurations and a fixed number of configurations. In the whole data set about 2.1% of the generated trajectories have maximum acceleration or minimum distances exceeding the respective thresholds by more than 150%, which is a clear sign of configurations stuck during optimization.




Fig. 7: Example of the elastic band being stuck in a local minimum. Two objects enclosing the band at both sides are shown at two different predicted states. The corresponding configurations of the elastic band are highlighted with a circle in the respective color. They do not move during further iterations of the optimization, as the cost function repelling them from the obstacles and the cost functions attracting them to each other (due to max acceleration and velocity) are in balance. The stuck configurations prevent the trajectory from slowing down to avoid the obstacles, as for such a maneuver it would be necessary to pull more configurations closer to the ego vehicle.

To examine the quality of the proposed initialization process, it was compared to initializing the band with configurations sampled on a straight line to the target with linear velocity adjustment. Figure 8 shows that the average distance, the configurations were moved during the optimization process, is clearly reduced, i.e. the final state of the band is reached with less iterations. On the other hand, when taking only the longitudinal position change into account (i.e. the optimization of the velocity), the simple initialization is closer to the final state. This is due to the proposed initialization taking into account the velocity of the target along its trajectory, which is then smoothed out again.

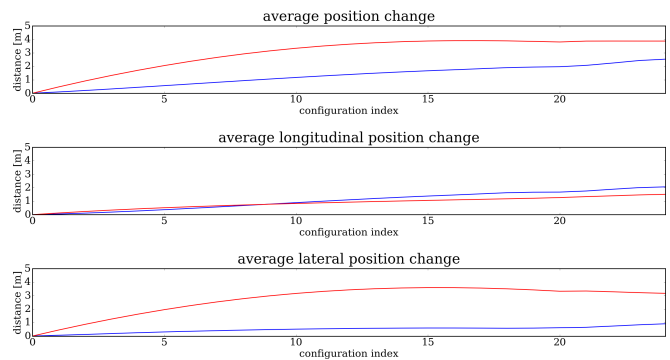


Fig. 8: Comparison of the proposed initialization (blue graphs) with a simple initialization sampled on a straight line to the target with linear velocity adjustment (red graphs). The graphs represent the average distance each configuration c_i was moved during the optimization process. The longitudinal and lateral distance were measured with respect to the orientation of the respective initial configuration.

V. CONCLUSIONS

The experimental results show that the approach of following other cars using elastic bands is valid to compensate for temporary unavailable maps. The proposed initialization process is superior to a naive initialization, as significantly

Scenario	cars observed (max/min/avg)	valid targets (max/min/avg)	target changes	velocity [$\frac{m}{s}$] (max/min/avg)	long. acc [$\frac{m}{s^2}$] (max/avg)	lat. acc [$\frac{m}{s^2}$] (max/avg)	dist to obst. [m] (max/min/avg)
1	18 / 6 / 11.2	10 / 1 / 5.5	2	13.9 / 0.0 / 8.8	4.3 / 1.2	2.1 / 0.8	10.4 / 0.2 / 7.0
human				10.8 / 0.0 / 7.5	3.2 / 0.7	2.8 / 0.7	11.4 / 1.8 / 4.0
2	10 / 4 / 7.1	7 / 1 / 4.1	0	13.5 / 0.0 / 9.0	3.4 / 0.6	2.5 / 1.0	10.5 / 0.6 / 4.8
human				14.0 / 0.0 / 8.4	3.0 / 0.6	5.4 / 1.0	5.5 / 0.4 / 3.1
3	9 / 5 / 6.8	5 / 2 / 2.4	0	12.5 / 7.5 / 11.3	2.2 / 0.5	1.7 / 0.7	14.9 / 1.7 / 8.1
human				13.4 / 4.2 / 10.7	1.9 / 0.5	2.8 / 1.0	17.0 / 3.2 / 9.4
4	14 / 4 / 5.6	6 / 1 / 2.7	0	13.8 / 6.9 / 10.9	2.0 / 0.7	1.8 / 0.7	6.8 / 1.1 / 5.1
human				13.8 / 4.7 / 10.6	3.6 / 0.8	3.7 / 0.6	7.6 / 0.5 / 3.9
5	18 / 2 / 9.4	12 / 1 / 4.4	4	12.3 / 0.0 / 7.9	3.5 / 0.8	2.9 / 0.8	11.5 / 0.9 / 6.5
human				12.3 / 0.0 / 8.6	3.9 / 0.7	3.8 / 1.1	11.2 / 1.2 / 5.8
6	15 / 5 / 11.1	9 / 1 / 6.1	3	13.5 / 0.0 / 8.7	4.8 / 0.9	1.9 / 0.7	10.1 / 0.1 / 5.3
human				11.7 / 0.0 / 7.5	3.8 / 0.9	2.6 / 0.9	8.6 / 0.5 / 4.8

TABLE II: COMPARING GENERATED TRAJECTORIES TO HUMAN DRIVER TRAJECTORIES

less iterations are required to achieve a locally optimal solution.

Nonetheless there are some shortcomings which need to be addressed in future work. The concept of elastic bands does not guarantee a collision free trajectory within the maximal bounds for vehicle dynamics. This problem is even amplified when using fixed time intervals, due to configurations stuck in local minima during the optimization. This could be mitigated to some degree by using more than one initial band and hierarchically selecting a subset of the bands for further optimization.

Also, more elaborated strategies for selecting the target vehicle need to be developed. In this paper a rather simple strategy was used, ignoring the global route or destination of the ego vehicle. The destination of the observed vehicles could be estimated or acquired with C2X communication.

Another important aspect is the prediction of other vehicles' trajectories. The naive approach of assuming constant velocity or acceleration is often inaccurate in complex urban scenarios. Obviously the prediction could be improved by directly communicating planned trajectories between cars. Alternatively better heuristics should be used, e.g. other drivers are also following observed vehicles' trajectories in most scenarios.

ACKNOWLEDGMENT

This work was supported by the Deutsche Forschungsgemeinschaft (DFG), SPP 1835 - Kooperativ interagierende Automobile and by the Bundesministerium für Bildung und Forschung (BMBF) - Intelligente und effiziente Elektromobilität der Zukunft (e-Mobilize).

REFERENCES

- [1] Teb local planner package for ros. http://wiki.ros.org/teb_local_planner. Accessed: 2017-01-10.
- [2] Stanfords Robotic Vehicle "Junior." Interim Report. Technical report, Stanford University, CA, USA, 2007.
- [3] G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g2o: A general framework for graph optimization. In *IEEE International Conference on Robotics and Automation*, 2011.
- [4] T. Gu, J. M. Dolan, and J.-W. Lee. Automated tactical maneuver discovery, reasoning and trajectory planning for autonomous driving. In *IEEE International Conference on Intelligent Robots and Systems*, pages 5474–5480, October 2016.
- [5] M. Keller, F. Hoffmann, C. Hass, T. Bertram, and A. Seewald. Planning of optimal collision avoidance trajectories with timed elastic bands. *IFAC Proceedings Volumes*, 47(3):9822–9827, 2014.
- [6] M. Likhachev, G. Gordon, and S. Thrun. Ara*: Anytime a* with provable bounds on sub-optimality. *Advances in Neural Information Processing Systems (NIPS)*, 16, 2016.
- [7] R. Linker and T. Blass. Optimal path planning for car-like off-road vehicles. In *Conference on Robotics, Automation and Mechatronics, 2008*. IEEE, 2008.
- [8] S. Quinlan and O. Khatib. Elastic bands: Connecting path planning and control. In *Robotics and Automation, 1993. Proceedings., 1993 IEEE International Conference on*, pages 802–807. IEEE, 1993.
- [9] Rasekhipour, Yadollah, Chen, Shih-Ken, Khajepour, Amir, and Litkouhi, Baktiar. A potential field-based model predictive path-planning controller for autonomous road vehicles. 2016.
- [10] C. Rösmann, W. Feiten, T. Wösch, F. Hoffmann, and T. Bertram. Trajectory modification considering dynamic constraints of autonomous robots. In *Robotics; Proceedings of ROBOTIK 2012; 7th German Conference on*, pages 1–6. VDE, 2012.
- [11] C. Rösmann, F. Hoffmann, and T. Bertram. Planning of multiple robot trajectories in distinctive topologies. In *Mobile Robots (ECMR), 2015 European Conference on*, pages 1–6. IEEE, 2015.
- [12] C. Rösmann, F. Hoffmann, and T. Bertram. Timed-elastic-bands for time-optimal point-to-point nonlinear model predictive control. In *Control Conference (ECC), 2015 European*, pages 3352–3357. IEEE, 2015.
- [13] S. Rotter. Swarm behaviour for path planning. Master's thesis, Freie Universität Berlin, 2014.
- [14] F. Ulbrich, D. Goehring, T. Langner, Z. Boroujeni, and R. Rojas. Stable timed elastic bands with loose ends. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 186–192, June 2017.
- [15] M. Wang. *A Cognitive Navigation Approach for Autonomous Vehicles*. mbv, 2012.
- [16] J. Ziegler and C. Stiller. Spatiotemporal state lattices for fast trajectory planning in dynamic on-road driving scenarios. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1879–1884, Oct 2009.