Long-Term Prediction of Vehicle Behavior Using Short-Term Uncertainty-Aware Trajectories and High-Definition Maps

Sai Yalamanchi, Tzu-Kuo Huang, Galen Clark Haynes, Nemanja Djuric

UBER ATG

The 23rd IEEE International Conference on Intelligent Transportation Systems
September 20 – 23, 2020

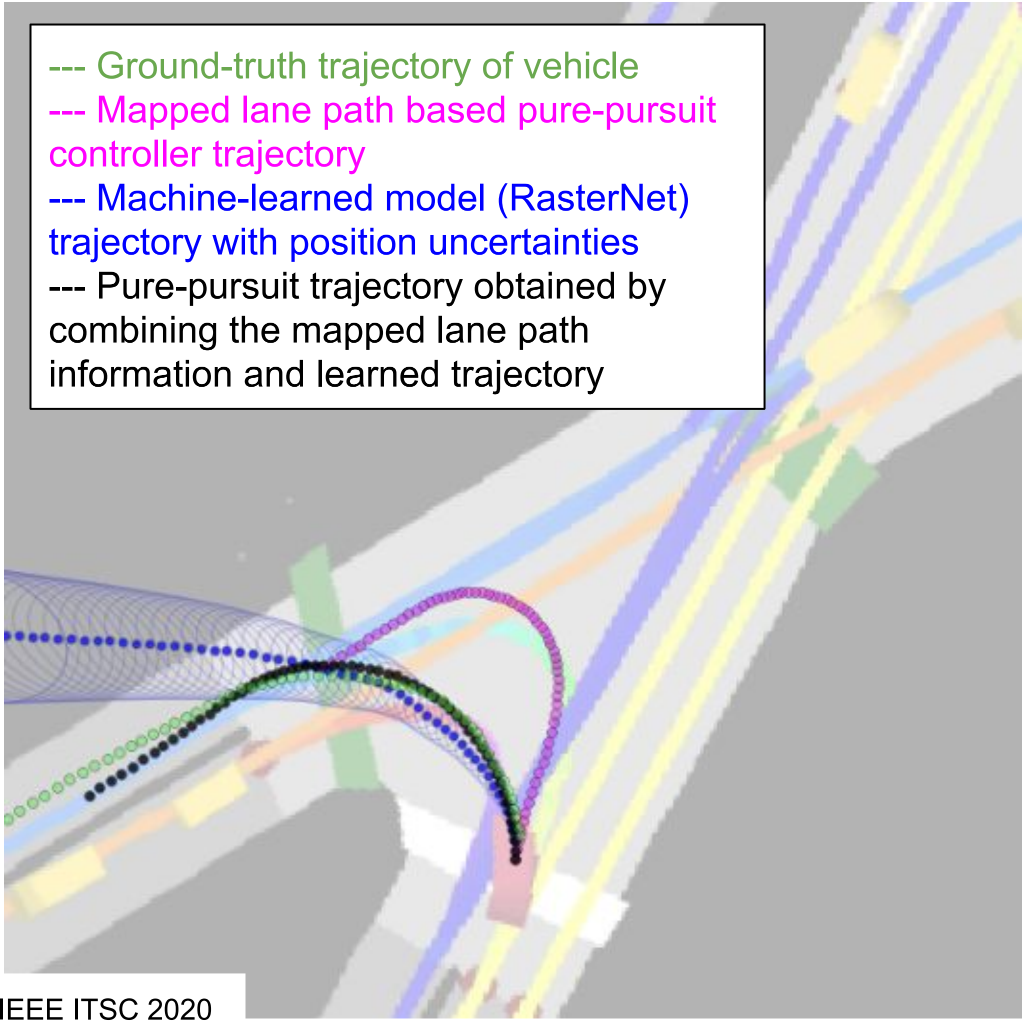

Trajectory Prediction For Autonomous Driving

Kinematic controller-based trajectories tracking lane paths tend to be less accurate in the short term (e.g., the trajectory cuts too wide in the figure).

Machine-learned trajectories tend to be less accurate in the long term (e.g., the trajectory overshoots the road in the figure).

In this work, we focus on combining mapped lane path information with machine-learned trajectories in an uncertainty-aware framework to produce a <u>spatial path</u> that more accurately represents motion patterns of vehicles.

Framework

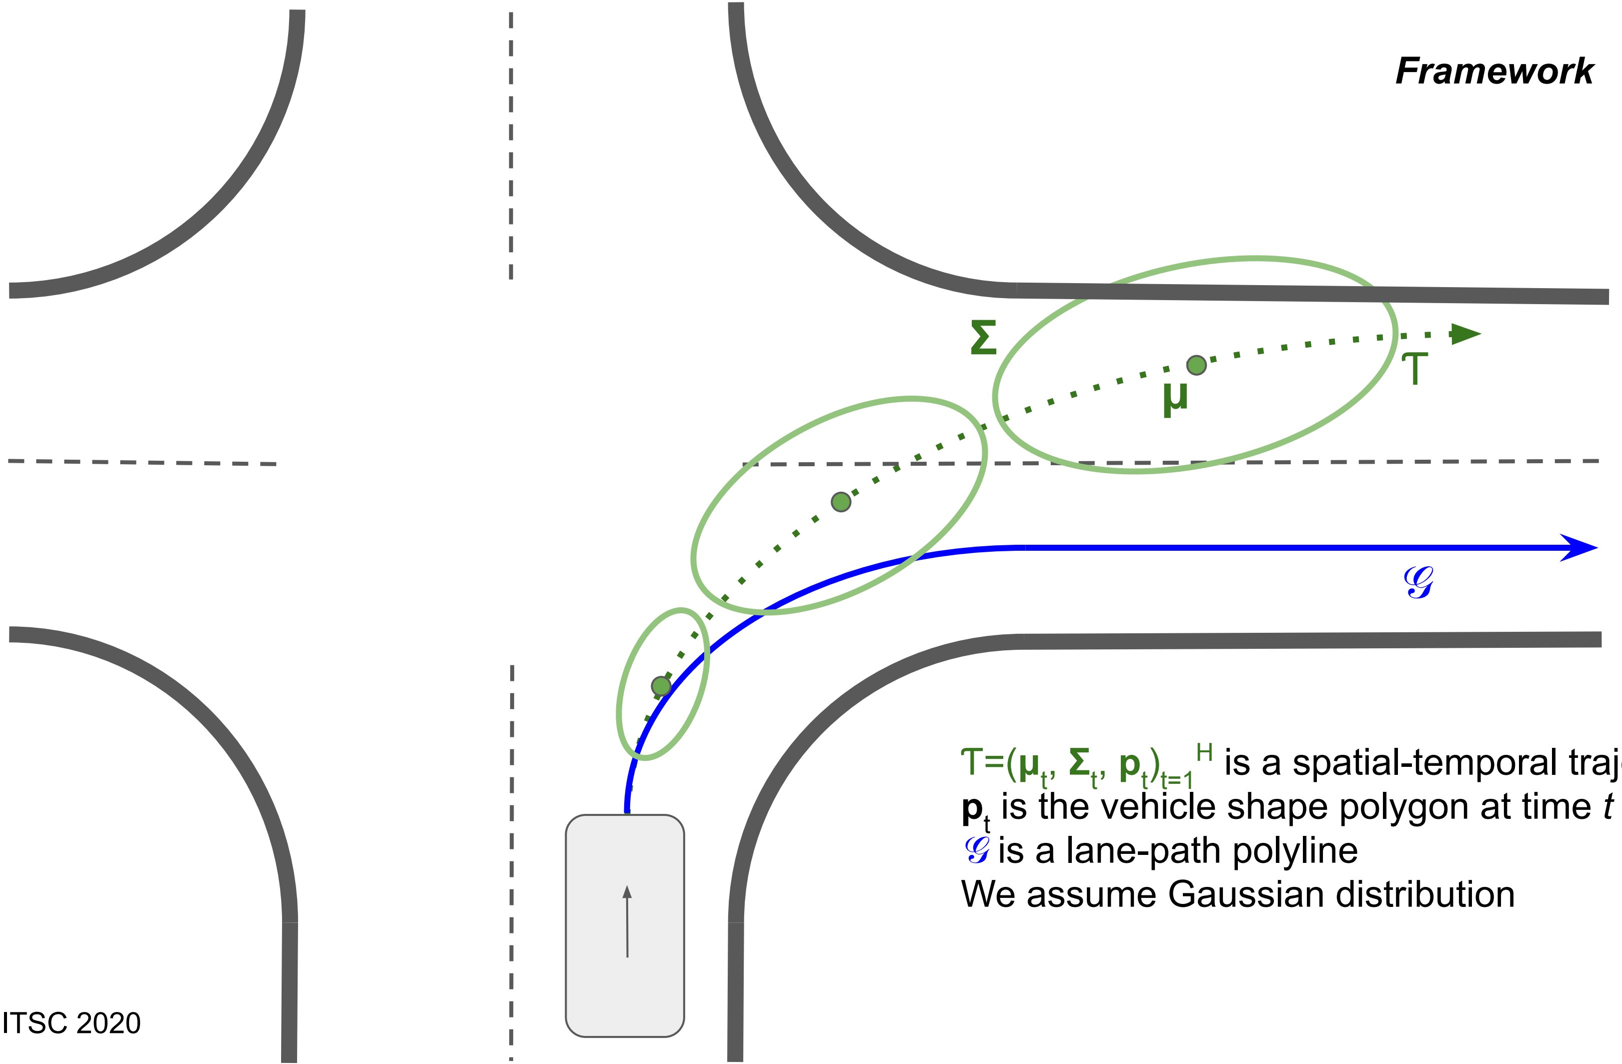

$\mathcal{T}=(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t, \mathbf{p}_t)_{t=1}^{H}$ is a spatial-temporal trajectory

$\mathbf{p}_t$ is the vehicle shape polygon at time t

$\mathcal{G}$ is a lane-path polyline

We assume Gaussian distribution

Motivation

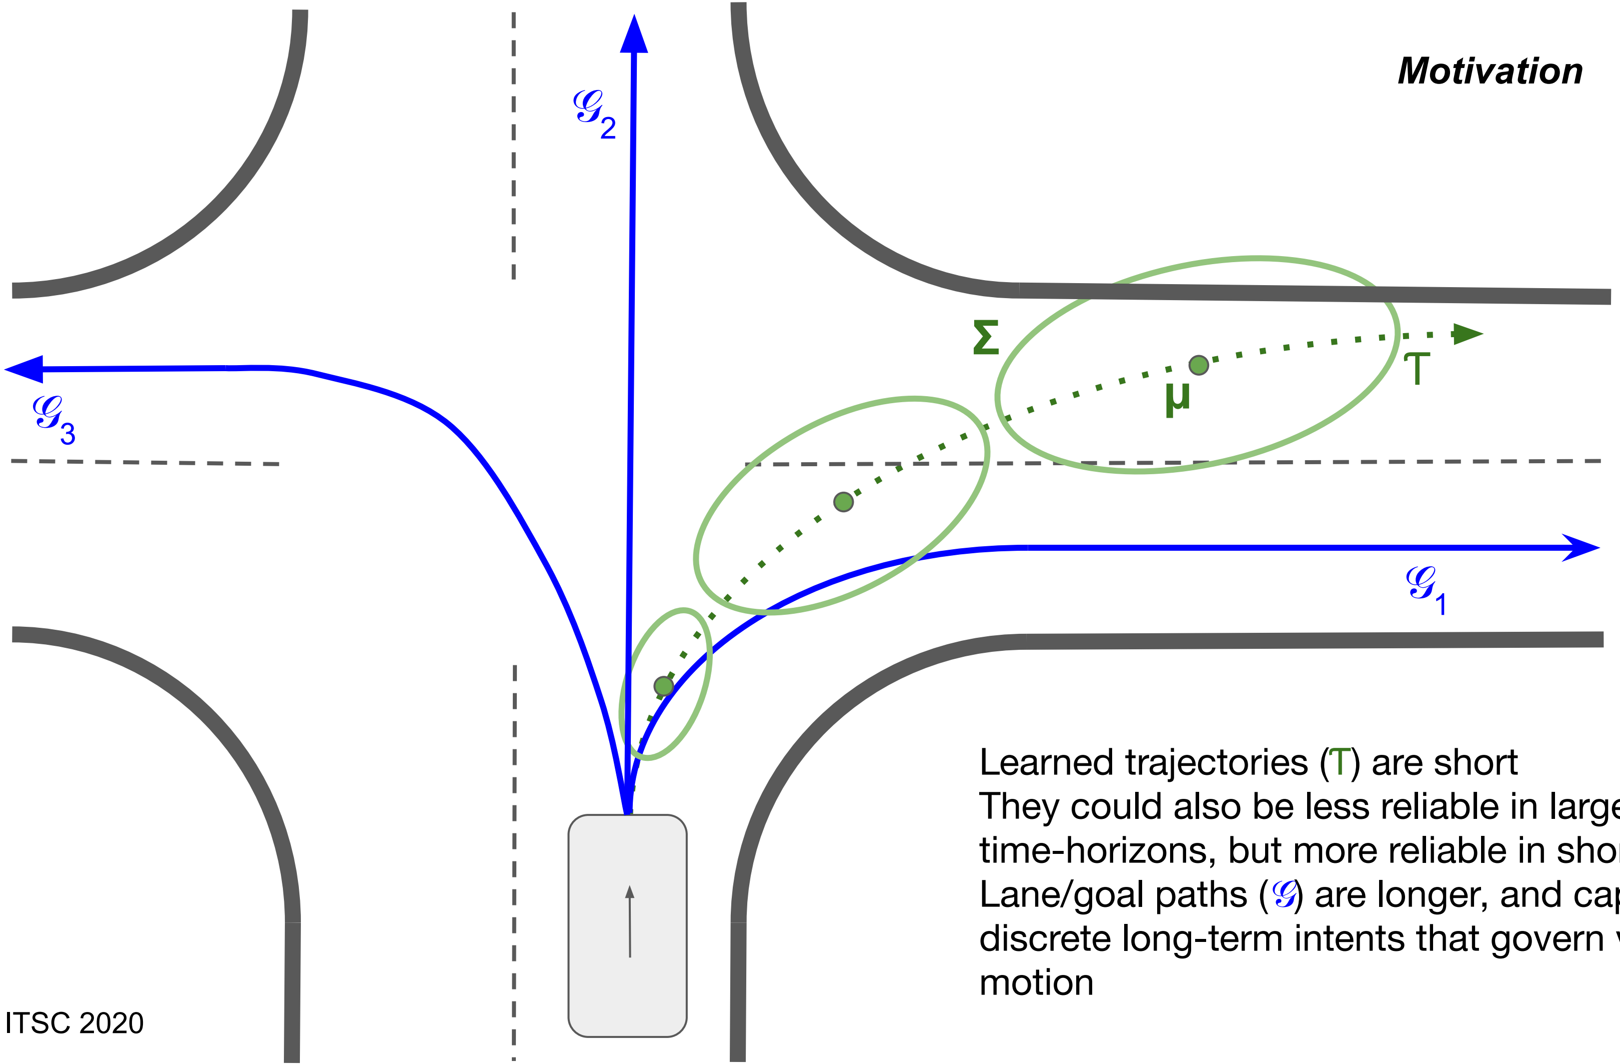

Learned trajectories ($\mathcal{T}$) are short
They could also be less reliable in large time-horizons, but more reliable in short-term
Lane/goal paths ($\mathcal{G}$) are longer, and capture discrete long-term intents that govern vehicle motion

Problem Description

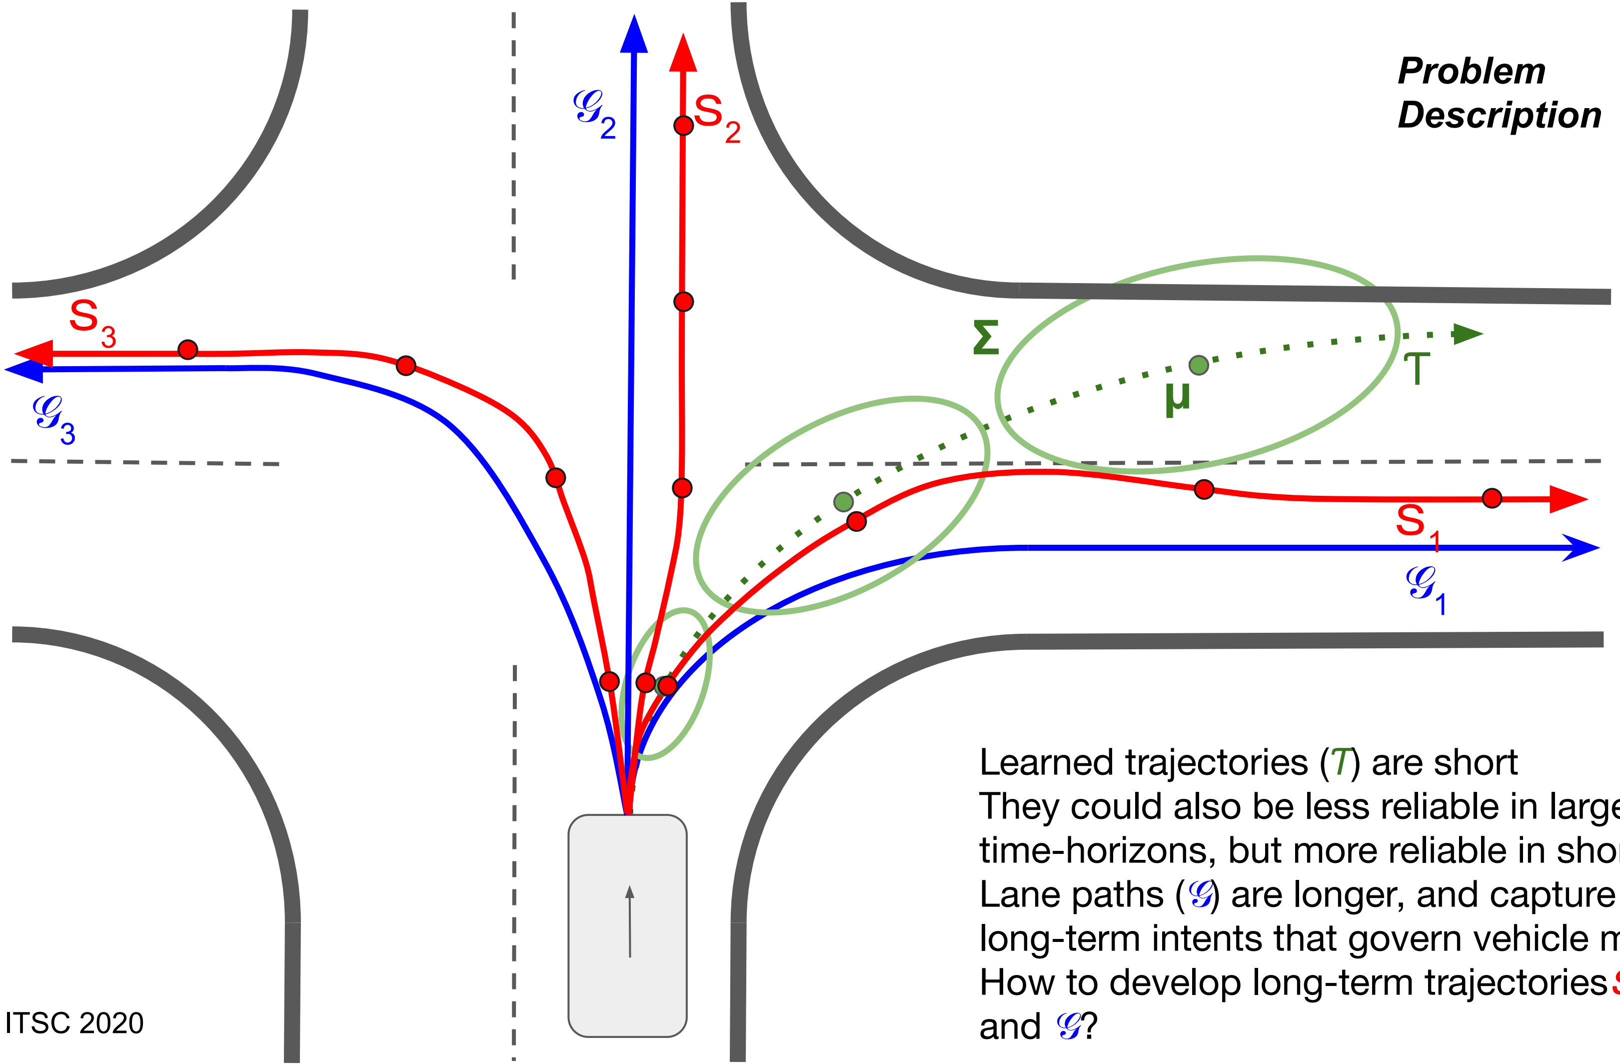

Learned trajectories ($\mathcal{T}$) are short
They could also be less reliable in large time-horizons, but more reliable in short-term
Lane paths ($\mathcal{G}$) are longer, and capture discrete long-term intents that govern vehicle motion
How to develop long-term trajectories $\boldsymbol{S}$ from $\mathcal{T}$ and $\mathcal{G}$?

Desired Properties of S

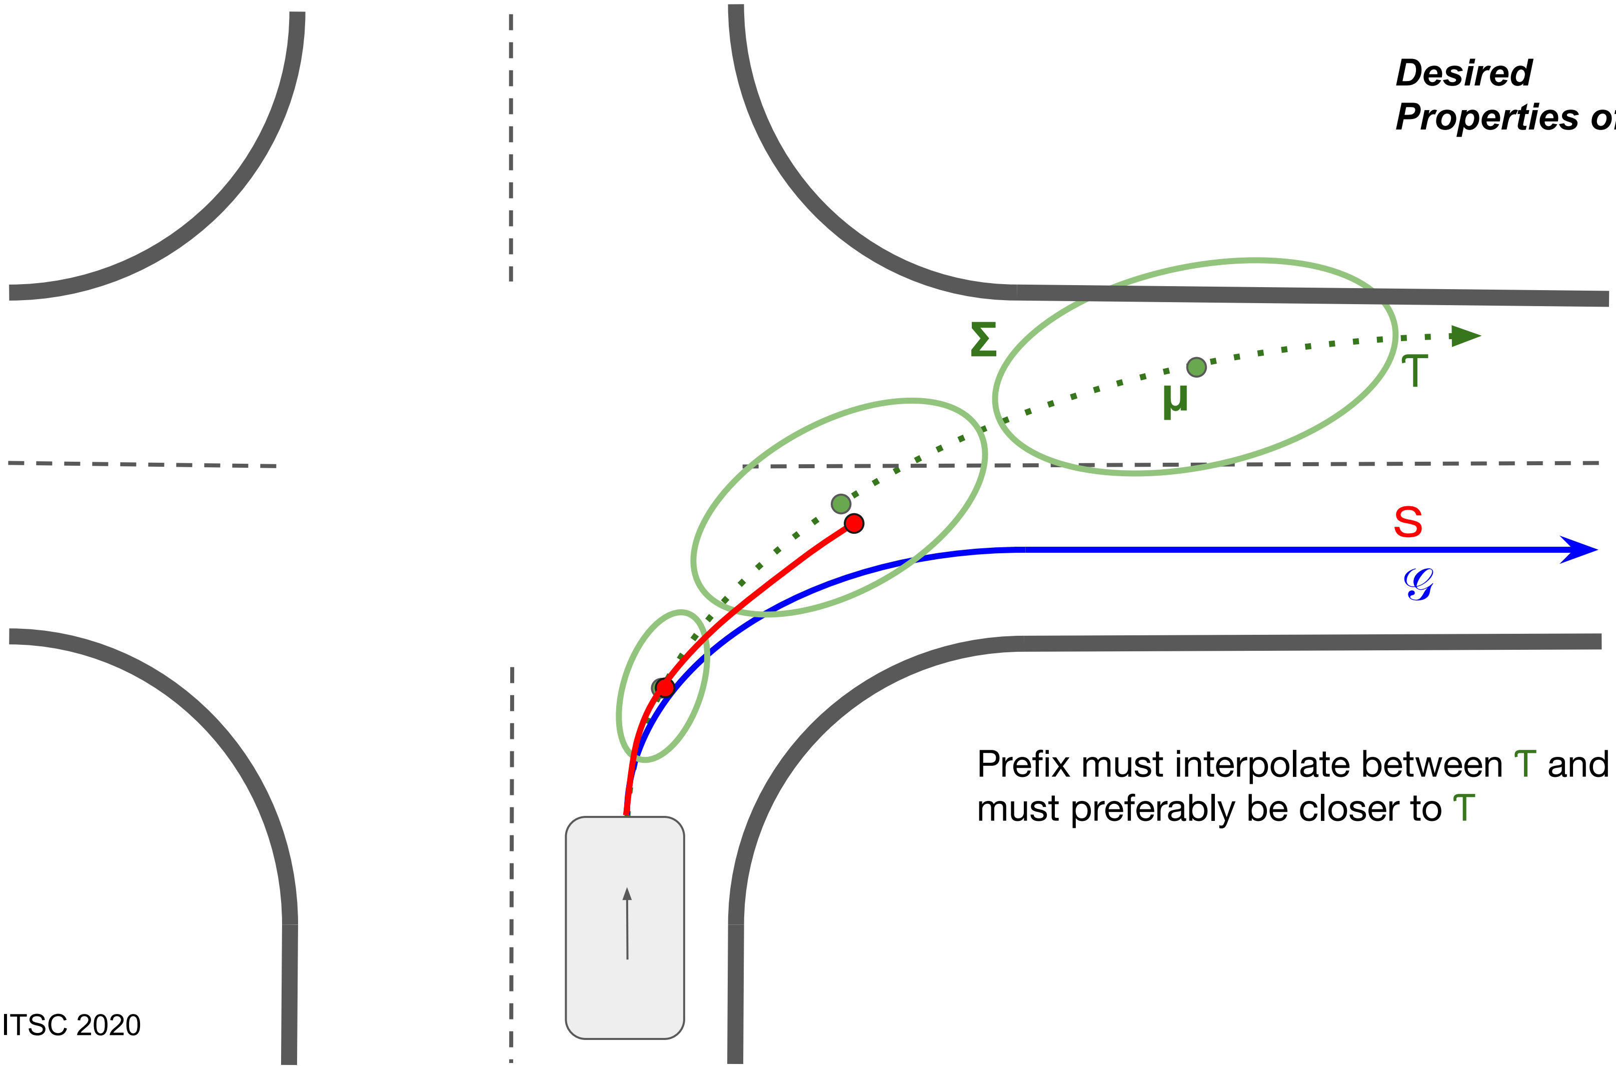

Prefix must interpolate between $\mathcal{T}$ and $\mathscr{G}$ but must preferably be closer to $\mathcal{T}$

Desired Properties of S

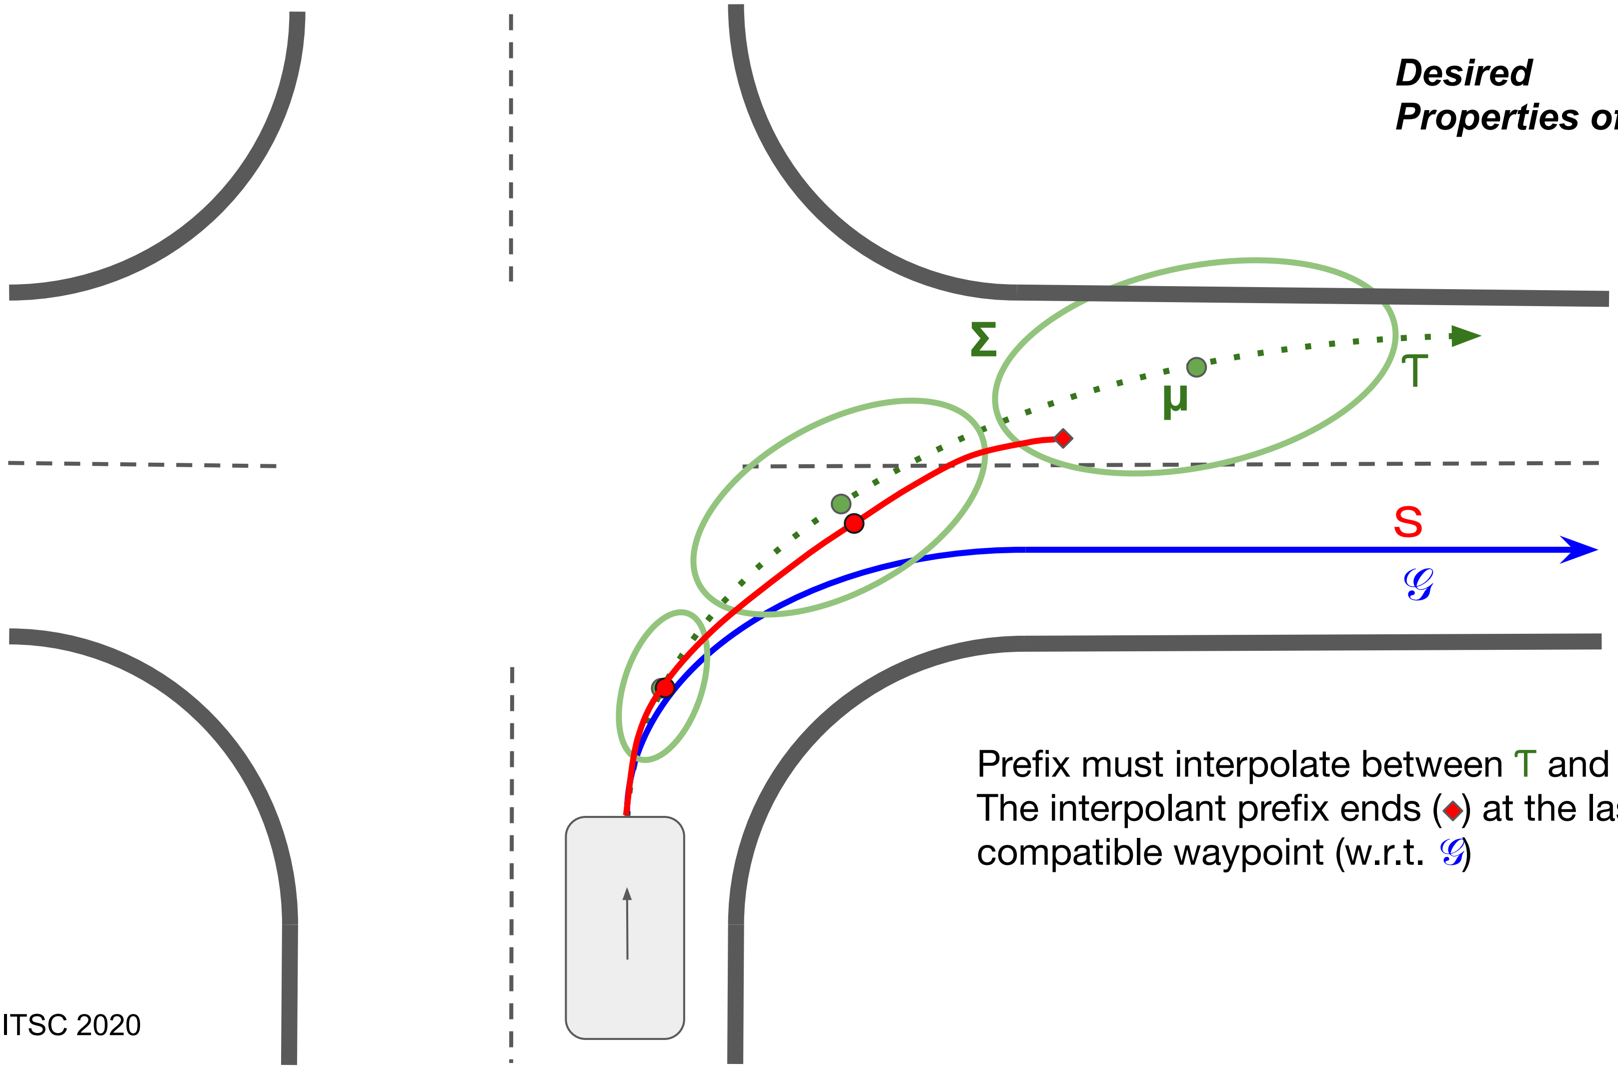

Prefix must interpolate between T and 𝒢
The interpolant prefix ends (◆) at the last compatible waypoint (w.r.t. 𝒢)

***Desired Properties of** S*

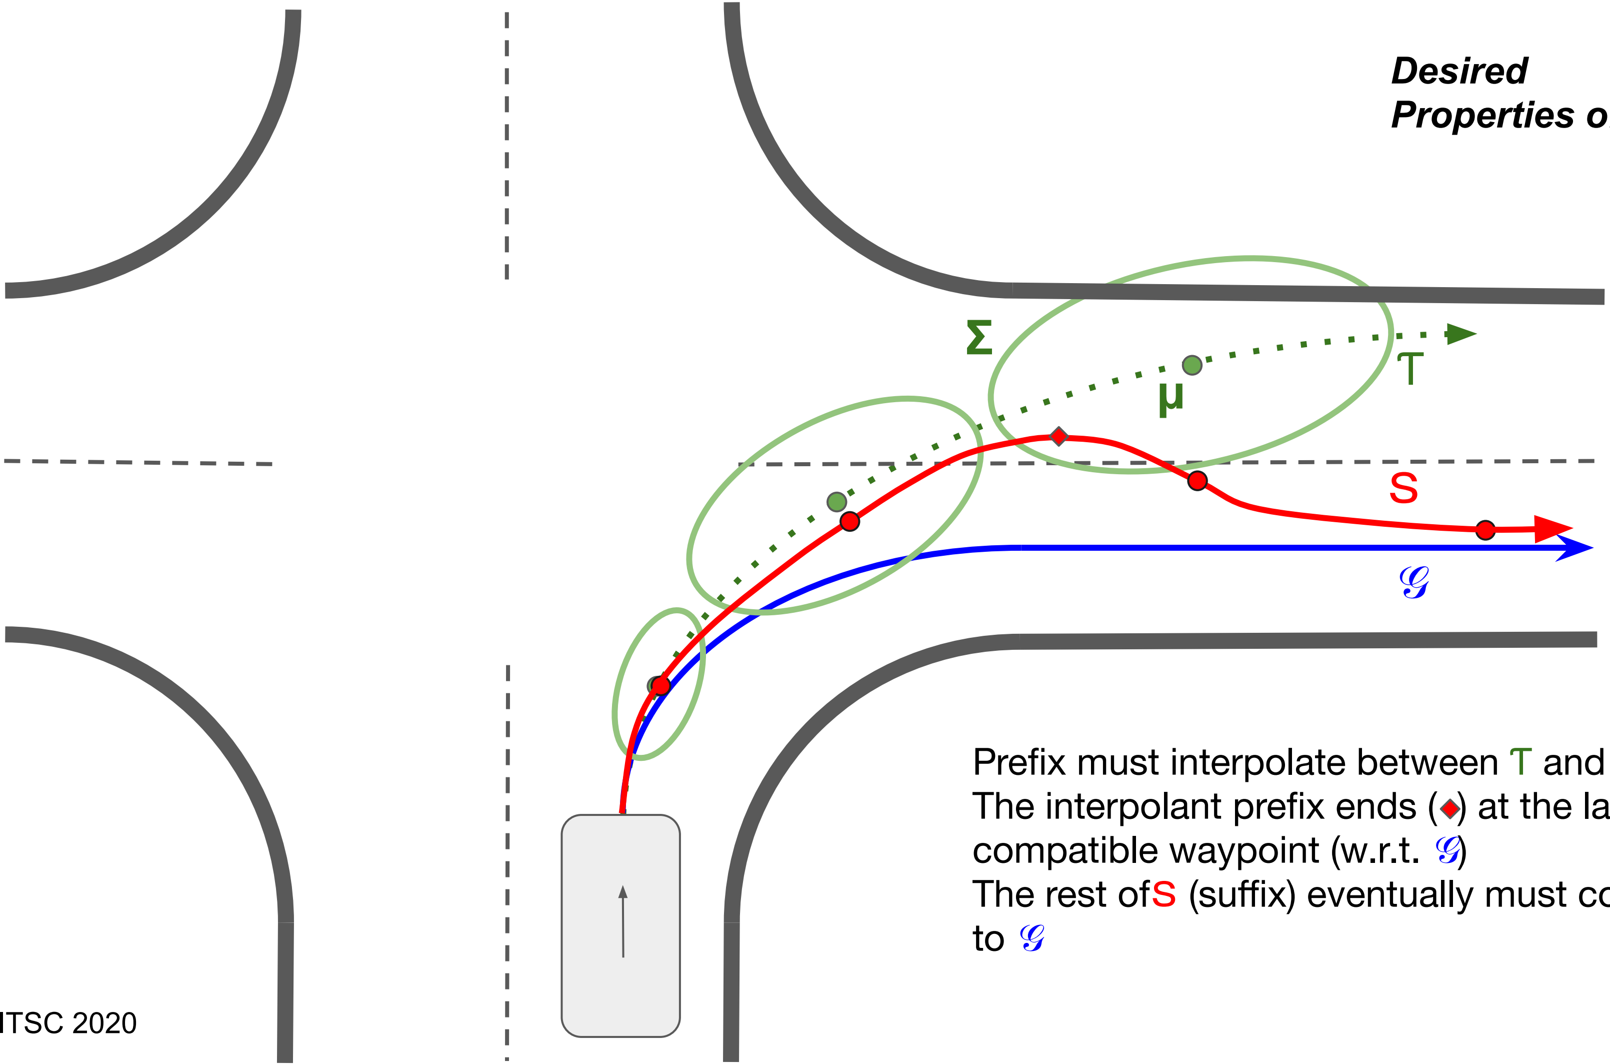

Prefix must interpolate between T and $\mathscr{G}$
The interpolant prefix ends (◆) at the last compatible waypoint (w.r.t. $\mathscr{G}$)
The rest of S (suffix) eventually must converge to $\mathscr{G}$

Desired Properties of S

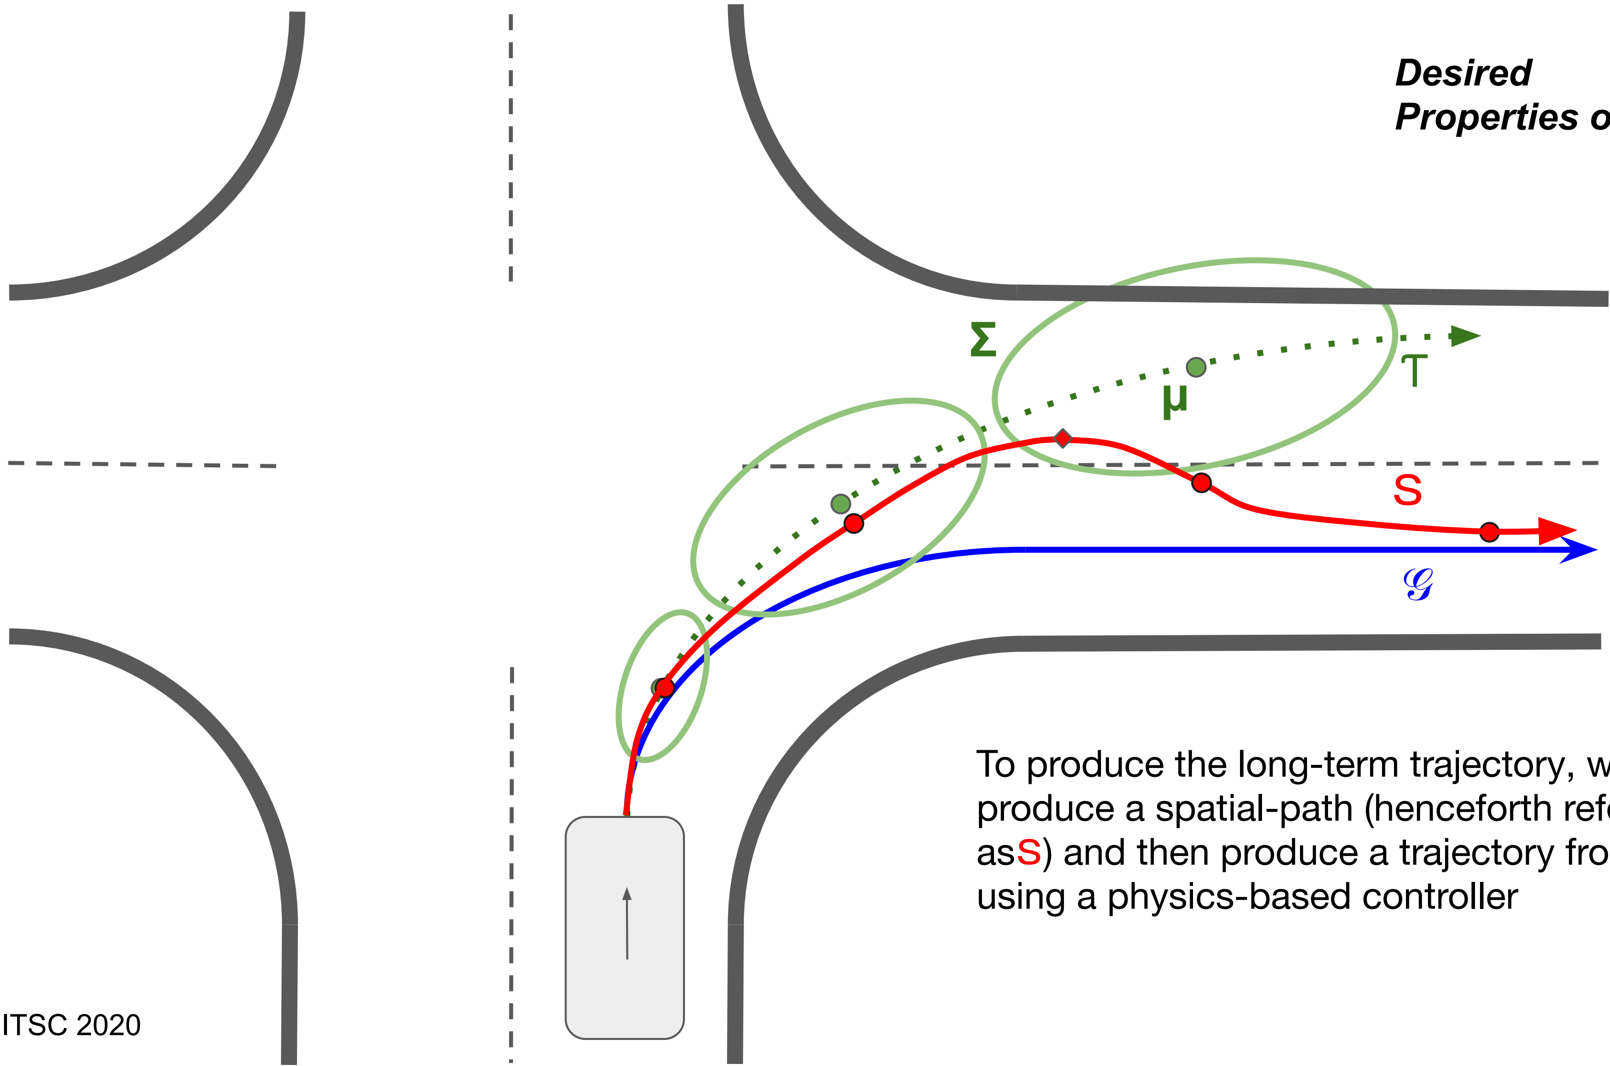

To produce the long-term trajectory, we first produce a spatial-path (henceforth referred to asS) and then produce a trajectory from it using a physics-based controller

Problem Formulation

Conditioned on a goal $\mathcal{G}$ and waypoint parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ of a trajectory $\mathcal{T}$, we would like to find the most likely position $\mathbf{y}$ of the resulting path prefix waypoint. To that end, we propose to solve the following optimization problem,

$$\arg\max_{\mathbf{y}} \quad \log \mathbb{P}(\mathbf{y} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) - \frac{\lambda}{2}\|\mathbf{y} - \Pi_{\mathcal{G}}(\mathbf{y})\|^2,$$

where $\Pi_{\mathcal{G}}(\mathbf{y})$ denotes the projection of $\mathbf{y}$ onto the goal path $\mathcal{G}$, and $\lambda \geq 0$ is a parameter that controls the cost of goal deviation.

We employ an alternating minimization procedure to solve this optimization problem. We alternate between finding the projection and finding the optimum of the objective using a fixed projection.

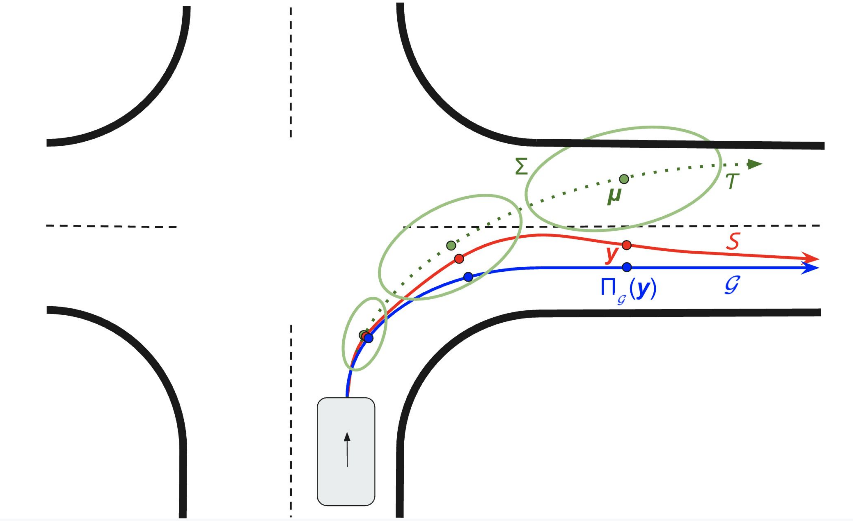

Resulting Solution

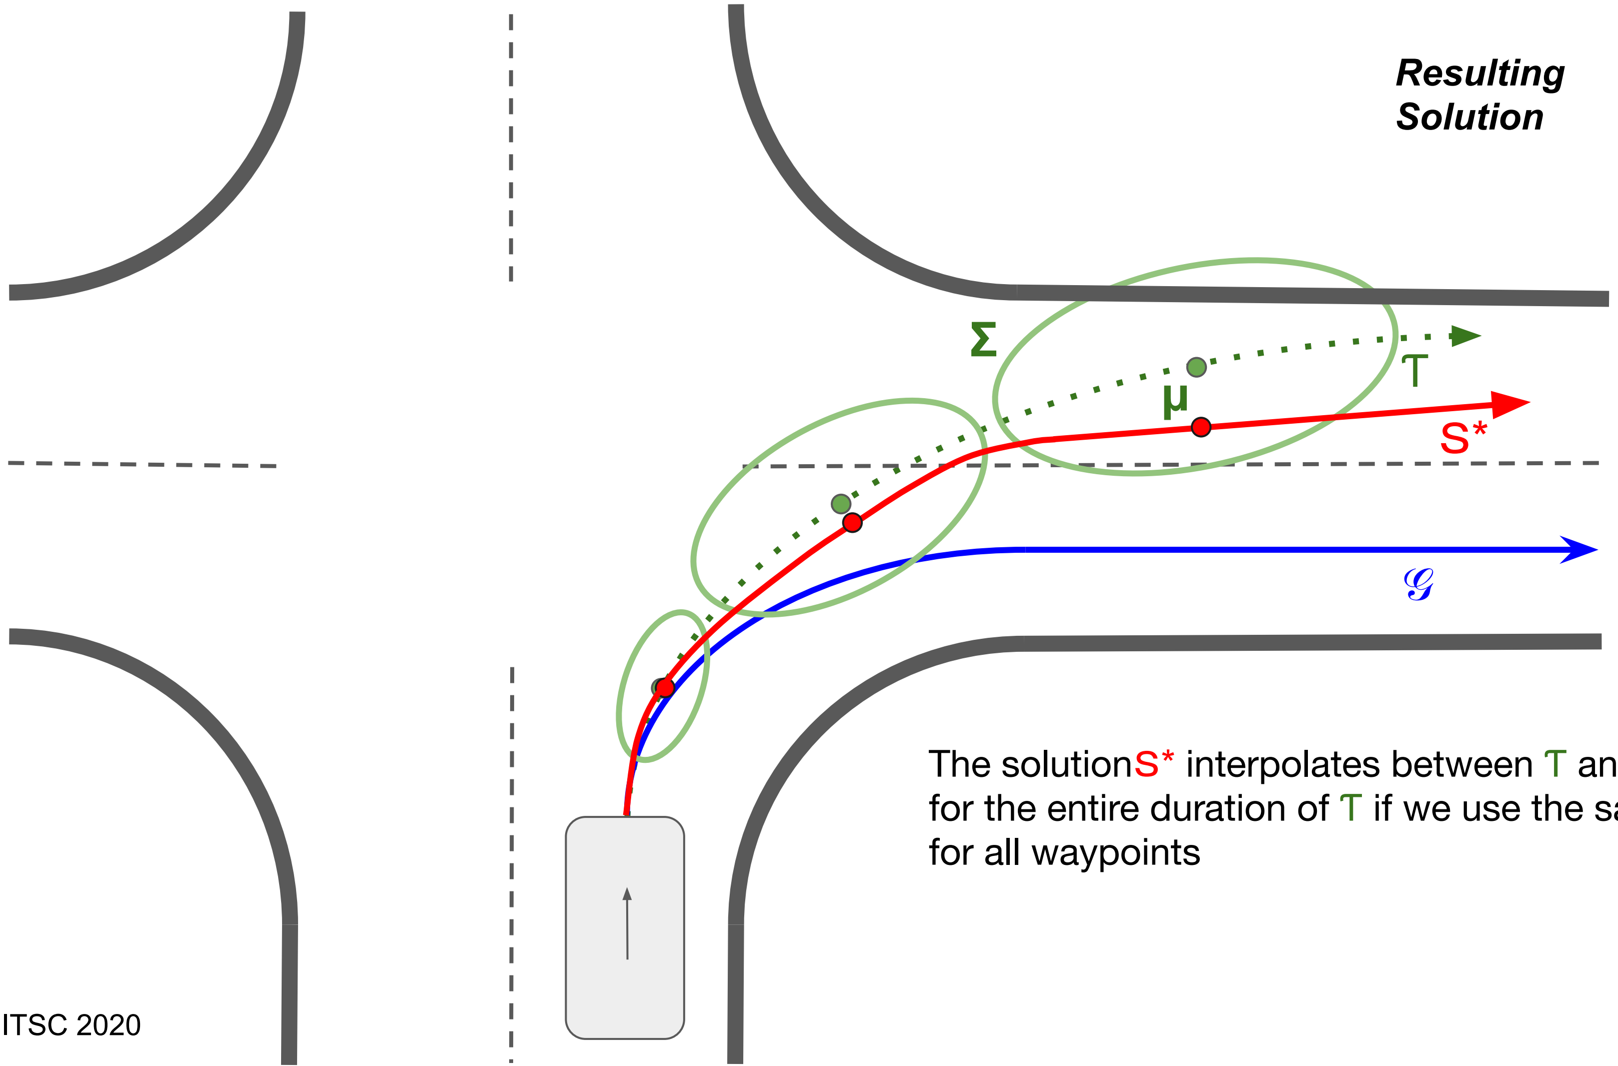

The solution S* interpolates between T and 𝒢 for the entire duration of T if we use the same λ for all waypoints

Resulting Solution

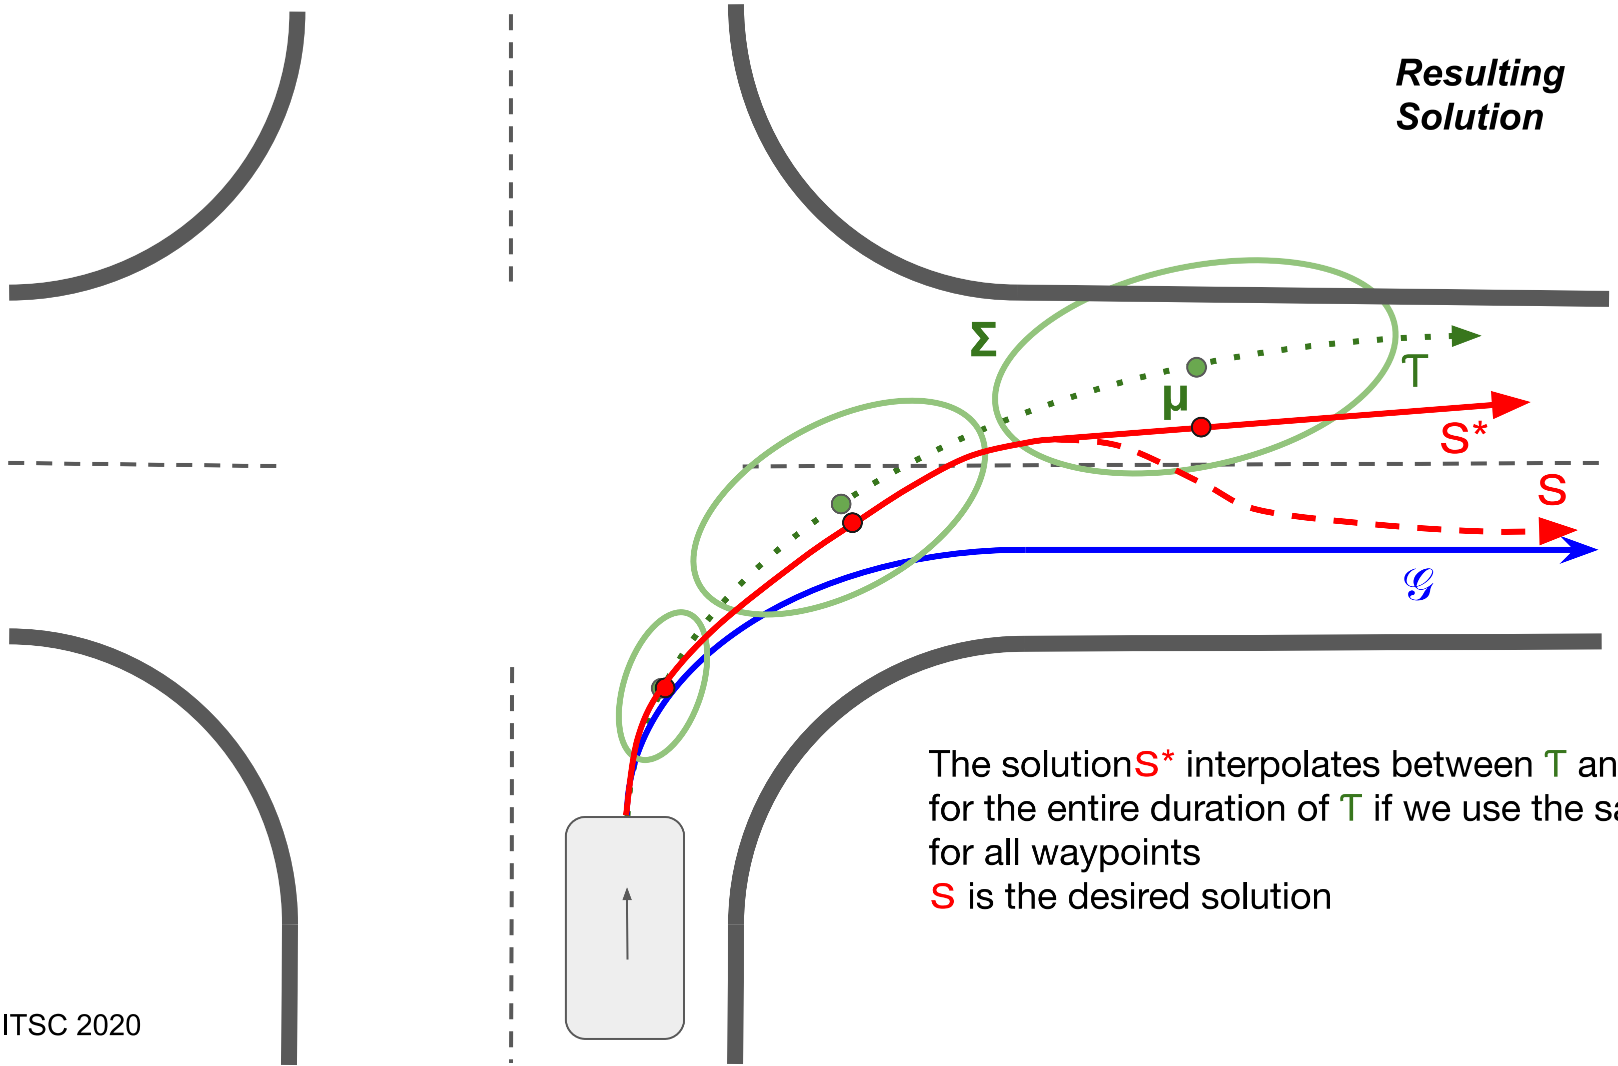

The solution S* interpolates between T and 𝒢 for the entire duration of T if we use the same λ for all waypoints

S is the desired solution

Resulting Solution

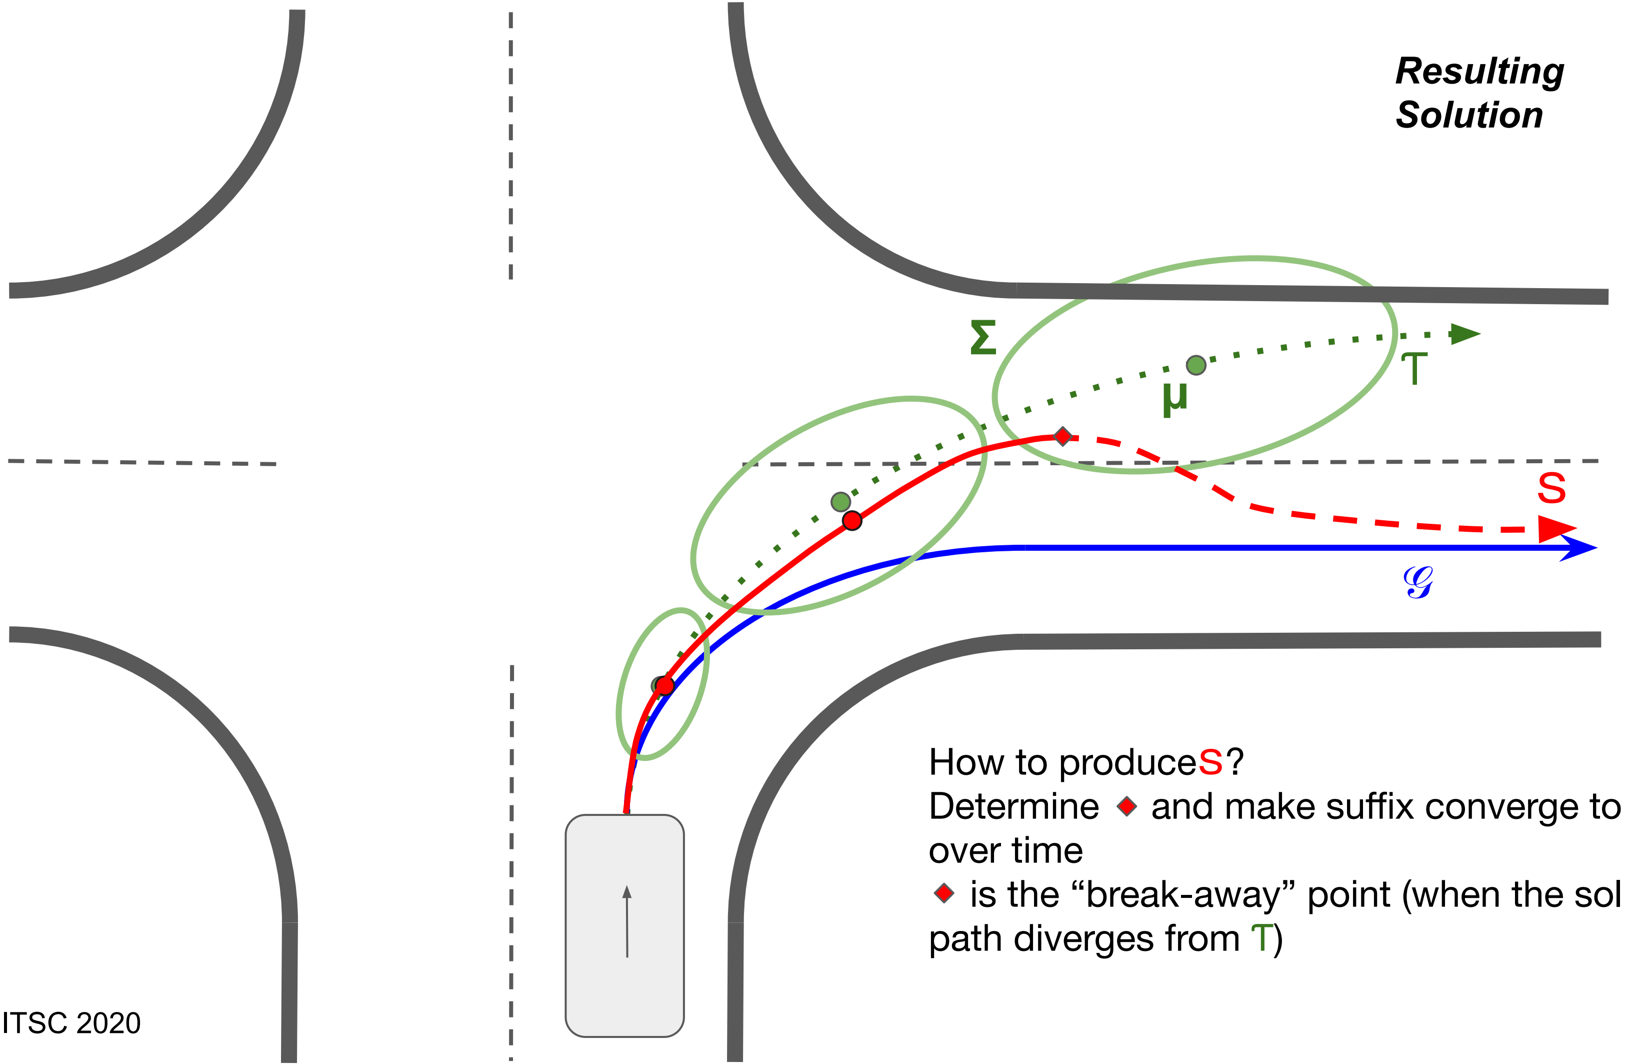

How to produce S?

Determine ◆ and make suffix converge to $\mathscr{G}$ over time

◆ is the “break-away” point (when the solution path diverges from T)

Resulting Solution

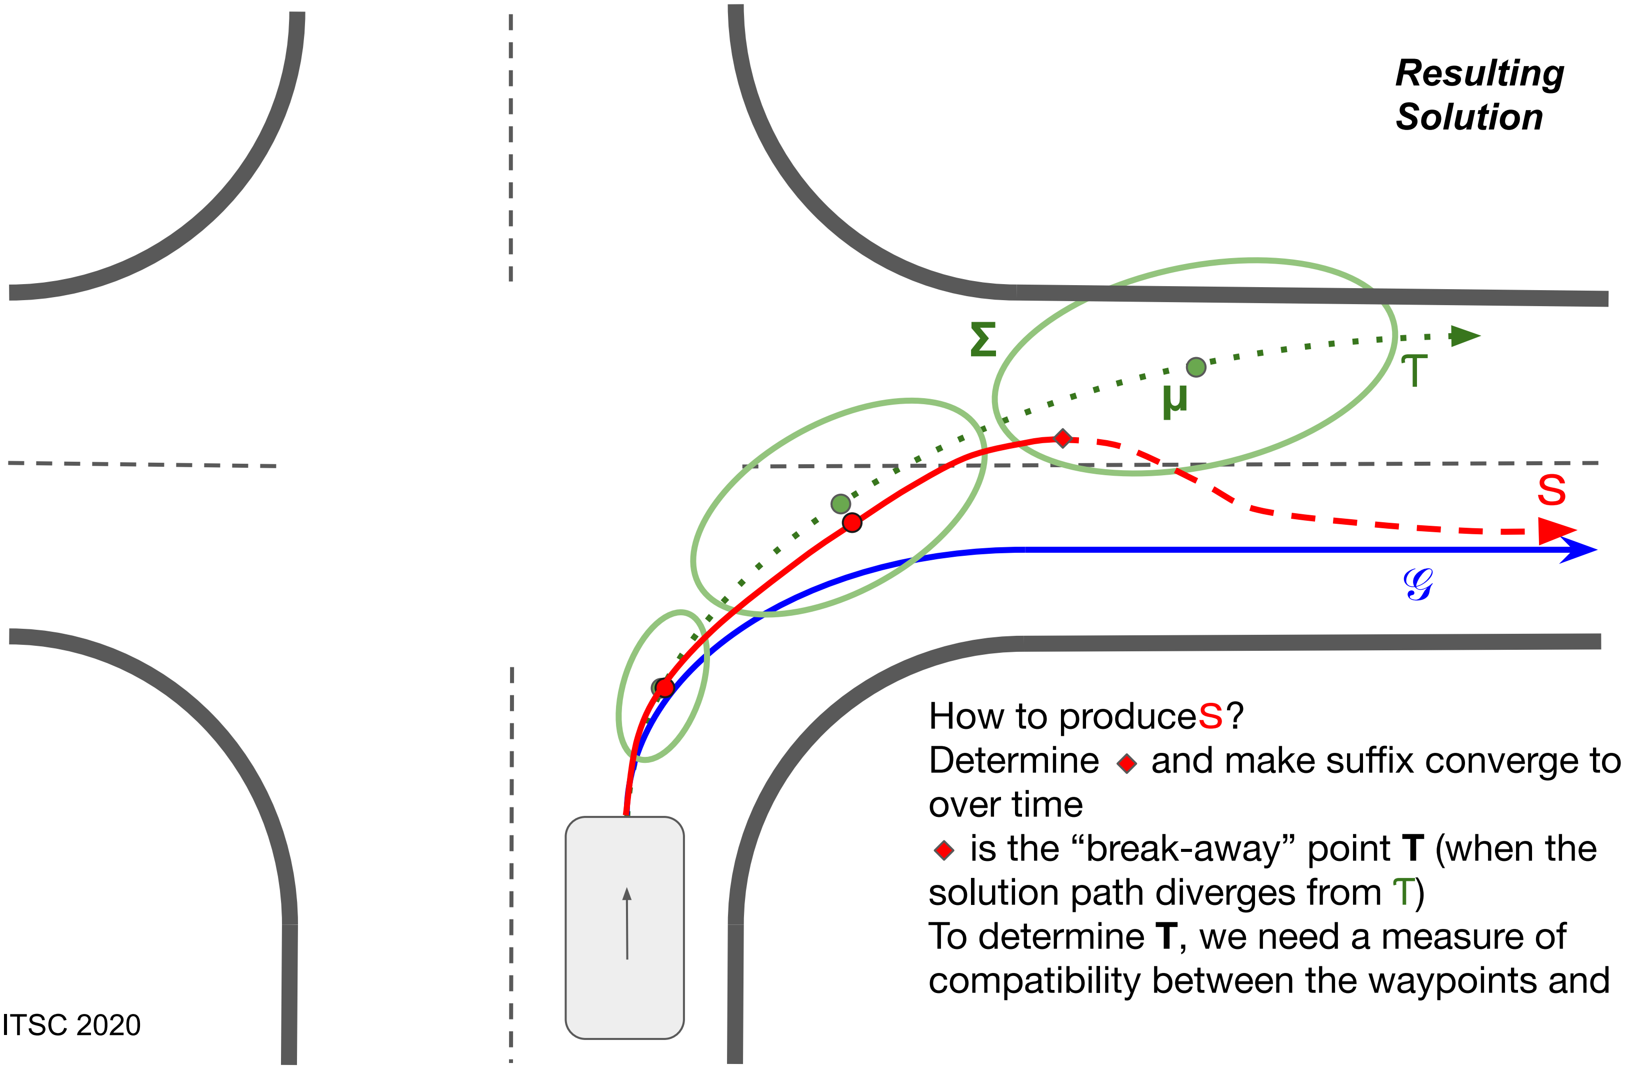

How to produce S?

Determine ◆ and make suffix converge to $\mathscr{G}$ over time

◆ is the "break-away" point **T** (when the solution path diverges from T)

To determine **T**, we need a measure of compatibility between the waypoints and $\mathscr{G}$

Waypoint compatibility score

1) Waypoint compatibility score: Let $S(\mathbf{w}_t, \mathcal{G}) \in [0, 1]$ denote a compatibility score of the waypoint $\mathbf{w}_t$ w.r.t. $\mathcal{G}$, determined using the actor position distribution $\mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$, a 2-D polygon representation of their shape at time step t (denoted by $\mathcal{P}_t$), and the goal $\mathcal{G}$. The waypoint $\mathbf{w}_t$ is considered close to $\mathcal{G}$ iff $S(\mathbf{w}_t, \mathcal{G}) \geq \alpha$, where $\alpha \in (0, 1)$ is a fixed compatibility threshold.

Waypoint compatibility score
Assumption #1

Actors are rigid body objects

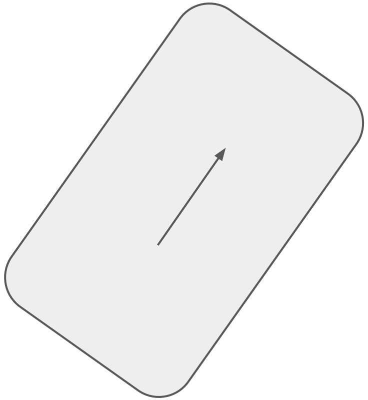

Waypoint compatibility score
Assumption #1

Actors are rigid body objects

- - - is the control point position uncertainty

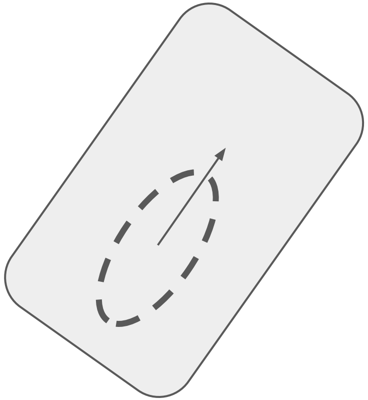

Waypoint compatibility score
Assumption #1

Actors are rigid body objects
- - - is the control point position uncertainty
It can be used to estimate the position uncertainty of any point on the actor’s body

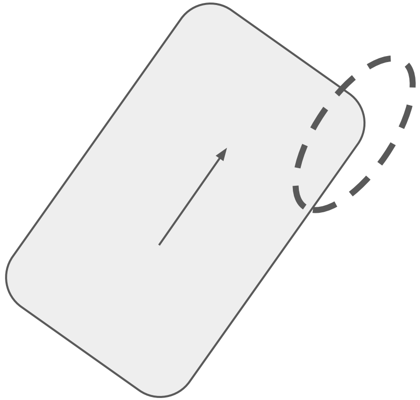

Waypoint compatibility score

Let $\mathbf{p}$ be any point on the actor's body, and $\mathcal{X}_\theta = \mathcal{N}(\mathbf{p}, \mathbf{\Sigma}_t)$ represent the position uncertainty distribution of $\mathbf{p}$, where θ are the distribution parameters. Let $D_{\mathcal{X}_\theta}(\mathbf{x}) = \|\mathbf{x} - \mathbf{p}\|_{\mathbf{\Sigma}_t^{-1}}$ represent the Mahalanobis distance of some point $\mathbf{x}$ from $\mathcal{X}_\theta$, and $\Pi_{\mathcal{G}}^{D_{\mathcal{X}_\theta}}(\mathbf{p}) = \arg\min_{\mathbf{g}\in\mathcal{G}} D_{\mathcal{X}_\theta}(\mathbf{g})$ denote a function that computes the closest point on $\mathcal{G}$ from $\mathbf{p}$ based on the Mahalanobis distance, where $\|\mathbf{v}\|_{\mathbf{Q}} := \sqrt{\mathbf{v}^\top \mathbf{Q} \mathbf{v}}$.

Waypoint compatibility score

A is the closest point on the blue goal path w.r.t. the Mahalanobis Distance evaluated using the point **p** and its uncertainty distribution (- -)

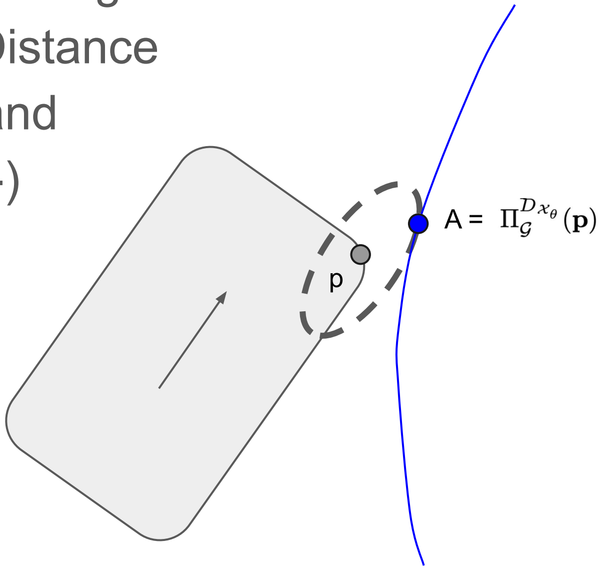

Waypoint compatibility score

B is the closest point from **p** w.r.t. the Euclidean L2 distance

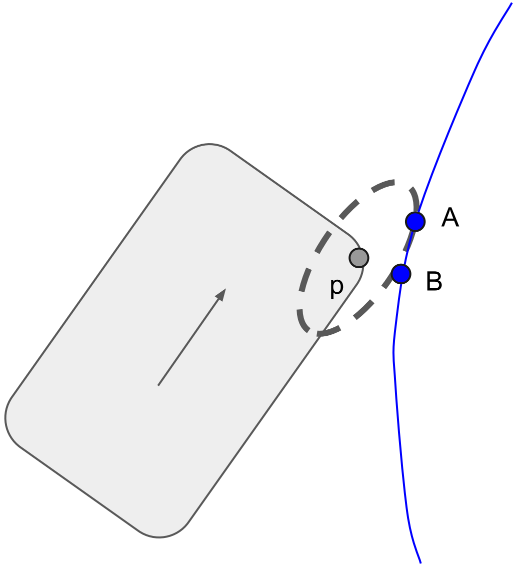

Waypoint compatibility score

$S(\mathcal{X}_\theta, \mathcal{G}) =$ 1 - prob. mass in shaded region

Null hypothesis: *A* is a sample drawn from the position uncertainty distribution of *p*

Compatibility score: p-value

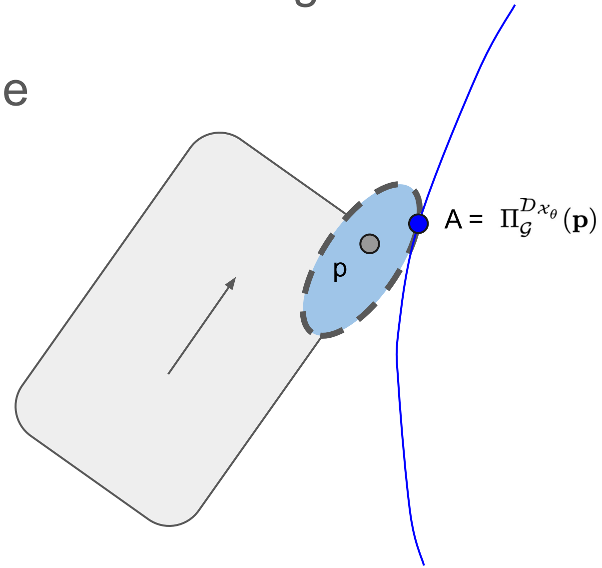

Waypoint compatibility score

Assumption #2

The wheelbase of the average vehicle is smaller than the average road curve radius

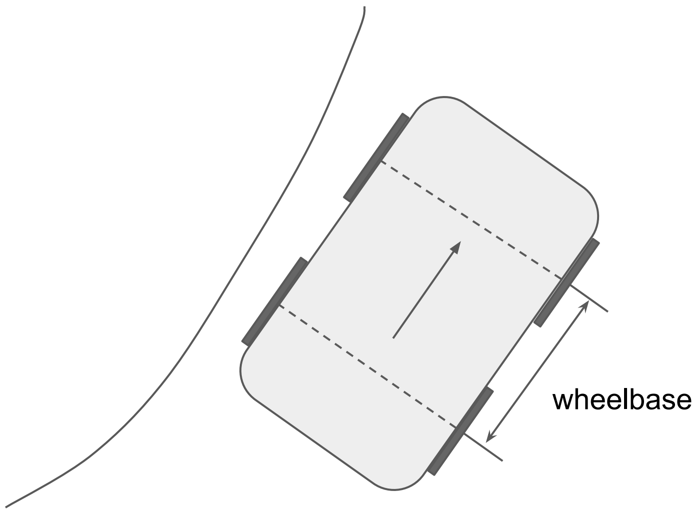

Waypoint compatibility score

Assumption #2

The wheelbase of the average vehicle is smaller than the average road curve radius

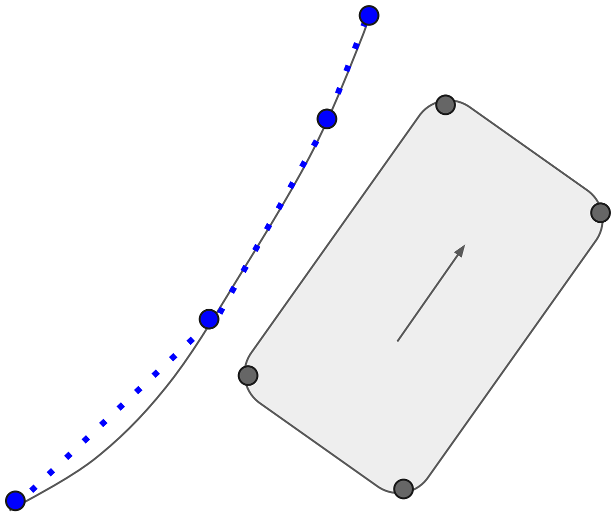

Waypoint compatibility score

$S(\mathbf{w}_t, \mathcal{G}) =$ 1

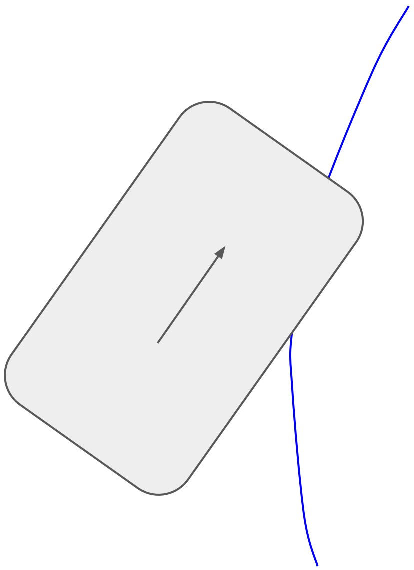

Waypoint compatibility score

$$S(\mathbf{w}_t, \mathcal{G}) = \max_{\mathbf{p} \in \mathcal{P}_t} S(\mathcal{X}_\theta, \mathcal{G})$$

= 1 - prob. mass in shaded region

Approximated by considering 4 corners of the shape polygon

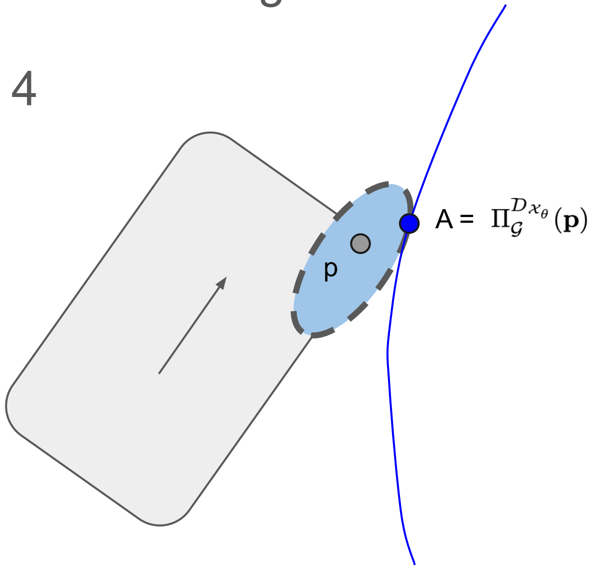

Waypoint compatibility score

We use $S(\mathcal{X}_\theta, \mathcal{G})$ to determine $S(\mathbf{w}_t, \mathcal{G})$. If $\mathcal{P}_t$ overlaps the goal path polyline, we set $S(\mathbf{w}_t, \mathcal{G}) = 1$. Otherwise, $S(\mathbf{w}_t, \mathcal{G}) = \max_{\mathbf{p} \in \mathcal{P}_t} S(\mathcal{X}_\theta, \mathcal{G})$. Assumption (ii) enables us to determine $S(\mathbf{w}_t, \mathcal{G})$ efficiently by considering the actor polygon vertices alone, without the need to account for the vertices of goal polyline. Then, we define T as the latest time horizon t at which $S(\mathbf{w}_t, \mathcal{G}) \geq \alpha$, as

$$T = \max\{t : S(\mathbf{w}_t, \mathcal{G}) \geq \alpha, t \in \{1, \ldots, n\}\}.$$

Convergence to $\mathscr{G}$

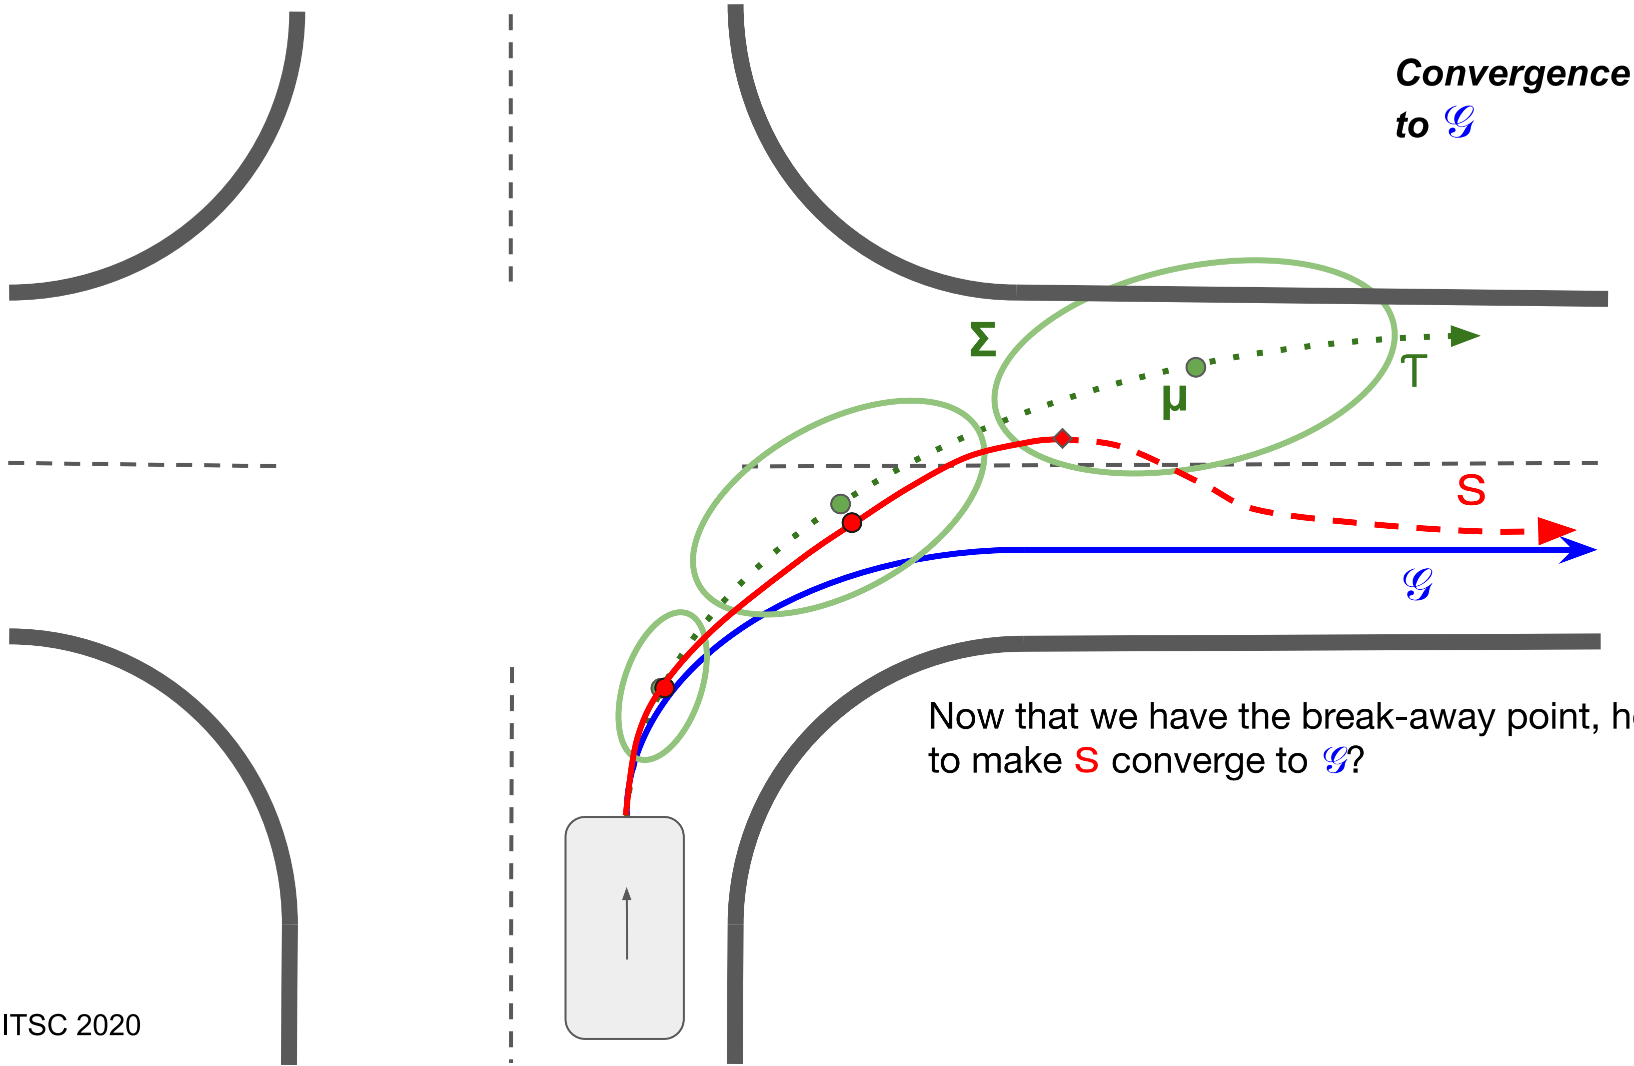

Now that we have the break-away point, how to make S converge to $\mathscr{G}$?

Convergence of solution path to the goal

Let $f(t)$ be a non-negative, monotonically decreasing function of time. For a fixed constant $c > 0$, we choose the following schedule for waypoints of the solution path,

$$\|\mathbf{y}_t - \mathbf{g}_t\| = cf(t),$$

for $t = T+1, \ldots, H$, where H is the prediction horizon and $\mathbf{g}_t = \Pi_{\mathcal{G}}(\mathbf{y}_t)$.

Convergence to $\mathscr{G}$

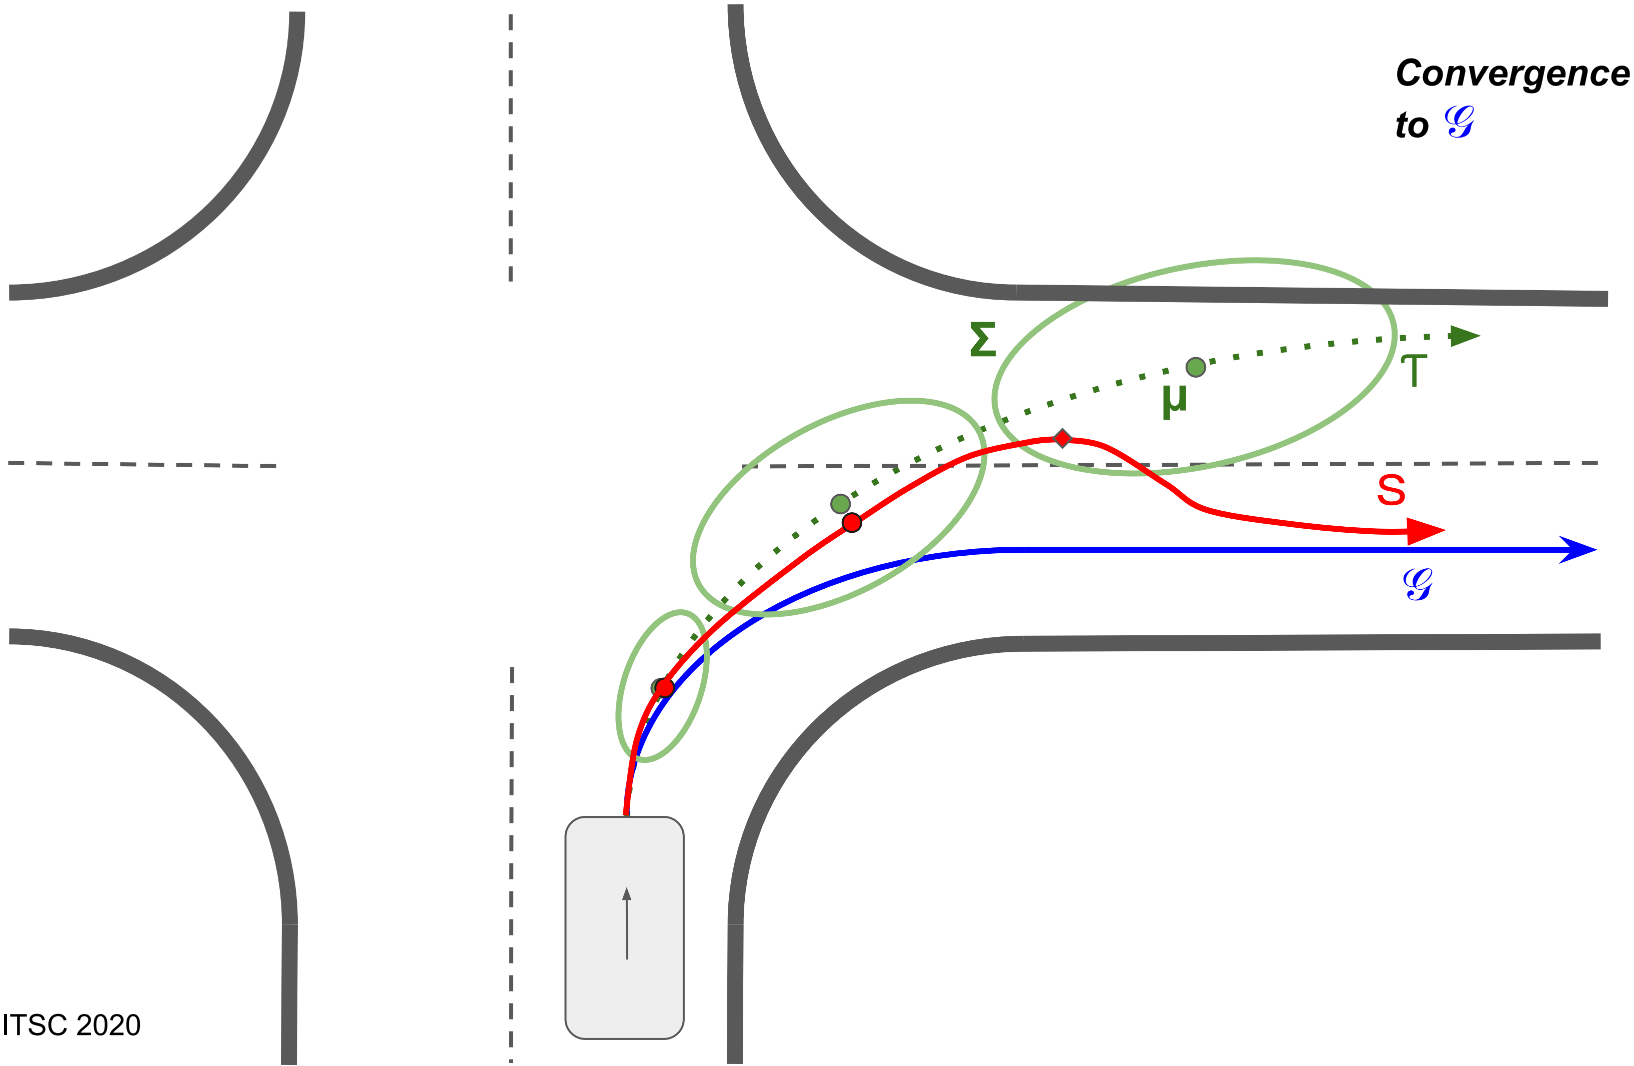

Developing a path beyond T

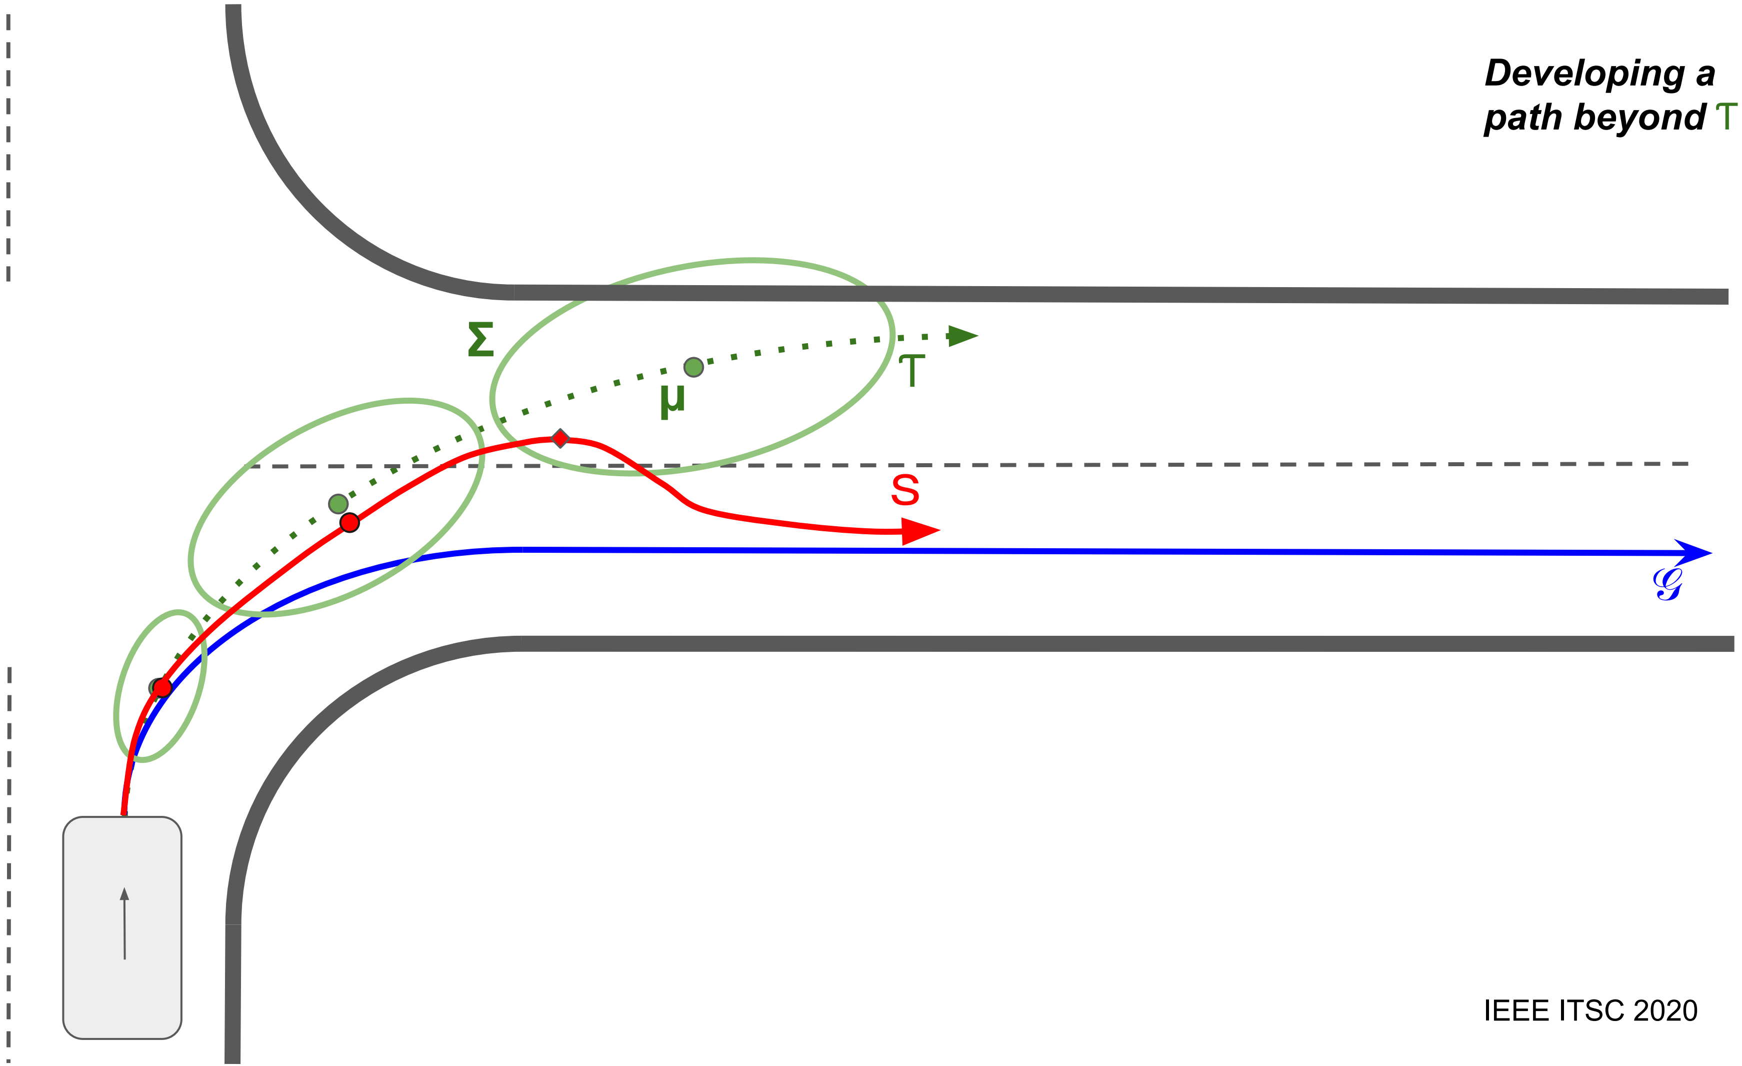

Developing a path beyond T

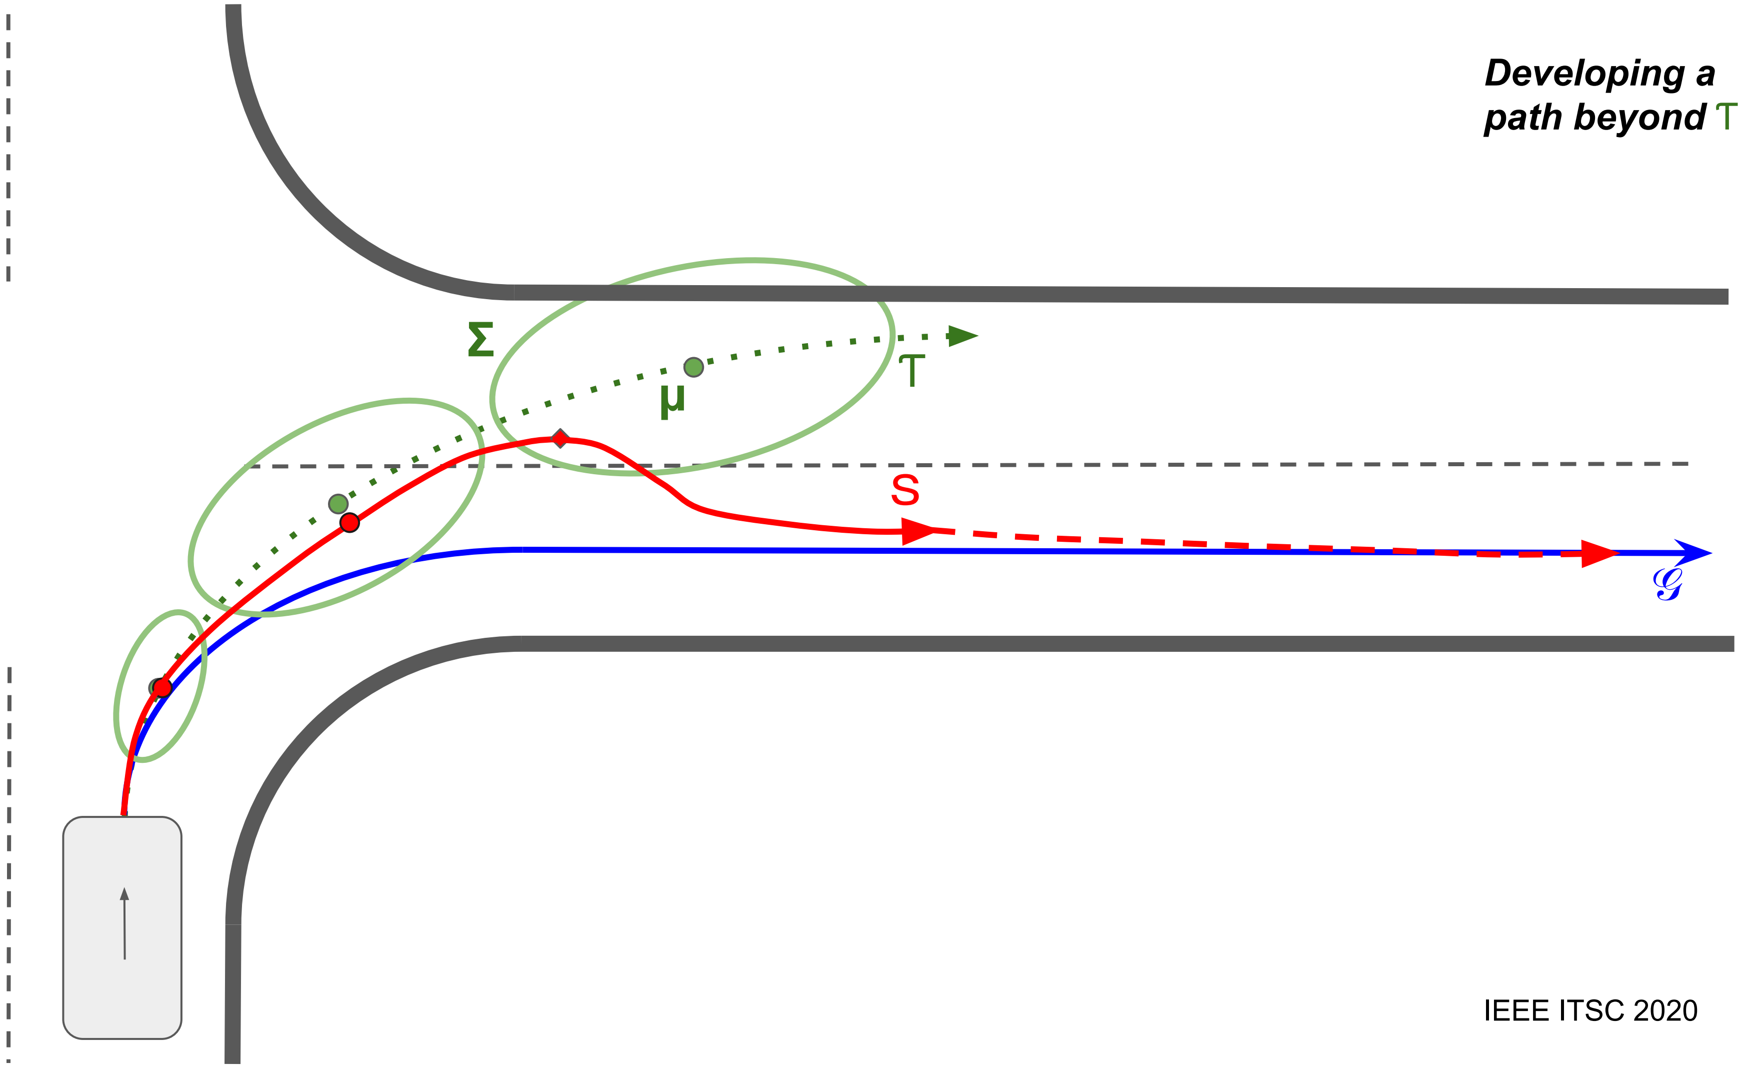

Developing a path beyond T

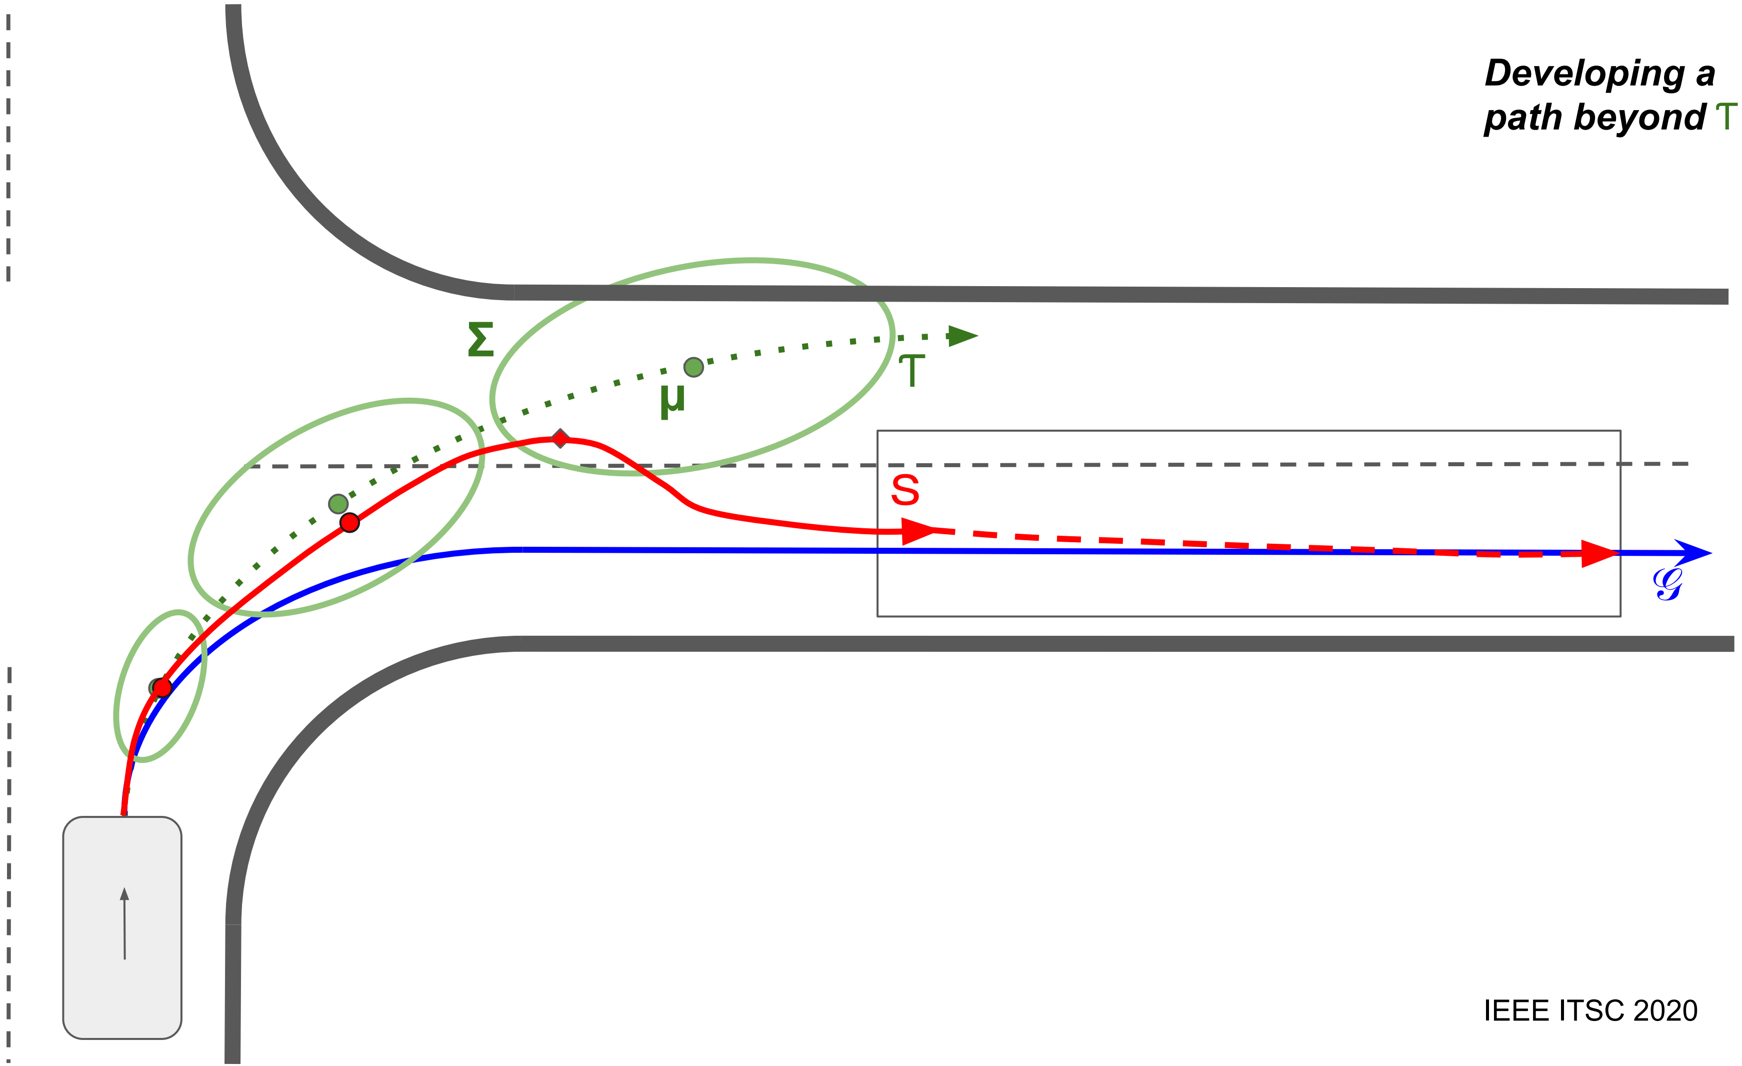

Developing a path beyond T

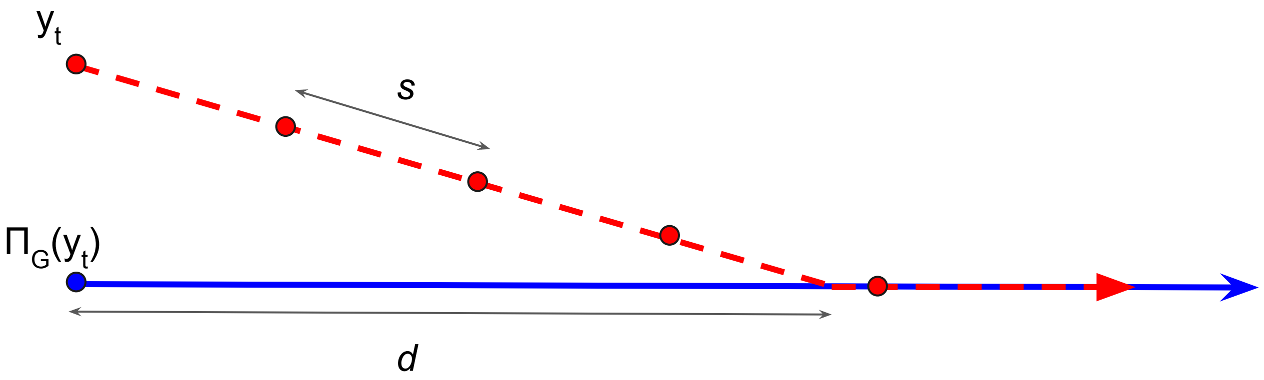

Initial offset is linearly decayed over distance *d*
with a spatial resolution of *s*

Experiments

We use RasterNet to produce 6s trajectories T

Lane association produces a list of goal paths $\mathcal{G}$ for every vehicle actor

Baseline models

- RasterNet (*RN*)
- Pure-pursuit trajectories (*PP*, also 6s) from $\mathcal{G}$
- Linear-decay stitching (*LS(n)*), which is an *RN* trajectory for the first *n* seconds followed by a *PP* trajectory tracking $\mathcal{G}$ from the vehicle position at $t=n$
- *Ballistic* trajectories (rollouts of initial state using physics)

We only consider (cross-track) CT errors

Our method: US (uncertainty-aware stitching)

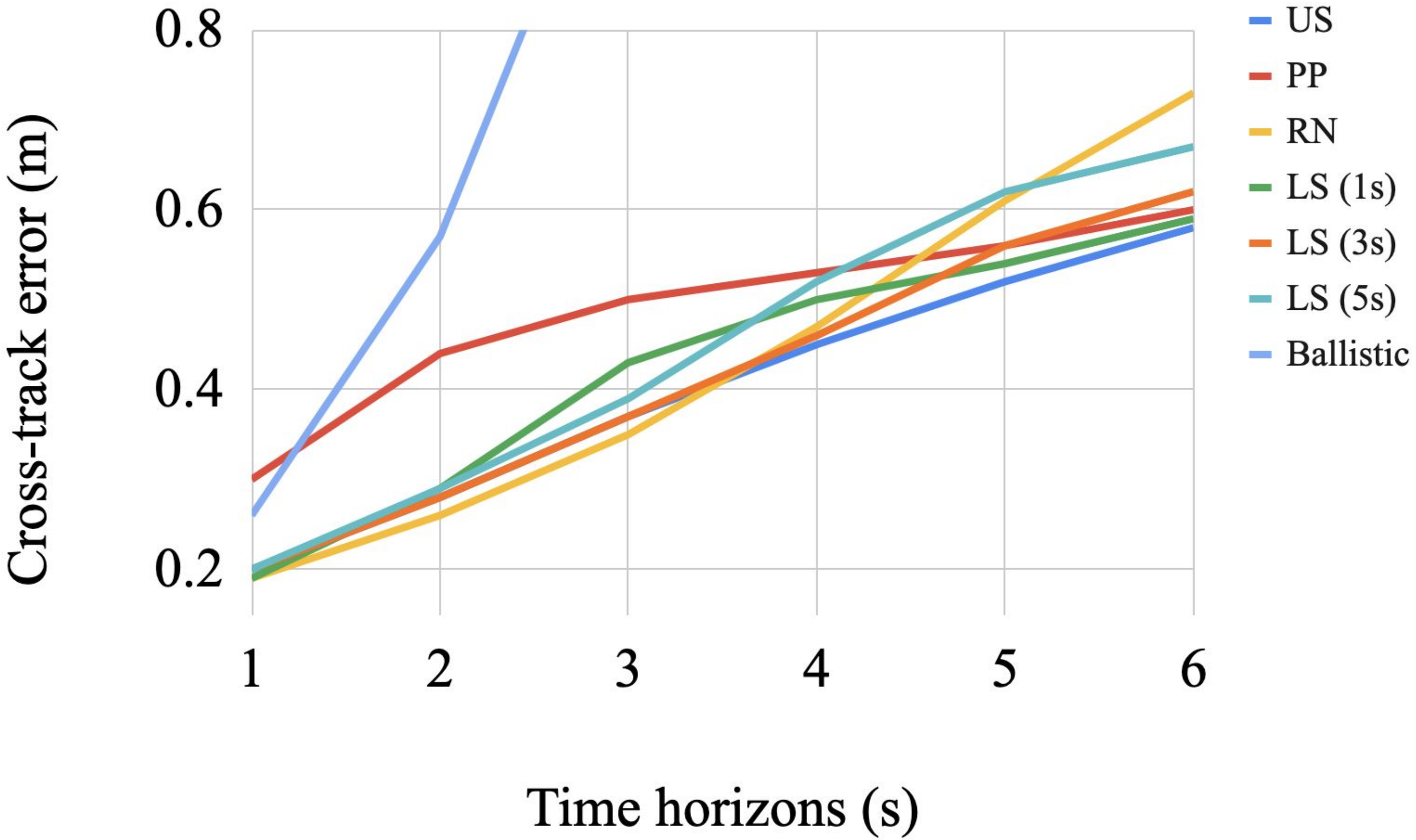

Comparison of CT error against baselines

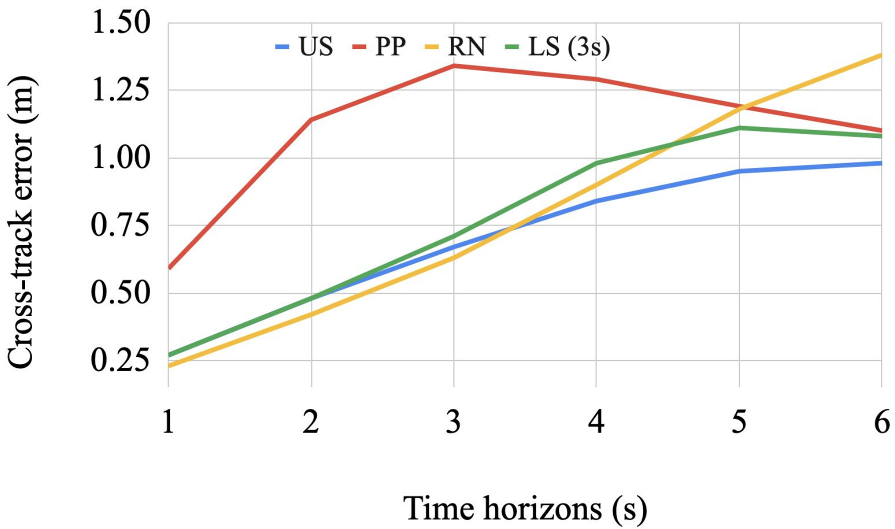

Comparison of CT error against baselines for left-turning tracks (6% of data)

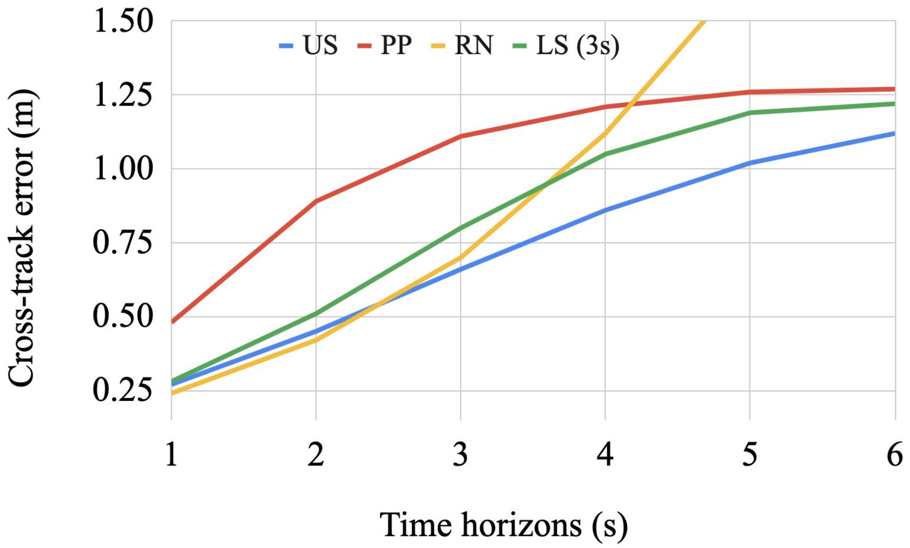

Comparison of CT error against baselines for right-turning tracks (5.3% of data)

Qualitative Examples And Conclusion

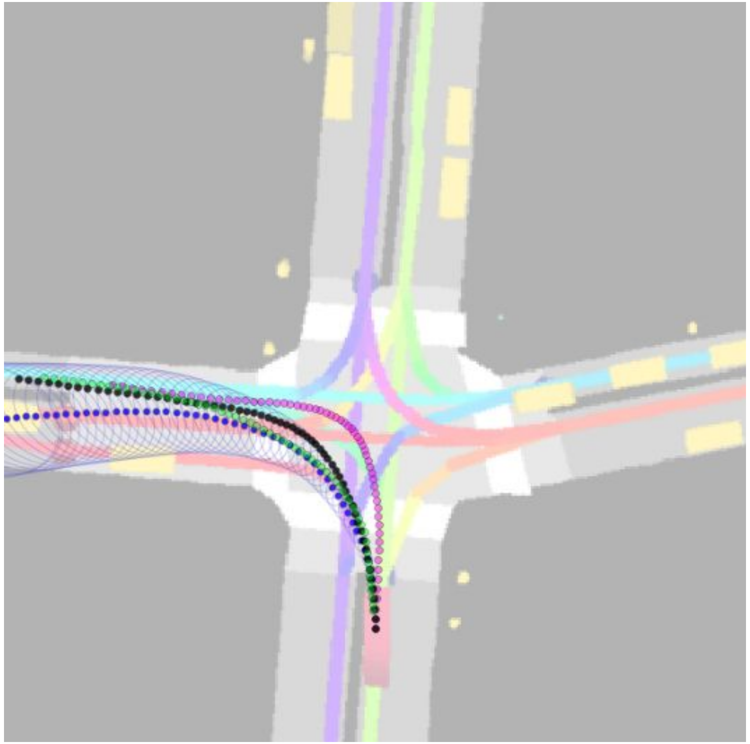

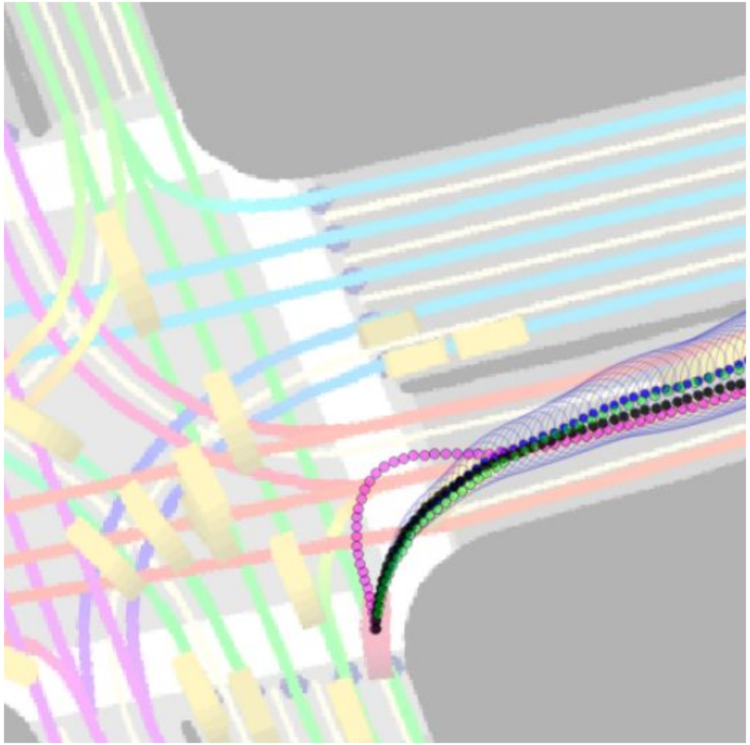

--- *PP*
--- Ground-truth
--- *RN*
--- *US*

Qualitative Examples And Conclusion

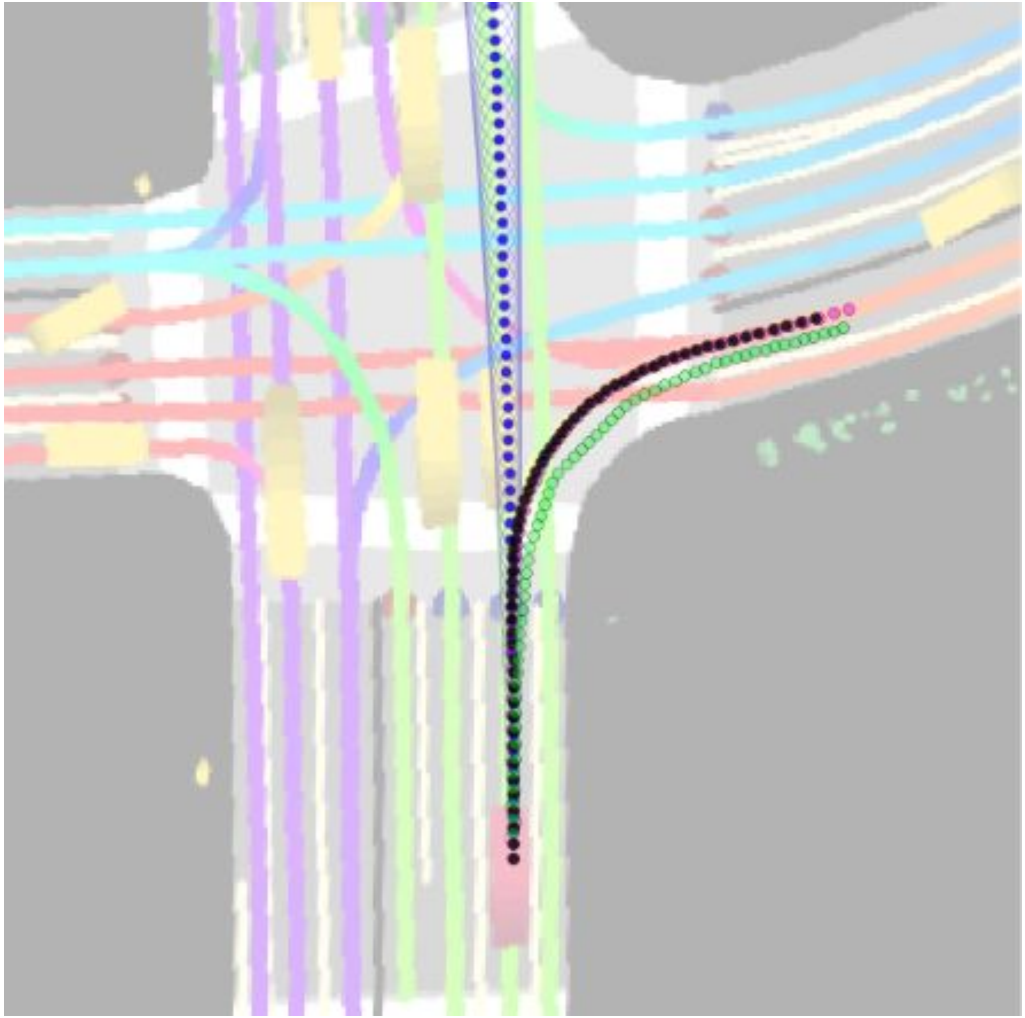

--- *PP*
--- Ground-truth
--- *RN*
--- *US*

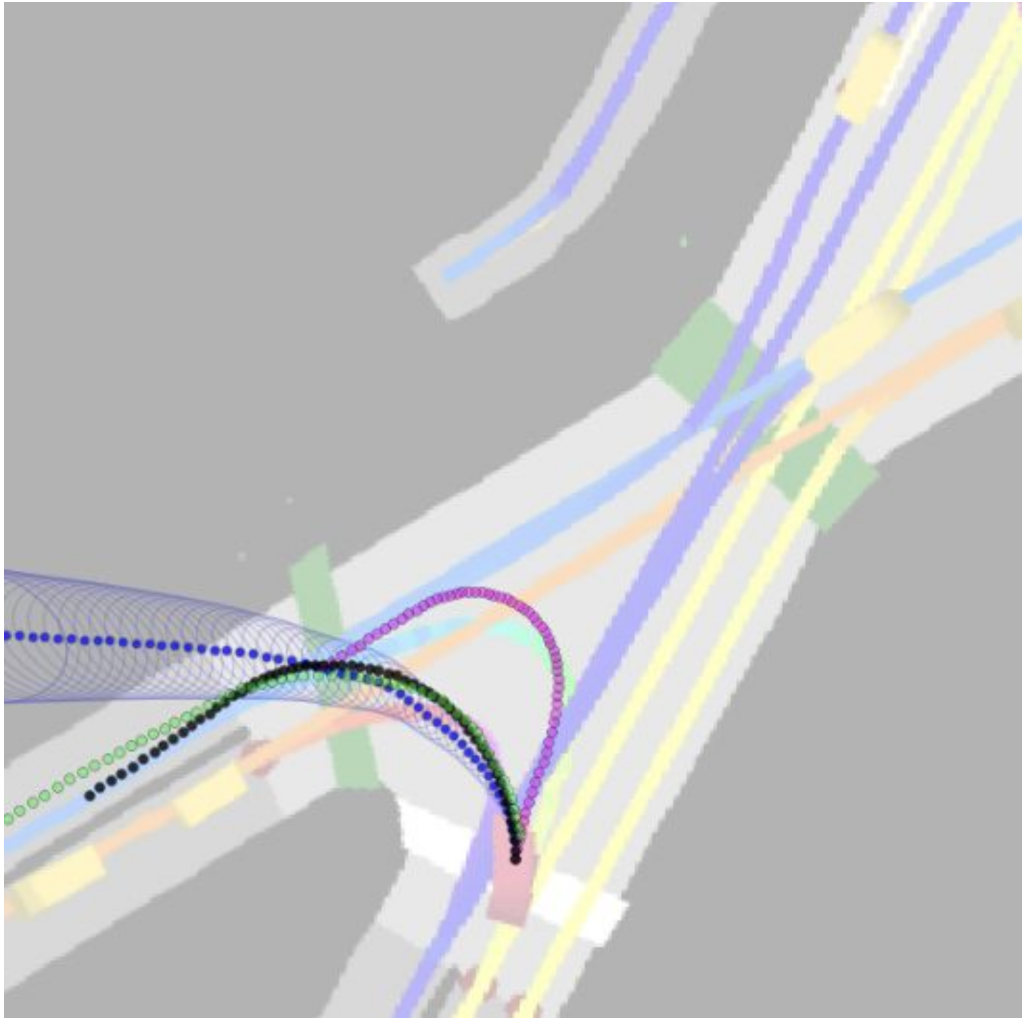

Qualitative Examples And Conclusion

--- *PP*
--- Ground-truth
--- *RN*
--- *US*