Versatile Scene-Consistent Traffic Scenario Generation as Optimization with Diffusion

Zixu Zhang, Ziyu Huang, Ameya Vaidya, Chen Lyu, Jaime F. Fisac

SAFE Robotics Laboratory · Princeton University · Nanyang Technological University Singapore

Traffic Modeling as Optimal Control

Optimization Objective → Joint Control Sequence ↓

$$\min_{\mathbf{u}} J_\theta(\mathbf{x}, \mathbf{u}; \mathbf{c}) \leftarrow \text{Scene}$$

↑ Joint Trajectory

$$\text{s.t} \quad \mathbf{x}_0 = x_0, \leftarrow \text{Initial State}$$

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t),\ \forall t \in \{0, \ldots, T-1\}$$

↖ Known Dynamics

What are *key features* of J_θ?

- *Realism & Scalability*
- *Scene-Consistency*
- *Interaction*
- *Controllability*

Under *Maximum Entropy IRL* formulation, ***Boltzmann Distribution*** of J_θ approximates $p(\mathbf{u}\,|\,\mathbf{c})$:

How to obtain J_θ?

$$p(\mathbf{u}\,|\,\mathbf{c}) \approx p_\theta(\mathbf{u}\,|\,\mathbf{c}) := \boxed{\frac{1}{Z_\theta}\exp(-J_\theta(\mathbf{x}(\mathbf{u}), \mathbf{u}; \mathbf{c}))}$$

Energy-Based Method

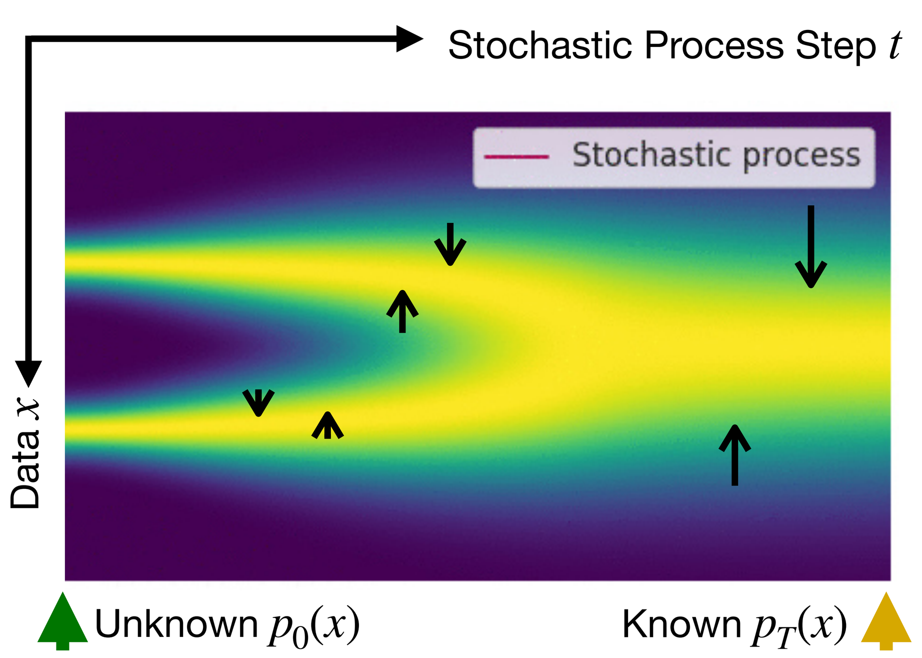

Diffusion model learns to estimate the gradient of $\log p_t$

$$\mathbf{s}_{t=0}(\mathbf{u}\,|\,\mathbf{c}) = \nabla_{\mathbf{u}} \log p_\theta(\mathbf{u}\,|\,\mathbf{c})$$

$$= -\nabla_{\mathbf{u}} J_\theta(\mathbf{x}(\mathbf{u}), \mathbf{u}; \mathbf{c}) - \underbrace{\nabla_{\mathbf{u}} \log Z_\theta}_{=0}$$

Diffusion generative allows us to learn ∇J from data, and conduct compositional optimization with user-defined objectives

*V*ersatile *B*ehavior *D*iffusion

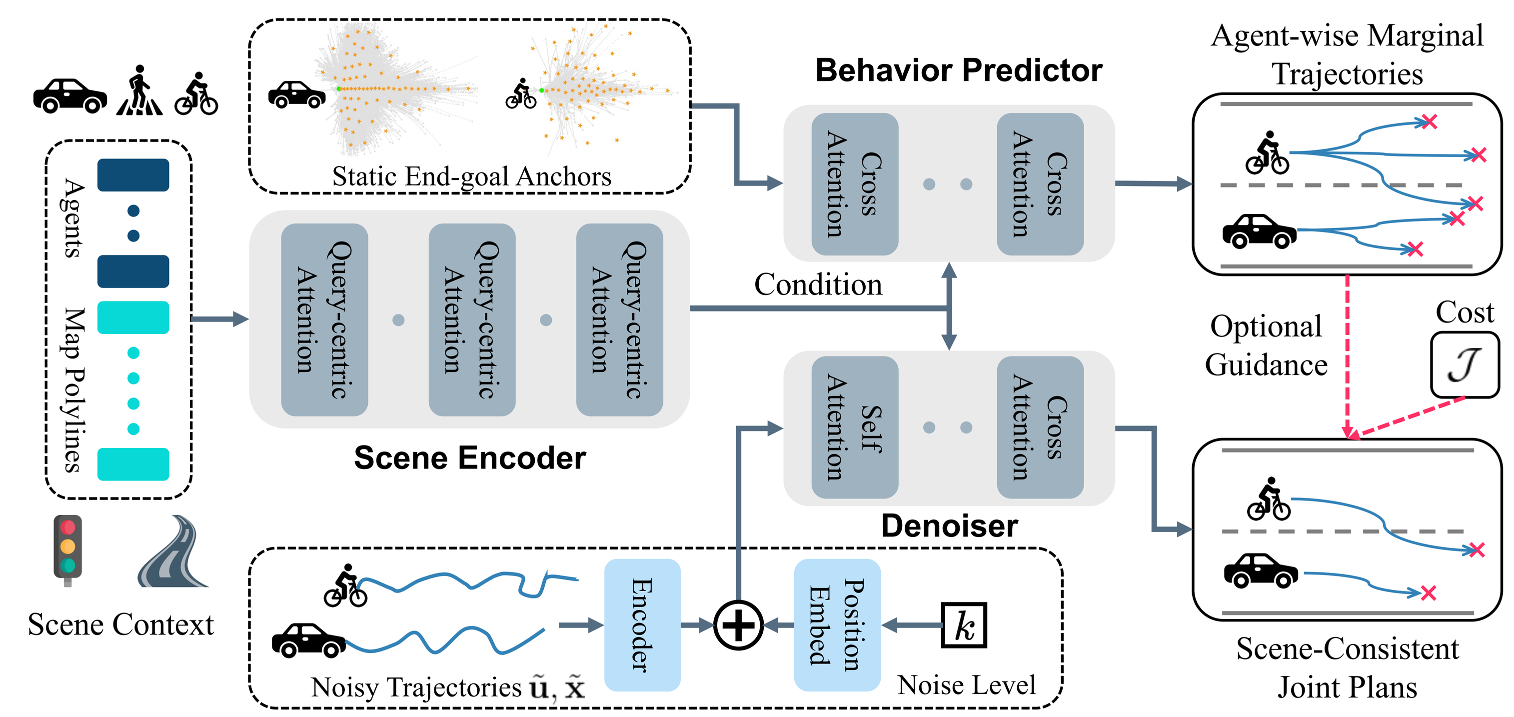

- *Query-centric transformer adapts to scenes with various number of agents*
- *Auxiliary prediction head improves training stability*
- *VBD achieve SOTA performance on Waymo Open Motion Dataset and won the 2nd place in 2024 SimAgent Challenge*
- *VBD generate scenarios both within and beyond training data distribution with simple guidance*

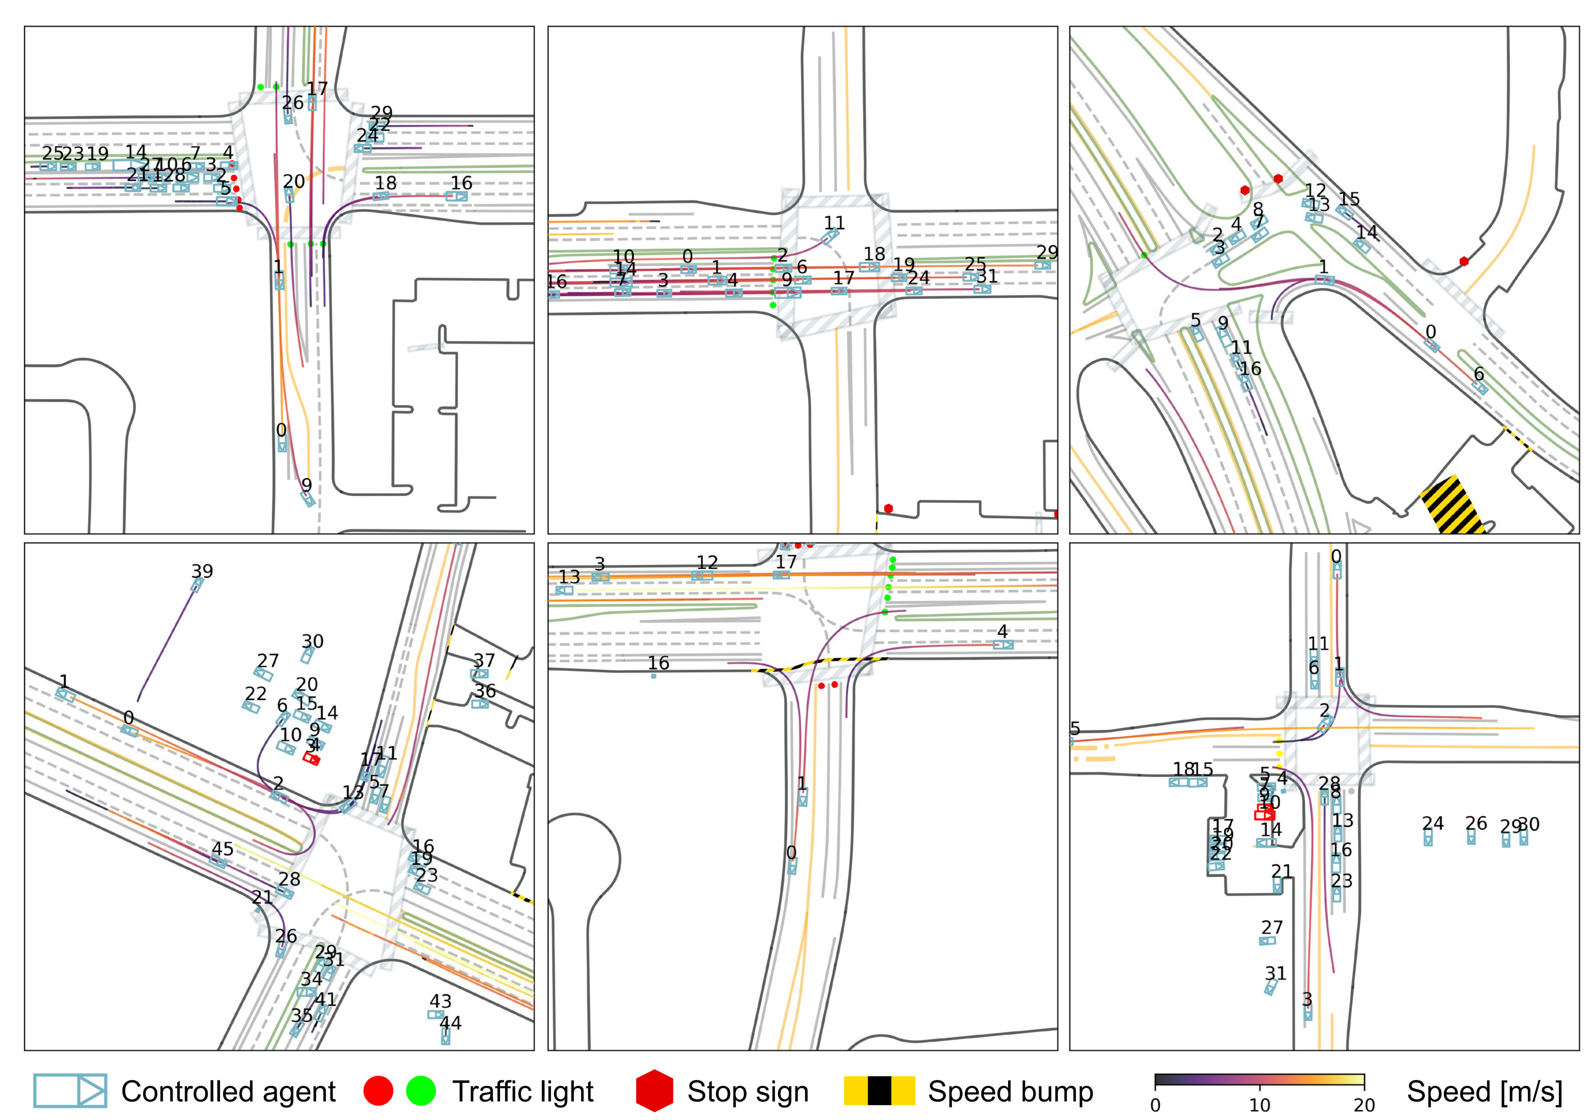

Improving Narrow Passage Problem by Enforcing Constraints

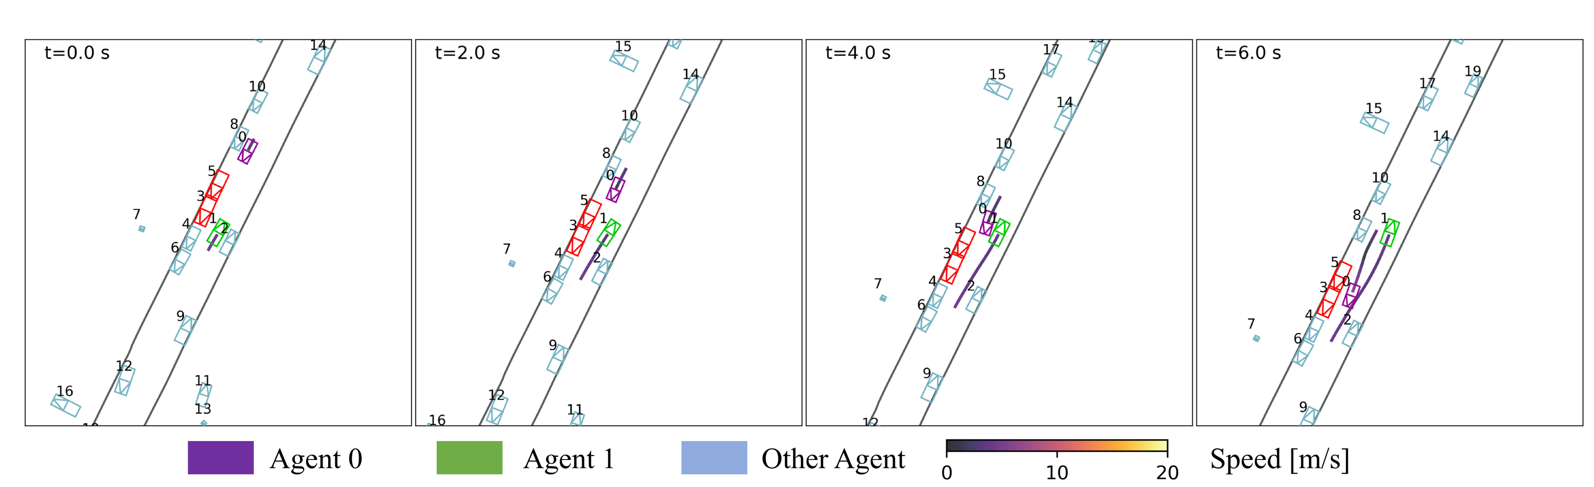

Scene-Consistency With Prior Condition

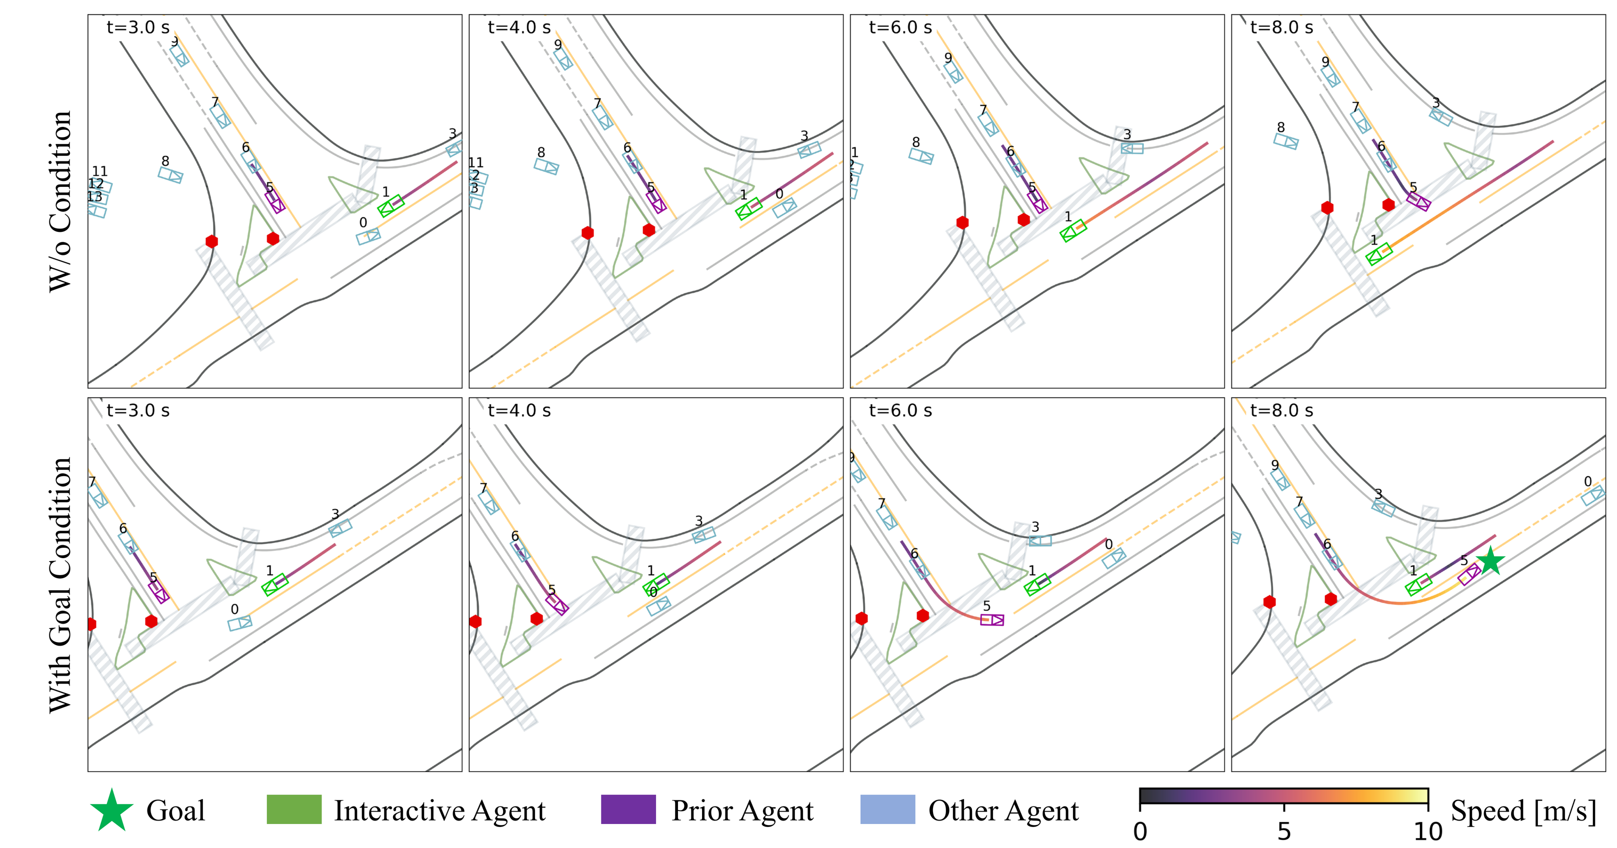

Generating Interactive Long-tail Scenarios via Game Theoretic Guidance

Iteratively guide the denoting process in a gradient descent-ascent fashion to generate interactive pursuit–evasion behaviors

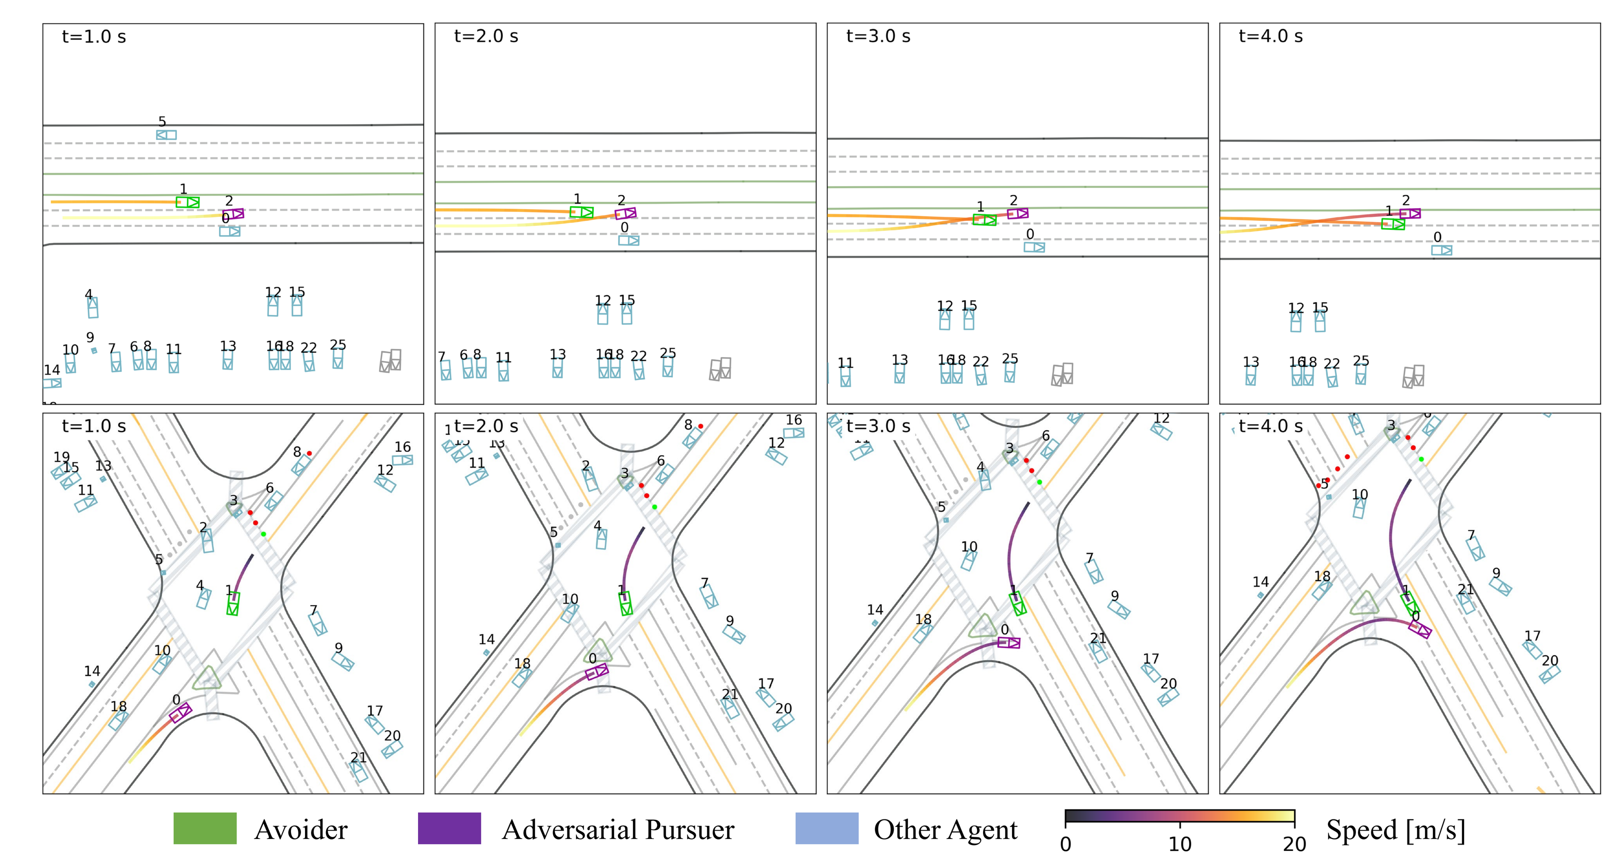

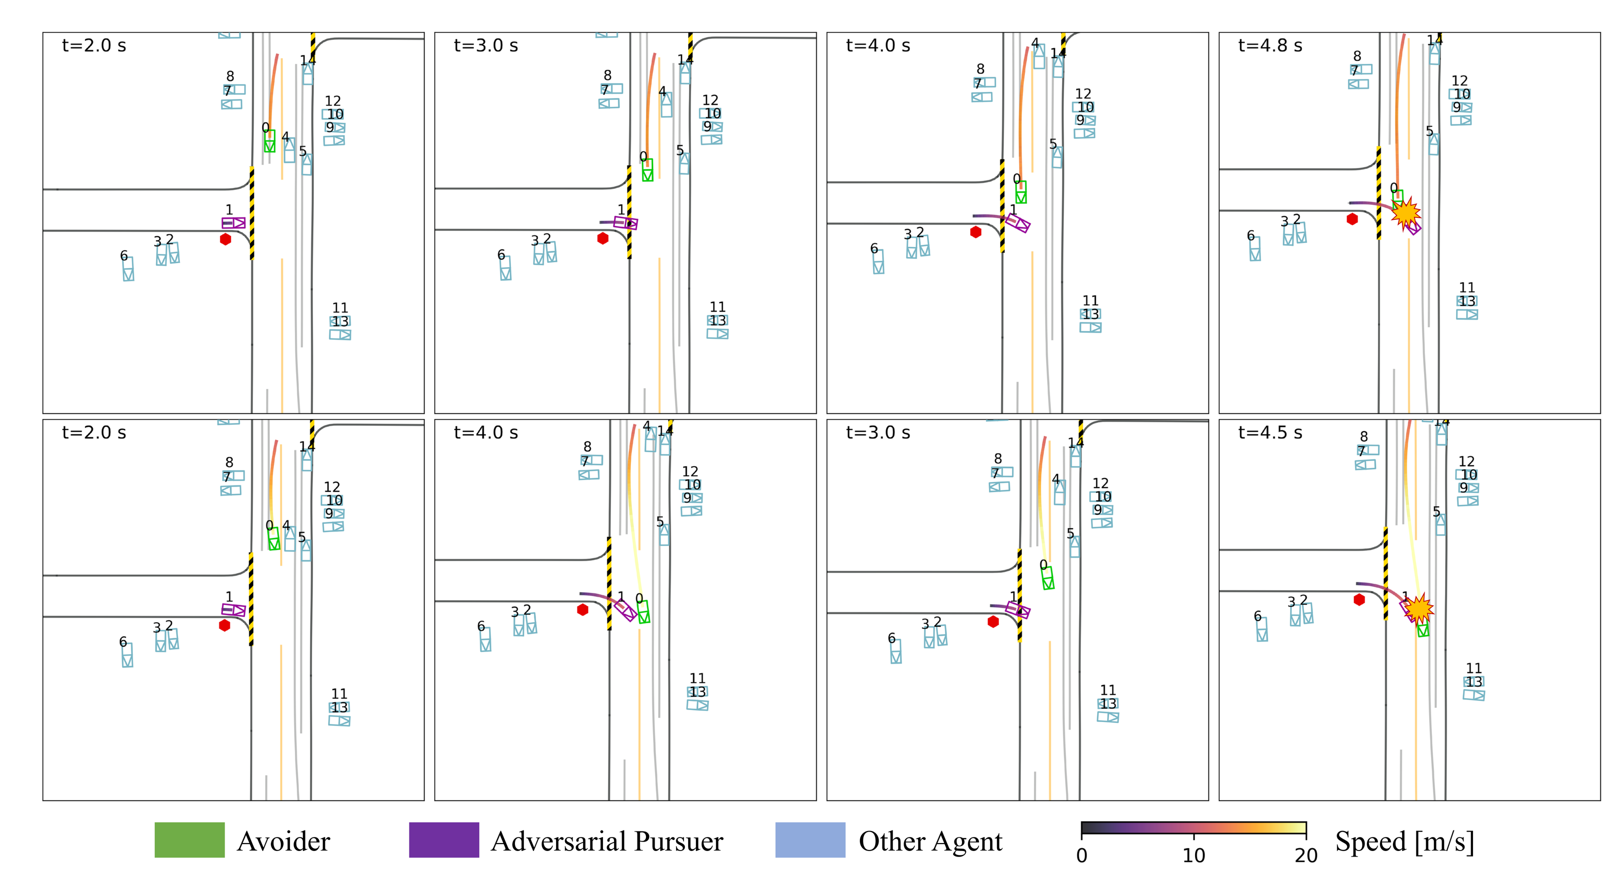

References

Ziebart, Brian D., et al. "**Maximum entropy inverse reinforcement learning**." *AAAI*. Vol. 8. 2008.

Song, Yang, et al. "**Score-Based Generative Modeling through Stochastic Differential Equations**." *ICLR 2021*