BIMI LAB
BIO-INSPIRED MACHINE INTELLIGENCE

DEARBORN

B-Morpheus: Perceptual-Motor Active Inference for Autonomous Driving

KMU 국민대학교

Jaerock Kwon, PhD
Assistant Professor
Electrical and Computer Engineering
University of Michigan-Dearborn

May 27, 2024

How to Make a Car Drive Itself?

Conventional Approach

Steering, Acceleration, Brake

- Design feature extractors and controllers with stochastic/deterministic models.

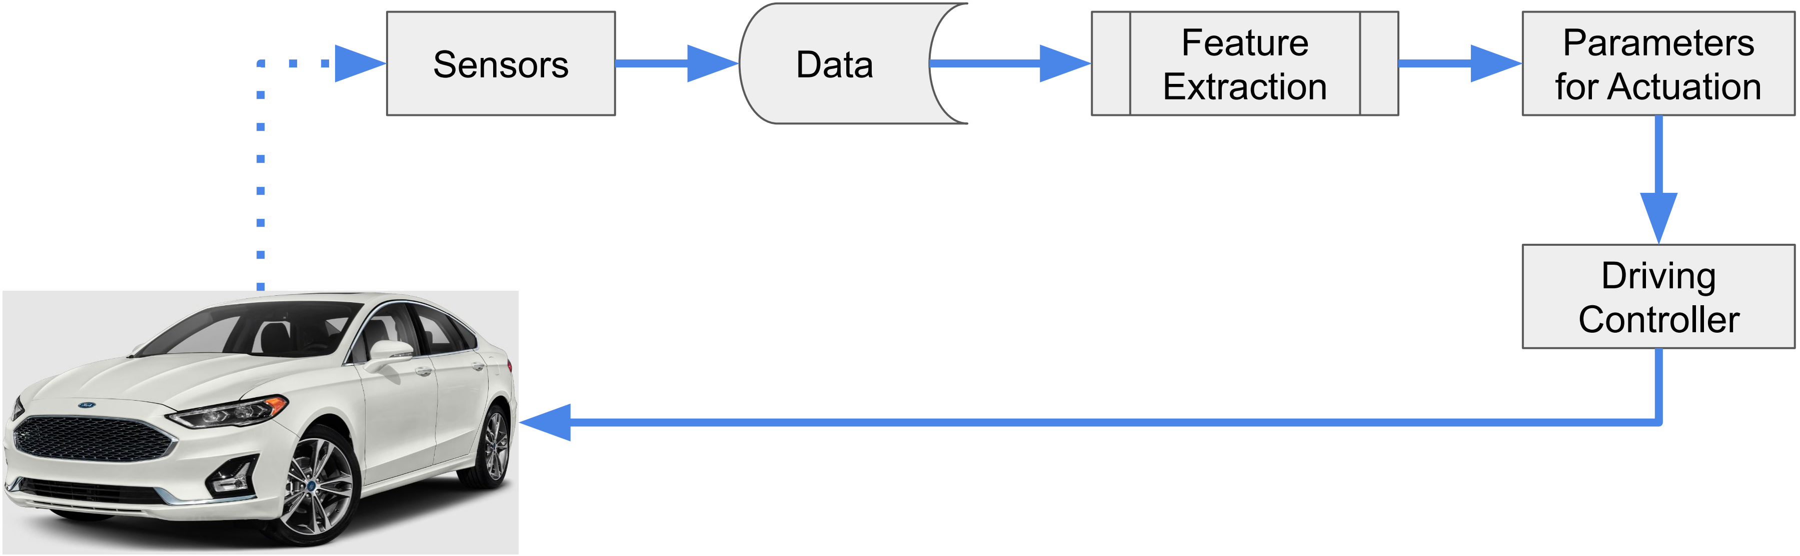

Behavior Cloning Approach

Steering, Acceleration, Brake

- Design feature extractors and controllers with stochastic/deterministic models.

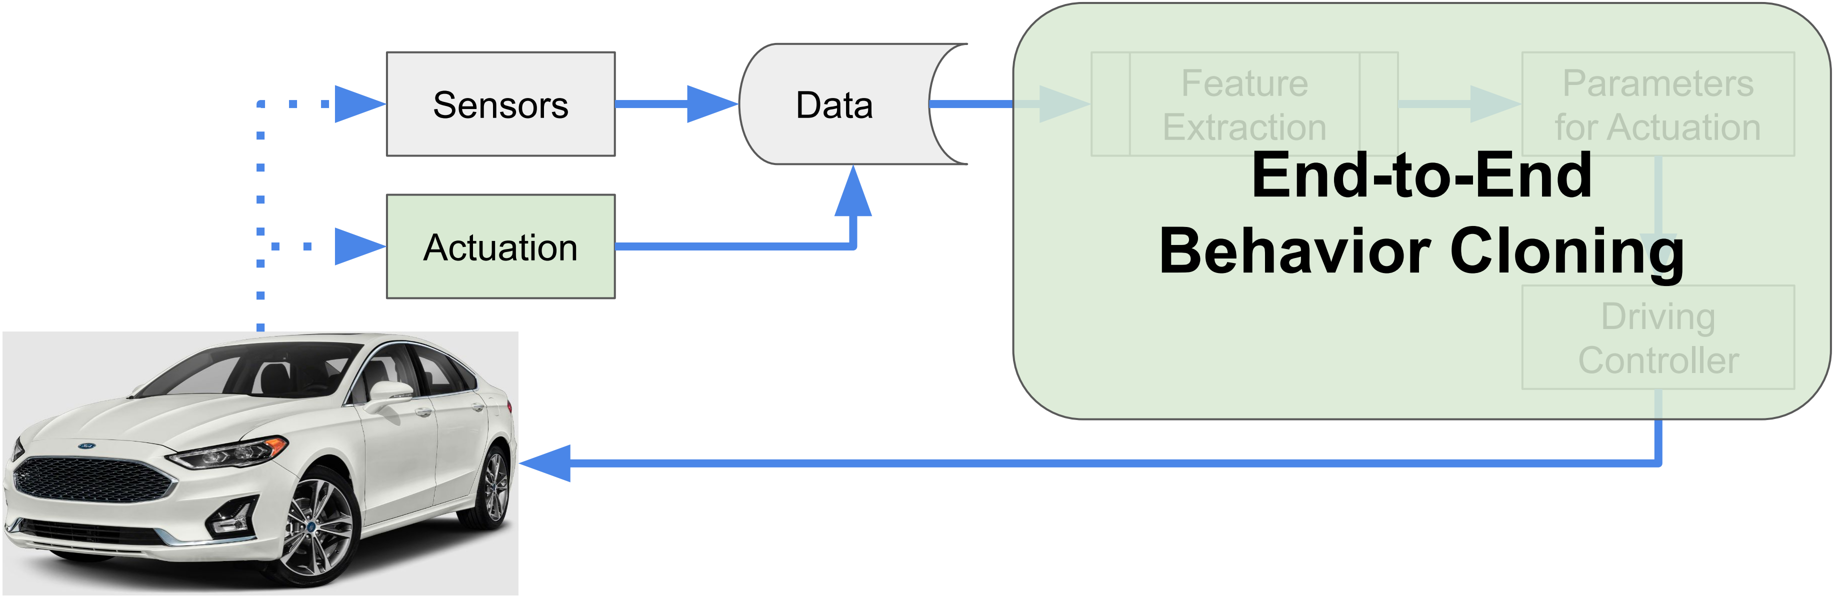

Why Behavior Cloning?

- In ideal conditions, conventional feature extraction and model based approaches work well.
- Yet, in practice, it is difficult to get proper actuation parameters for a model in bad road conditions and/or adversarial weather conditions.

Behavior Cloning

Behavior Cloning

- Try to mimic an expert behavior on a task.

- Sensory input: x → Action output: y
- Collect an expert behavior in a tuple format: (input: x, action y).
    - input → [expert] → output

- If we can find a nonlinear function to map the input x into the output y,
- We can infer a proper output without modeling a complex internal state.

ALVINN

In 1989 by Dean A. Pomerleau

Steer like I do

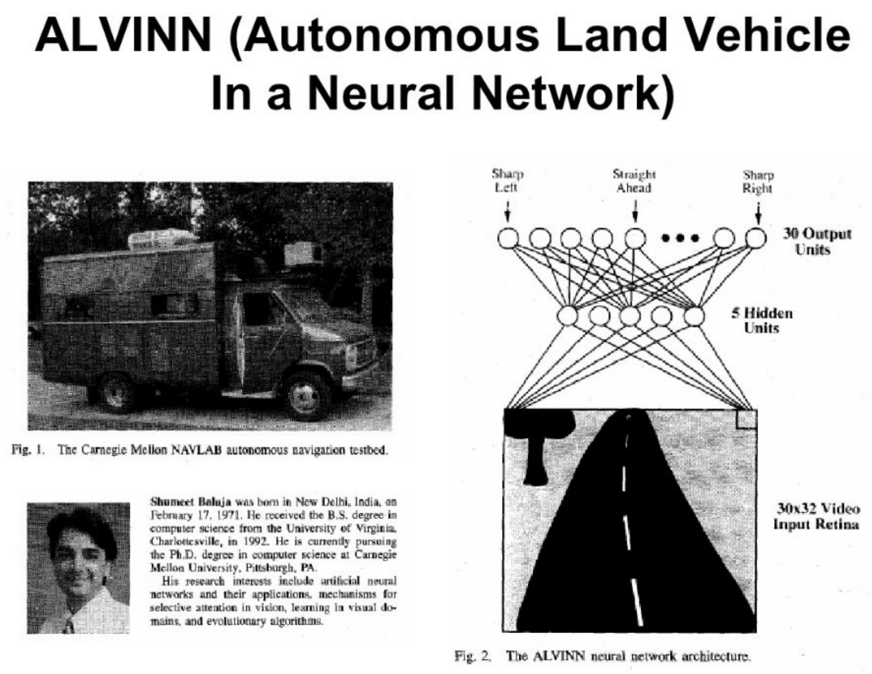

DAVE:

In 2003

- DARPA Autonomous Vehicle Experiment (DAVE)
- A Convolutional Neural Network (CNN: Proposed by LeCun in 1998) was used.
  - Six layers
    - 3,148,776 connections
    - 71,886 independent parameters
- Training data: 225K samples from 3,200 runs

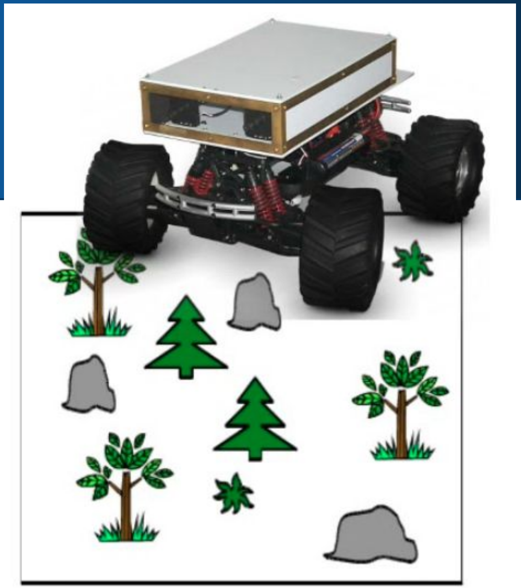

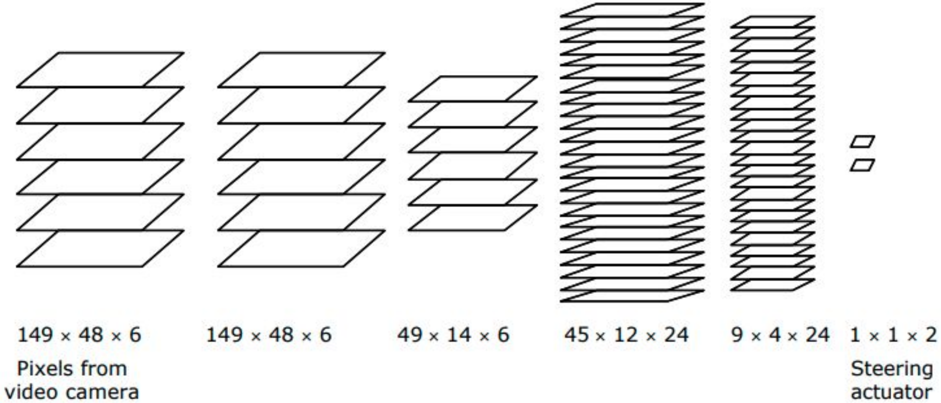

NVIDIA - PilotNet (*originally DaveNet*)

In 2016

- Neural Network Architecture:
  - PilotNet:
    - CNN
    - 27 million connections
- Hardware: NVIDIA PX2
  - 12 CPU cores, Pascal GPUs

NVIDIA DRIVE PX Specification Comparison		
	DRIVE PX	DRIVE PX 2
SoCs	2x Tegra X1	2x Tegra "Parker"
Discrete GPUs	N/A	2x Unknown Pascal
CPU Cores	8x ARM Cortex-A57 + 8x ARM Cortex-53	4x NVIDIA Denver + 8x ARM Cortex-A57
GPU Cores	2x Tegra X1 (Maxwell)	2x Tegra "Parker" (Pascal) + 2x Unknown Pascal
FP32 TFLOPS	> 1 TFLOPS	8 TFLOPS
FP16 TFLOPS	> 2 TFLOPS	16 TFLOPS?
TDP	N/A	250W

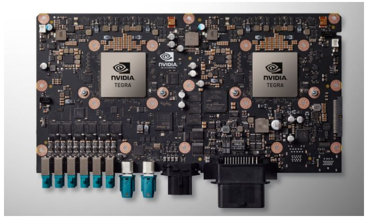

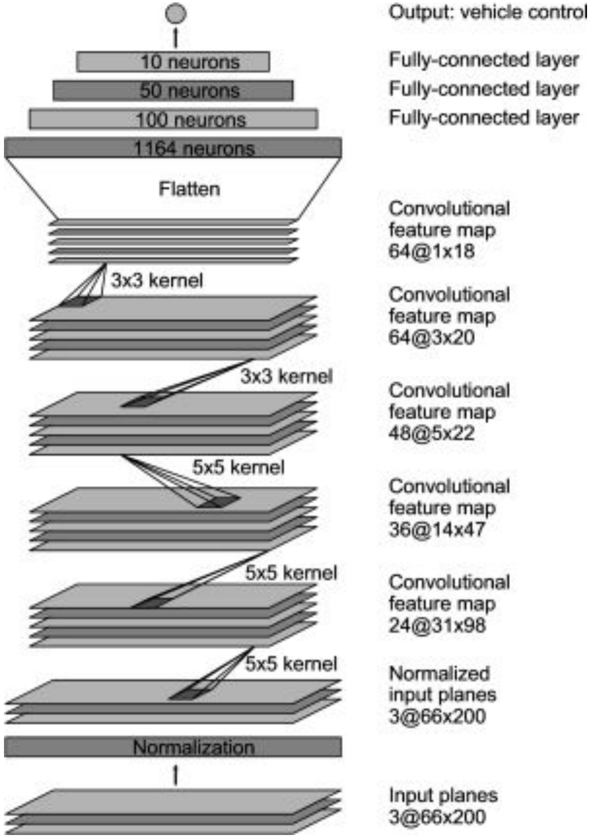

NVIDIA - PilotNet (*originally DaveNet*)

In 2016

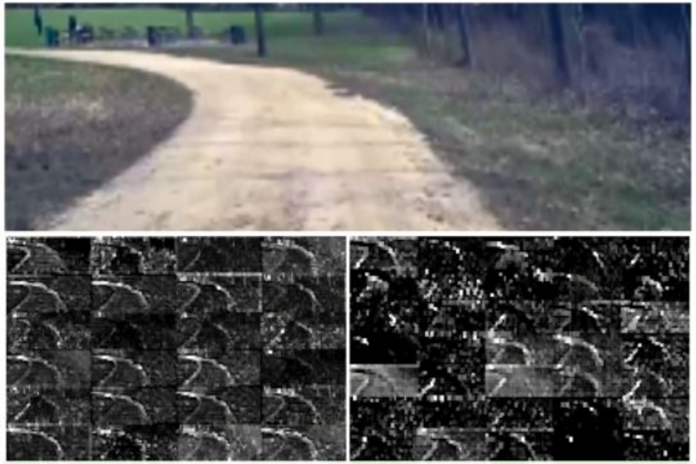

Behavior Cloning

Training

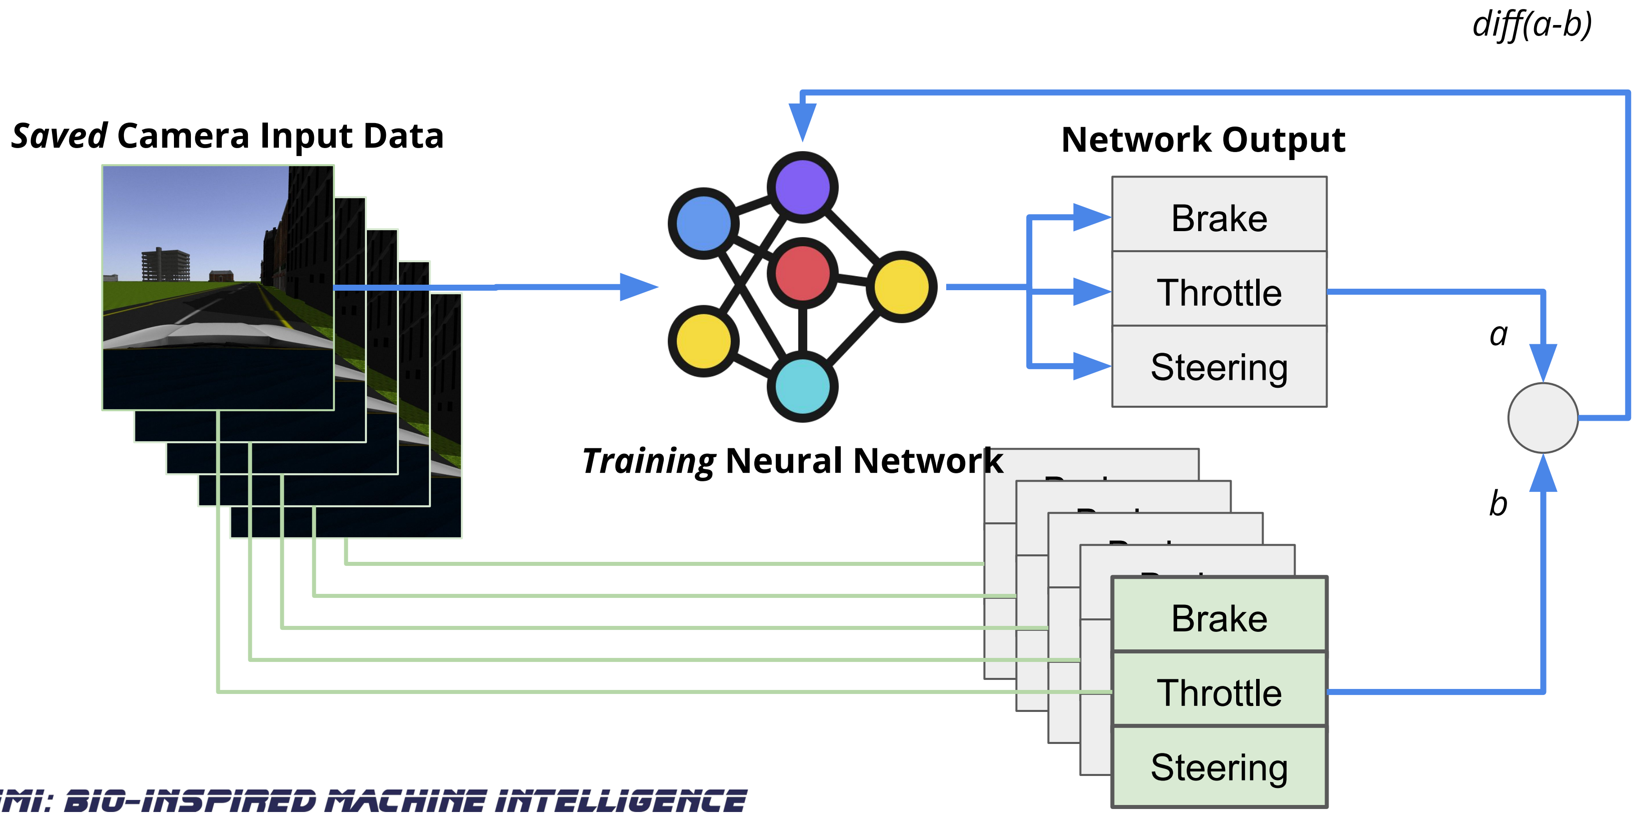

Behavior Cloning

Inference

- Learn a policy to **map** an input image to control signals (steering, throttle, brake)

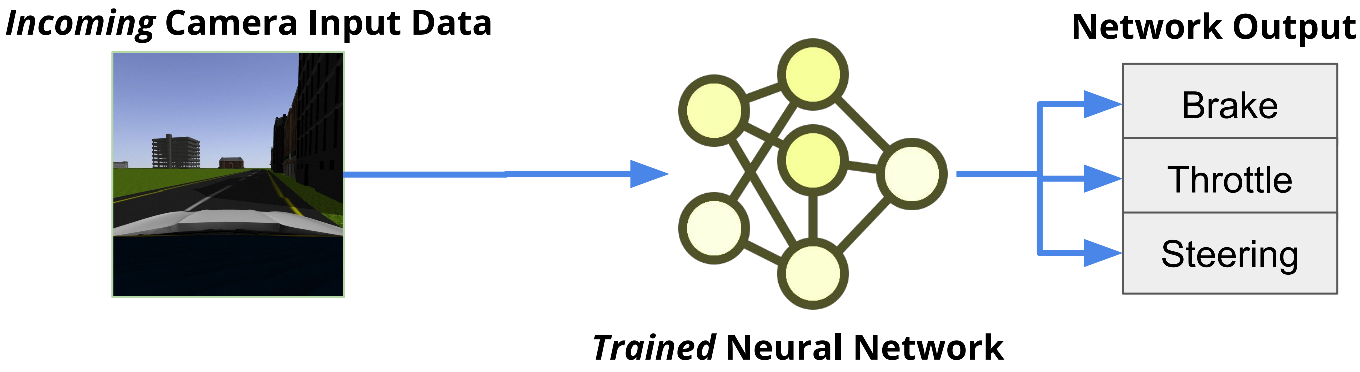

But, Wait…
Does a Human Learn How to Drive in This Way?

Problems in Behavior Cloning

Behavior Cloning

- Try to mimic an expert behavior on a task.

- Sensory input: x → Action output: y
- Collect an expert behavior in a tuple format: (input: x, action y).
    - input → [expert] → output

- If we can find a nonlinear function to map the input x into the output y,
- We can infer a proper output without modeling a complex internal state.

Behavior Cloning

- Try to mimic an expert behavior on a task.
  ← Mapping without knowing the reason of the behavior.
- Sensory input: x → Action output: y
- Collect an expert behavior in a tuple format: (input: x, action y).
  - input → [expert] → output
    ← (1) Lack of Important but not frequent (e.g., restore from errors) data.
    ← (2) Test in an environment with different probabilistic distribution from the training data (distribution shift)

- If we can find a nonlinear function to map the input x into the output y,
- We can infer a proper output without modeling a complex internal state.

Human Learning for Driving

- Humans **do not learn** the mapping: sensory input ⇒ motor output.

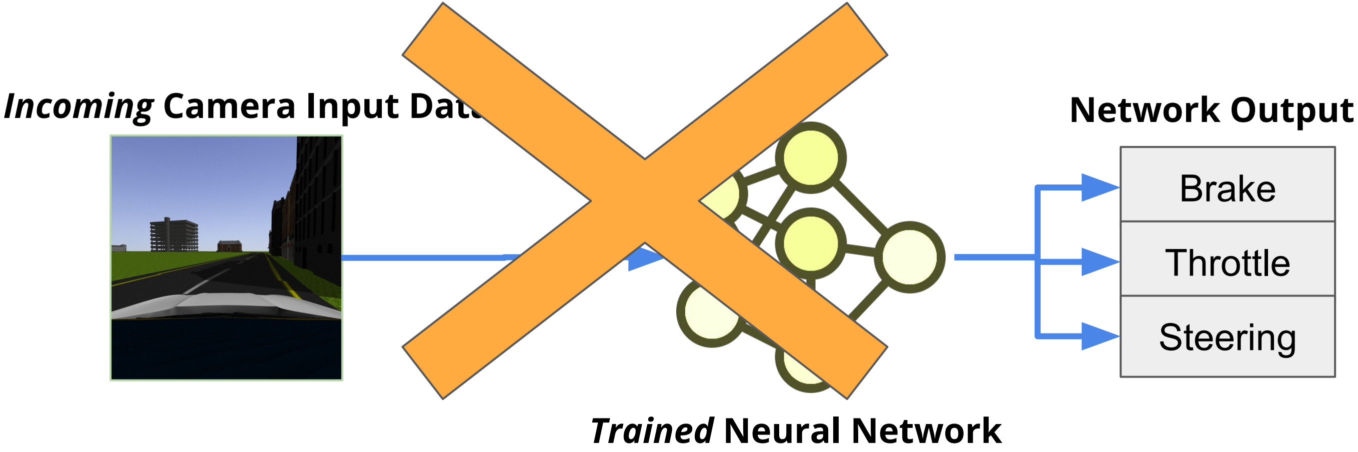

Human Learning for Driving

Learn by **Active Motor Babbling**

- *Causally* associating sensory feedback with motor actuation.

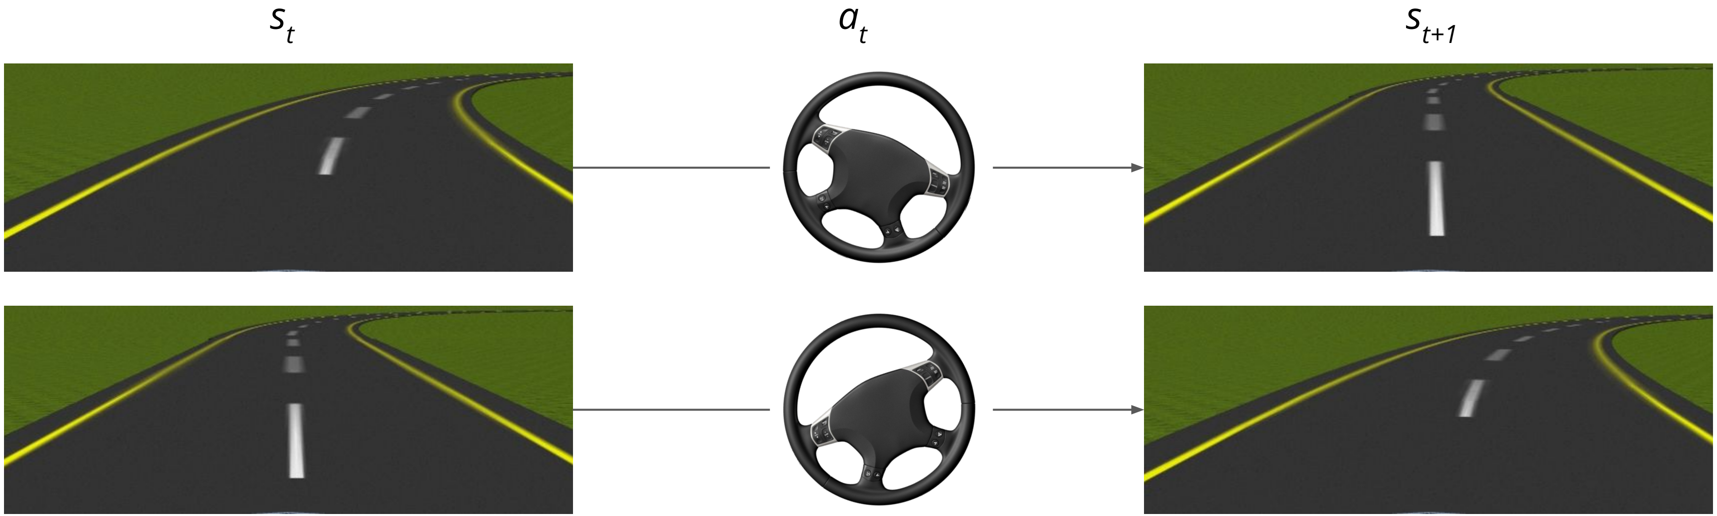

How to Learn Driving Like Humans Do?

Conventional Methods

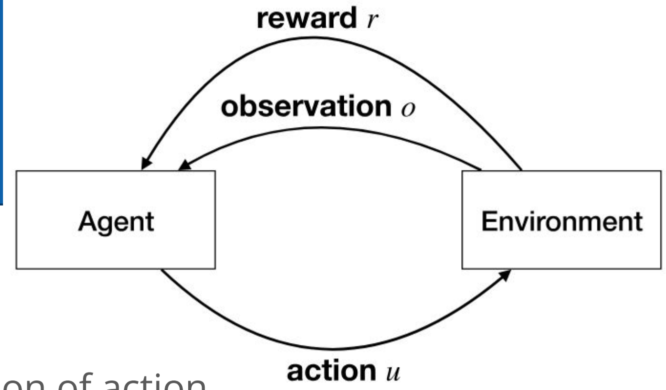

- Model Predictive Control (MPC)
  - See control problems as optimization.
  - Select an action that has the minimum cost from a horizon of action sequences in a dynamic physics environment.
  - **Challenging to define an exact model and cost functions.**
- Reinforcement Learning
  - Model-free: through trial-and-error-based search an action sequences that are expected the maximum reward.
  - Model-based: through mathematical models, search and action sequences that are expected the maximum reward.
  - **Challenging to define reward functions in complex real-world environment.**

Active Inference

DEARBORN

- Integrated model of perception and control based on the brain's perception model.
- Active Inference based on Free Energy Principle
  - (1) Deep neural generative models
  - (2) Motor imagery
- Find optimal control signals to minimize Free Energy.
- Advantages:
  - No need to define system models
  - No need to define cost or reward functions
  - Strong to distribution shift issues because it is not simple mapping from input to output.
  - Possible to explain the reason of an action because it is not trial and error-based but a cause-and-effect-based method.

Active Inference

Free Energy Principle

Free Energy Principle

- Proposed by Karl Friston.
- Originally from thermodynamics in a closed system with constant temperature and volume.
  - Helmholtz free energy: F = U - TS (U: internal energy, T: temperature, S: entropy)
  - Gibbs free energy: G = U - TS + PV (P: pressure, V: volume)
- Friston's free energy principle:
  - An information theory quantity that bounds the evidence for a model (encoded by the brain) of data (sensory input).
  - Free energy is greater than the negative log-evidence or 'surprise' in sensory data.
  - Under simplifying assumptions, it is just the **amount of prediction error**.
- "Systems change to decrease their free energy."

Variational Free Energy

- VFE is an information-theoretic quantity that is minimized during variational inference.
- Variational (or Bayesian) Inference
  - Approximation of complex problems into a simpler one.
  - Make an inference problem to an optimization problem by adding variational parameters.
- The brain posses hierarchical generative models capable of generating expected sensory data, which learn by minimizing the prediction error between the predicted and observed sensory data.
- **Active Inference extends this idea by applying it to action.**

Active Inference

Perception System

Probabilistic Inference Engine

- **Perception Model**: to infer probability distributions about a cause of sensory information from outside.
- **Generative Model**: prior probability distribution established via perception model

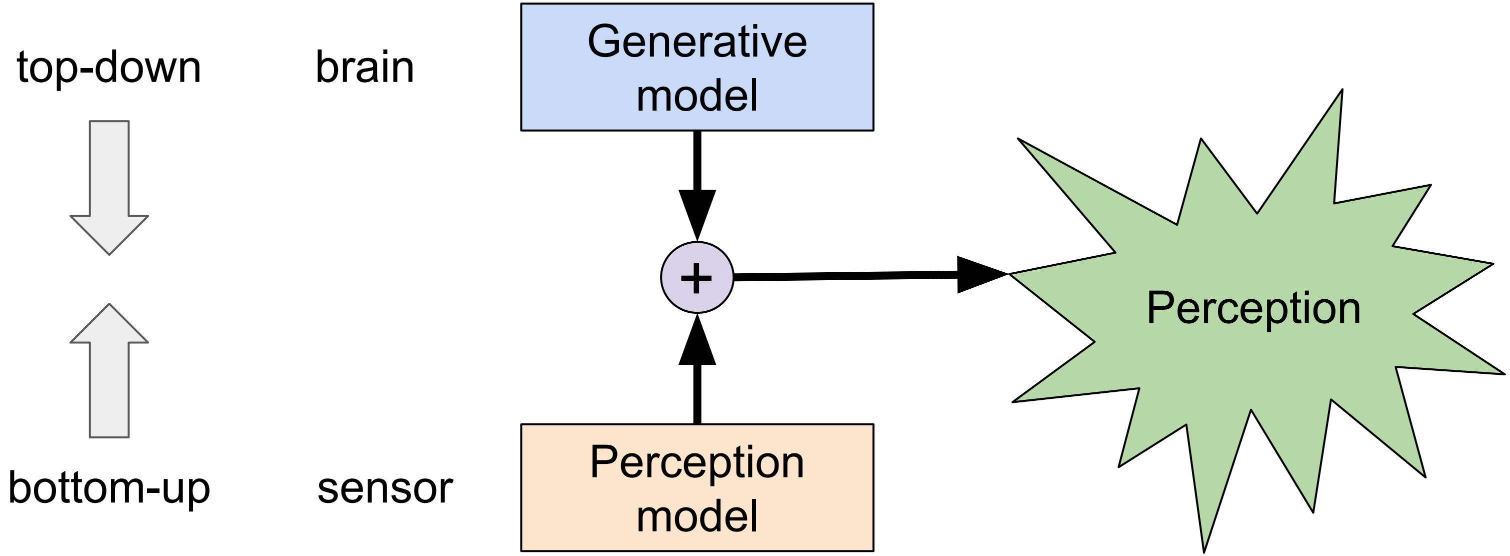

How to perceive the world (1/2)

- Information about the world → [sensory apparatus] → compressed and abstracted.

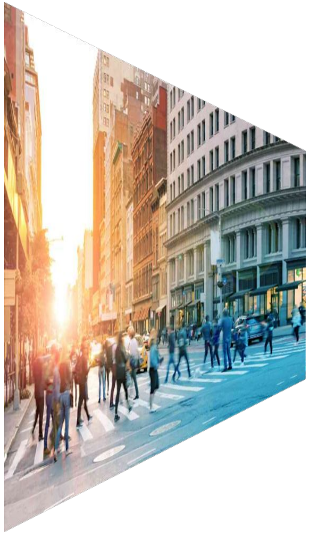

How to perceive the world (2/2)

- How to infer the meaning of the stream of compressed electrical signals
  - (1) Make changes on the world by **actions**.
  - (2) Sensory organs receive the **changes** caused by the action.
  - (3) Find the pattern of sensory information and action that causes the change.

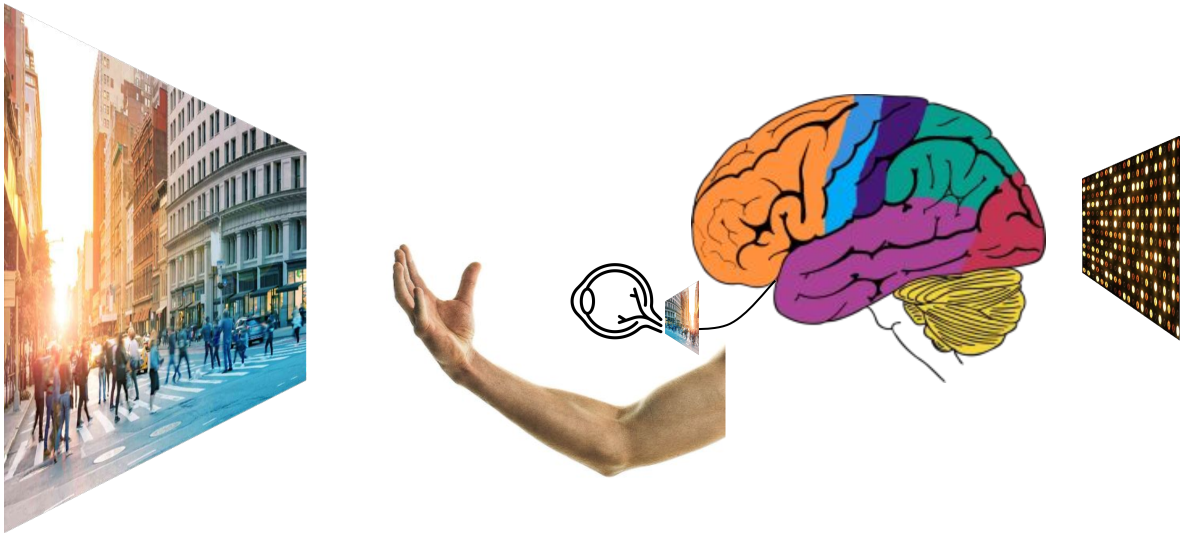

Perception and Action in Active Inference

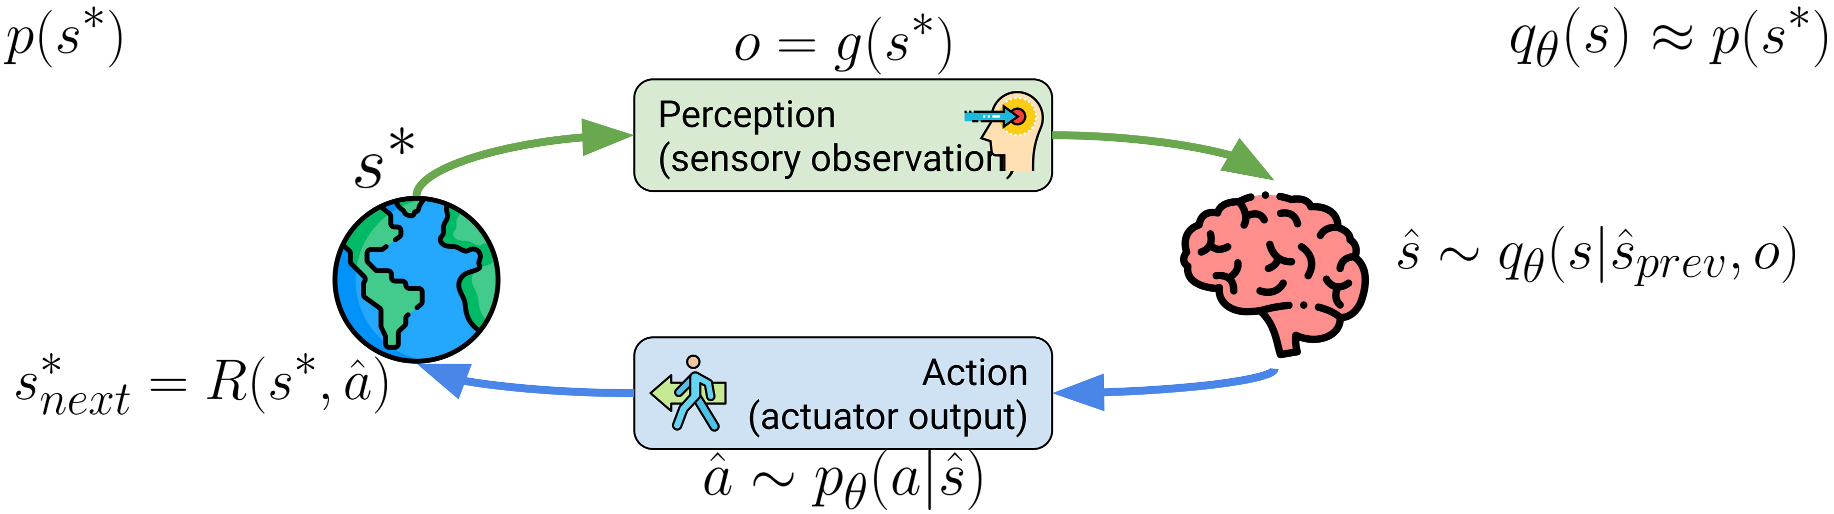

Perception and Action in Free Energy Principle

F = the ***negative*** variational free energy

Evidence Divergence

$$F = \log P(o_t) - KL[Q(s_t)\,|\,P(s_t\,|\,o_t)]$$

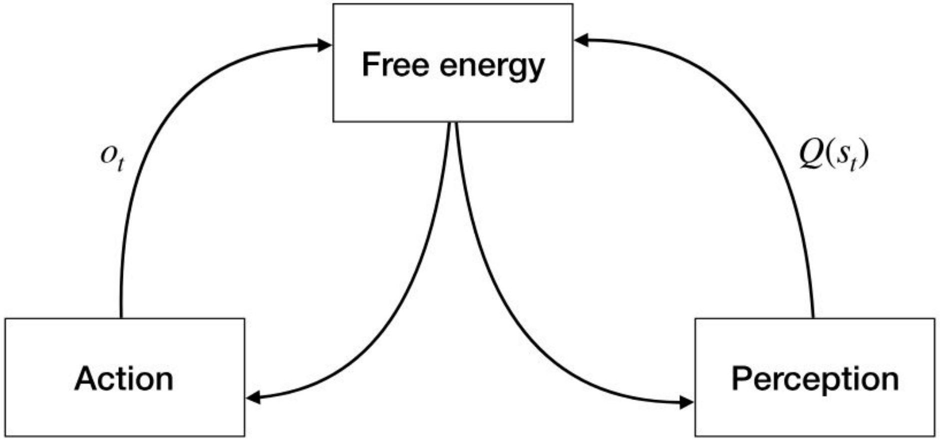

$$Q(\pi) = \sigma\left(\sum_{\rho \le t} F(\pi, \rho) + \sum_{\tau > t} G(\pi, \tau)\right)$$

$$Q(s_t) = \arg\max_{Q} F$$

Changing sensations through action to maximize evidence

Changing beliefs to minimize divergence

Perception and Action in Free Energy Principle

- To reduce prediction errors of generative models, either
  - (1) Improve perception model
  - **(2) Find an action policy (or action sequences) to minimize prediction errors.**

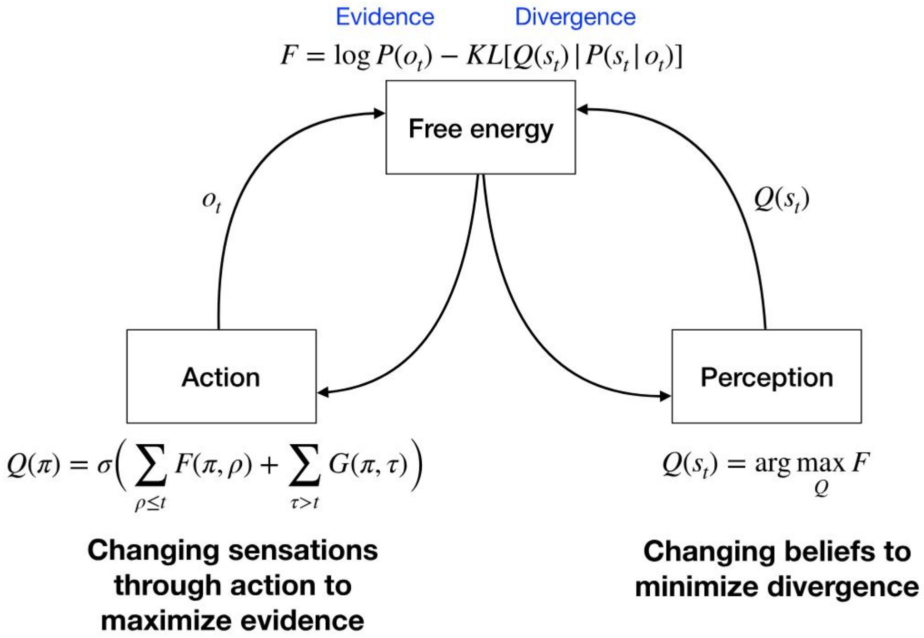

How to find action policies

Motor Imagery - Simulation Theory

- The same area is activated when
  - When you move your muscle and
  - When you imagine (or simulate) the motion.
- The brain can simulate the sensory information changes caused by an action without actually doing the action.

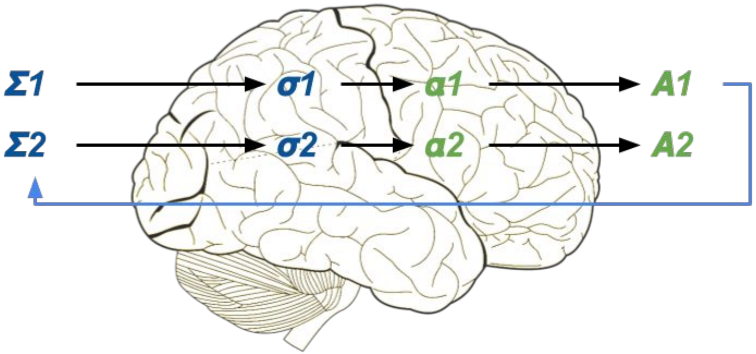

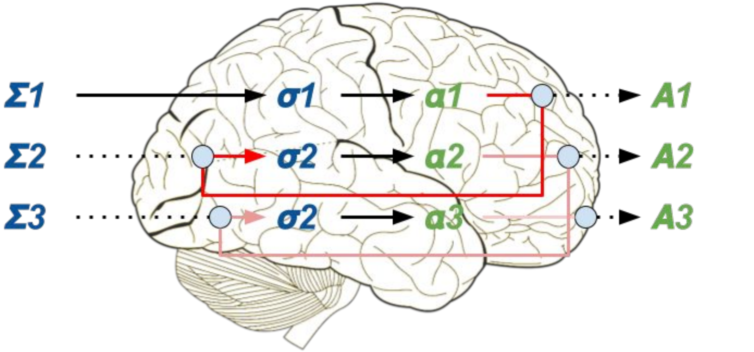

Forward Model - *Mental Imagery*

- Internal sensorimotor models that predict how the sensory situation changes as a result of an agent actions.

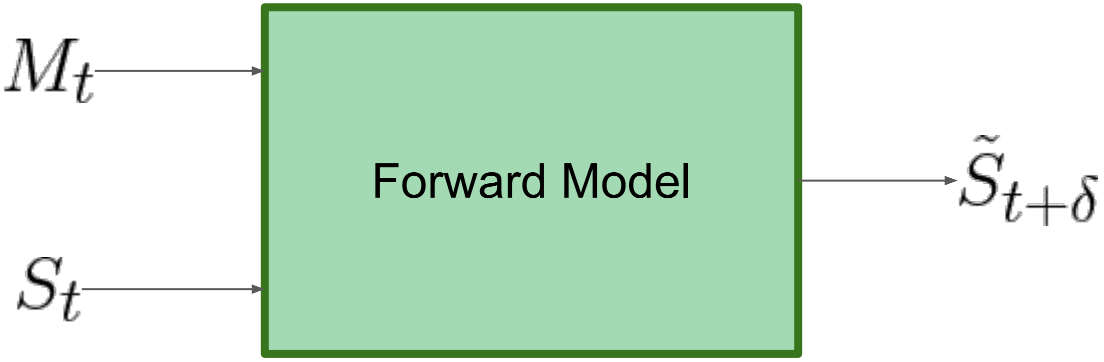

Forward Model In Vehicle

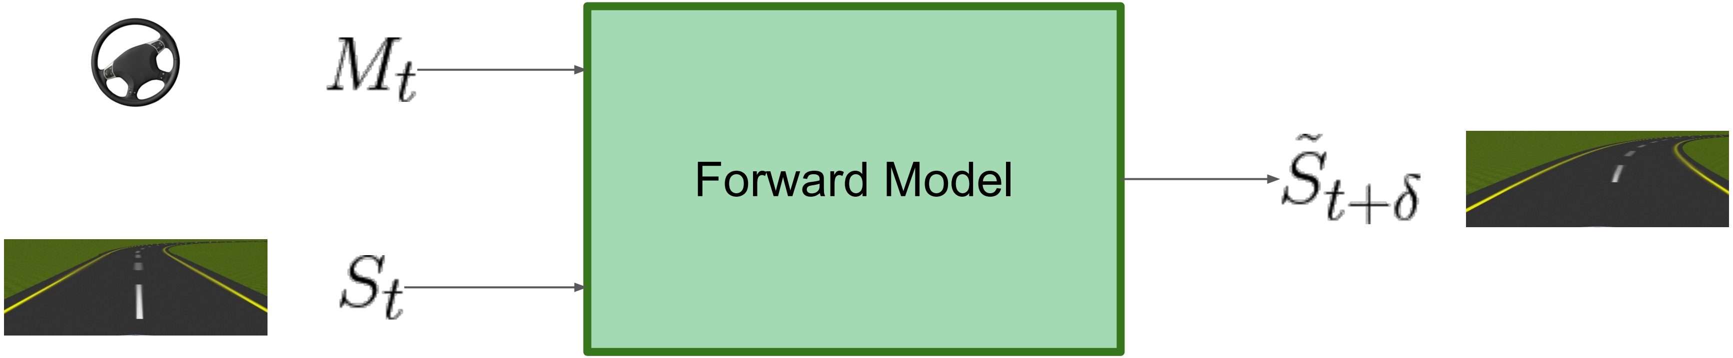

Inverse Model - *Mental Transformation*

Given the current state and next state, estimate the action that could cause the next state. Optimization problem.

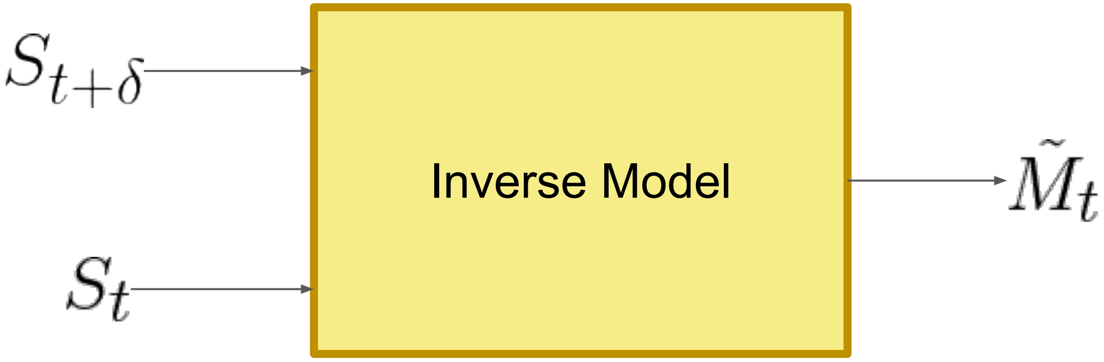

Inverse Model In Vehicle

Given the current state and next state, estimate the action that could cause the next state. Optimization problem.

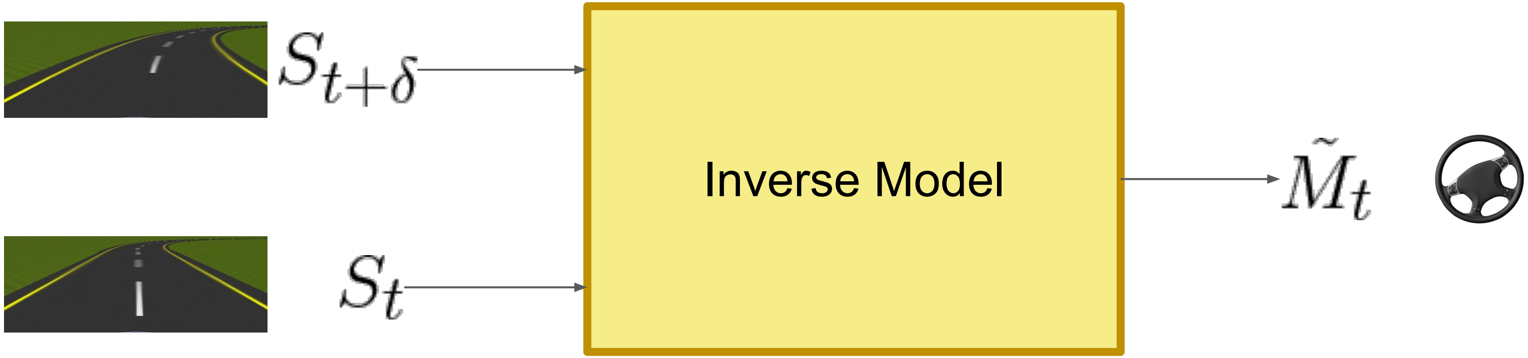

Training Forward Model

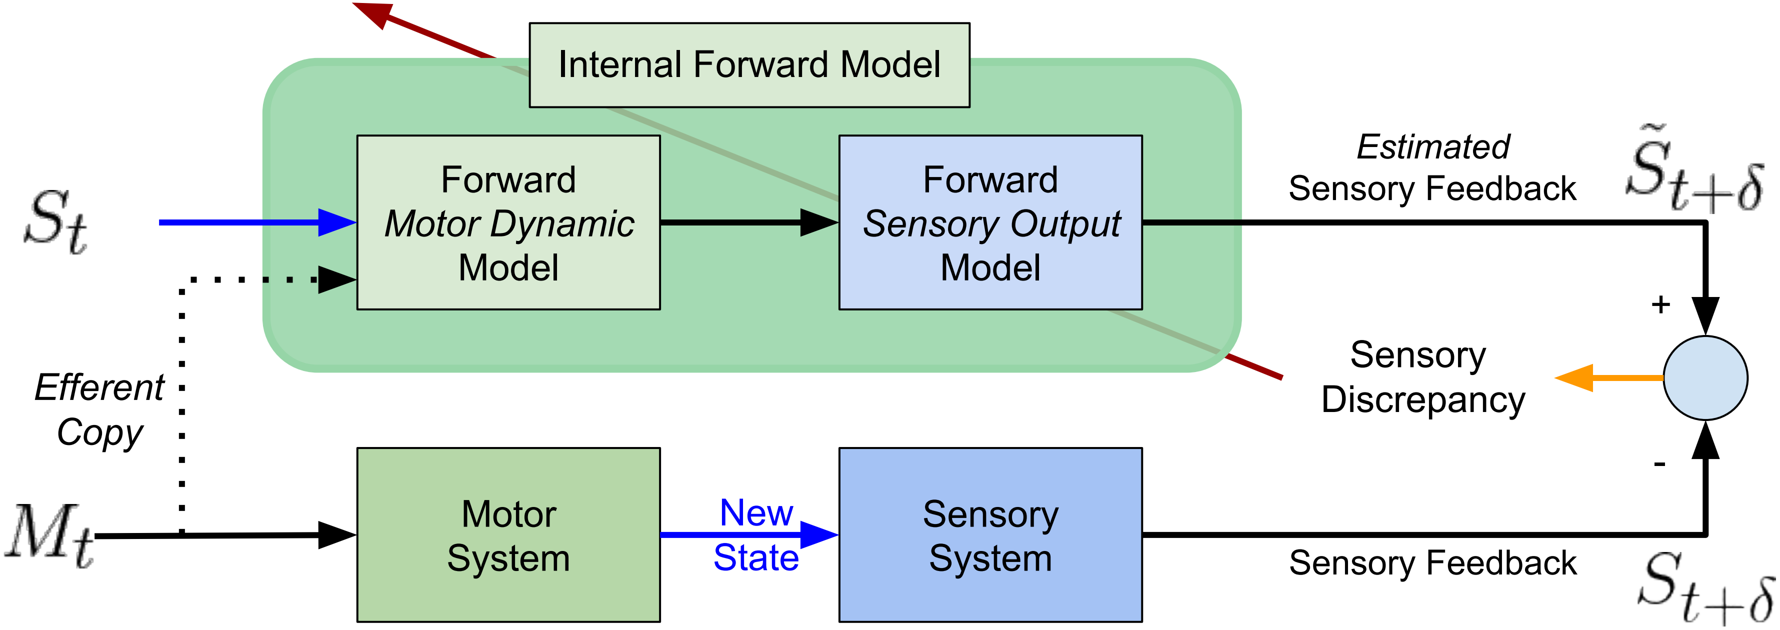

Forward Model with Variational Autoencoder

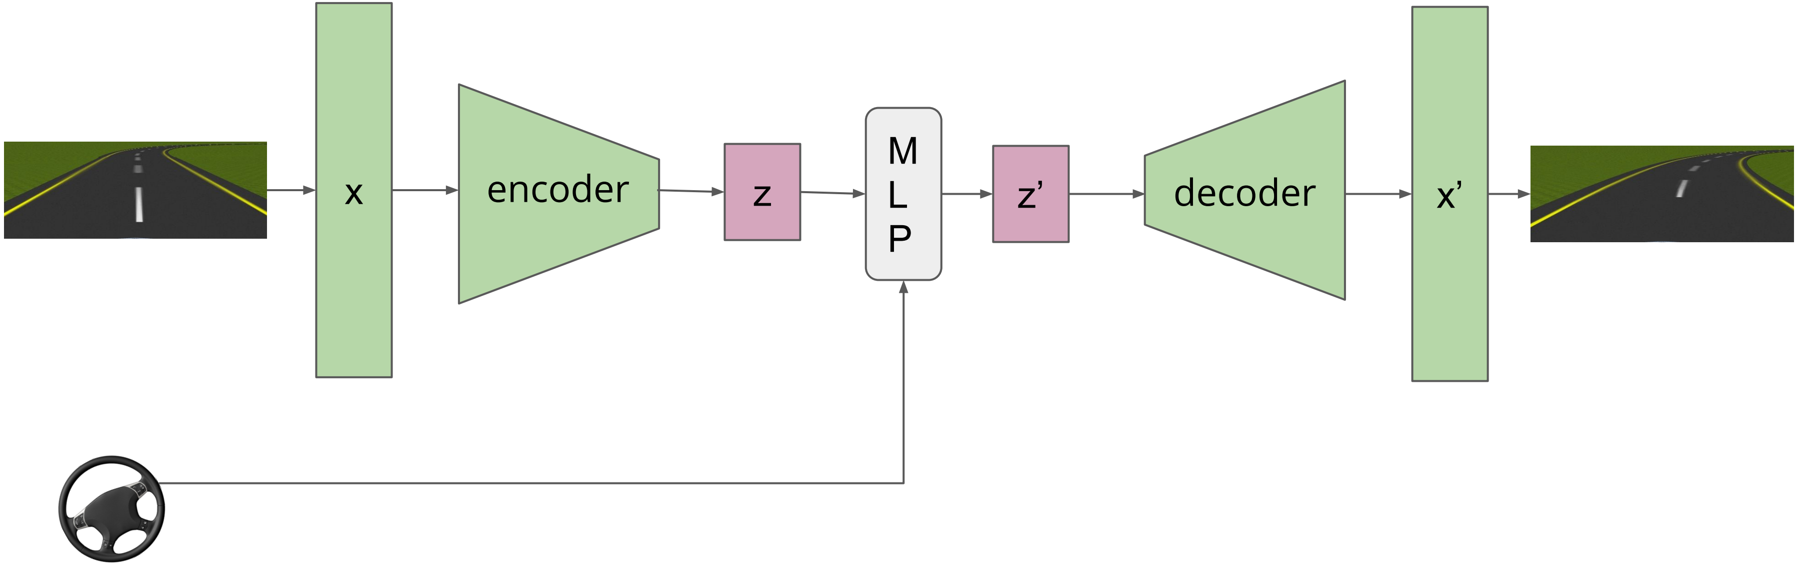

Forward Model with Variational Autoencoder

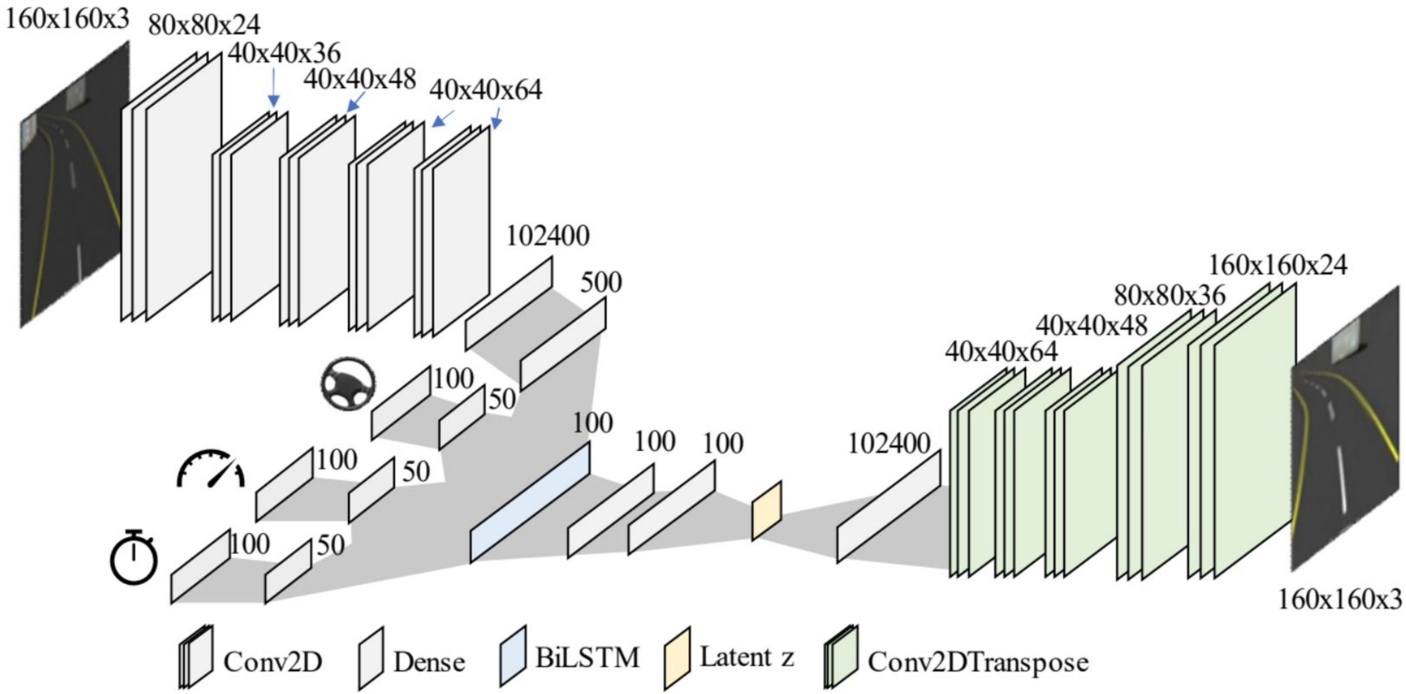

Experimental Environments

Forward Model using VAE

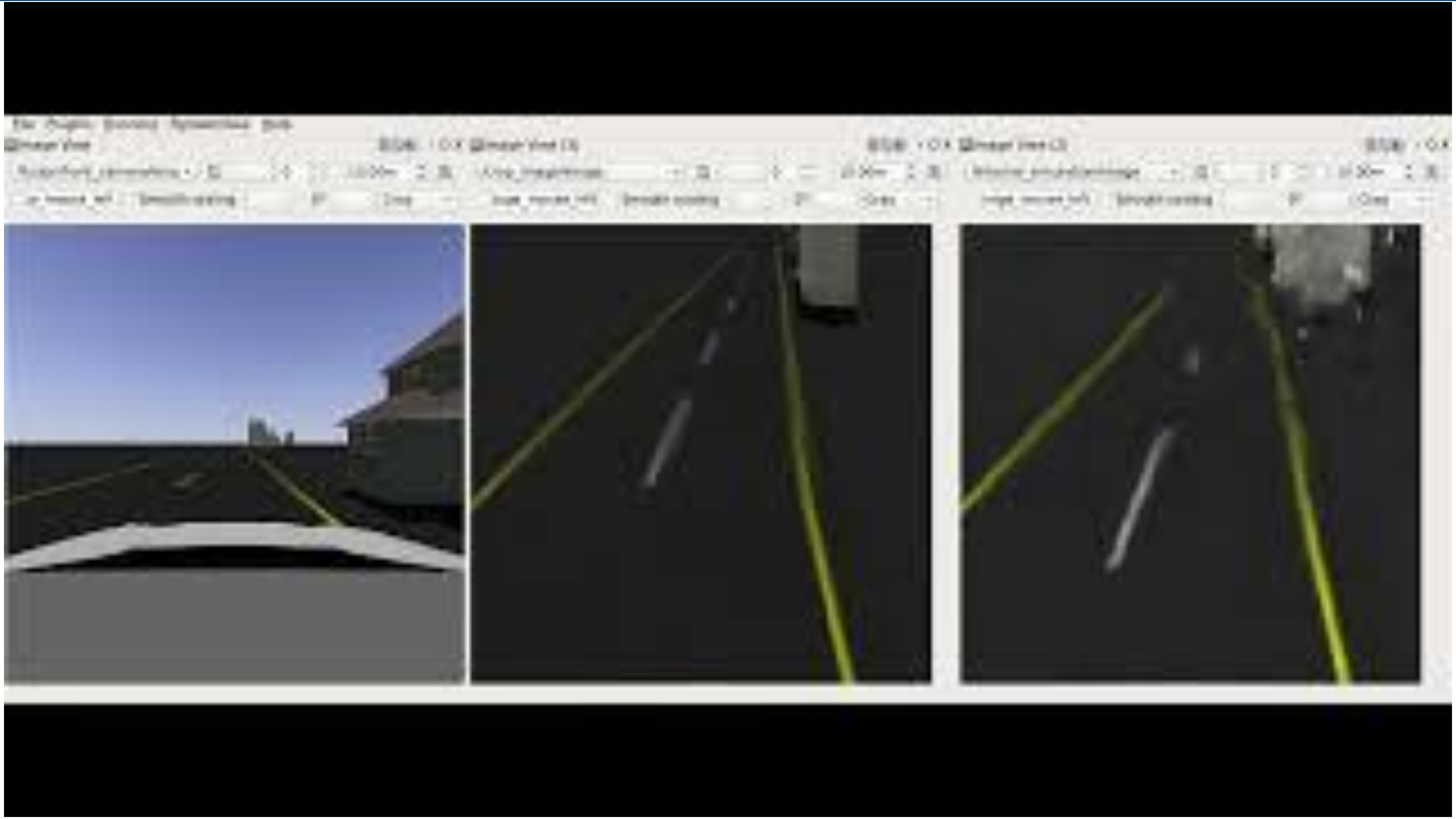

Driving with Images Generated by VAE

Training Inverse Model

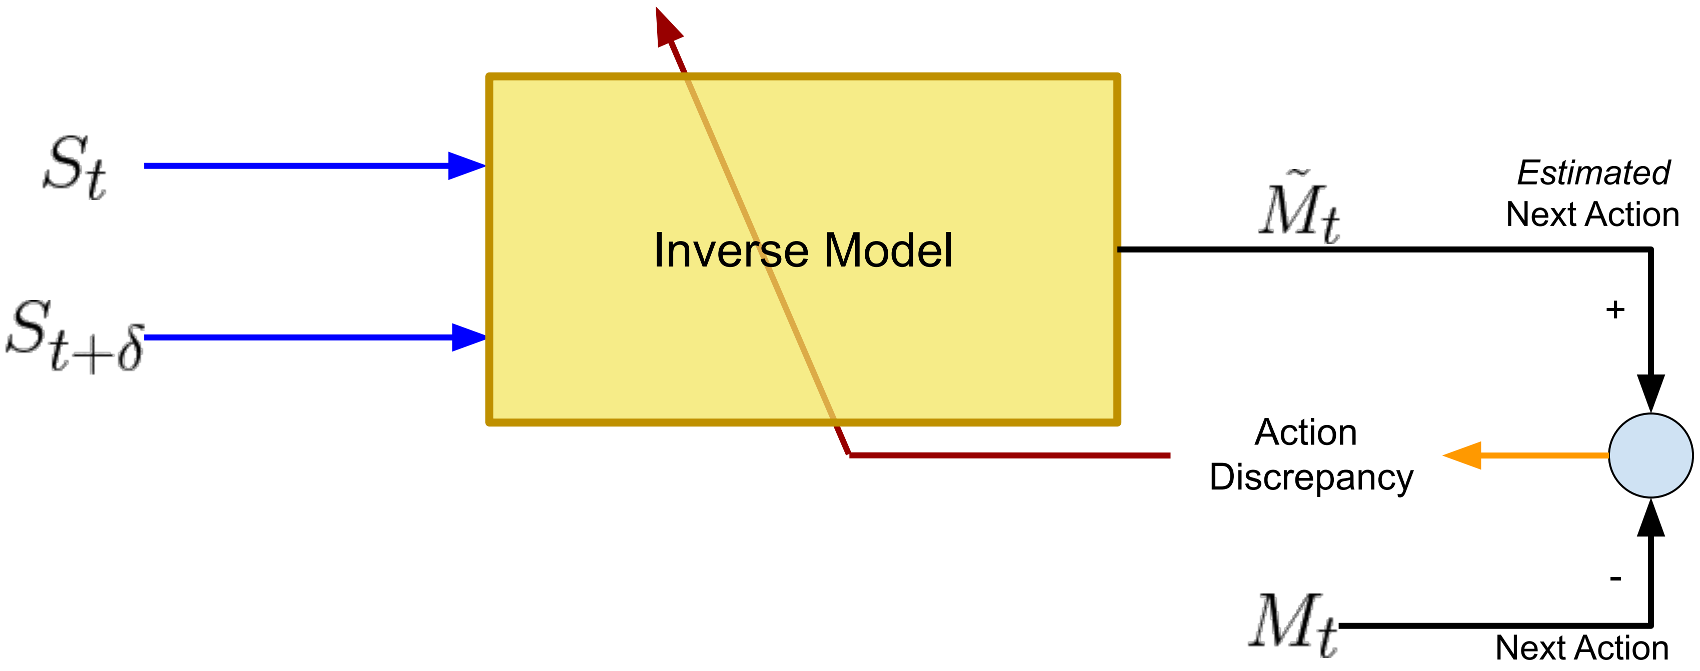

Training Inverse Model - *Problem*

Training dataset: challenging to collect all state transition caused by an action.

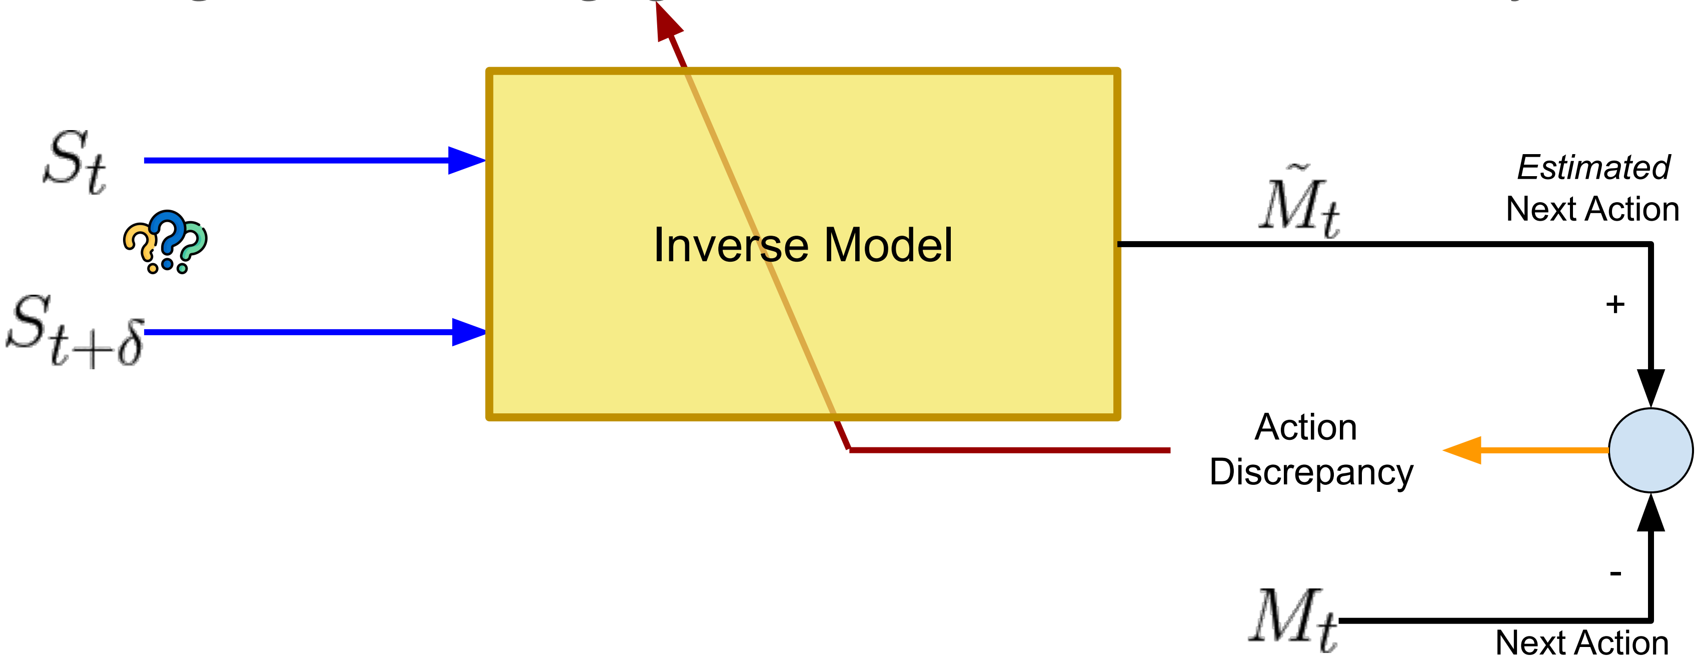

Mental Simulation Theory

Mental Simulation Theory

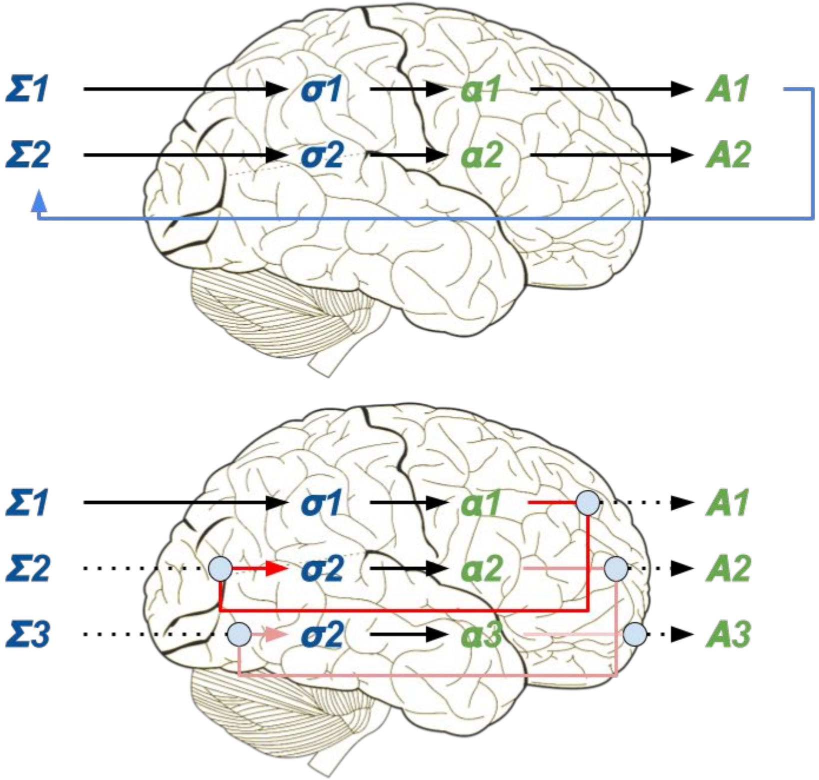

Mental Simulation Theory

- Simulation theory can be implemented through the Inverse Model.
- Motor actions can be inferred from the current sensory information and the future sensory situation caused by the the motor actions.

Action Policy through Mental Imagery

- First, define a task. (e.g., lane keeping)
- Second, define a preferred state (desired state or preference).
- Internally simulate possible action policies without actual overt action.
- Calculate the Expected Free Energy (EFE) from the policies.
- Choose an action policy that has a minimum EFE!

Applications

Lane Keeping using Active Inference

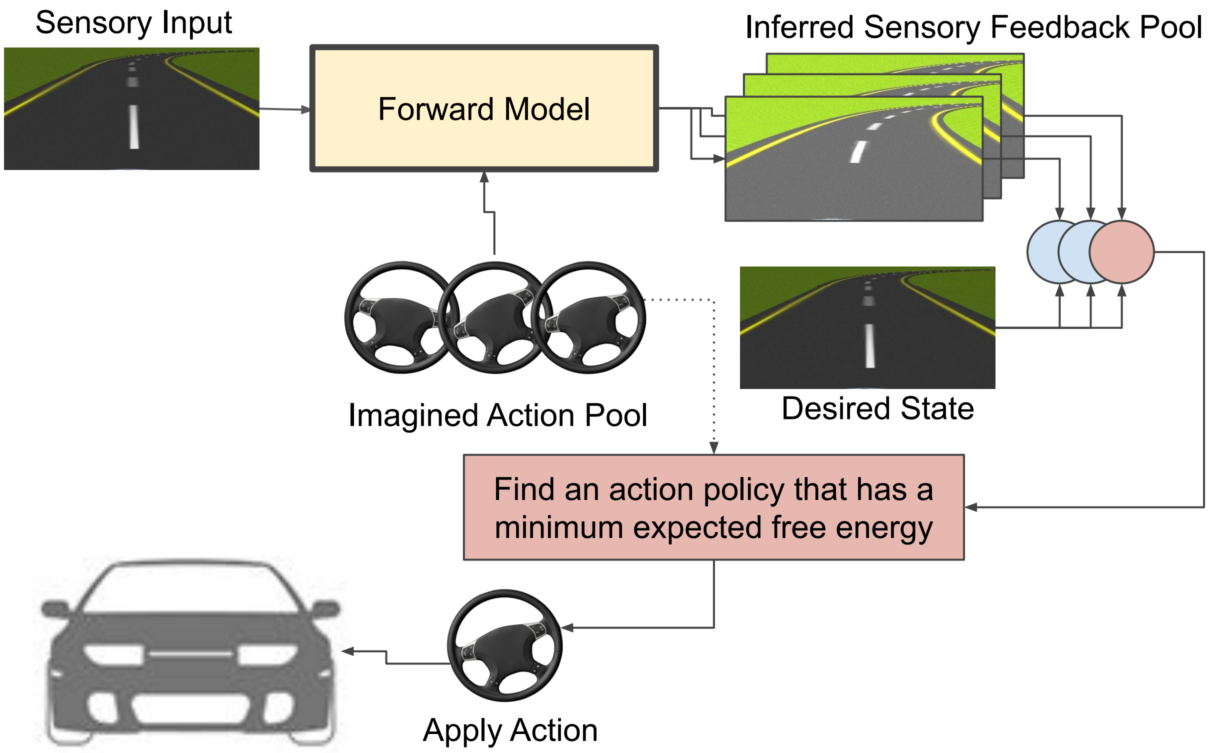

- To achieve a desired state (preference) from an actuation, possible actuations can be internally simulated.
- Find an actuation that has the least expected free energy.
- Apply the actuation to the actual system.

Demo in Carla Simulator

Birds Eye View



Demo

View from Rear

Another Demo in Carla

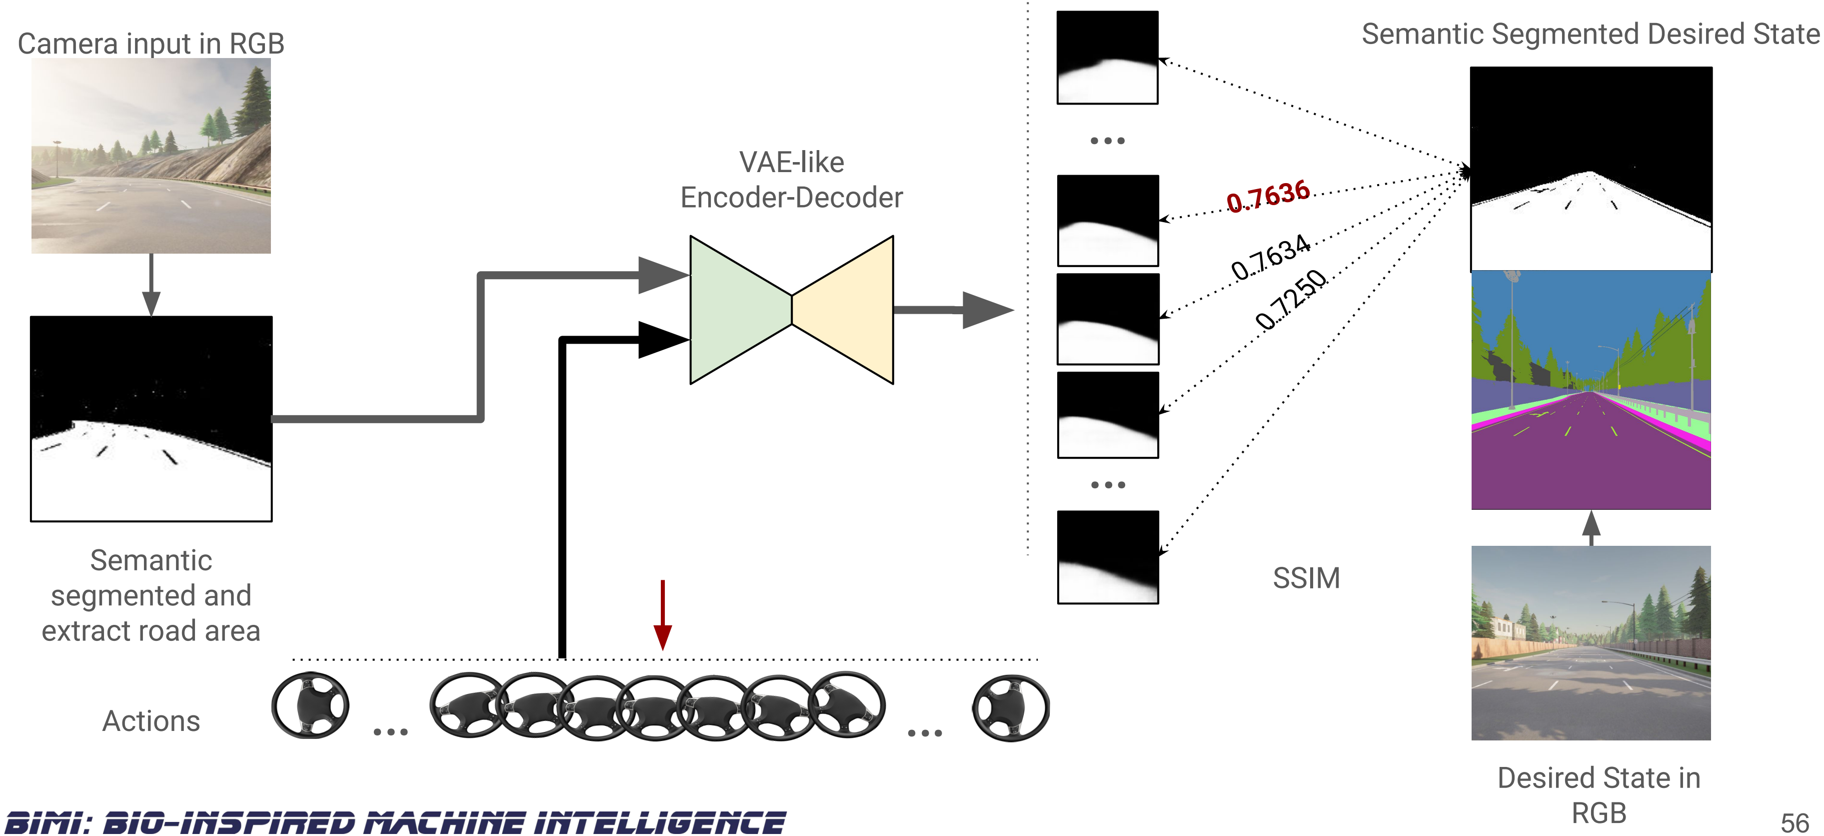

Another Demo in Carla

Lane Change

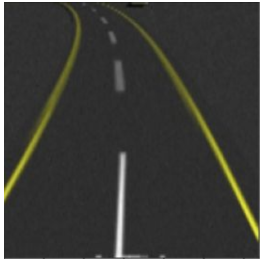

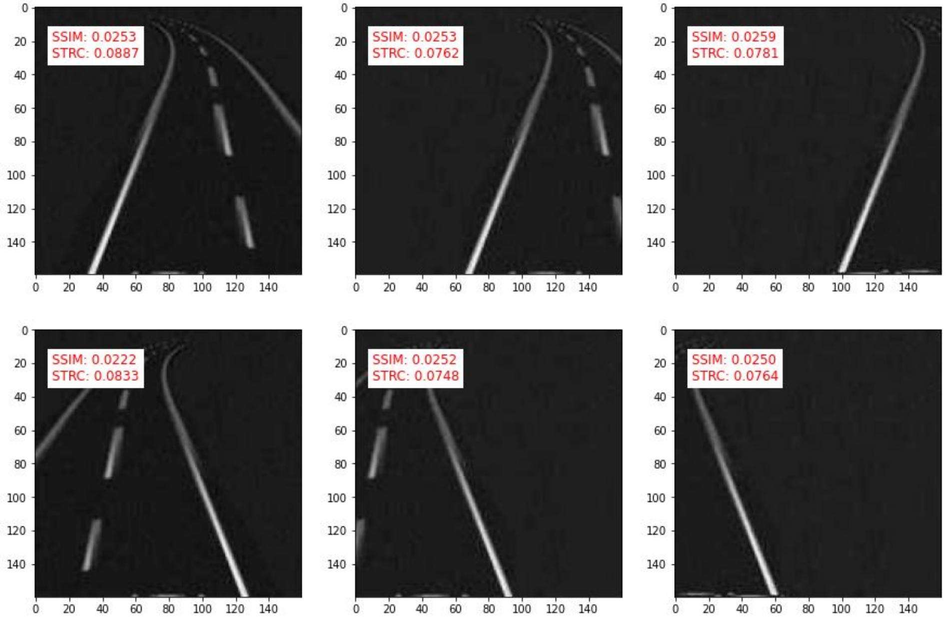
SSIM: 0.0253
STRC: 0.0887
SSIM: 0.0253
STRC: 0.0762
SSIM: 0.0259
STRC: 0.0781
SSIM: 0.0222
STRC: 0.0833
SSIM: 0.0252
STRC: 0.0748
SSIM: 0.0250
STRC: 0.0764

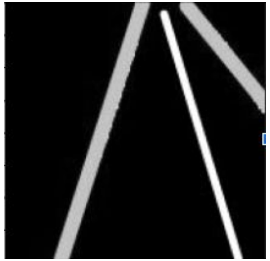

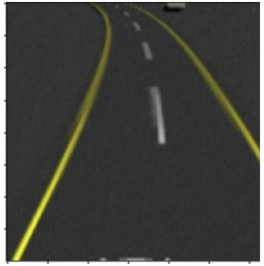

Lane Change Demo

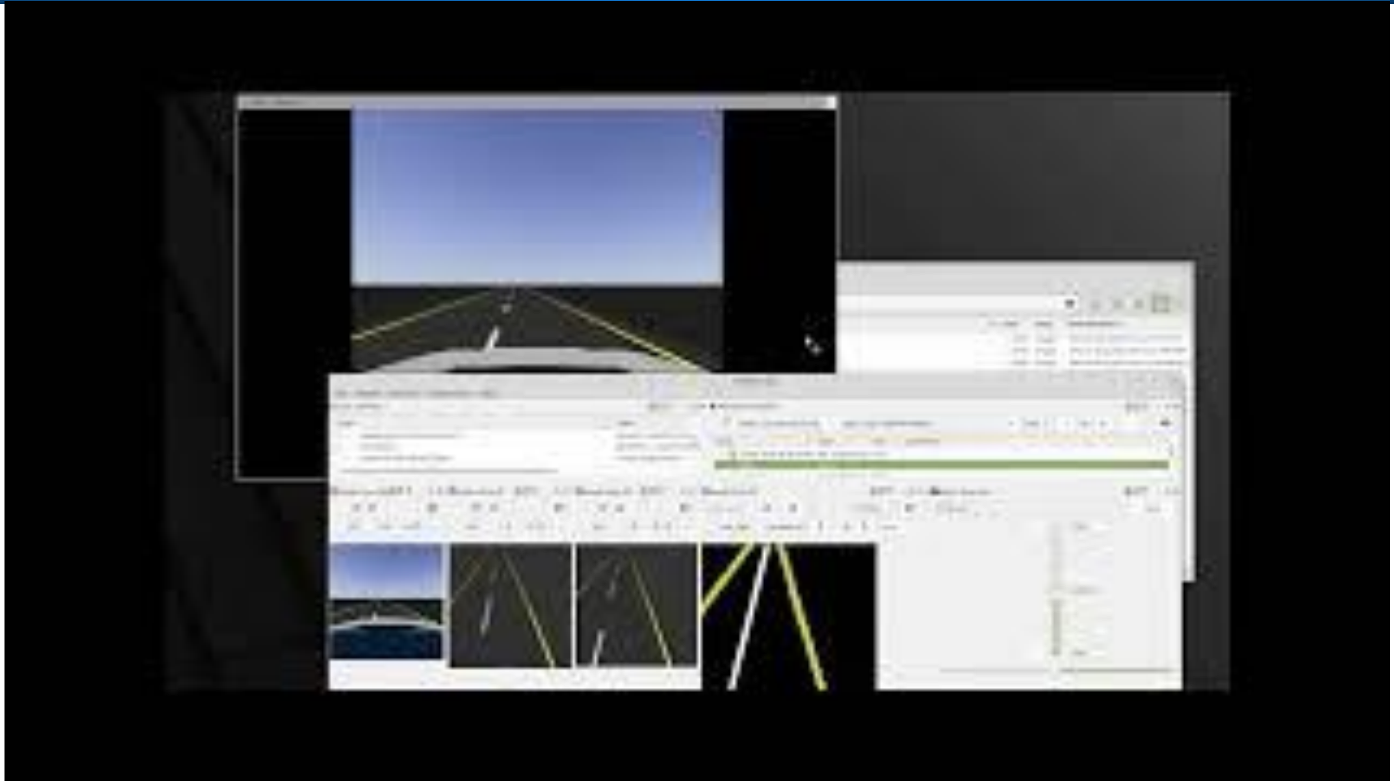

Summary

Contributions:

DEARBORN

- Learn driving through action selection through mental transformation.
- Use DNN-based Internal Forward and Inverse Model
  - Robustness and adaptability.
- Utilize *perceptual-motor* active inference framework to plan goal-directed behavior
  - This makes the decision making process more transparent because *we* can explain why a certain steering angle was applied to a certain condition.
  - Free from distribution shift problems.
    - Not learning the mapping sensory states to control signals.
    - Learning *causal* association between actions and sensory situations induced by the actions.

Questions and Discussion

End of Presentation