nuro

March 26, 2023

# Autonomy: Open Questions


Aleksandr Petiushko

**Autonomy Interaction Research**


Keynote Talk at 2023
Spring BDD Retreat

Berkeley DeepDrive

# Content

# Part I: Automation

01 Automation

02 Development

03 Evaluation

04 Conclusion

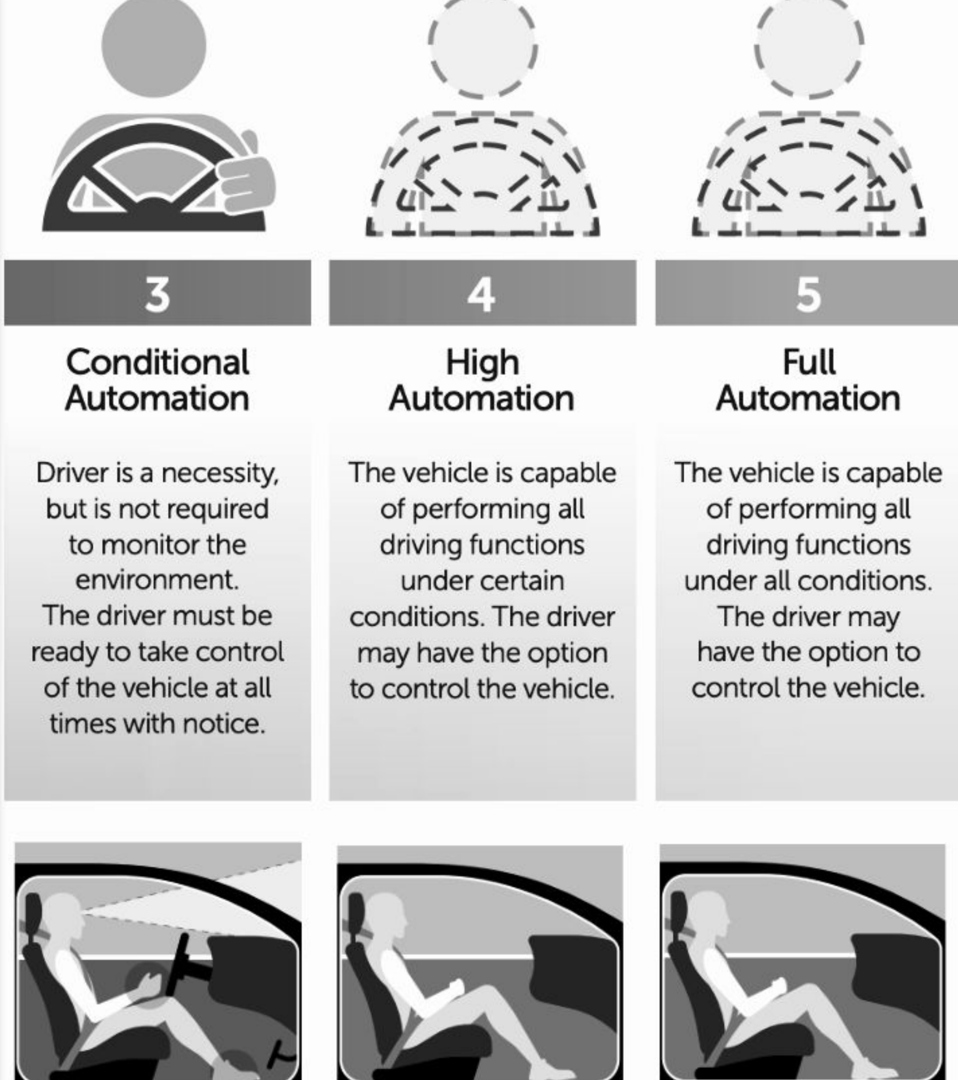

# Do we have the clear understanding / roadmap for introducing high Automation levels?

# Levels of Automation

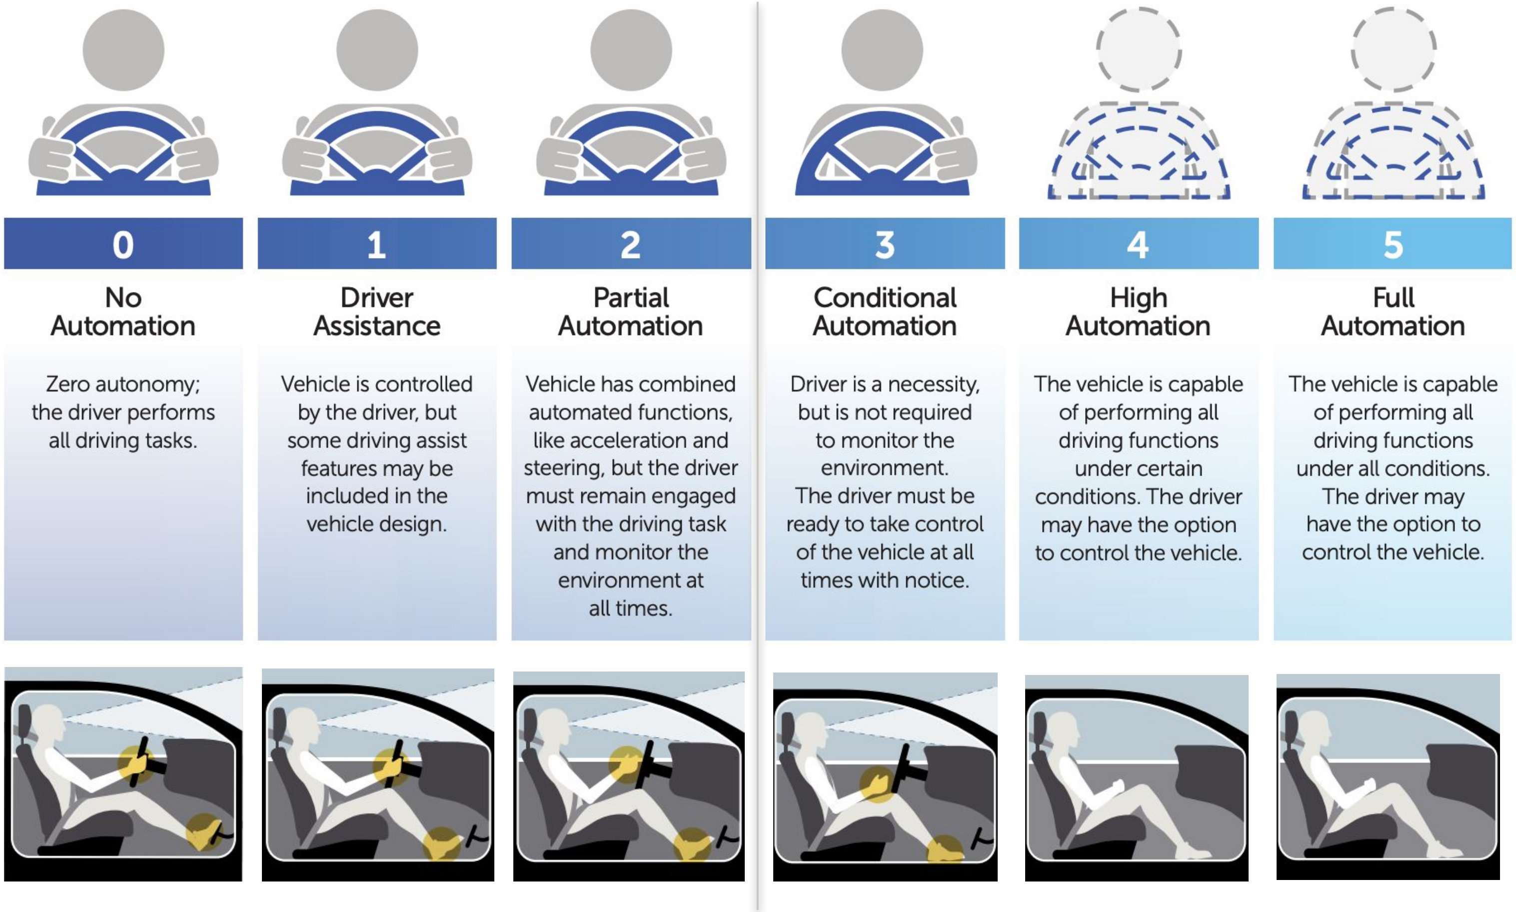


NHTSA: 1, 2 + SAE (Society of Automotive Engineers) J3016

# Conditional Automation

**Q**: how to make **notice** for driver ***in advance***?
Is it **realistically** doable and useful?

**Problem**:
- Example: **collision avoidance signal**[1]
- **Time of human reaction**: 1-2 seconds[2]
- **False** positives avoidance **vs true** positives coverage

W/ and **w/o** waiting for the human **feedback**:
- **Automatic Emergency Braking**
  - Pros: greatly *reduces rear-end collisions* (by 40-50%)
  - Cons: still not ideal (have *hundreds per year accidents* caused by drivers placing too much confidence in automatic brakes)

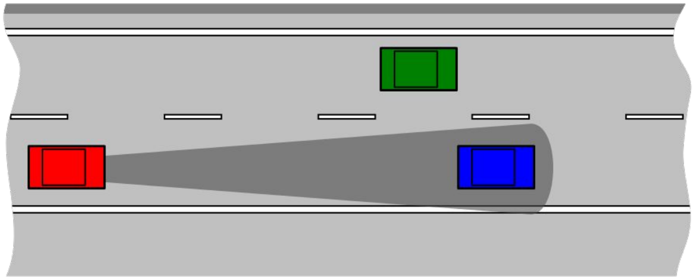

0.7 sec -- about as fast as it gets
1.0 sec -- old standard
1.5 sec -- common use
2.0 sec -- common use
**2.3 sec -- AVERAGE**
2.5 sec -- used in a few states
3.0 sec -- NSC and UK Standard

**Driver reaction times**

*Wiki on [Collision Avoidance System]()*
*McGehee, Daniel. et al. "[Driver reaction time in crash avoidance research: Validation of a driving simulator study on a test track]()." 2000. + [copradar.com]()*

# High vs Full Automation

**Q**: how to understand that we are **in** or **out** of our “**certain** conditions”?

**Problem**:
- need to understand the input **distribution shift**
- need to understand it for **every single module** inside the Autonomy Stack (e.g., Perception, Prediction, Planning, etc)

Possible **solution**:
- (Variational) **Autoencoders**[1]
  - Cons: How to behave if *OOD/Anomaly* (see “*Conditional Automation*”)?

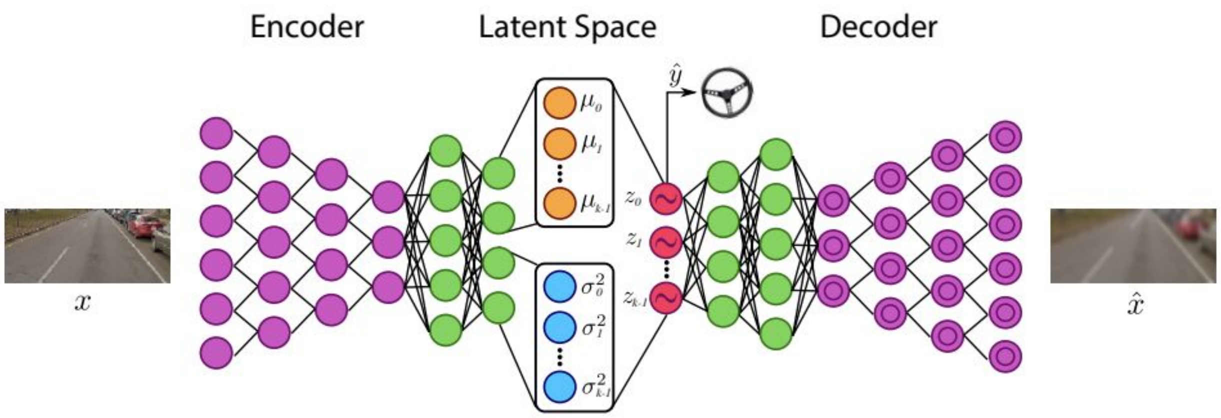


*Amini, Alexander, et al. "Variational autoencoder for end-to-end control of autonomous driving with novelty detection and training de-biasing." 2018.*

# Full Automation

**Q**: how to make the model **working** for **all input** (even weird) conditions?

**Problem**:
- **known unknowns**: specific adversarial RL agents for the specifically designed scenario
- **unknown unknowns**: some physically plausible input providing “bad” outputs (e.g., collisions)

Possible **solutions**:
- **Adversarial RL** agents
  - Cons: *limited* by scenario generation and RL engine capabilities
- **Backpropagation**[1] w.r.t. Input
  - Cons: full-stack usually *hardly backpropagatable, constraint*s on Input

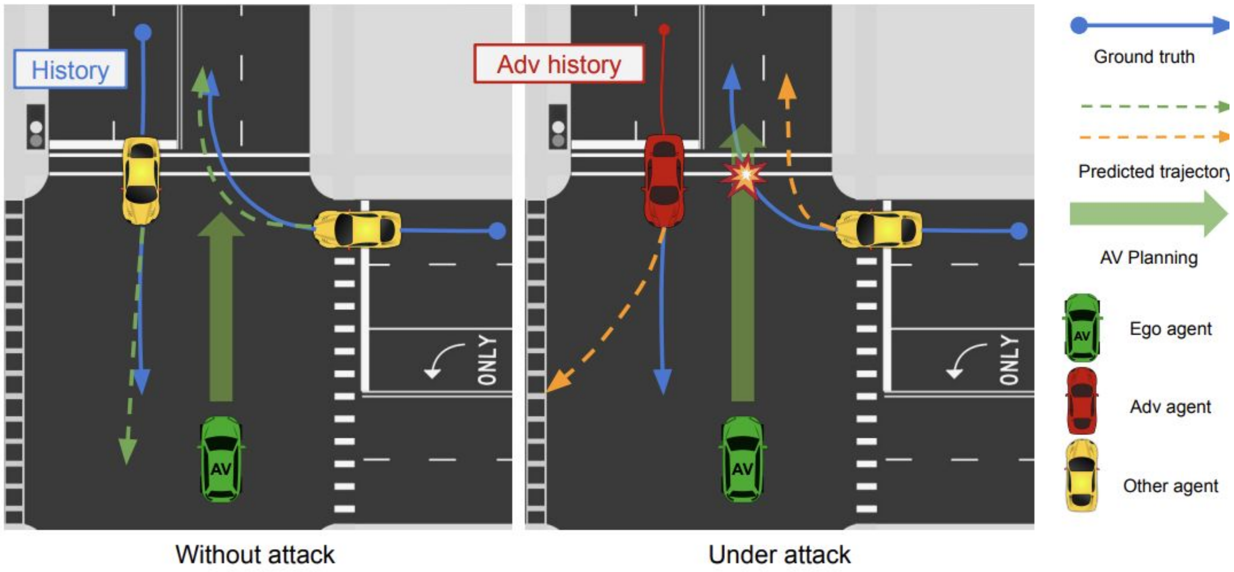


*Cao, Yulong, et al. "[Advdo: Realistic adversarial attacks for trajectory prediction](). " 2022.*

# Part II: Development

01 Automation

02 Development

03 Evaluation

04 Conclusion

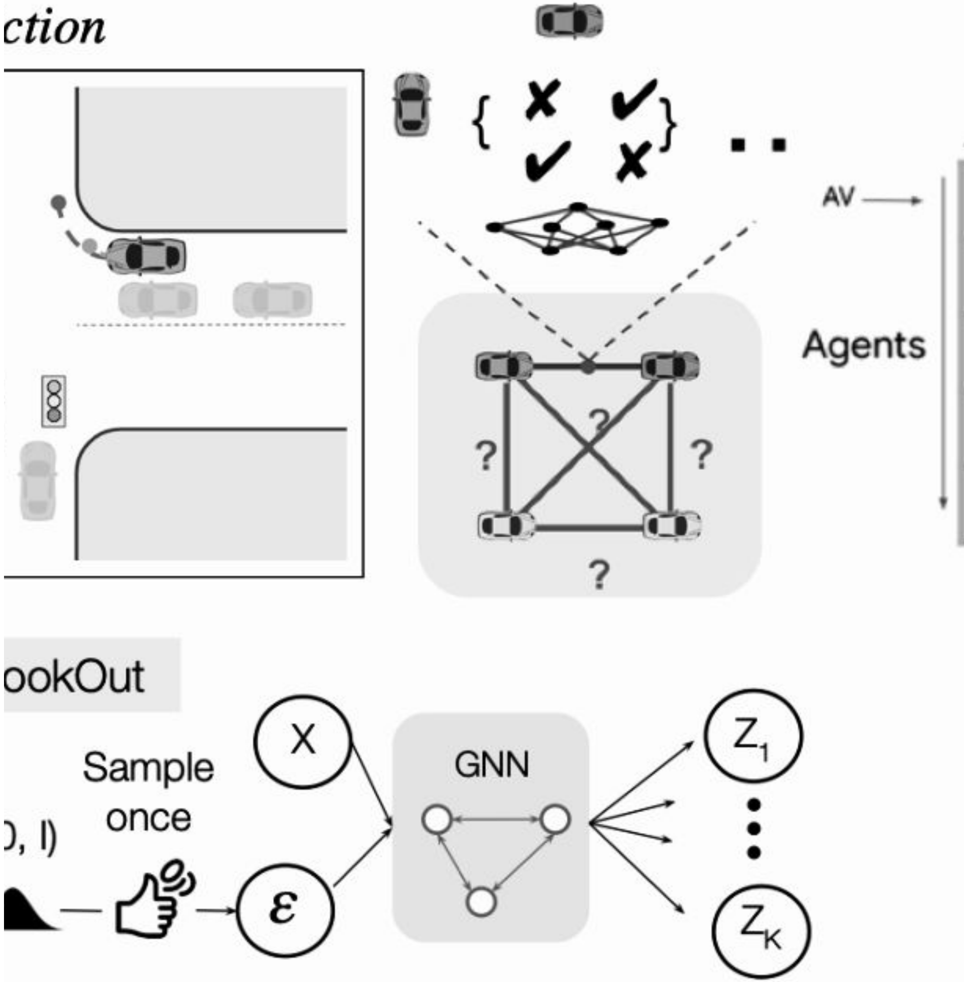

# What could be the **development** stepping stones for reaching the self-driving?

# Differentiability

**Q**: how to propagate the learning signal (and uncertainty estimations) through the whole stack?

**Problem**:
- avoid **end2end** approach like *Behavior Cloning*
- **re-use** the existing modules and *expert* knowledge

Possible **solutions**:
- **Approximation** of non-differentiable modules by:
  - differentiable **wrapping**[1]
  - differentiable **approximation**[2]
  - Cons:
    - *constraints* on modules inside wrapping
    - *hard* / *slow* to approximate some existing modules (iLQR, sampling)

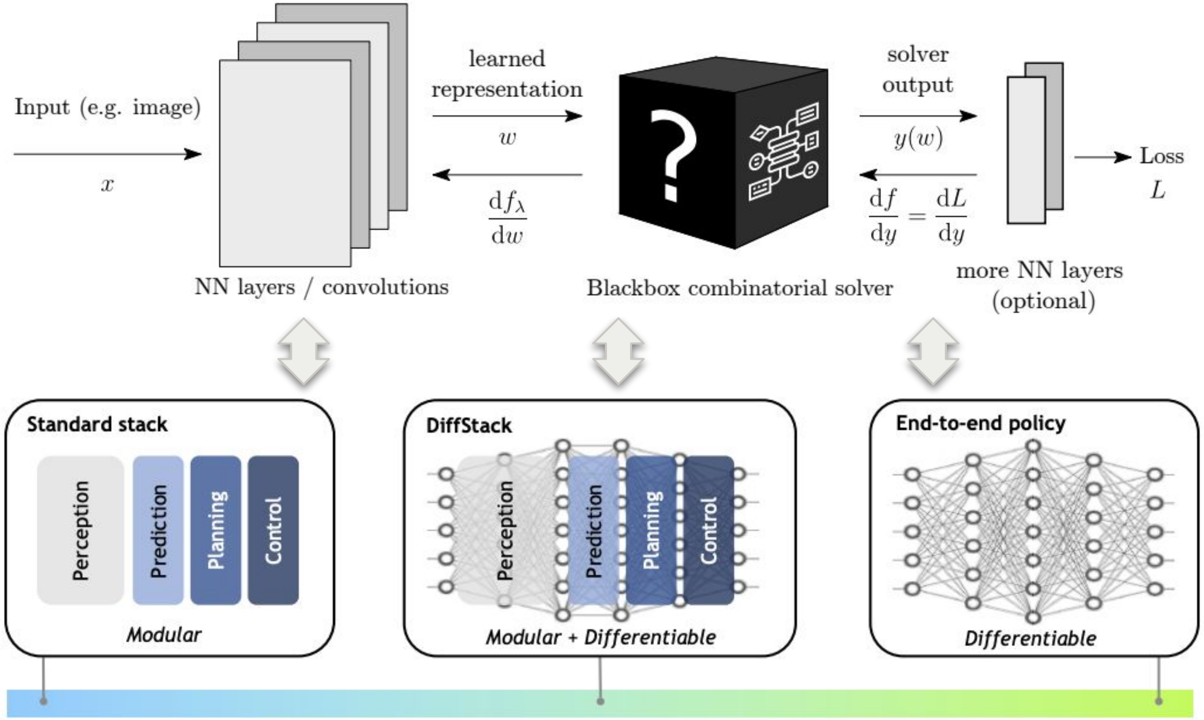


*Vlastelica, Marin, et al. "Differentiation of blackbox combinatorial solvers." 2019*
*Karkus, Peter, et al. "DiffStack: A Differentiable and Modular Control Stack for Autonomous Vehicles." 2022.*

# Jointness I

**Q**: how to **ensure consistency** between:
- **between prediction and planning,**
- different predictions,
and how to evaluate it?

**Problem**:
- **feedback loop** between the robot future and other road agents futures
- mining of **interactivity** scenes

Possible **solutions**:
- **Heuristically** (e.g., by distance) defining the interactive scenes/agents
- Conditional Behavior Prediction by the **new model input** (robot planned future)
- Conditioning in the **autoregressive** way

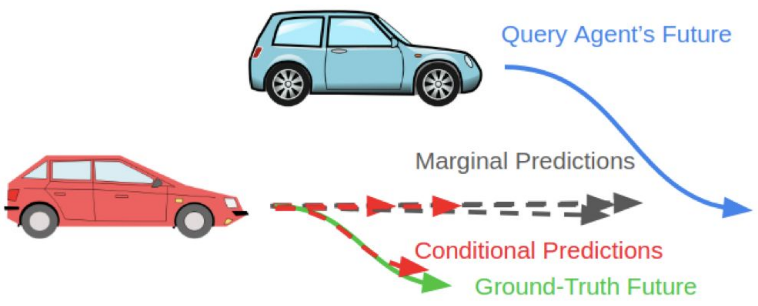


**Conditional Behavior Prediction**[1]

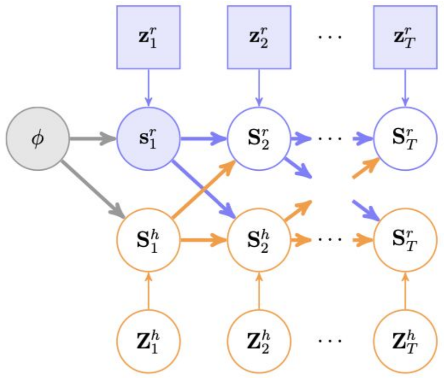


**PRECog**[2]

*Tolstaya, Ekaterina, et al. "Identifying driver interactions via conditional behavior prediction." 2021*
*Rhinehart, Nicholas, et al. "Precog: Prediction conditioned on goals in visual multi-agent settings." 2019*

# Jointness II

**Q**: how to **ensure consistency** between:
- between prediction and planning,
**- different predictions,**
and how to evaluate it?

**Problem**:
- working on top of **marginals** is **error-prone**
- considering all the combinations of agents leads to a **combinatorial** complexity **explosion**

Possible **solutions**:
- Different mitigations:
  - Joint pairwise by **message passing**[1]
  - Jointness by **transformer decoder**[2]
  - Jointness by the **unified latent**[3]
- These are still mitigations

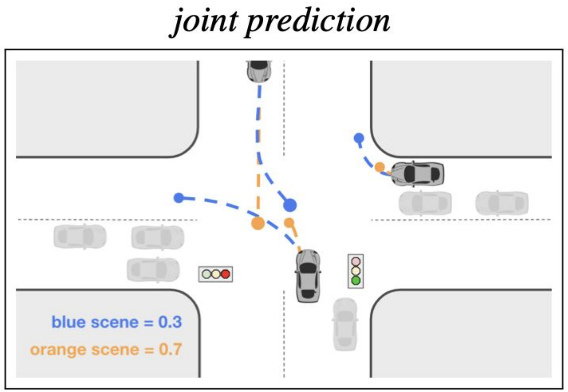


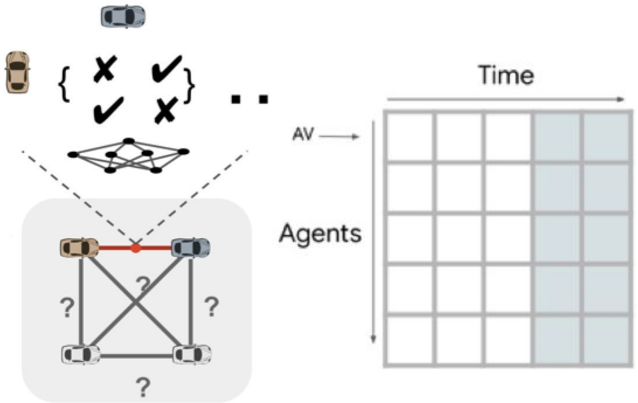


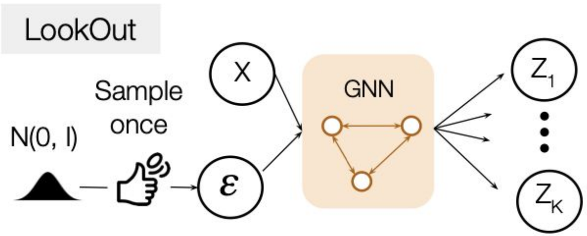


*Luo, Wenjie, et al. "[JFP: Joint Future Prediction with Interactive Multi-Agent Modeling for Autonomous Driving]." 2023*
*Ngiam, Jiquan, et al. "[Scene Transformer: A unified architecture for predicting multiple agent trajectories]." 2021*
*Cui, Alexander, et al. "[Lookout: Diverse multi-future prediction and planning for self-driving.]" 2021*

# Jointness III

**Q**: how to **ensure consistency** between:
- between prediction and planning,
- different predictions,
and **how to evaluate it?**

**Problem**:
- need new **joint metrics**
- need public **datasets** and **challenges** supporting it

Possible **solutions**:
- **Scene-level** analogs of marginals
  - minSADE vs minADE
- **Waymo**[1] (pairwise joint) and **Interaction**[2] (pairwise and fully joint conditional) datasets

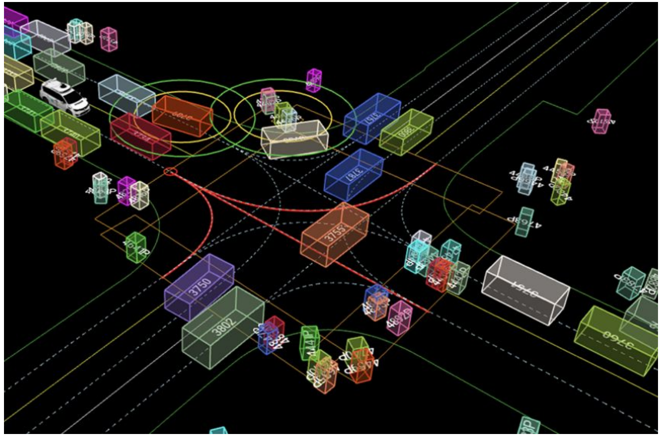

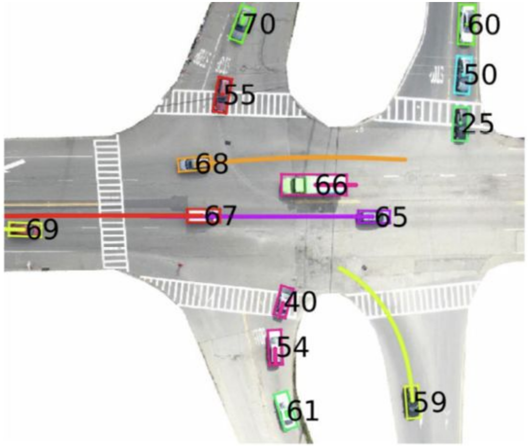


$$minADE = \frac{1}{l}\sum_{i=1}^{l} \min_{k} ||x_i^k - x_i^{gt}||$$

$$minSADE = \frac{1}{l}\min_{k} \sum_{i=1}^{l} ||x_{scene,i}^k - x_i^{gt}||$$

*Ettinger, Scott, et al.* "[Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset](#)." *2021*
*Zhan, Wei, et al.* "[Interaction dataset: An international, adversarial and cooperative motion dataset in interactive driving scenarios with semantic maps](#)." *2019*

# RL for AV

**Q**: how to **incorporate Reinforcement Learning** (RL) into the Autonomy Stack taking into account safety requirements?

**Problem**:
- Explicit Planning by RL is unstable / unreliable
- Hard to balance and optimize multiple safety constraints

Possible **solutions**:
- Instead of explicit Planning by RL, **fine-tuning by RL rollouts**
  - Cons: having the good model is a *chicken-egg* problem
- Usage of **Human Preference**[2] labels (RL from Human Feedback (HF)): ChatGPT[1]-like approach
  - Cons: 1) *absence* of a good *foundation* model for AD; 2) *hard* to get *lots of HF labels* for AV
- Still unknown what is the best way to **inject safety constraints** (and is it needed explicitly?)

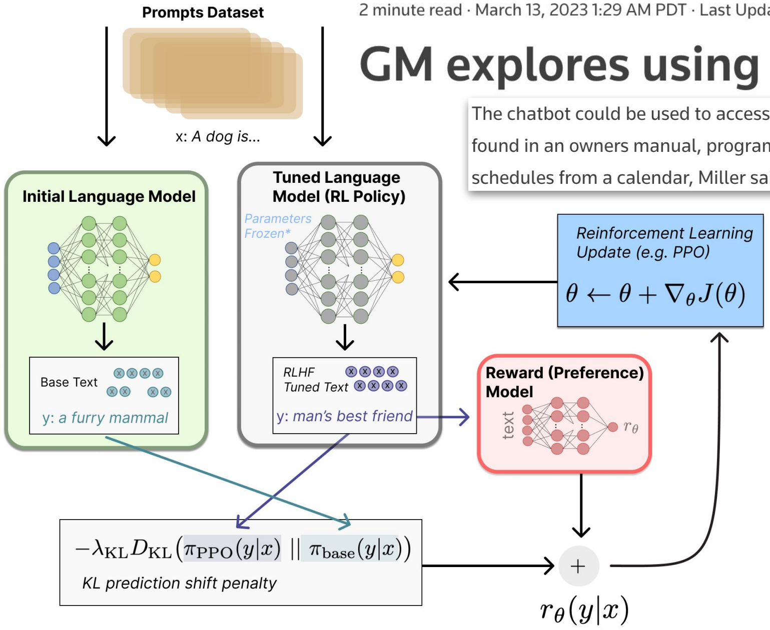


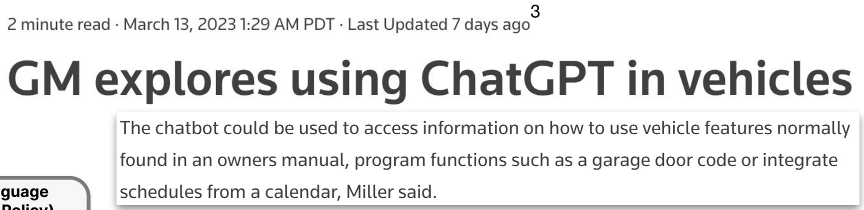
2 minute read · March 13, 2023 1:29 AM PDT · Last Updated 7 days ago [3]

GM explores using ChatGPT in vehicles

The chatbot could be used to access information on how to use vehicle features normally found in an owners manual, program functions such as a garage door code or integrate schedules from a calendar, Miller said.

1. OpenAI: [ChatGPT]
2. Hugginface: [RL from HF]
3. Reuters: [GM explores using ChatGPT in vehicles]

# Part III: Evaluation



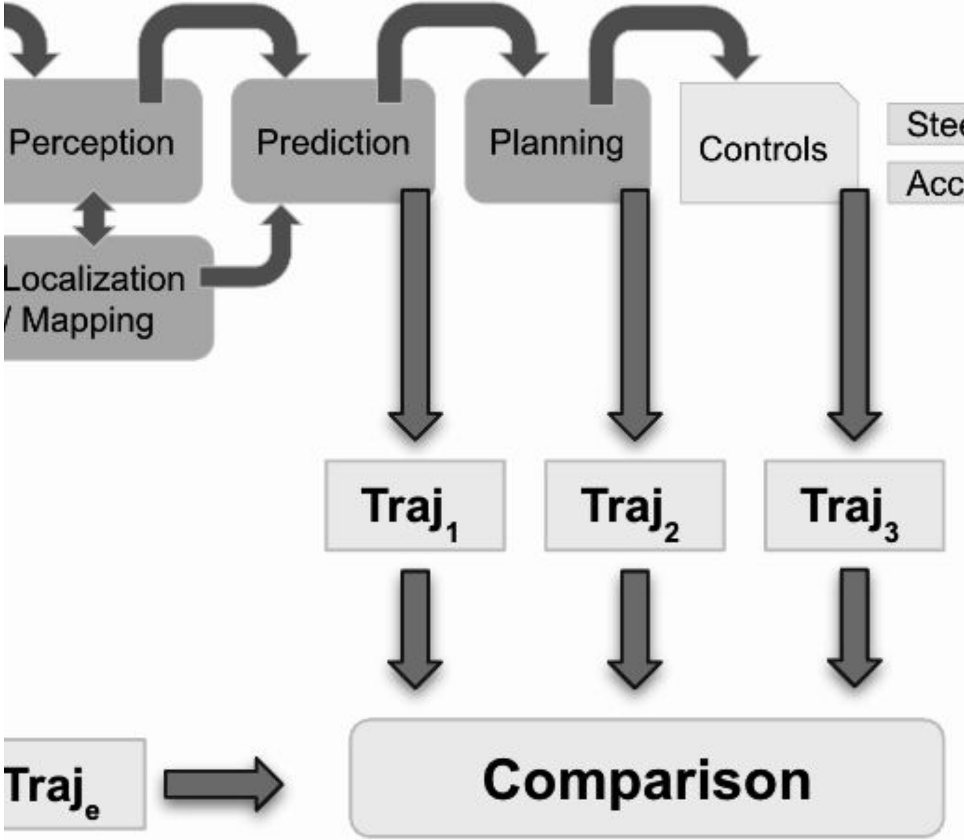

# How to **evaluate** our progress being engineers?

# Evaluation

**Q**: how to make the evaluation process be **less costly** and **faster**?

**Problem**:
- **how** (metrics) and **where** (modular vs end2end) to evaluate?
- need in **submodular** eval?

Possible **solutions**:
- **End2end comparison** with the human expert
  - Cons: it is only Imitation Learning-like metric
- **Submodular comparison** with the human expert
  - Cons: need to produce the robot trajectory *as soon as possible*
- *Necessity* vs *sufficiency*

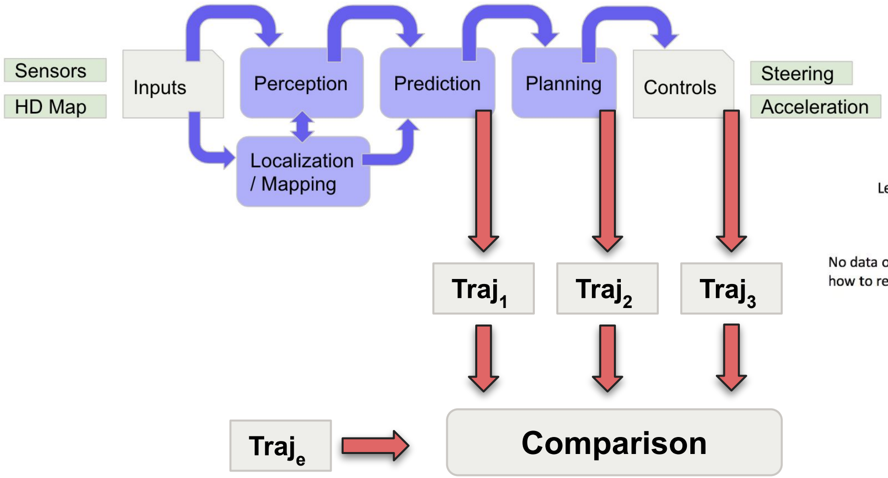


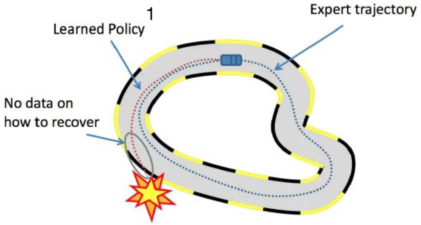


*Medium: [Imitation Learning](), 2019*

# Part IV: Conclusion

01 Automation

02 Development

03 Evaluation

04 Conclusion

# Conclusion

- Formal Automation Levels definition are not clarifying the possible approaches to reach them
- Stepping stones towards the full self-driving are reasonable but not set in stone
- Consistency in a model output is going to be a trend; but need deeper support from datasets/metrics/challenges
- Evaluation is painful
- “*ADGPT*” to the rescue?


Thank
You.
nuro