Driving Policy Transfer via Modularity and Abstraction

Presenter: Yasaswini Konthala

October 26, 2023

High-Level Problem Description

- This paper addresses the challenge of transferring driving policies from simulation to reality.

- Autonomous driving is a crucial component of general-purpose robot autonomy, with broad implications for various industries.

- Solving this problem can revolutionize transportation, making it safer, more efficient, and environmentally friendly.

Importance and Impact

- ❖ Autonomous driving systems have the potential to transform industries, from logistics to public transportation, making them more efficient and reducing accidents.
- ❖ Improved autonomous driving can save lives, reduce traffic congestion, and lower carbon emissions.
- ❖ Autonomous vehicles could revolutionize the automotive and transportation industries.

Technical Challenges

- Complex Transfer: Transferring control policies from simulation to reality involves addressing differences in sensor readings, vehicle dynamics, and environmental context.

- Safety Concerns: Safety is paramount; understanding and mitigating risks are critical due to the blackbox nature of end-to-end models.

- Diversity of Scenarios: Realistic urban driving scenarios require a huge amount of data to cover the full spectrum of driving conditions.

Key Insights

- The proposed approach leverages modularity and abstraction to encapsulate the driving policy, enabling direct transfer from simulation to reality.

- The system architecture involves three stages: **perception, driving policy**, and **low-level control**, improving adaptability and robustness.

- Training the driving policy on real-world perception data, rather than perfect ground-truth data, enables it to adapt to real-world imperfections.

Problem Setting

❖ The authors of the paper aim to transfer driving policies from simulation to reality effectively.

❖ Objective: Achieve direct transfer of driving policies without the need for retraining or fine-tuning in real-world conditions.

❖ Challenge: Overcoming the reality gap, encompassing differences in sensor data, dynamics, and environmental context between simulation and reality.

Context / Related Work

- **Transfer from Simulation to Reality:** Previous studies have extensively explored transferring knowledge from simulation to real-world scenarios in computer vision and robotics.

  - Use of synthetic data for training and evaluation in indoor and driving scenarios.

  - Notable challenges persist in direct transfer despite high-fidelity simulation.

Context / Related Work

- **Transfer of Sensorimotor Control Policies:** Prior work primarily focused on manual grasping and manipulation, utilizing specialized learning techniques and network architectures.
  - Techniques include domain adaptation, depth maps, and domain randomization.

- **Transfer of Driving Policies:** Transfer of driving policies has historical significance, with early attempts focusing on rudimentary lane following.
  - More recent efforts have addressed obstacle avoidance and collision prevention in smaller robotic vehicles.

Limitations of Prior Work

- **Complexity and Robustness:** Existing approaches often lack the complexity and robustness required for outdoor urban driving.

- **Nuisance Factors:** Some approaches use domain randomization to handle the differences between simulation and reality, which may not be directly applicable to urban driving.

- **Modularity**: Previous modular designs result in error accumulation and require meticulous engineering in each component.

Proposed Approach

- The proposed autonomous driving system comprises three components: a perception module, a driving policy, and a low-level controller.

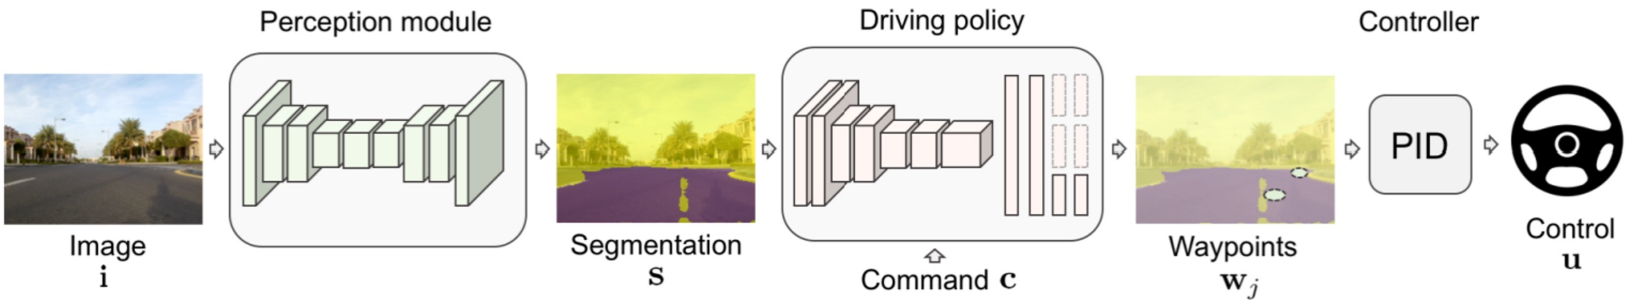

Figure : System

architecture

Perception Module

- The **Perception Module** processes a raw RGB image, generating a binary segmentation map that abstracts away unnecessary information and retains essential data.

- A convolutional neural network, based on the ERFNet architecture, is utilized for binary road segmentation, trained on the real-world Cityscapes dataset.

- The perception module's role is to filter out nuisance factors, such as texture, lighting, and weather, while preserving critical information.

Table 3: ErfNet-Fast architecture, used as perception module in our method.

Layer	Type	out channels	out resolution
1	Downsampler block	16	100×44
2-6	$5 \times$ Non-bt-1D	16	100×44
7	Downsampler block	64	50×22
8	Non-bt-1D (dilated 2)	64	50×22
9	Non-bt-1D (dilated 4)	64	50×22
10	Non-bt-1D (dilated 8)	64	50×22
11	Non-bt-1D (dilated 16)	64	50×22
12	Deconvolution (upsampling)	16	100×44
13-14	$2 \times$ Non-bt-1D	16	100×44
15	Deconvolution (upsampling)	2	200×88

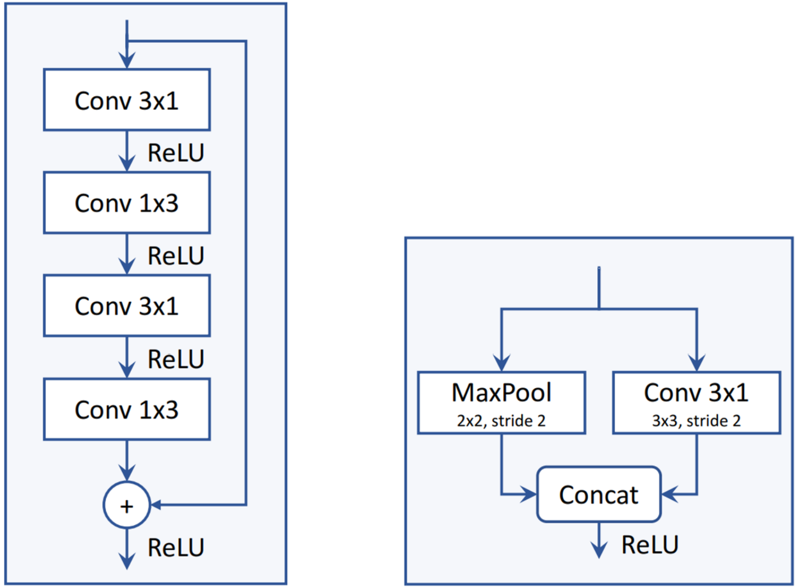

Non-bottleneck 1D

Downsampler

Figure 8: The ERFNet [39] building blocks, used in our architecture.

Driving Policy

- The **Driving Policy** receives the segmentation map and outputs a local trajectory plan represented by waypoints.

- The waypoints w_j are encoded by the distance r_j and the (oriented) angle ϕ_j with respect to the heading direction v of the car.

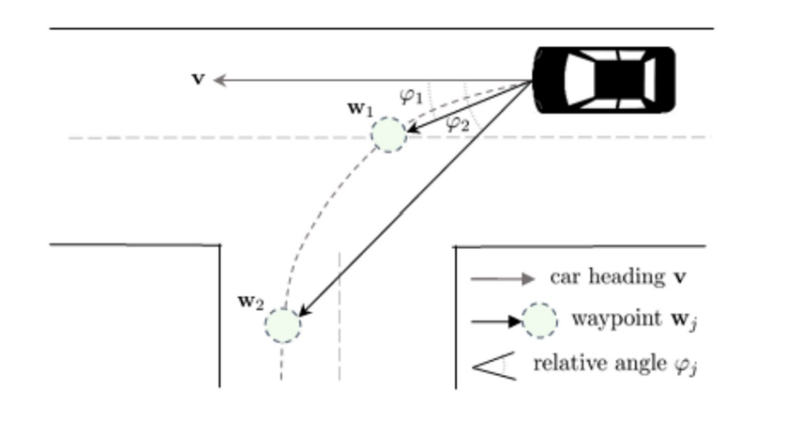

- Trained in simulation using conditional imitation learning (CIL), allowing conditioning on high-level commands for navigation.

- Dataset: $\{o_i , c_i , a_i\}^N_{i=1}$ of observation-command-action tuples, from trajectories of an expert driving policy
  - An **observation** can be an image or a segmentation map;
  - The **action** can be either vehicle controls (steering, throttle) or waypoints;
  - The **command** is a categorical variable indicating one of three high-level navigation instructions (left, straight, right)

❖ Used a deep network as the function approximator and adopt the branched architecture of Codevilla et al. [8], with a shared convolutional encoder and a small fully-connected specialist network for each of the commands.

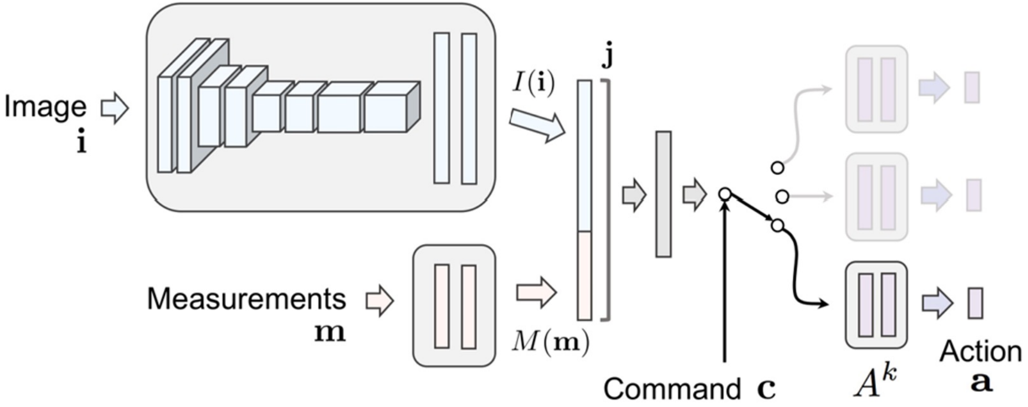

Branched architecture of Codevilla et al. [8]

Control

- **Control** is achieved using PID controllers for throttle and steering angle.
- Throttle PID control uses the difference between target and current speed as the error.
- Steering angle PID control employs the oriented angle between the viewing direction and the direction to the first waypoint (ϕ_1) as the error.

Experimental Setup

- **Two domains:** simulation and a physical 1/5-scale truck

- **Simulation**: CARLA, an open-source simulator for urban driving, which provides two towns (Town 1 and Town 2) with different layouts, sizes, and visual styles.

- **Physical System:** A modified 1/5-scale Traxxas Maxx truck equipped with hardware, the onboard computer, an Nvidia TX2, is responsible for most of the computation.

The authors compare their approach with several baselines:

1. **Image to Control (img2ctrl):** Predicting low-level control directly from color images.
2. **Image to Waypoint (img2wp)**: Predicting waypoints directly from color images.
3. **Segmentation to Control (seg2ctrl)**: A policy that predicts low-level control from segmentation maps.
4. **Segmentation to Waypoint:** The proposed full model that predicts waypoints from segmentation maps.
5. Variants of these models trained with data augmentation.

Evaluation Metrics

- ❖ The success of their approach is evaluated using a protocol where 25 start-goal pairs in each town are tested, and the percentage of successfully completed episodes is measured.

- ❖ Success is defined as the vehicle reaching the goal point given high-level commands.

- ❖ Additional metrics related to road - including the success rate over different navigation trials with varying distances to be driven.

Experimental Results

Driving in Simulation:

- ❖ The proposed approach outperforms baselines in both training and generalization to the test environment and weather conditions.
- ❖ Data augmentation further enhances performance, particularly in the challenging Town 2/Weather 2 scenario.

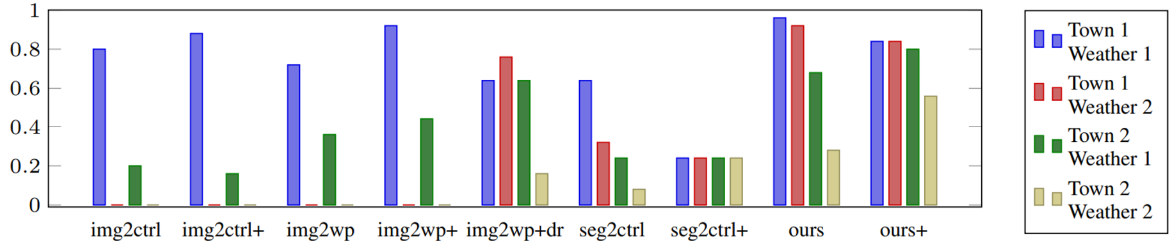

Experimental Results

Driving in the Physical World

- ❖ The proposed modular approach achieves 100% success in the real world with data augmentation.

- ❖ Surprisingly, **img2wp** performs better in some real-world trials.

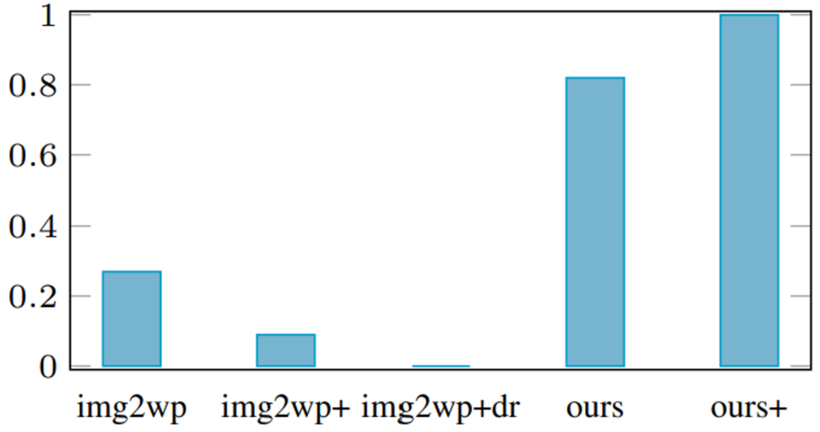

Discussion of Results

- ❖ The experimental results indicate the robustness of the proposed modular approach for autonomous urban driving.
- ❖ The approach outperforms baseline methods in both simulation and the real-world physical driving environment.
- ❖ Remarkably, it demonstrates strong generalization to complex real-world conditions, without the use of real-world data during training.

Discussion of Results

- ❖ The results confirm the effectiveness of the proposed perception-to-action pipeline, with waypoint prediction based on segmentation data.
- ❖ The waypoint-based approach provides significant benefits for generalization to diverse conditions, outperforming control-based and image-to-waypoint baselines.
- ❖ Data augmentation enhances performance, especially in challenging scenarios, such as Town 2 and Weather 2.

Key limitations of the approach

- Data Intensive Training: The approach relies heavily on simulation data
- Segmentation Reliability: The accuracy of the perception module's segmentation may be sensitive to real-world variability
- Real-world Generalization: While the method shows promise, there is still room for improvement in real-world conditions, especially when dealing with complex, unexpected scenarios.

Practical challenges in the Real-World deployment

- Sensor Realism: Ensuring the fidelity and robustness of onboard sensors and real-world sensor data collection can be challenging.

- Safety Concerns: Ensuring that the system behaves safely in unforeseen situations is a significant challenge.

- Dynamic Environments: Handling interactions with pedestrians, cyclists, and other road users, especially in urban settings, is a complex real-world challenge.

Future Work

- ❖ Develop methods for ensuring safety in real-world deployments. This includes handling unforeseen situations, adhering to traffic rules, and understanding and predicting the behavior of other road users.
- ❖ Enhance the system's capability to predict and respond to the intentions of pedestrians, cyclists, and other drivers, considering human behavior models.
- ❖ Investigate how autonomous vehicles can collaborate with each other and with non-autonomous vehicles to optimize traffic flow and safety.

Extended Readings

[1] M. Bojarski, D. D. Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, et al. End to end learning for self-driving cars.

[2] F. Codevilla, M. Muller, A. Dosovitskiy, A. L ¨ opez, and V. Koltun. End-to-end driving via conditional ´ imitation learning.

[3] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun. CARLA: An open urban driving simulator.

[4] A. Paszke, A. Chaurasia, S. Kim, and E. Culurciello. ENet: A deep neural network architecture for real-time semantic segmentation.

Summary

❖ The paper focuses on achieving autonomous urban driving, with the central challenge of enabling AI systems to generalize from simulation to real-world scenarios effectively.

❖ The paper introduces a modular approach that decouples perception, planning, and control.

❖ The proposed approach showcases remarkable generalization from simulation to both Town 1 and Town 2 environments, marking a significant step towards the practical deployment of autonomous vehicles in urban settings.