Hochschule
Bonn-Rhein-Sieg
University of Applied Sciences

b-it Bonn-Aachen International Center for Information Technology

Institute for AI and Autonomous Systems

# Learning for Robot Navigation

## An Overview

Dr. Alex Mitrevski
Master of Autonomous Systems

# Structure

- Why learning for robot navigation
- Navigation learning methods
- The case of autonomous driving

Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8

**Motion planning and control for mobile robot navigation using machine learning: a survey**

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS, VOL. 23, NO. 6, JUNE 2022 4909

Deep Reinforcement Learning for Autonomous Driving: A Survey

B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Yogamani, and Patrick Pérez

**Learning Robotic Navigation from Experience: Principles, Methods, and Recent Results**

**Sergey Levine, Dhruv Shah**
UC Berkeley
{svlevine, shah}@eecs.berkeley.edu

# Why Learning for Robot Navigation

# Typical Robot Navigation Workflow

- A typical robot navigation scenario starts by **creating a map of the environment in which the robot needs to navigate**
- In that map, the robot then has to:
  - **localise**
  - **find path plans** that bring it from its current location to a goal location
  - **find and apply low-level motion commands** that will bring the robot to the goal
- Under this framework, a robot can only navigate in a known environment with a (more or less) static environment (so that the map does not change)

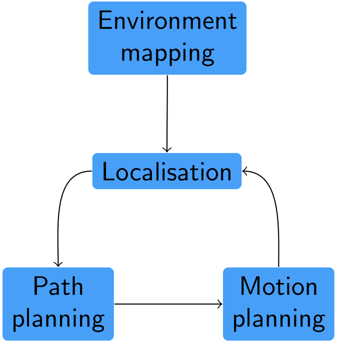

# Continuous Mapping Using SLAM

- A more generic workflow is one in which **simultaneous localisation and mapping (SLAM)** is performed, which allows the map to be updated (in principle)
- SLAM also makes it possible to navigate in unknown environment, as **a robot can create a map on the fly while using it for navigation**
- This eliminates the limitation of only being able to navigate in a known environment. **Problem solved?**

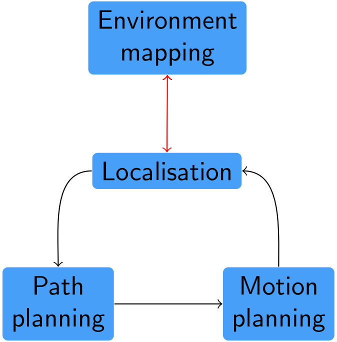

# Global and Local Planning

Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8

**Motion planning and control for mobile robot navigation using machine learning: a survey**

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

- As mentioned on the previous slides, **the actual navigation act involves path and motion planning**

# Global and Local Planning

Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8

Motion planning and control for mobile robot navigation using machine learning: a survey

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

- As mentioned on the previous slides, **the actual navigation act involves path and motion planning**
- Path (aka global) planning is the problem of **finding a viable (typically collision-free) path from the robot's current location to its destination**
  - A path plan is **created within a known map**
  - A path is **usually decomposed into a sequence of waypoints** through which the robot should pass

# Global and Local Planning


Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8

Motion planning and control for mobile robot navigation using machine learning: a survey

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

- As mentioned on the previous slides, **the actual navigation act involves path and motion planning**
- Path (aka global) planning is the problem of **finding a viable (typically collision-free) path from the robot's current location to its destination**
  - A path plan is **created within a known map**
  - A path is **usually decomposed into a sequence of waypoints** through which the robot should pass
- Motion (aka local) planning is concerned with **finding appropriate robot motion commands**
  - Unlike global planners, **local planners consider current sensor measurements**
  - Local planners typically ensure that the robot passes through the waypoints, **potentially based on certain motion constraints**
  - Local planning requires a **motion model of the robot**

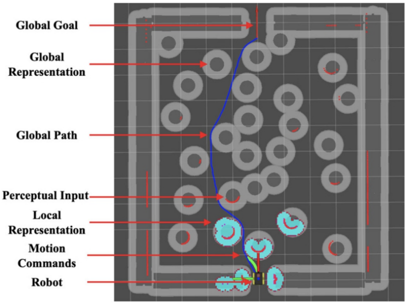

# Challenges with the Traditional Navigation Approach

## (Too) Geometric navigation

Typical navigation algorithms use only the geometry of the environment, but **important semantic aspects and cues are ignored**

# Challenges with the Traditional Navigation Approach

## (Too) Geometric navigation

Typical navigation algorithms use only the geometry of the environment, but **important semantic aspects and cues are ignored**

## Planar navigation

Many motion planners **assume that a robot navigates on a plane — this does not always hold, particularly in outdoor environments**

# Challenges with the Traditional Navigation Approach

### (Too) Geometric navigation

Typical navigation algorithms use only the geometry of the environment, but **important semantic aspects and cues are ignored**

### Planar navigation

Many motion planners **assume that a robot navigates on a plane — this does not always hold, particularly in outdoor environments**

### Hand-designed motion planning

Robot motion planning algorithms are usually hand-tuned, and **reusability between robot platforms is usually not guaranteed**

# Challenges with the Traditional Navigation Approach

### (Too) Geometric navigation

Typical navigation algorithms use only the geometry of the environment, but **important semantic aspects and cues are ignored**

### Planar navigation

Many motion planners **assume that a robot navigates on a plane — this does not always hold, particularly in outdoor environments**

### Hand-designed motion planning

Robot motion planning algorithms are usually hand-tuned, and **reusability between robot platforms is usually not guaranteed**

### Inability to use prior experiences

**Every navigation trial is treated independently of prior trials**, which makes it impossible to improve based on experience

# Challenges with the Traditional Navigation Approach

## (Too) Geometric navigation

Typical navigation algorithms use only the geometry of the environment, but **important semantic aspects and cues are ignored**

## Planar navigation

Many motion planners **assume that a robot navigates on a plane — this does not always hold, particularly in outdoor environments**

## Hand-designed motion planning

Robot motion planning algorithms are usually hand-tuned, and **reusability between robot platforms is usually not guaranteed**

## Inability to use prior experiences

**Every navigation trial is treated independently of prior trials**, which makes it impossible to improve based on experience

## Challenges with new environments

Most standard navigation frameworks are **designed for known environments**; adapting them for new environments can be challenging

# Indoor vs. Outdoor Navigation

## Indoor navigation

- Structured environments overall, but with a large diversity in how this structure manifests itself
- Environments typically change dynamically
  - Although that depends on the application domain (e.g. a package delivery center may be rather static)
- Often navigation among people necessary

# Indoor vs. Outdoor Navigation

## Indoor navigation

- Structured environments overall, but with a large diversity in how this structure manifests itself
- Environments typically change dynamically
  - Although that depends on the application domain (e.g. a package delivery center may be rather static)
- Often navigation among people necessary

## Outdoor navigation

- A large variety of environments and terrains (including extraterrestrial environments)
- Environments usually more static
  - But again, that depends on the application domain (e.g. roads are not very static)
- In some domains, navigation needs to follow well-defined rules (e.g. in autonomous driving)

# Navigation Learning Strategies

# Formal Framework

- From a formal point of view, the overall problem of applying learning for navigation **does not differ much from that of using learning for manipulation**

# Formal Framework

- From a formal point of view, the overall problem of applying learning for navigation **does not differ much from that of using learning for manipulation**
- Thus, we can also observe navigation learning in terms of a task family of $\mathrm{t}$ tasks, which can be represented as **a collection of Markov Decision Processes (MDPs)**:

$$P(\mathcal{M}) = \{(S_i, A, \mathcal{T}_i, R_i, C_i, \gamma) \mid 1 \leq i \leq \mathrm{t}\}$$

# Formal Framework

- From a formal point of view, the overall problem of applying learning for navigation **does not differ much from that of using learning for manipulation**

- Thus, we can also observe navigation learning in terms of a task family of $\mathrm{t}$ tasks, which can be represented as **a collection of Markov Decision Processes (MDPs)**:

$$P(\mathcal{M}) = \{(S_i, A, \mathcal{T}_i, R_i, C_i, \gamma) \mid 1 \leq i \leq \mathrm{t}\}$$

- The overall state representation $S_i$ also the same: $S_i = S_r \cup S_e$
  - But navigation is not concerned with handling objects, so a factored object-centric representation would typically not be useful for $S_e$

# State Representations

- The state for learning-based navigation can be represented in a variety of ways
- Some of these are illustrated in the figure below

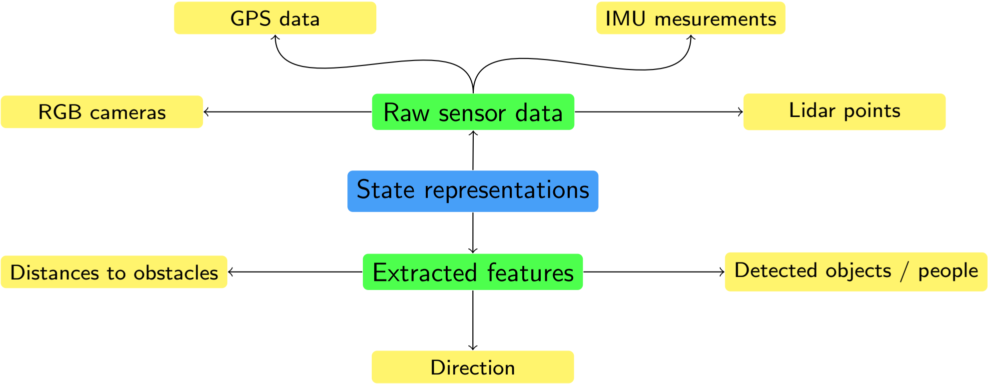

# What to Learn for Navigation?

Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8

**Motion planning and control for mobile robot navigation using machine learning: a survey**

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

End-to-end navigation learning

Learning for navigation

Learning dedicated navigation components

Parameter learning

# Parameter Learning

Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8

Motion planning and control for mobile robot navigation using machine learning: a survey

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

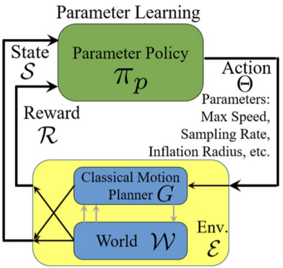


- Learning can be used within the established navigation framework, where **parameters that are used by components can be learned** (e.g. parameters of local planners)

# Parameter Learning

Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8
Motion planning and control for mobile robot navigation using machine learning: a survey
Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

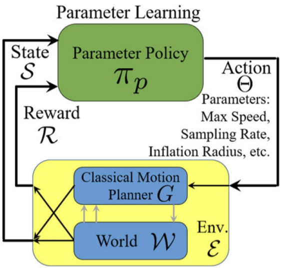


- Learning can be used within the established navigation framework, where **parameters that are used by components can be learned** (e.g. parameters of local planners)
- Learning can be useful for a large variety of parameters, e.g. weights of cost functions used during planning, motion model parameters, parameters of predictive models, etc.

# Parameter Learning


Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8

Motion planning and control for mobile robot navigation using machine learning: a survey

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

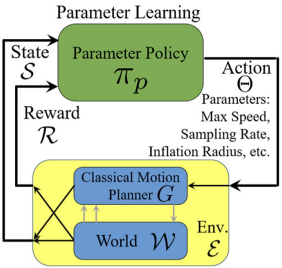


- Learning can be used within the established navigation framework, where **parameters that are used by components can be learned** (e.g. parameters of local planners)
- Learning can be useful for a large variety of parameters, e.g. weights of cost functions used during planning, motion model parameters, parameters of predictive models, etc.
- This makes it possible to both **preserve the properties of existing algorithms** and **improve those based on the experiences of a robot**

# Learning Dedicated Navigation Components


Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8

Motion planning and control for mobile robot navigation using machine learning: a survey

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

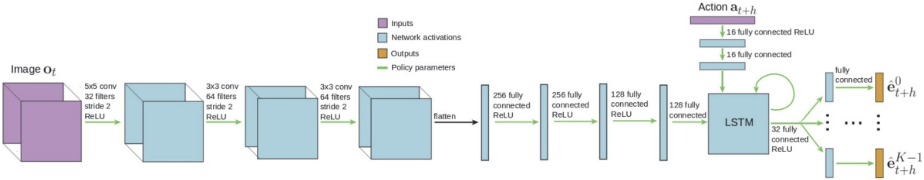


G. Kahn, P. Abbeel and S. Levine, "BADGR: An Autonomous Self-Supervised Learning-Based Navigation System," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1312–1319, 2021.

- Another alternative to incorporate learning into a robot navigation framework is to **replace complete components by a learning-based version**

# Learning Dedicated Navigation Components


Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8
Motion planning and control for mobile robot navigation using machine learning: a survey
Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

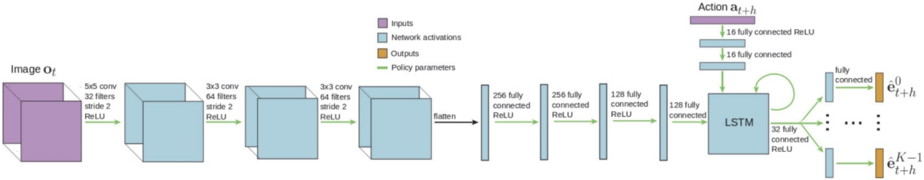


G. Kahn, P. Abbeel and S. Levine, “BADGR: An Autonomous Self-Supervised Learning-Based Navigation System,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1312–1319, 2021.

- Another alternative to incorporate learning into a robot navigation framework is to **replace complete components by a learning-based version**
- Different components can be replaced in this manner, for instance:

# Learning Dedicated Navigation Components


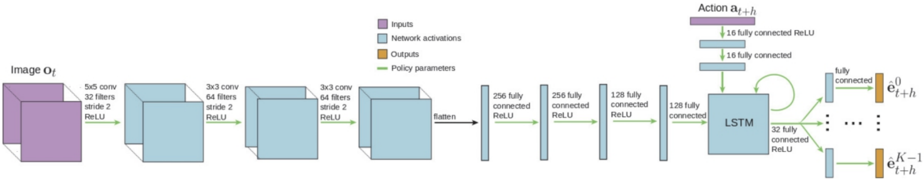


G. Kahn, P. Abbeel and S. Levine, "BADGR: An Autonomous Self-Supervised Learning-Based Navigation System," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1312–1319, 2021.

- Another alternative to incorporate learning into a robot navigation framework is to **replace complete components by a learning-based version**
- Different components can be replaced in this manner, for instance:
  - a trained policy can be used instead of a local planner


Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8

Motion planning and control for mobile robot navigation using machine learning: a survey

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

# Learning Dedicated Navigation Components

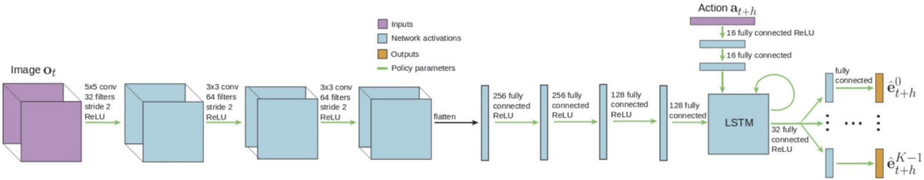


G. Kahn, P. Abbeel and S. Levine, "BADGR: An Autonomous Self-Supervised Learning-Based Navigation System," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1312–1319, 2021.

- Another alternative to incorporate learning into a robot navigation framework is to **replace complete components by a learning-based version**
- Different components can be replaced in this manner, for instance:
  - a trained policy can be used instead of a local planner
  - predictions produced by a learned model can be used for motion planning, e.g. as in BADRG (a diagram of the component learned there is shown above)

# End-To-End Navigation Learning

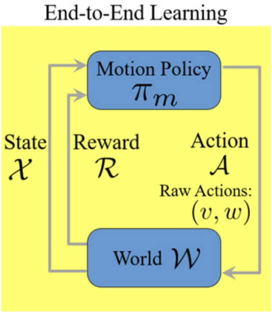


- A more radical approach is to **replace the complete navigation pipeline with learning-based components that are trained in an end-to-end fashion** (i.e. where the learned model outputs motion commands directly based on sensory input)

# End-To-End Navigation Learning


Learning Robotic Navigation from Experience: Principles, Methods, and Recent Results

Sergey Levine, Dhruv Shah

UC Berkeley

{svlevine, shah}@eecs.berkeley.edu


Autonomous Robots (2022) 46:569–597

https://doi.org/10.1007/s10514-022-10039-8

Motion planning and control for mobile robot navigation using machine learning: a survey

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

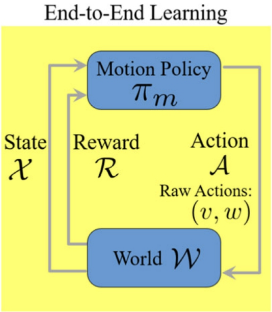


- A more radical approach is to **replace the complete navigation pipeline with learning-based components that are trained in an end-to-end fashion** (i.e. where the learned model outputs motion commands directly based on sensory input)
- The attractive prospect of this is that **the engineering effort required by the traditional navigation framework is, in principle, reduced**
  - Although whether that holds depends on how learning is exactly performed

# End-To-End Navigation Learning

**Learning Robotic Navigation from Experience: Principles, Methods, and Recent Results**

Sergey Levine, Dhruv Shah
UC Berkeley
{svlevine, shah}@eecs.berkeley.edu

Autonomous Robots (2022) 46:569–597
https://doi.org/10.1007/s10514-022-10039-8

**Motion planning and control for mobile robot navigation using machine learning: a survey**

Xuesu Xiao[1] · Bo Liu[1] · Garrett Warnell[3] · Peter Stone[1,2]

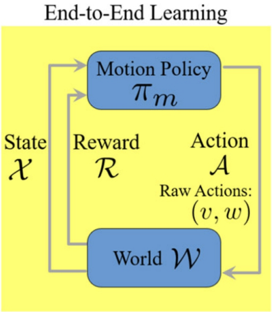


- A more radical approach is to **replace the complete navigation pipeline with learning-based components that are trained in an end-to-end fashion** (i.e. where the learned model outputs motion commands directly based on sensory input)
- The attractive prospect of this is that **the engineering effort required by the traditional navigation framework is, in principle, reduced**
  - Although whether that holds depends on how learning is exactly performed
- End-to-end learning leads to **black-box models whose decision-making process can be difficult to analyse**
  - This is in contrast to components used in traditional navigation, which are usually based on mathematical models that are easier to understand

# Action Spaces for Navigation Policies

- Depending on what is learned, navigation policies can have a variety of action spaces:

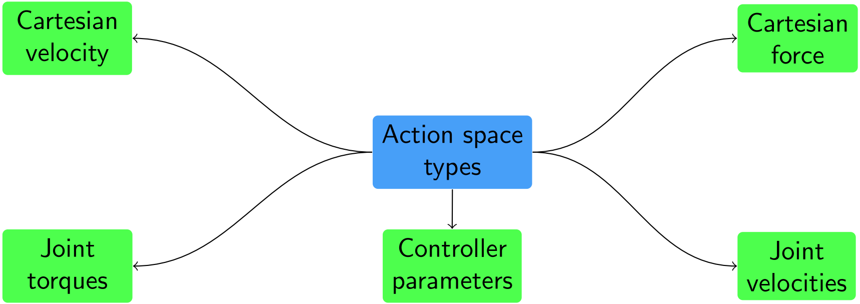

# Action Spaces for Navigation Policies

- Depending on what is learned, navigation policies can have a variety of action spaces:

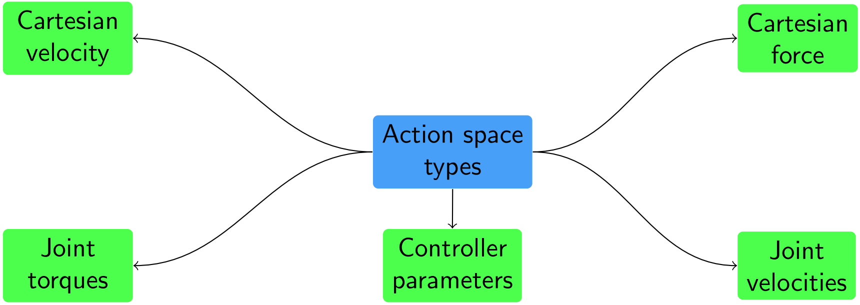


- The figure looks familiar? It is exactly the same as the one we had in the case of robot manipulation (last lecture)!

# The Case of Autonomous Driving

# Learning in Autonomous Driving

- Autonomous driving is **a very challenging problem due to the fact that roads are very dynamic and involve many other traffic participants**
  - Pure model-based methods can quickly reach their limit here because there are simply too many relevant aspects to consider

# Learning in Autonomous Driving

- Autonomous driving is **a very challenging problem due to the fact that roads are very dynamic and involve many other traffic participants**
  - Pure model-based methods can quickly reach their limit here because there are simply too many relevant aspects to consider
- Autonomous driving is an important area where the benefits of learning have been directly visible
  - Particularly **advances in computer vision have made it possible to process complex input** and, as a result, produce more complex driving behaviours

# Driving Tasks Where Learning Has Been Applied

IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS, VOL. 23, NO. 6, JUNE 2022

Deep Reinforcement Learning for Autonomous Driving: A Survey

B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Yogamani, and Patrick Pérez

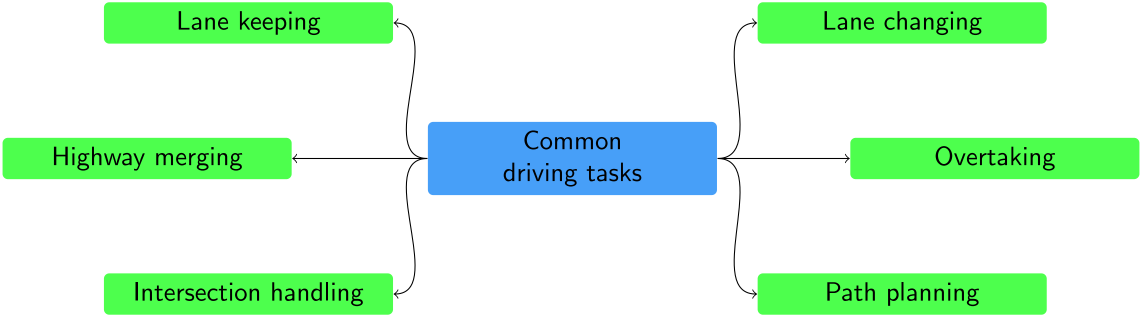

# Autonomous Driving Framework


Deep Reinforcement Learning for Autonomous Driving: A Survey

B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Yogamani, and Patrick Pérez

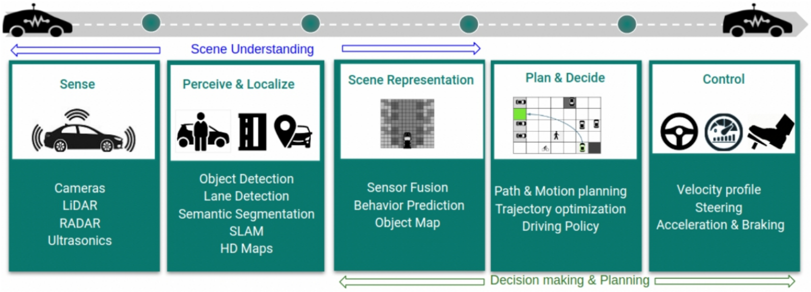


- From a general point of view, **the navigation framework for an autonomous vehicle is the same as that for any other mobile robot** (see the figure below)
  - But the necessity to follow traffic rules is a major difference with other navigation domains

# Autonomous Driving Framework

IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS, VOL. 23, NO. 6, JUNE 2022

Deep Reinforcement Learning for Autonomous Driving: A Survey

B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad A. Al Sallab, Senthil Yogamani, and Patrick Pérez

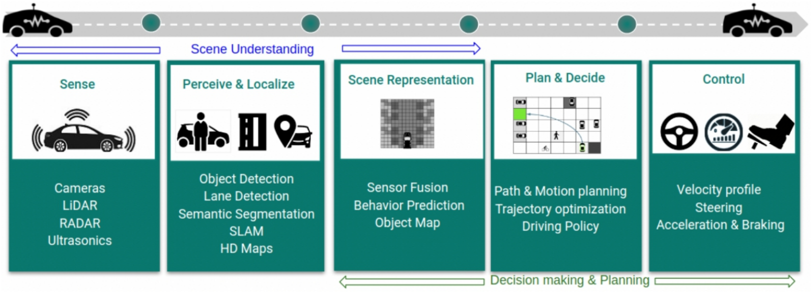


- From a general point of view, **the navigation framework for an autonomous vehicle is the same as that for any other mobile robot** (see the figure below)
  - But the necessity to follow traffic rules is a major difference with other navigation domains
- An autonomous vehicle navigates in an outdoor environment, so **it can benefit from additional signals (e.g. GPS) that cannot be used in indoor environments**

# Autonomous Driving Based on Deep Learning

- Due to its ability to process high-dimensional data, **deep learning has been applied to a variety of autonomous driving problems**

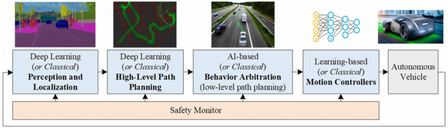


S. Grigorescu et al., "A survey of deep learning techniques for autonomous driving," *Journal of Field Robotics*, vol. 37, no. 3, pp. 362–386, Apr. 2020.

# Autonomous Driving Based on Deep Learning

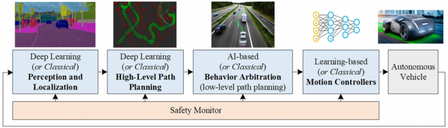


S. Grigorescu et al., "A survey of deep learning techniques for autonomous driving," *Journal of Field Robotics*, vol. 37, no. 3, pp. 362–386, Apr. 2020.

- Due to its ability to process high-dimensional data, **deep learning has been applied to a variety of autonomous driving problems**
- The most common approach is to **use deep learning for specific aspects of the driving pipeline**, for instance:

# Autonomous Driving Based on Deep Learning

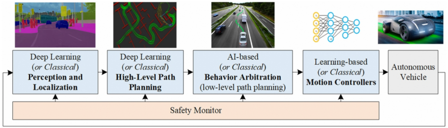


S. Grigorescu et al., "A survey of deep learning techniques for autonomous driving," *Journal of Field Robotics*, vol. 37, no. 3, pp. 362–386, Apr. 2020.

- Due to its ability to process high-dimensional data, **deep learning has been applied to a variety of autonomous driving problems**
- The most common approach is to **use deep learning for specific aspects of the driving pipeline**, for instance:
  - detecting people and other vehicles (scene understanding)

# Autonomous Driving Based on Deep Learning

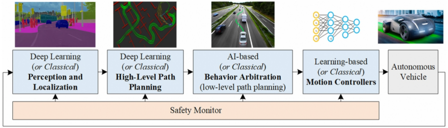


S. Grigorescu et al., "A survey of deep learning techniques for autonomous driving," *Journal of Field Robotics*, vol. 37, no. 3, pp. 362–386, Apr. 2020.

- Due to its ability to process high-dimensional data, **deep learning has been applied to a variety of autonomous driving problems**
- The most common approach is to **use deep learning for specific aspects of the driving pipeline**, for instance:
  - detecting people and other vehicles (scene understanding)
  - controlling particular driving behaviours, such as lane keeping, using a learned policy (motion planning)

# Autonomous Driving Based on Deep Learning

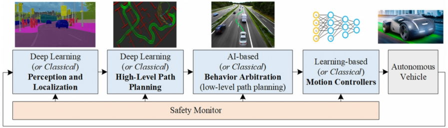


S. Grigorescu et al., "A survey of deep learning techniques for autonomous driving," *Journal of Field Robotics*, vol. 37, no. 3, pp. 362–386, Apr. 2020.

- Due to its ability to process high-dimensional data, **deep learning has been applied to a variety of autonomous driving problems**
- The most common approach is to **use deep learning for specific aspects of the driving pipeline**, for instance:
  - detecting people and other vehicles (scene understanding)
  - controlling particular driving behaviours, such as lane keeping, using a learned policy (motion planning)
- Attempts to use end-to-end deep learning also exist

# Autonomous Driving Based on Deep Learning

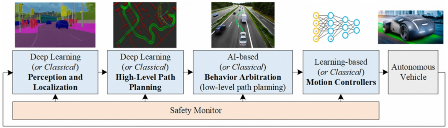


S. Grigorescu et al., "A survey of deep learning techniques for autonomous driving," *Journal of Field Robotics*, vol. 37, no. 3, pp. 362–386, Apr. 2020.

- Due to its ability to process high-dimensional data, **deep learning has been applied to a variety of autonomous driving problems**
- The most common approach is to **use deep learning for specific aspects of the driving pipeline**, for instance:
  - detecting people and other vehicles (scene understanding)
  - controlling particular driving behaviours, such as lane keeping, using a learned policy (motion planning)
- Attempts to use end-to-end deep learning also exist
- Deep neural networks **need dedicated hardware** (e.g. graphic cards) **to run efficiently**
  - Autonomous driving requires real-time control, so applying deep learning on autonomous vehicles requires **implementations that support efficient inference**

# Explainability

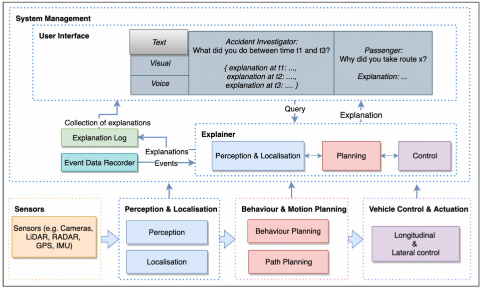


D. Omeiza et al., "Explanations in Autonomous Driving: A Survey," *IEEE Trans. Intelligent Transportation Systems*, vol. 23, no. 8, pp. 10142–10162, Aug. 2022.

- When using (deep) learning in robot navigation (but particularly in autonomous driving), **it is not always easy to understand why a decision has been made**

# Explainability

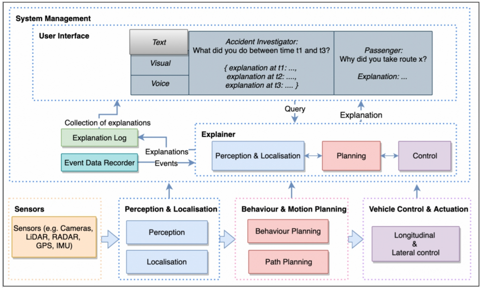


D. Omeiza et al., "Explanations in Autonomous Driving: A Survey," *IEEE Trans. Intelligent Transportation Systems*, vol. 23, no. 8, pp. 10142–10162, Aug. 2022.

- When using (deep) learning in robot navigation (but particularly in autonomous driving), **it is not always easy to understand why a decision has been made**
- The property of being able to understand an autonomous decision-making process is referred to as **explainability**
  - Particularly relevant in case of failures that lead to accidents

# Explainability

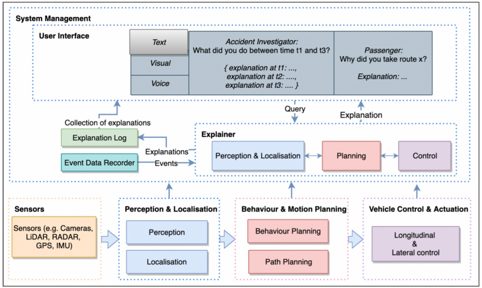


D. Omeiza et al., "Explanations in Autonomous Driving: A Survey," *IEEE Trans. Intelligent Transportation Systems*, vol. 23, no. 8, pp. 10142–10162, Aug. 2022.

- When using (deep) learning in robot navigation (but particularly in autonomous driving), **it is not always easy to understand why a decision has been made**
- The property of being able to understand an autonomous decision-making process is referred to as **explainability**
  - Particularly relevant in case of failures that lead to accidents
- Certifying autonomous driving behaviours is **generally difficult if explainability cannot be guaranteed**
  - Various ideas for how explainability should be achieved exist in the literature (one example is shown on the left)

# Explainability

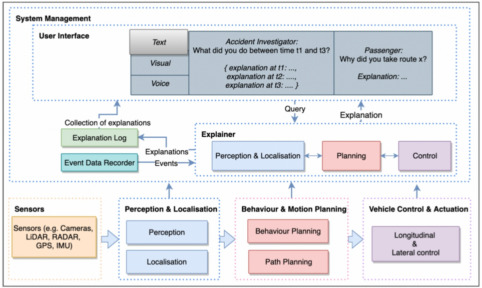


D. Omeiza et al., "Explanations in Autonomous Driving: A Survey," *IEEE Trans. Intelligent Transportation Systems*, vol. 23, no. 8, pp. 10142–10162, Aug. 2022.

- When using (deep) learning in robot navigation (but particularly in autonomous driving), **it is not always easy to understand why a decision has been made**
- The property of being able to understand an autonomous decision-making process is referred to as **explainability**
  - Particularly relevant in case of failures that lead to accidents
- Certifying autonomous driving behaviours is **generally difficult if explainability cannot be guaranteed**
  - Various ideas for how explainability should be achieved exist in the literature (one example is shown on the left)
- We will discuss techniques for explainability later in the course