Centre *for* Assuring Autonomy

The challenges of autonomous vehicle decision-making in a real-world scenario

Consider a scenario such as that shown in figure 1 below.

Decision Matrix

Action	Safety	Performance	Deadline
Accelerate	0.78 – 0.82	0.96 – 0.98	2 sec
Mt. Speed	0.86 – 0.98	0.78 – 0.95	N/A
Stop	0.90 – 0.99	0.76 – 0.81	3 sec

Scenes

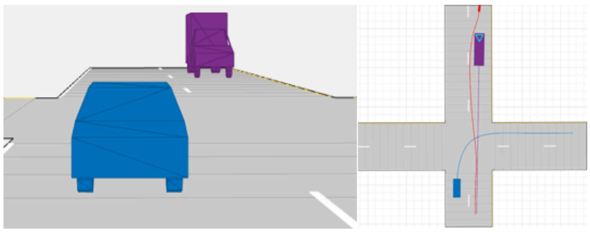

Figure 1: This is an illustration of an autonomous vehicle (blue) intended to take a right turn at an intersection. The scene on the right depicts the position of the car (4 seconds away from the middle of the junction) and the truck as well as their intended trajectory. The table on the left shows a decision matrix containing all valid actions for our car to choose from considering the situation. Each of these actions are associated with values in terms of safety and performance reflecting the car's assessment of risk and efficiency. Some actions have a time ("deadline") attached to them which informs how long that particular action will remain valid.

Imaging a fully autonomous car (blue) approaching a junction, intending to make a right turn. Let's call it our "ego vehicle". While approaching the junction, it detects a truck (purple) coming from the opposite direction. The car is equipped with a stereo camera for environmental perception and GPS for localization and global path planning. Its available actions are: accelerate, decelerate, stop, maintain speed, and steer left or right. The challenge lies in deciding which action to take to ensure a safe and smooth turn. To make this decision, the car must predict the outcome of each possible action and select the one that provides the best result in terms of safety and performance. Let's break this problem into smaller parts to better understand the decision-making process of the ego vehicle.

Understanding its own position

The ego vehicle first needs to monitor its own state, including speed, position, and distance to the junction. Since it cannot make the turn until it reaches the middle of the junction, it has some time to decide and adjust its velocity. For example, attempting a sharp turn at 30 mph without slowing down risks losing control, something the ego vehicle must avoid by understanding its dynamics and geometry of the roads and junctions. Therefore, required by the vehicle dynamics and the manoeuvre, the ego vehicle reduces its speed to 10 mph to take the turn as safely and smoothly as possible.

Next, the ego vehicle must analyze its surroundings. Detecting the oncoming truck, it identifies the truck as a potential risk because their paths are likely to intersect when the ego vehicle makes the turn at the junction. By estimating its own time to reach the middle of the junction (let's say 4 seconds) and the truck's position and velocity, the ego vehicle predicts where both vehicles will be. This prediction allows it to classify the truck as an object of interest for tracking.

Now, based on the understanding of the situation and prediction of how it will evolve in the near future, our ego vehicle generates a list of possible actions that are reasonable and valid at that moment. In this case, three options stand out: accelerate, maintain speed, and stop.

Now, the challenge is to select the best option that maximizes our ego vehicle's overall objectives. The primary goal of our ego vehicle is to avoid collisions and minimize near misses (i.e. situations where a collision or accident is narrowly avoided), thus ensuring maximum safety. At the same time, improving performance by reducing travel time and enhancing fuel efficiency remains an important objective.

To evaluate safety and performance, the ego vehicle relies on predefined metrics. **Time to Collision (TTC)** estimates how long it will take for a collision to occur if both vehicles maintain their current paths. **Minimum Distance to Collision (DTC)** measures the closest distance the car gets to other vehicles during its manoeuvre. **Collision Probability** provides a probabilistic estimate of the crash risk, factoring in uncertainties in other vehicles' behavior and prediction errors. Performance, on the other hand, may consider factors like manoeuvre completion time, fuel efficiency and comfort of the passenger inside.

Let's take a look at how our ego vehicle may reason over the three options we mentioned earlier.

Option 1: Accelerate

This is a valid option because, if our car can accelerate early enough, it may clear the junction before the truck arrives. This option minimizes travel time and improves performance. However, it comes with a risk: if the truck suddenly speeds up, the car may not have enough time to stop safely. While acceleration offers decent performance, its safety value is slightly lower due to the possibility and effect of sudden behavioural change of the surrounding vehicles.

Option 2: Maintain Speed

Another option for the car is to maintain its current speed of 10 mph. Based on the truck's position and speed, the car estimates that by the time it reaches the turning point, the truck will have already crossed the junction. This means the car can make the turn without needing to stop completely.

Keeping the current speed is also a safer choice compared to accelerating because it allows the car to either speed up or stop easily if the situation changes. This also gives the car more time to gather extra information and monitor how the situation develops. However, this option is not as fast as accelerating, leading to a slight increase in travel time. As a result, its performance value is slightly lower than the previous option.

Option 3: Stop

The car could choose to stop at the middle of the junction and give way to the truck. This option prioritizes safety by eliminating the risk of collision. However, it significantly impacts performance by increasing travel time and disrupting the flow of traffic.

The safety and performance assessments depend on sensor data and predictions. Both of these have inherent uncertainties. To account for this, the safety and performance values in the decision matrix are expressed as ranges between 0 and 1. As the car gathers more data, these ranges narrow, reflecting the improvement of confidence in its predictions.

While gathering more information can improve decision-making, the car cannot wait indefinitely. Time constraints are critical. Each action has a **"deadline"**—the time by which it must be enacted to remain effective. For instance:

- **Accelerate:** Must begin within 2 seconds to ensure the vehicle will clear the junction before the truck arrives. If this decision is enacted after that time, the manoeuvre will no longer be considered safe.
- **Stop:** Must begin within 3 seconds to halt smoothly at the middle of the junction. If this decision is made after that time then excessive

braking will be required for the vehicle to stop which may endanger the passenger inside or vehicle following behind.

- **Maintain Speed:** Requires no immediate action since it's already in progress and therefore, no deadline.

In this scenario, let's assume the car chooses to maintain speed, offering a good trade-off between safety and performance.

Adapting to changing environments

Since self-driving cars are designed to operate in real-world environments, they face unique challenges in decision-making. The real world is dynamic and constantly changing—what is true in one moment may no longer apply the next. This is precisely the case for our ego vehicle. One second ago, when the ego vehicle chose to maintain its speed, it was focused solely on the truck ahead. However, a second later, it detects a motorbike previously hidden behind the truck. This new information prompts the ego vehicle to reevaluate the situation and come up with another set of possible actions along with their updated assessments for safety and performance as shown in the decision matrix in figure 2 below. As the velocity of our ego vehicle is quite low and it is still 3 seconds away from the middle of the junction, it decides to maintain the same speed, allowing itself to gather more data to estimate the motorbike's speed, position, and intention.

Decision Matrix

Action	Safety	Performance	Deadline
Accelerate	0.84 – 0.88	0.96 – 0.98	1 sec
Mt. Speed	0.72 – 0.80	0.78 – 0.95	N/A
Stop	0.92 – 0.96	0.75 – 0.80	2 sec

Scenes

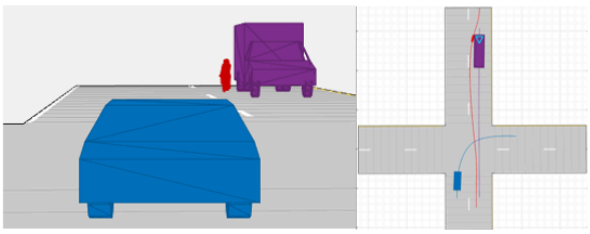

Figure 2: This is an illustration of how the scene described in figure 1 has evolved after 1 second. This is the moment when our ego vehicle first detects the motorbike (red) previously hidden behind the truck with information about its intention. With this new information, the ego vehicle re-assesses the situation and updates the safety and performance measurements for each viable option in the decision matrix.

After another second, the car observes that the motorbike has accelerated and is attempting to overtake the truck. At this point, accelerating or maintaining speed becomes unsafe. Therefore, it chooses to opt for the only valid option at that moment: stopping in the middle of the junction as it offers the highest safety guarantee by allowing the motorbike and truck to clear the area first.

Decision Matrix

Action	Safety	Performance	Deadline
Accelerate	Invalid	Invalid	N/A
Mt. Speed	Invalid	Invalid	N/A
Stop	0.96 – 0.98	0.73 – 0.77	1 sec

Scenes

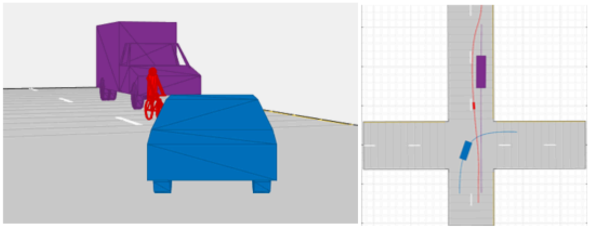

Figure 3: This illustrates the evolution of our first scene (figure 1) after 2 seconds. Our ego vehicle has collected enough camera images of the motorbike since it was first detected and observed that the motorbike is trying to overtake the truck. Based on this new observation and understanding of the situation, our car updates its decision matrix which shows that first two actions are no longer valid leaving "Stop" as the only viable and safe action to enact.

While this scenario is relatively simple compared to the complex situations an autonomous vehicle will encounter in the real world, it serves as a baseline case for understanding key decision-making challenges. By analysing this controlled example, we can identify crucial factors affecting autonomous decisions and explore how they can be managed. Future experiments can be built upon this scenario by introducing additional complexities, but this baseline provides a solid starting point for refining decision-making strategies in dynamic environments.