KING: Generating Safety-Critical Driving Scenarios for Robust Imitation via Kinematics Gradients

Niklas Hanselmann Katrin Renz Kashyap Chitta
Apratim Bhattacharyya Andreas Geiger

Mercedes-Benz
MAX PLANCK INSTITUTE FOR INTELLIGENT SYSTEMS
EBERHARD KARLS UNIVERSITÄT TÜBINGEN

ECCV TEL AVIV 2022

Motivation

Problem: Critical driving scenarios are extremely rare and *underrepresented* in real world data. Simulation is a promising solution but limited by *hand-crafted behavior* that does not permit the necessary *diversity*.

Goal: *Efficiently* generate safety-critical driving scenarios via optimization and use them to *augment* regular data to improve brittle driving policies.

Limited CARLA Scenarios

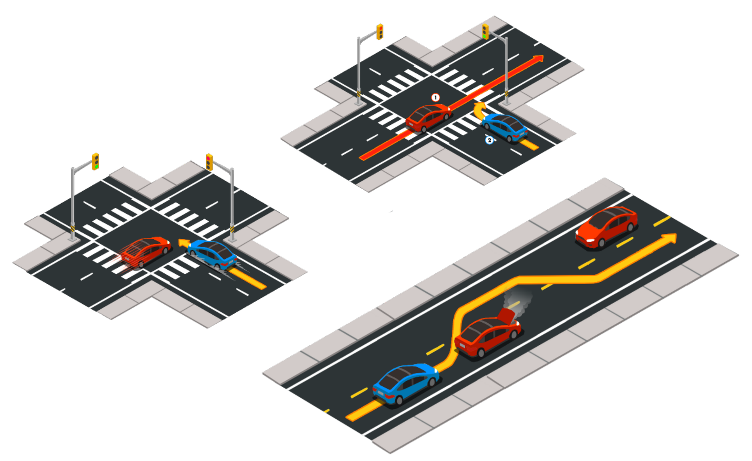

Long-Tail Distribution

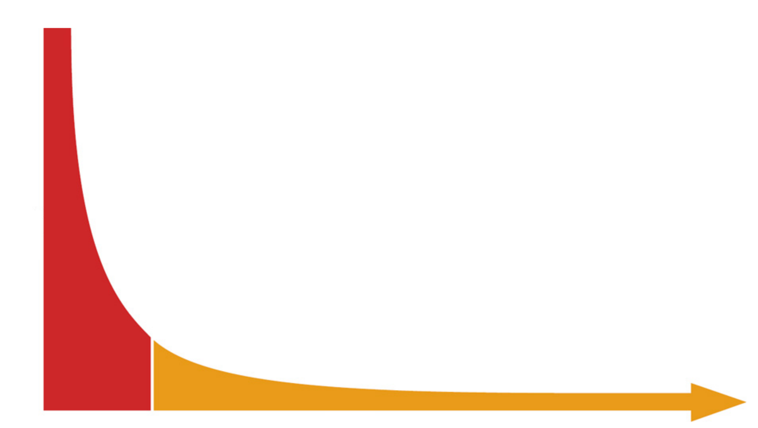

Optimization via Kinematics Gradients

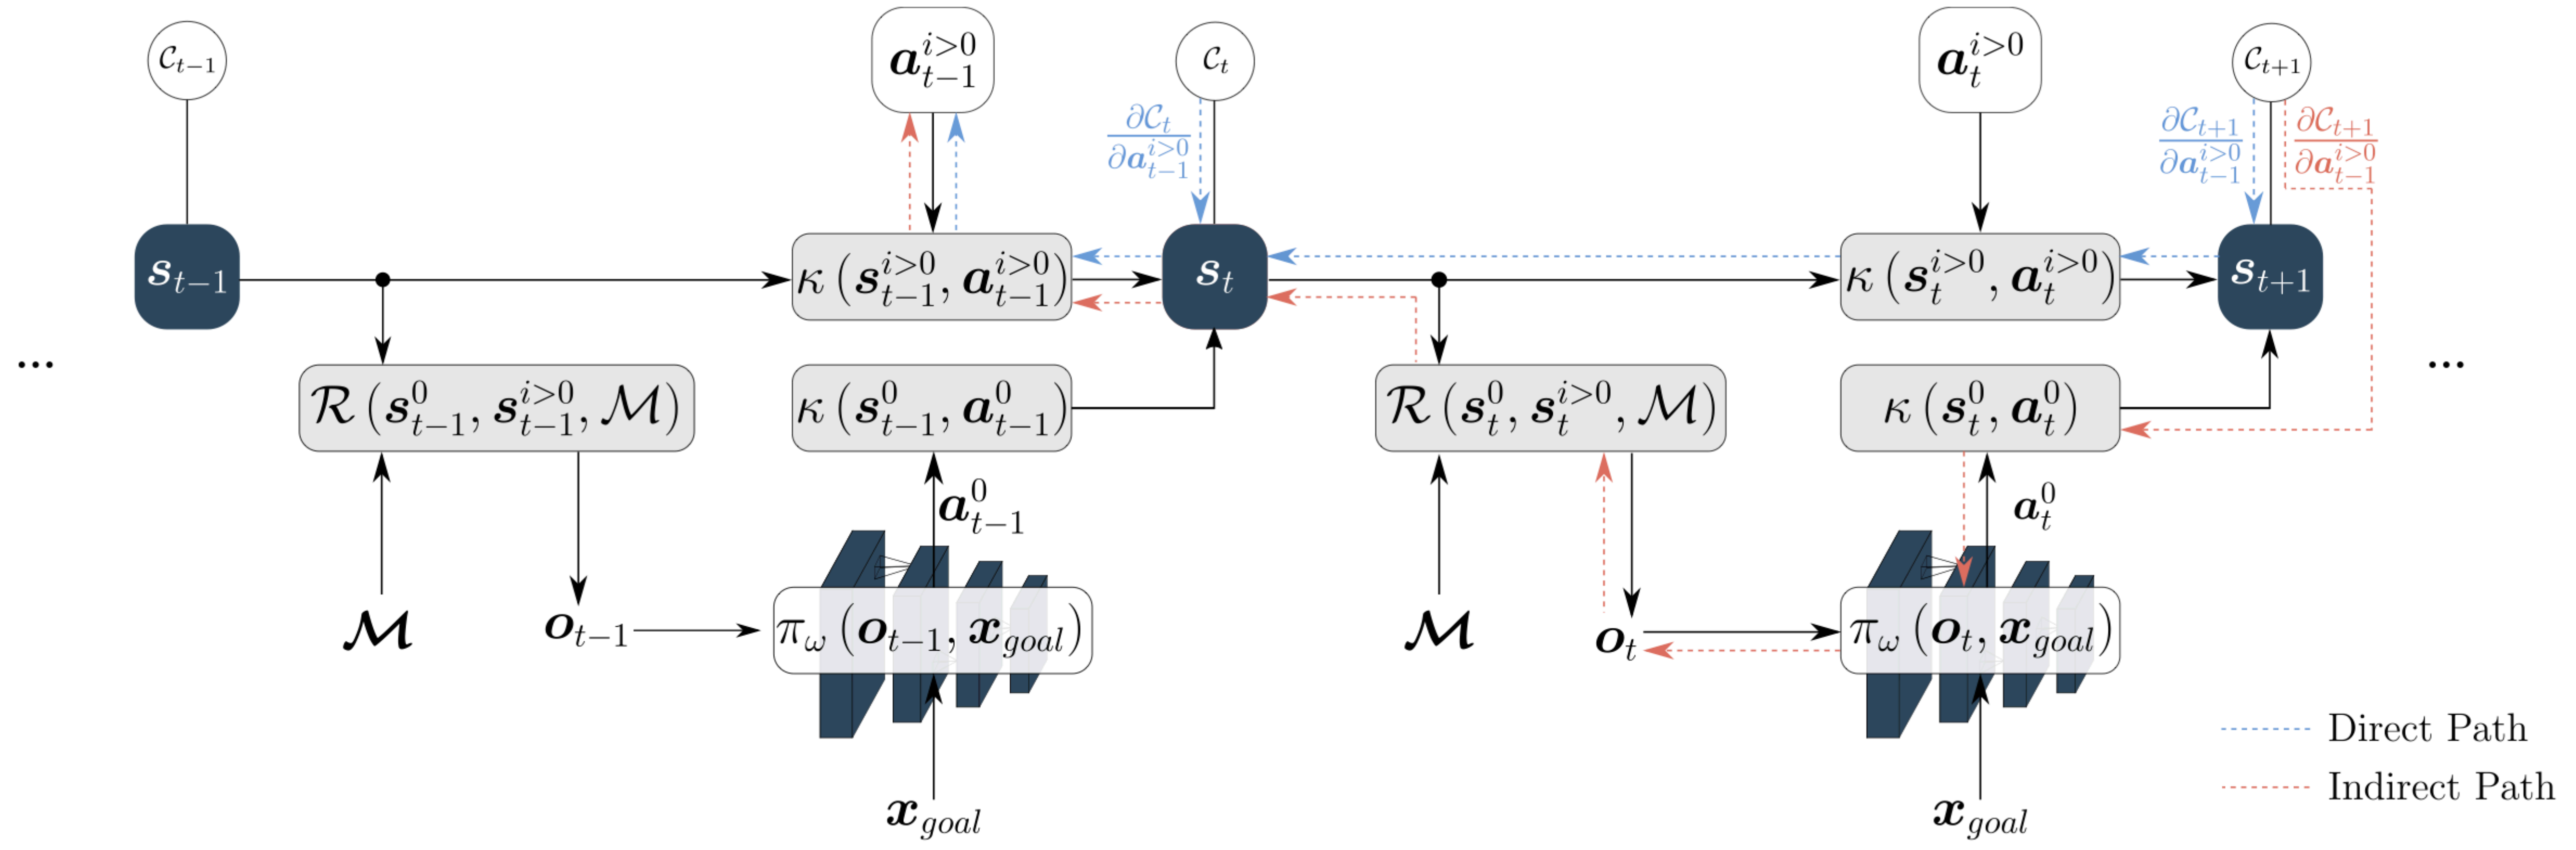

Key Idea: Consider partial derivatives through the direct path only. This has two main advantages:

- Renderer and driving model can be non-differentiable.
- The backward pass is significantly less expensive.

Improved Robustness

Before After

Dataset	Held-out KING scenarios CR ↓	CARLA scenarios DS ↑	CARLA scenarios CR ↓
No Fine-tuning	100.00±0.00	86.74±0.67	17.48±1.86
$\mathcal{D}_{reg}$	57.14±0.00	86.85±0.62	19.51±0.00
$\mathcal{D}_{crit} \cup \mathcal{D}_{reg}$	**28.57±0.00**	**90.20±0.00**	**8.13±0.70**

Procedure

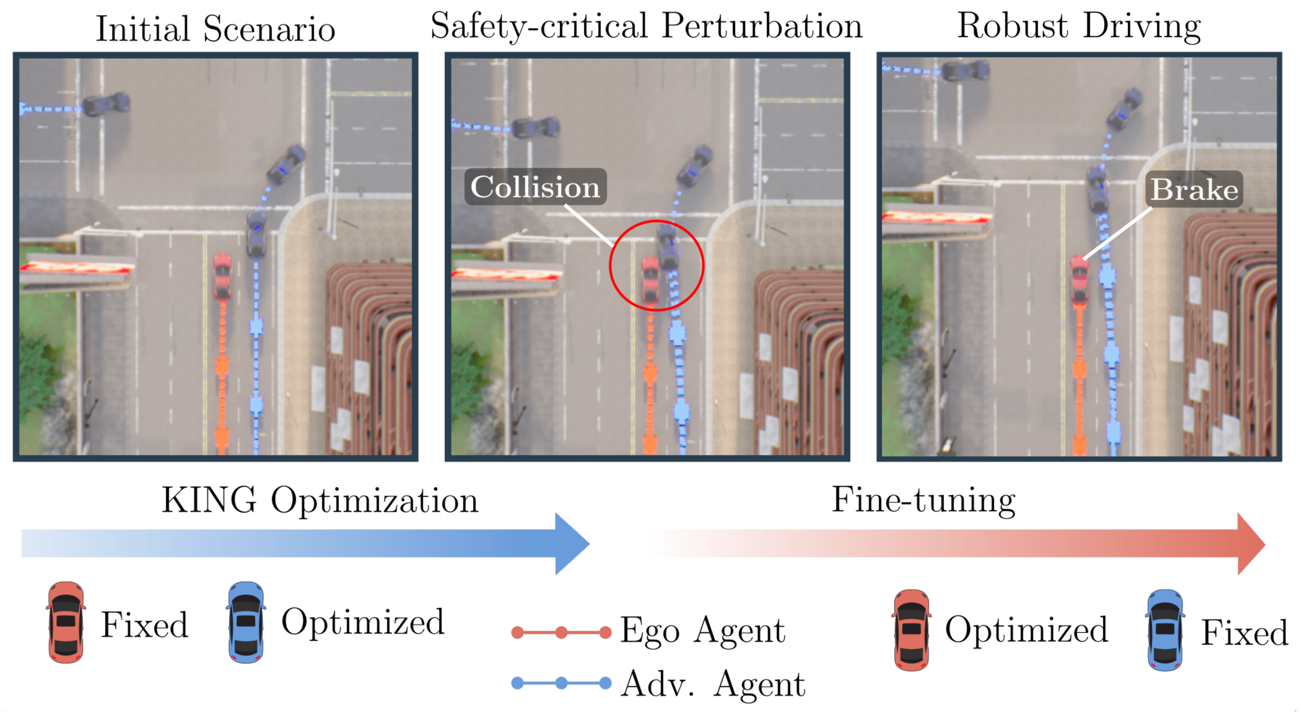

Generation of Diverse, Safety-Critical Scenarios

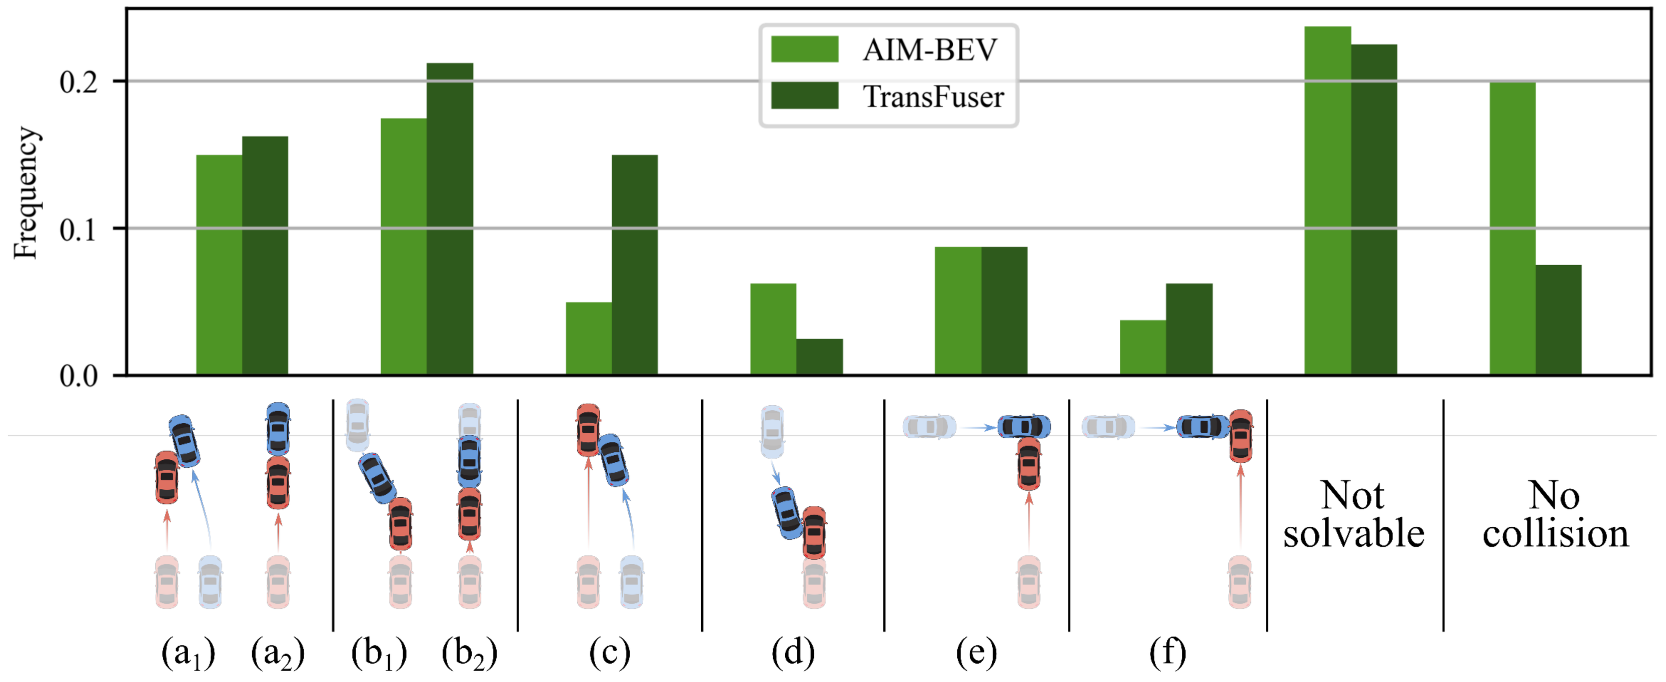

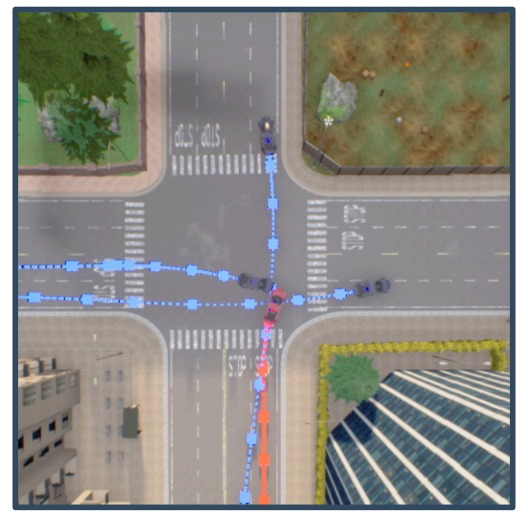

Cluster (e)

Cluster (f)

Better and Faster

Method	4 Agents CR ↑	4 Agents $t_{50\%}$ ↓	4 Agents s/it ↓	Avg. {1,2,4} Agents CR ↑	Avg. {1,2,4} Agents $t_{50\%}$ ↓	Avg. {1,2,4} Agents s/it ↓
Random Search	68.75	15.22	1.48	66.67	9.66	1.38
Bayesian Optimization	63.75	22.12	2.06	65.00	14.34	1.73
SimBA [1]	61.25	19.68	1.48	64.17	15.84	1.38
CMA-ES [2]	62.50	9.39	1.52	68.33	8.17	1.40
Bandit-TD [3]	21.25	-	5.02	29.58	-	4.43
KING (ours)	**78.75**	**6.40**	2.03	**82.50**	**7.78**	1.90

[1] Guo et al., "Simple black-box adversarial attacks", ICML 2019
[2] Hansen et al., "Completely derandomized self-adaptation in evolution strategies", Evolutionary Computing 2001
[3] Ilyas et al., "Prior convictions: Black-box adversarial attacks with bandits and priors", ICLR 2019

https://lasnik.github.io/king/
niklas.hanselmann@mercedes-benz.com