Conflicting Moral Codes for Self-Driving Cars: The Single Car Case

Ahmed Ibrahim , Yasmin Mansy, Nourhan Ehab and Amr Elmougy
Computer Science and Engineering Department, German University in Cairo, Cairo, Egypt

Keywords: Conflict Resolution, Thomas Kilmann Model, Self-Driving Cars.

Abstract: There has been incremental advances in the development of self-driving cars. However, there are still several gaps that must be filled before a fully autonomous vehicle can be achieved. The ability to resolve conflicts in the event of an unavoidable accident is one of the most prominent and crucial aspects of a self-driving car that is currently absent. To address this gap, this paper aims to resolve moral conflicts in self-driving cars in the case of an unavoidable accidents. Assuming we have a predefined rule set that specifies how a car should morally react, any clash between these rules could result in a critical conflict. In this paper, we propose a novel procedure to resolve such conflicts by combining the Thomas Kilmann conflict resolution model together with decision trees. Evaluation results showcase that our proposed procedure excels in distinct ways, enabling the self-driving car to make a decision that will yield the best moral outcome in conflicting scenarios.

1 INTRODUCTION

Driving a vehicle with minimum human intervention is no longer science fiction (Cowger Jr, 2018). Currently, the automotive industry is undergoing a quantum shift towards a future in which the driver's role in operating his or her vehicle will diminish to the point of being eliminated (Barabas et al., 2017). For a self-driving car to reach a stage where it would be fully autonomous, it must be able to make faster decisions under a wide range of situations, which could include moral dilemmas such as choosing which potential victims to avoid. Such situations include unavoidable accidents where life and death decisions have to be made.

Although autonomous vehicles will eliminate the major cause of human error, unavoidable deadly accidents, where the vehicle must make a life-or-death decision in a fraction of a second, cannot be prevented (Borenstein et al., 2017). In situations where moral decisions need to be made by the vehicle, an ethical framework must be implemented, where a predefined set of rules would allow the car to react. However, if these rules contradict, the car would be unable to determine a proper action, leading to a conflict. As discussed later in this study, this must be resolved to prevent catastrophic losses from indecision.

Accordingly, the objective of this research which is to apply different conflict resolution techniques to resolve conflicts that may arise due to various factors that would be discussed in further details later in this study, allowing the self-driving car to react in a timely manner. As such, the ability to resolve conflicts that may arise within a single self-driving vehicle has not been previously addressed. In addition, the Thomas Kilmann Model has primarily been used to resolve conflicts between humans. Since we are attempting to imitate a human-like trait, ethics, it stands to reason that in order to imitate the ways in which conflicts in ethics should be resolved, we would need to employ another negotiation method used by humans, which has not been attempted before to resolve conflicts within self-driving vehicles using the Thomas Kilmann Model in previous literature. Thus, the aim is to provide a fast and accurate conflict resolution technique that would allow self-driving cars to resolve the conflict and react on time in case of an unavoidable accident. With the help of the Carla simulator and ScenarioRunner, it is possible to build complex conflicting scenarios where different resolution techniques can be applied to imitate a real-world scenario as closely as possible.

The rest of this paper is structured as follows. Section 2 will present necessary background. Section 3 describes the proposed conflict resolution model. Section 4 presents the evaluation of the proposed model. Finally, Section 5 outlines some concluding remarks.

Ibrahim, A., Mansy, Y., Ehab, N. and Elmougy, A.
Conflicting Moral Codes for Self-Driving Cars: The Single Car Case.
DOI: 10.5220/0011677500003393
In *Proceedings of the 15th International Conference on Agents and Artificial Intelligence (ICAART 2023) - Volume 1*, pages 249-256
ISBN: 978-989-758-623-1; ISSN: 2184-433X

2 BACKGROUND

Unavoidable accidents are described as crucial situations in which no solutions can be developed in the time available (to a human or self-driving car) to completely avoid the accident. The focus of this work is on unavoidable accidents that result in fatalities, or catastrophes in which human lives are lost. A self-driving car must decide what action to take in the event of an unavoidable accident, which includes deciding on potential casualties. This is only partially today investigated legally and ethically (Karnouskos, 2018).

2.1 The Need for Ethics

In a research paper by (Lin, 2016), the author provides a basic scenario demonstrating the need for ethics in self driving vehicles. When a self-driving car is faced with a terrible choice in the distant future: where it would hit an eight-year-old girl if it swerved to the left, or an 80-year-old grandma if it swerved to the right. Given the speed of the car, both victims would almost probably die upon impact. If you do not swerve, you will hit and kill both victims. Therefore, there are sufficient reasons to believe that you should swerve in one of two directions (Lin, 2016). But how should the vehicle react in such a scenario? There may be reasons to pick one over the other, regardless of how unappealing or distressing those reasons may be. This is a dilemma that is not easily solvable, highlighting the importance of ethics in the development of self-driving cars.

2.2 Conflicts in Decisions

Referring to the same case given in a research paper by (Patrick, 2016), some may claim that hitting the grandma is the lesser of two evils, at least in some people's eyes. As the girl has her entire life ahead of her, whereas the grandmother has already lived a full life and had her fair share of experiences, there are reasons that seem to weigh in favor of saving the little girl over the grandmother, if an accident is unavoidable.

But what if there were two girls of the same age on each side of the road? Even if an autonomous car with an ethical framework can select between the two evils. Both sides would have equal priority according to the car's framework, having the same age, gender, and any other aspect the car could compare to. Resulting in a conflict since the car can no longer make a decision. So even if we develop an ethical framework capable of resolving such complex scenarios in an attempt to replicate a human trait, conflicts will still arise. Conflicts can arise for a number of reasons, some of which are predictable and could be avoided, while others we would not even know about.

Some might argue that future cars will not have to make difficult ethical decisions, that stopping the car or handing control to the human driver is the easy way around ethics. However, braking and handing over control will not always be enough. These solutions may be the best we have today, but if automated cars are ever to be used more widely outside of limited highway environments, they will need more response capabilities (Lin, 2016).

In future autonomous cars, crash-avoidance features alone will not be sufficient. Sometimes an accident will be unavoidable as a matter of physics, for myriad reasons–such as insufficient time to press the brakes, technological errors, misaligned sensors, terrible weather conditions, and just pure bad luck. Therefore, self-driving cars will require crash-optimization strategies.

2.3 Thomas-Kilmann Conflict Model

In 1974, Kenneth W. Thomas and Ralph H. Kilmann introduced the ThomasKilmann Conflict Mode Instrument, which is presented as an alternative approach to the arising conflicts. The Thomas-Kilmann model identifies five distinct conflict resolution behaviours and how they effect problem resolution. As illustrated in Figure 1, the various types of behaviour reflect various degrees of assertiveness and cooperation (Fahy et al., 2021).

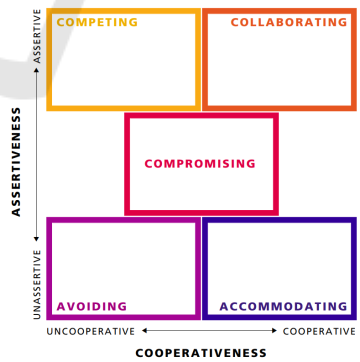

Figure 1: Thomas-Kilmann Conflict Model.

The Thomas-Kilmann technique has primarily been used to resolve conflicts in human-to-human negotiations. It has previously been used to resolve conflicts between medical staff in pediatric surgery. (Fahy et al., 2021). This is a real-time, critical, and time-sensitive situation in which a decision made could mean the difference between life and death.

And must be made within a matter of seconds. This is very similar to the situations we would face in the event of an unavoidable accident, where the conflicts must be resolved within very short deadlines. Furthermore, since we are attempting to mimic a human-like trait, ethics, it stands to reason that in order to mimic the way conflicts in ethics should be resolved, we would need to employ another human negotiation method.

Next, we will analyse the five distinct types of behaviour and how they can be utilised to resolve an ethical conflict in the event of an unavoidable accident.

1. **Competing:** Power-oriented, assertive, and uncooperative. Having rules/decisions of equal importance. Each decision uses whatever power seems appropriate to win that position, by competing on any aspect favoring one side over the other.
2. **Collaborating:** assertive and cooperative. Usually involves two or more entities trying to find a solution that meets and satisfies both concerns.
3. **Compromising:** assertive and cooperative. Compromising splits the difference or exchanges concessions to find a quick middle-ground position.
4. **Avoiding:** unassertive and uncooperative. Knowing one of two conflicting decisions is wrong allowing for a better decision to be considered.
5. **Accommodating:** is unassertive and cooperativethe opposite of competingand involves an element of self-sacrifice. During conflict resolution, one side may realise its decision was wrong, thus backing up allowing for a better solution.

3 CONFLICT RESOLUTION

Our proposed conflict resolution model consists of a pipeline of three main stages: (1) Collecting diverse data using a variety of sensors. Along with the formulation of rules that would serve as the car's ethical foundation in the event of an unavoidable accident; (2) Conflict resolution in case moral conflicts arise by utilizing decision trees; and (3) Deciding which actions to take after the conflict resolution. The three stages are described in this section.

3.1 Data & Rule Set

First we start by data collection, CARLA simulator (Dosovitskiy et al., 2017) was used which is an open-source autonomous driving simulator that utilize lot of sensors that vehicles rely on to obtain data from their surroundings, ranging from cameras to Radar, LIDAR, and many more.

After we are done with data collected using different and various types of sensors, we then referred to the MIT Moral Machine Experiment (Awad and Dsouza, 2018) as our rule set. Allowing us to create a 18 priority based rule set as shown in Figure 2.

Having the top 9 rules indicated in blue as the highest priority rules and the other nine indicated in red as the lower priority rules. These priorities were established in the following manner according to a research paper by (Anbar, 1983) on estimating the difference between two probabilities.

$$(nd * (1 - (highest\Delta P - current\Delta P)) + k$$

Where n is the number of the rules multiplied by a factor d and Δ P is the difference between the probability of sparing characters having the attribute in blue and the probability of sparing characters having the opposing attribute in red (for example Elderly/Young) (Awad and Dsouza, 2018). And k is an offset that is larger than the smallest expected negative number in our rule set.

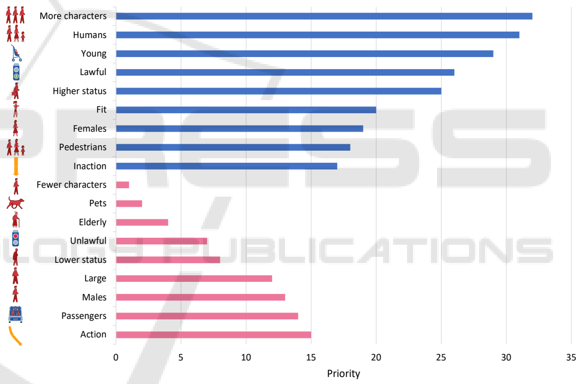

Figure 2: Priority Based Rule Set.

3.2 Conflict Resolution via Decision Trees

In this work, we use decision trees as means to resolve conflicts. A decision tree is an optimal choice for data structuring since it allow for computational efficiency and accuracy of classification according to a research paper by (Safavian et al, 1991) (Safavian and Landgrebe, 1991).

3.2.1 Types of Conflicts

Decision trees classify any object, in our case a scenario by moving it down the tree from the root to a leaf node, with the leaf node providing the scenario's evaluation. Each node in the tree represents a test case for some attribute, and each edge descending from that node represents one of the test case's possible answers. This recursive approach is repeated for each sub-tree rooted at the new nodes. A priority score can be assigned at each level moving down the decision tree and adds up until a leaf node is reached. The situation with the highest final priority score will be given the top priority. There are two types of possible conflicts that can occur:

1. **Conflicting Nodes:** this form of conflict would occur within the decision tree nodes themselves. We classified such types of conflicts into three main categories.
   (a) **Comparable Branching:** Where an attribute depends on comparing its value to the other lanes, such as more (versus less) characters, having both lanes with the same number of characters would cause a conflict and prevent branching to either node.
   (b) **Mix of Characters:** Having a mix of different characters standing on the same lane, such as having young and elderly ages standing on the same lane in certain scenarios, prevents branching to either node, as illustrated in Figure 3.
   (c) **Attributes with no Value:** Finally, not having the value for a certain attribute available at a certain level, which could be due to various factors such as a lack of time or inaccurate sensors.
2. **Conflicting Priority Scores:** occurs when the decision tree has been entirely traversed and the assigned scores for each lane are exactly equal.

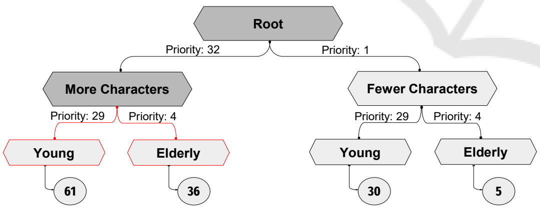

Figure 3: Conflicting decision tree.

3.2.2 Resolution Techniques

We use the Thomas Kilmann model techniques: **competing**, **avoiding**, and **accommodating**. The remaining two techniques (collaborating and compromising) are not considered since they deal with multi-agent cases, which is not the focus in this paper. Alongside the above mentioned techniques, the **verifying nodes** technique is used. Since the time required to traverse the entire tree may exceed the time available before the collision, the resolution strategies make use of a **deadline** by which we must stop traversing the decision tree and assign the current scores, leaving extra time to resolve any conflicts if both lanes have the same priority scores. In what follows, we describe the operation of the different resolution techniques on decision trees.

Competing: implies contending over one or more aspects that might give one lane an edge over the other. In our case the competition would be based on two critical factors, **time** and **highest priority score**.

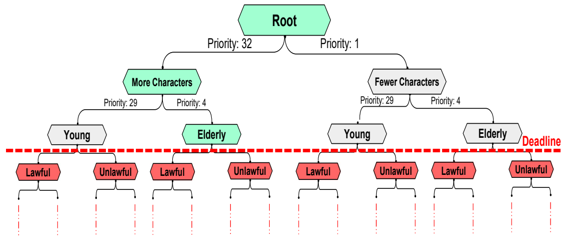

Figure 4: Deadline met before traversing entire tree.

Figure 6 illustrates a scenario in which we were unable to traverse the entire decision tree as we hit the deadline. Although we ended up with the same priority score, resulting in a conflict, it is uncertain whether either of the lanes might have achieved a higher priority if we had enough time to complete the entire decision tree.

As there is insufficient time to traverse the entire tree, the competing technique scans all of the leaf nodes for each of the two lanes in search of the highest possible priority score that could have been obtained. And the one that have the highest possible score is assigned a higher priority.

Competing is utilized in situations where time is a prime concern, so time is considered to be a vital factor. Each lane's time to reach a final node and provide a score is recorded, regardless of whether this final node is a leaf or if the lane did not complete the entire tree. The quicker a lane was able to receive a score, the fewer conflicting nodes it encountered while traversing the tree, resulting in it being faster and having an advantage over the other lane by completing the task in less time. And because an unavoidable accident is an extremely time-sensitive situation, it would be essential to have a lane with fewer conflicts and a shorter computation time.

Avoiding: implies admitting that you were mistaken or knowing that the other side has a better approach to this circumstance (Schaubhut, 2007).

As shown in Figure 5, lane A contained all Pets, while lane B contained humans, which should be prioritised according to our rule set. However, lane A received a higher Priority Score due to the fact that the traversal of the decision tree was interrupted after just two levels, resulting in both lanes having equal or even higher priority to lane A.

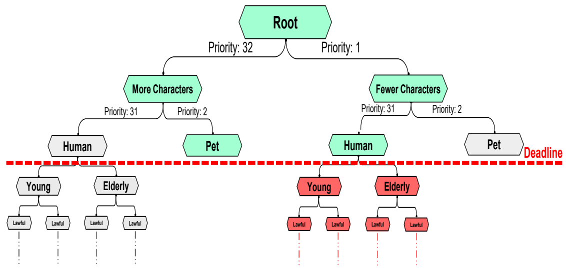

Figure 5: One of the two lanes is a leaf node.

In this case, the decision tree was not entirely completed. The avoidance technique would be applied by determining whether or not the final node for each lane has any children. If a lane came to a stop at a leaf node. It should declare so and thus avoid the conflict in the first place, saving us more time and resulting in a better outcome since the decision would be made earlier.

Accommodating: is realizing that its own decision was incorrect, allowing for a better solution to be considered (Schaubhut, 2007). It is also possible that the current choice or node has been chosen most of the time so it would accommodate to allow the other nodes to contribute to the process of decision making as well, given that both lanes have equal priority.

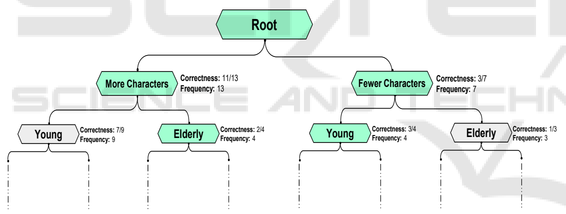

Figure 6: Correctness and frequency logs.

The main idea behind the accommodating technique is to record a history log, which is in charge of two main things:

- **Correctness Log.** The correctness log stores information on comparable instances in the same street or neighbourhood and the ratio at which choosing this leaf or node was correct.
- **Frequency Log.** The Frequency log records the number of times that leaf or node was chosen, allowing for a balance in the frequency at which each node is chose, given that both nodes have equal priority.

In the event of a conflict if we have prior knowledge about the accident location or the conflicting decision tree nodes. The accommodating technique examines the node's history for each processing lane, as shown in 6. Starting with the Correctness ratio, if a node's correctness ratio is low, it allows for the other lane to be selected. Similarly, if a node's frequency is higher, indicating that it has been selected more frequently, it allows for the other lane to be selected.

Verifying Nodes: is a method inspired from (Bazan et al., 2016) establishing a procedure for generating trees with verifications of cuts defined on numerous attributes. The distribution of objects based on the optimal cut is confirmed in each node of a tree formed using this method by subsequent cuts on different attributes that separate objects in a similar fashion.

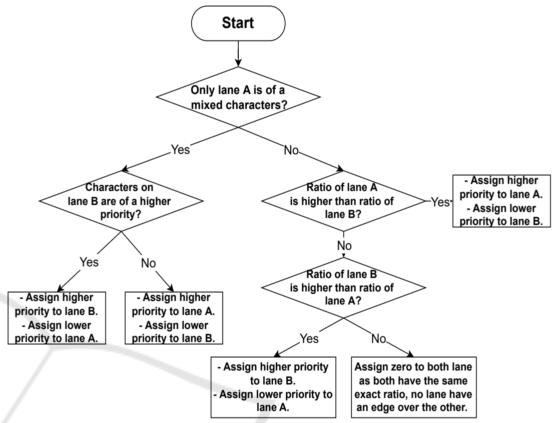

Figure 7: Verifying node flowchart.

The verifying node technique resolve conflicting nodes where we have a mix of characters standing within the same lane preventing us from branching as illustrated earlier in Figure 3.

As shown in Figure 7, we start by comparing one of the two lanes, let's call it lane A, to the other, lane B. If only one of the two lanes contains a mix of characters, we check to see if lane B contains characters with a higher priority; if it does, lane B is assigned the higher priority.

In the event that both lanes contained a mix of characters, the lane with the higher ratio of high-priority characters is assigned a higher priority.

For the two other conflicting nodes scenarios mentioned earlier: **comparable branching** and having **attributes with no value**, as the values for both lanes are equally identically or unknown, we skip that level of the decision tree and assign zero score for both lanes at that level, preventing it from affecting the accuracy and final priority scores of either lane.

3.3 Decision Making

The process of decision making depends on two main factors:

- **Time Remaining Before Collision.** The time remaining determines whether we will be able to traverse the whole tree, and whether we will be able to run all the conflict resolution techniques presented within the remaining time.

- **The Complexity of the Scenario.** Complexity of the scenario, number of characters, lighting, and other factors determine the time needed to process this data and whether missing data will cause node conflicts that must be resolved, taking more time.

Having one conflict resolution technique resolve the issue is sufficient to terminate and move forward with this decision. Different resolution techniques may be suitable in different situations, as some are faster and more certain than others. In the next section, we'll examine all the different techniques and determine which one excels in which aspects.

4 RESULTS & DISCUSSION

To ensure that the presented conflict resolution technique will be effective in a real-world scenario, and to analyze and evaluate each of the techniques. This section presents the method used to evaluate this project and the results we obtained.

Carla simulation was used combined with ScenarioRunner To simulate a real-world scenario as closely as possible, generating conflicting scenarios in which the proposed resolution techniques can be evaluated.

4.1 Evaluation Technique

Three different metrics were used to evaluate the efficiency of our conflict resolution techniques, that take into account a range of relevant factors:

- **Decidability.** In this analysis, we adopt the decidability index d (Williams, 1996), as a measurement index of whether the technique is decidable and would eventually result in a resolution.
- **Correctness.** Correctness is the accuracy of a decision, and verifying a technique's correctness can be done by comparing its outcome to the top nine priorities in Figure 2, and calculating a ratio of the number of rules met against these top nine priorities, which was adopted based on a related line of research includes robust detection (Huber, 1965).
- **Time.** The time to reaction has been considered in the literature (Wagner et al., 2018; Tamke et al., 2011; Junietz et al., 2018) for which the time needed by each technique to resolve a conflict would be measured in a variety of complex scenarios, allowing us to determine whether a particular technique would be able to meet the deadline and, if not, by how much it missed the deadline.

4.2 Testing Scenarios

As shown in Figure 8, all scenarios were ran three different times at three different distances to simulate a shorter time remaining for reaction and conflict resolution at each point. In each scenario, a random set of characters were present on lanes A and B, where the car was faced with a conflicting situation and must react within the time available.

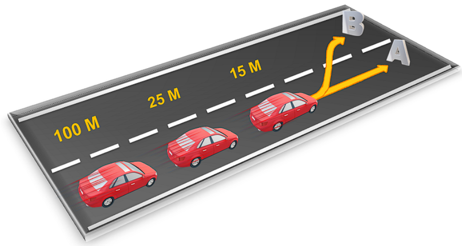

Figure 8: Evaluation Technique.

To ensure accurate results, a random set of 80 different conflicting scenarios were generated. Having fifty scenarios represent the case where the entire decision tree could be traversed, while the remaining thirty scenarios represent the case where there is no time to traverse the entire decision tree before the deadline. The results of the two cases are deeply analyzed in the following section.

4.3 Results

Figure 9 show an overall analysis after running 50 different scenarios that succeeded in **traversing the whole decision tree**.

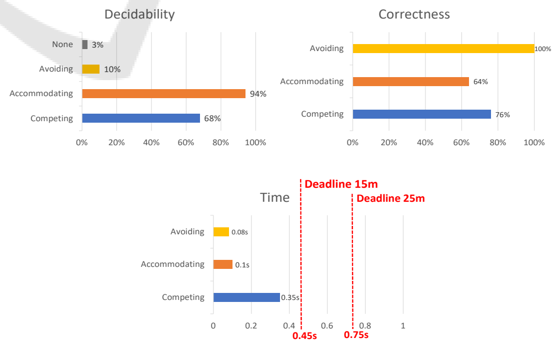

Figure 9: Results of traversing the whole tree.

Where each of the graphs represented in Figure 9 is discussed in more details below:

- **Decidability.** In 97% of the cases, we reached a resolution, with the Accommodating technique being the most successful, reaching a resolution in 94% of the scenarios, followed by the Competing technique, with a resolution of 68%, and finally the Avoiding technique, which only succeeded in reaching a resolution in 10% of the cases.

- **Correctness.** Although the Accommodating technique was the most effective in resolving most scenarios, it was the least accurate, achieving only 64% accuracy, followed by the Competing technique with 76% accuracy, and finally the Avoiding technique, which was the most accurate with 100% accuracy in an event of a conflict.
- **Time.** The Avoiding and Accommodating techniques were the fastest in just under 0.1s, with the Avoiding technique achieving faster averages by 0.08s. The competing technique was the slowest with 0.35s, yet it managed to finish before the deadline of the shortest distance (0.45s).

4.3.1 Deadline met Analysis

Figure 10 show an overall analysis after running 30 different scenarios that were terminated before traversing the whole decision tree.

Where each of the graphs represented in Figure 10 is discussed in more details below:

- **Decidability.** The results this time were rather interesting, with a resolution being reached in 100% of the cases. The Competing technique was the most successful, resolving 96% of the scenarios, followed by the Accommodating technique, with a resolution of 92%, and finally the Avoiding technique, with a resolution of only 12%.
- **Correctness.** The Avoiding technique continues to be the most accurate with a highest score of 100%, followed by the Competing technique, which achieved 78% accuracy, and then the Accommodating technique, which was once again the least accurate with 67%.
- **Time.** The Avoiding and Accommodating techniques still achieved nearly identical averages, with the Avoiding technique remaining faster and averaging 0.08s, while the competing technique became even slower with 0.48s, missing the deadline for the shortest distance (0.45s) and therefore failing to resolve the conflict in the third and closest run at 15m.

5 DISCUSSION

The purpose of this research was to reach a resolution in the first place in the event of a conflict. The results in section 4.3 show that in either case, whether the

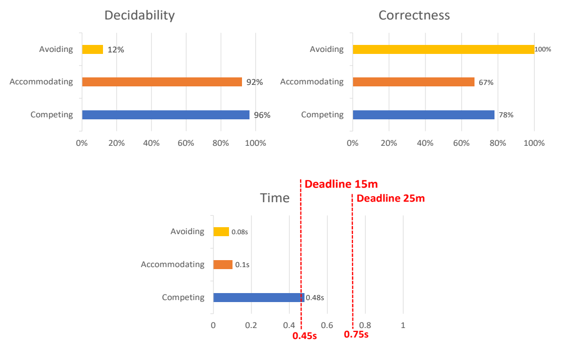

Figure 10: Deadline met before traversing whole tree.

algorithm managed to fully traverse the decision tree or not, it managed to resolve the conflict in 97% of all the scenarios generated.

Leaving 3% of scenarios in which the scenarios were near identical and there were basically no differences between the two lanes; at this point, choosing one or the other randomly would not make a difference, given that the analyses conducted were unable to distinguish between them.

5.1 Comparing Different Techniques

The Competing technique provides the best result of all techniques, reaching a resolution in most cases while excelling in situations where the algorithm couldn't traverse the entire decision tree. Having a decidability of 96% and an accuracy of greater than 75%, covers situations where there was insufficient time to traverse the entire tree. The main drawback is that it is the slowest, as it takes longer to check all the leaves of the decision tree to determine which lane could have earned the maximum potential score.

The Avoiding technique provided the most accurate result, with an accuracy of 100% in both cases, As if the decision tree traversal had not been terminated, the decision tree scores would not have been conflicting in the first place, with one of the lanes having a higher priority than the other. Furthermore, It's also the fastest among them all. The main downside is that, only 11% of scenarios are decidable on average.

The Accommodating technique had the highest decidability in both cases, at 93%, ensuring a resolution if the others failed. However, it had the lowest correctness, with an average of 66%; this percentage would increase as the self-driving car encountered more roads and scenarios, as it is the only technique that relies on the history log. Therefore, the more data that is available, the more precise the results will be.

In conclusion, the three techniques excel in distinct aspects. Therefore, integrating the results of all three techniques to enable the self-driving car to make a decision would yield the best possible results.

6 CONCLUSIONS & FUTURE WORK

The aim of the paper was to resolve conflicts that could arise in a self-driving car in the event of unavoidable accidents. A conflict resolution technique was implemented using the Thomas-Kilmann Conflict Model, along with the Verifying Nodes technique to resolve the conflicts that may arise. Both techniques were applied over a decision tree that was constructed based on the MIT Moral Machine Experiment results (Awad and Dsouza, 2018). We managed to create a system where we tested 80 different conflicting scenarios achieving a decidability of 97% in the scenarios, and an average accuracy of 80%. Overall, we managed to resolve 97% of the conflicting scenarios generated, in an average of just 0.2 seconds, leaving out only 3% where the scenarios were fairly identical that randomly choosing one over the other would have made no difference.

Future research can go in several directions. First, the presented testing scenarios can be enriched by gathered more data about the environment using image processing techniques. The behaviour of the proposed conflict-resolution model can then be further verified. It is also worth investigating the combination of the three different resolution techniques which could potentially result in faster and even more accurate results. Extending the work presented in this paper beyond the single car case to resolve conflicts among a swarm of cars is also a natural next step.

REFERENCES

Anbar, D. (1983). On estimating the difference between two probabilities, with special reference to clinical trials. *Biometrics*, pages 257–262.

Awad, E. and Dsouza (2018). The moral machine experiment. *Nature*, 563(7729):59–64.

Barabas, I., Todoruţ, A., Cordoş, N., and Molea, A. (2017). Current challenges in autonomous driving. In *IOP conference series: materials science and engineering*, volume 252, page 012096. IOP Publishing.

Bazan, J. G., Bazan-Socha, S., Buregwa-Czuma, S., Dydo, L., Rzasa, W., and Skowron, A. (2016). A classifier based on a decision tree with verifying cuts. *Fundamenta Informaticae*, 143(1-2):1–18.

Borenstein, J., Herkert, J., and Miller, K. (2017). Self-driving cars: Ethical responsibilities of design engineers. *IEEE Technology and Society Magazine*, 36(2):67–75.

Cowger Jr, A. R. (2018). Liability considerations when autonomous vehicles choose the accident victim. *J. High Tech. L.*, 19:1.

Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., and Koltun, V. (2017). Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR.

Fahy, A. S., Mueller, C., and Fecteau, A. (2021). Conflict resolution and negotiation in pediatric surgery. In *Seminars in Pediatric Surgery*, volume 30, page 151100. Elsevier.

Huber, P. J. (1965). A robust version of the probability ratio test. *The Annals of Mathematical Statistics*, pages 1753–1758.

Junietz, P., Bonakdar, F., Klamann, B., and Winner, H. (2018). Criticality metric for the safety validation of automated driving using model predictive trajectory optimization. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 60–65. IEEE.

Karnouskos, S. (2018). Self-driving car acceptance and the role of ethics. *IEEE Transactions on Engineering Management*, 67(2):252–265.

Lin, P. (2016). Why ethics matters for autonomous cars. In *Autonomous driving*, pages 69–85. Springer, Berlin, Heidelberg.

Safavian, S. R. and Landgrebe, D. (1991). A survey of decision tree classifier methodology. *IEEE transactions on systems, man, and cybernetics*, 21(3):660–674.

Schaubhut, N. A. (2007). Thomas-kilmann conflict mode instrument. *CPP Research Department*.

Tamke, A., Dang, T., and Breuel, G. (2011). A flexible method for criticality assessment in driver assistance systems. In *2011 IEEE Intelligent Vehicles Symposium (IV)*, pages 697–702. IEEE.

Wagner, S., Groh, K., Kuhbeck, T., Dorfel, M., and Knoll, A. (2018). Using time-to-react based on naturalistic traffic object behavior for scenario-based risk assessment of automated driving. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1521–1528. IEEE.

Williams, G. O. (1996). The use of d'as a decidability index. In *1996 30th Annual International Carnahan Conference on Security Technology*, pages 65–71. IEEE.