Advice Provision in Teleoperation of Autonomous Vehicles

YOHAI TRABELSI , Department of Computer Science, Bar-Ilan University, Israel
OR SHABAT , Department of Computer Science, Bar-Ilan University, Israel
JOEL LANIR , Department of Information Systems, University of Haifa, Israel
OLEG MAKSIMOV , Department of Computer Science, Bar-Ilan University, Israel
SARIT KRAUS , Department of Computer Science, Bar-Ilan University, Israel

Teleoperation of autonomous vehicles has been gaining a lot of attention recently and is expected to play an important role in helping autonomous vehicles handle difficult situations which they cannot handle on their own. In such cases, a remote driver located in a teleoperation center can remotely drive the vehicle until the situation is resolved. However, teledriving is a challenging task and requires many cognitive resources from the teleoperator. Our goal is to assist the remote driver in some complex situations by giving the driver appropriate advice. The advice is displayed on the driver's screen to help her make the right decision. To this end, we introduce the TeleOperator Advisor (TOA), an adaptive agent that provides assisting advice to a remote driver. We evaluate the TOA in a simulation-based setting in two scenarios: overtaking a slow vehicle and passing through a traffic light. Results indicate that our advice helps to reduce the cognitive load of the remote driver and improve driving performance.

CCS Concepts: • **Human-centered computing → Human computer interaction (HCI)**.

Additional Key Words and Phrases: Teleoperation, Autonomous vehicles, Remote driving, Teleoperation challenges

ACM Reference Format:
Yohai Trabelsi, Or Shabat, Joel Lanir, Oleg Maksimov, and Sarit Kraus . 2023. Advice Provision in Teleoperation of Autonomous Vehicles. In *28th International Conference on Intelligent User Interfaces (IUI '23), March 27–31, 2023, Sydney, NSW, Australia.* ACM, New York, NY, USA, 20 pages. https://doi.org/10.1145/3581641.3584068

1 INTRODUCTION

Autonomous vehicle (AV) technologies are improving dramatically over time (see, e.g., Li et al. [38]). Many companies, such as Waymo [63] and Mobileye [40], are developing cutting-edge technologies to improve AV capabilities. However, much remains to be done to make autonomous cars safe and efficient and to overcome regulatory hurdles. Many researchers, as well as industry leaders, believe that AVs will not be able to handle all driving situations, at least in the near future (see for example [9, 13, 26, 29, 34, 42, 53]). Some experts, such as former Weimo CEO John Krafcik, believe autonomous cars will never be able to drive in "all conditions" [60]. One solution already in use in industry [28, 42] is teleoperation, in which a human assistant monitors and operates the car remotely. The common model is that when an AV faces a situation it cannot handle, it sends a request for intervention to a teleoperation center, where a human operator is assigned to remotely drive the vehicle until the problematic situation is resolved.

As described by Tener and Lanir [59], AV teleoperation presents several major challenges. For example, the operator cannot feel the forces acting on her or hear the ambient sounds from the vehicle's environment. Another major challenge is cognitive load. The remote driver needs to quickly take control of the autonomous vehicle, assess its situation, and drive the vehicle remotely, taking into account the remote environment and the vehicle's responses to resolve the problem quickly. This taxing task requires many cognitive resources [59]. Such a heavy cognitive load can increase the likelihood of a collision, as described by Pawar and Velaga [45].

In this research study, we develop the TeleOperation Advisor(TOA), an assisting agent to help teleoperators with decision-making for challenging tasks. The idea is that an assisting agent, suggesting driving actions, can help reduce cognitive load as well as improve driving performance. In the literature, there are several types of systems that assist drivers in decision-making. Some of these systems are designed to warn the driver of approaching hazards, such as collision avoidance systems (see, e.g., Wang et al., [61]), or warn drivers when they intend to take an action that may result in a hazard, such as changing lanes (see, e.g., Sun et al., [58]). Another type of system is one that suggests actions to the driver, such as an overtaking attempt [46]. Although the former type is more common in driving, it is expected that the importance of suggesting actions to drivers will be more significant in teleoperation because of the difficulties mentioned above.

TOA repeatedly tries to identify the best action to be taken by the teleoperator in a given situation, taking into account the teleoperator's capabilities aiming at satisfying a set of ranked goals.

We have selected two scenarios with challenging tasks where a decision must be made and implemented a version of TOA that assist the teleoperators in these scenarios. The first scenario is the overtaking of a slow-moving vehicle on a highway. This scenario requires complex state awareness since the positions and speeds of nearby vehicles have a crucial influence on the decision of whether to overtake or not. The other scenario is going through a traffic light before turning red. The decision to be made here is whether it is better to stop at the traffic light or run it. Note that these challenging tasks may occur at any time during remote driving and not necessarily immediately when the driver takes control.

To evaluate the usefulness of our method, we conducted a user study, having human participants act as remote operators using the CARLA simulation [16]. The CARLA environment is an open-source driving simulation that has been used extensively in research on autonomous vehicles (see, e.g., [35]). Our results indicate that adding driving recommendations given by TOA successfully managed to reduce the cognitive load of the remote drivers. We have also shown that TOA significantly reduced the number of red light crossings in the traffic light scenario and reduced the use of brakes and improved the measured efficiency in the passing scenario. Finally, the participants indicated that they were satisfied with TOA and would want to use such an agent when remotely driving a vehicle.

Our work makes the following contributions:

(1) A general method for providing automatic advice in teleoperation when a human operator controls the vehicles remotely.
(2) The implementation of this method in two different scenarios: overtaking on a highway and crossing a traffic light.
(3) An evaluation of the method in an experiment with human participants.

2 RELATED WORK

2.1 Advising Agents

Automated agents' advice could be very beneficial to people when they need to make decisions in complex settings. There are many applications in which such advice is needed (see [6], [4],[50] for some examples). Supplying advice involves multiple challenges. The main challenge is to predict people’s behavior and decisions. See for example Rosenfeld et al., [51], Kraus [36] and Rosenfeld and Kraus [52]. However, a major difficulty is that many factors can influence human decision-making. Dietrich [14] enumerated past experience (Juliusson et al., [33]), cognitive biases (Stanovich and West, [56]), age and individual experiences (Bruine de Bruin et al., [10]), belief in personal relevance (Acevedo and Krueger, [1]) and an escalation of commitment as relevant factors for making a decision. Another challenge arises in settings that involve repeated interactions between an agent and a human who uses the agent’s advice. In such cases, the agent needs to consider how their actions in the present influence people’s future actions (Azaria et al [5]). Giving advice that will be accepted by people is also a challenge. Elmalech et al. [18] showed that giving intuitive sub-optimal advice is more beneficial than giving less intuitive optimal advice. Many Advanced Driver Assistance Systems (ADAS) warn drivers of certain dangerous events. Sun et al. [58]) developed a personalized system that warns drivers of dangerous lane changes. Wang et al. [61] developed a personalized system that warns drivers when a potential collision is expected. However, these works focus on warning drivers of various hazards. In contrast, our system focuses on helping drivers make decisions.

2.2 Teleoperation of Vehicles

Teleoperation is currently used in various types of ground vehicles. A famous example is teleoperated rovers, in particular those traveling on Mars (see, e.g., [22]). Another current application is mining trucks [55]. However, these sorts of applications are very different from driving in everyday’s traffic [43]. Few works show how teleoperation should work in everyday’s traffic (see, e.g., [25]). Several works try to overcome some of the various challenges of AV teleoperation, such as communication delays [27], situation awareness [24], and forces acting on the driver, e.g., giving the driver an artificial steering wheel feel [57]. However, these do not deal with driving in specific difficult tasks nor with the frequent context switch which is required in our settings. Mukhopadhyay et al [41] investigated the effect of Extended Reality (XR) ADAS systems on Take Over response and task completion times and found that their approach (using "augmented reality") resulted in significantly faster completion times. The goal of their approach was to help vehicles follow ego lanes which could be a difficult task in India, but it is different from higher-level decision-making processes, as in our work.

2.3 AV overtaking a slow vehicle

Although there are relatively few collisions caused by lane-changing maneuvers (4%), and although fatalities from these collisions account for only about 0.5% of total fatalities, lane-changing behavior is responsible for about 10% of traffic congestion and frequently causes traffic jams. This has implications for the economy [12, 32]. In addition, failure to find an appropriate time window to start an overtaking attempt can lead to inefficient driving and can also delay other vehicles.

Making the decision as to when it is appropriate to overtake is a challenge that is relevant for fully autonomous vehicles [44]. There are several methods that have been developed over the years to make an overtaking decision. The first is a rule-based approach. Its main advantage is that it is easy to implement [7, 20]. Another method is a utility-based approach. It is better suited for complex scenarios, but requires more effort to fine-tune the parameters [23, 64]. A more recent method, used primarily for fully autonomous vehicles, is reinforcement learning (see, e.g., [11, 39, 65]). As suggested by Armağan and Kumbasar, Fuzzy logic can also be used to decide when to start an overtaking attempt [3]. Our method is based on a combination of a set of rules and a set of utility functions. The easier-to-implement rules are used whenever possible, while the utility functions are used at the core of our approach to better handle the more complex decisions. Reinforcement learning is not used in our settings because it requires larger data sets.

2.4 Passing through a traffic light

Running red lights is considered one of the most important factors leading to collisions at signalized intersections [67]. A rule-based system that suggests to drivers whether to stop or proceed before a traffic light has been proposed by Bar-Gera et al. [8]. In their work, they showed that their advisory system actually significantly reduced the number of red light violations. However, they knew in advance what the traffic light would be in the future, while we do not, which may have helped them improve their advice. For example, if they knew that the traffic light would turn red in one second, they could suggest to the driver to stop, while we did not have this information and therefore need a more advanced method to give appropriate advice.

An important difference between their setting and ours is that in their setting, participants had to drive a long distance in a single car, so they did not have to switch between different environments. In our setting, on the other hand, they had to take over the car in different environments, drive different cars in different cities, and that is a much more difficult task.

3 THE TELEOPERATOR ADVISOR (TOA)

In this section, we describe a method for developing a TeleOperator Advisory agent (TOA) for teledriving in complex driving scenarios. Our method runs iteratively, generating appropriate recommendations for the teleoperator at each small time interval. The time units in which the iterations occur are called time ticks. We first describe the problem of providing advice in a single time tick, and then describe how the whole system works.

A problem π is defined by a teleoperator τ, a current time t, an ego-vehicle ϵ, and a set of objects Ω. The set of objects consists of, for example, vehicles, pedestrians and traffic signs. The ego-vehicle and the objects in Ω have properties relevant to the problem. The properties of the ego-vehicle for problem π are denoted as $B_{\pi,\epsilon}$ and the properties of each object $\omega \in \Omega$ are denoted as $B_{\pi,\omega}$. Possible properties of the ego-vehicle are its speed, acceleration, braking and acceleration capacity. Possible properties of the other objects are their type (e.g., car, motorcycle, truck, stop sign), their speed (which can be positive, zero, or negative), their distance from the ego-vehicle, a direction with respect to the ego-vehicle (e.g., front, back, left), and their order in the corresponding direction (e.g., the third car from the front).

Moreover, a problem π has a set of possible actions $A = \{\alpha_1, ..., \alpha_n\}$ that the teleoperator τ can perform on the vehicle ϵ such as accelerating, stopping, or starting an overtaking maneuver. The problem π is associated with a set of goals $\Xi = \{\xi_1, ..., \xi_n\}$ such as avoiding collisions or driving efficiently. Each goal ξ is associated with a weight $w_\xi \in [0, 1]$ that determines the importance of that goal. For example, avoiding collisions receives the highest possible value of 1 while driving efficiently may have a lower value of 0.7. We also define a personal weighting parameter, $\rho_\tau \in [0, 1]$, which represents the teleoperator's personal reaction and performance time. This weighting parameter is initially generated by running the *PVT* test [2, 15], a general, widely used, test that measures individual reaction time to visual stimuli. It is updated after each completed action of the teleoperator (see subsection 3.5 for more details). A parameter value of 0 represents a teleoperator that responds very quickly, and a value of 1 represents a very slow teleoperator.

Also, for each problem π, an action α, a goal ξ and the teleoperator's personal weight parameter ρ_τ we define a function $\mu(\pi, \alpha, \xi, \rho_\tau)$ that computes a score for achieving the goal ξ given the problem π and an action α. The goal of TOA is to optimize the function $argmax_{\alpha \in A_p} \sum_\xi w_\xi * \mu(\pi, \alpha, \xi, \rho_\tau)$.

The properties of the ego and the other vehicles as well as the relevant goals, their weighting, and the corresponding actions are initiated by an experienced human driver and verified by driving in the simulation.

3.1 Limiting the set of relevant objects and actions

Computing $\mu(\cdot)$ efficiency is necessary since it is needed to run this function repeatedly in real-time. Therefore, we propose to reduce the number of objects considered by the function. For example, a vehicle that is far from the ego-vehicle is irrelevant. For that, we define a set of rules created by a human expert. This function depends heavily on the problem. For example, for the traffic light scenario, it may be sufficient to consider only one car in front of the ego-vehicle and one car behind the ego-vehicle and ignore the other cars.

Additionally, there is a need to limit the set of possible actions. For example, we cannot switch between different recommended actions for each time tick, as this can cause a lot of confusion. Another reason is to avoid actions that are technically impossible. For example, if the ego-vehicle is in the left lane, which is the faster lane, it is impossible to try to overtake and move to a better lane. In addition, an action may not be applicable if the estimated time is not acceptable (see subsection 3.2 for details on calculating estimated time). For example, if there is not enough time to stop before a traffic light, a recommendation to stop is not appropriate. Finally, an action may not be applicable if the reason for taking the action is not yet relevant. For example, an action that starts an overtaking attempt may not be appropriate before the ego-vehicle approaches a slow vehicle.

We define w as the number of historical time ticks we are considering and A_w as the queue of actions recommended in the last w time ticks. The function $\lambda(\alpha, \pi, A_w) = \{True|False\}$ is given an action α, a problem π and the queue A_w and returns True if and only if α is applicable to the problem π.

3.2 Assigning a score for achieving a goal

In order to estimate the score of achieving the various goals, an estimation of the time required for τ to perform each action is required. This is computed by using the properties of the problem together with ρ_τ. The expected time for an action α by an operator τ can be measured by performing the various actions multiple times (e.g., 10 times per action), calculating an average, and multiplying the average by ρ_τ. If more data is available, using machine learning methods could be beneficial.

The function $\mu(\pi, \alpha, \xi, \rho_\tau)$ is based on a set of rules established by a human expert. The rules vary for different problems and goals. In general, the time interval from the current time until the estimated end of the action is considered, and the estimated locations of other relevant vehicles are taken into account. For example, in the overtaking scenario, a low collision score is assigned when a vehicle is expected to be very close to the ego vehicle. On the other hand, this case may result in a high-efficiency score because the task is expected to be completed faster. A low score for the goal of avoiding emergency braking may also be assigned in this case. Another example related to the traffic light scenario is the score for the action of stopping at a red light when the intersection is empty, which may increase the probability of a collision if there is a vehicle behind the ego-vehicle that is very closed to it and moving too fast. If the ego-vehicle stop, it may bumped into it. However, there may be a high score for the goal of obeying the traffic light.

If the scores are generated by a human expert, we allow only a few (e.g., 3) possible discrete scores per goal. In addition, an experienced human driver should drive in a simulation after the initial scores are created and later be asked for his or her opinion on the balance between the different goals. Based on this information, we should adjust the scoring rules accordingly. If enough data could be collected, the score could be estimated using machine-learning methods which will yield a more accurate estimation of the value of a recommended action.

3.3 The advising flow

The main loop of algorithm 1 iterates over the time ticks. For each time tick t in which there is at least one action for which $\lambda(\cdot)$ gives a positive value, these actions are considered and a weighted score is calculated for each. The action with the highest score is then recommended to the teleoperator. Finally, when an action has been performed, the personal weighting parameter is updated according to the teleoperator's performance.

Algorithm 1 The procedure of TOA

$A_w \leftarrow Queue()$
$A_w.initialize()$
for all $t \in$ Time ticks **do**
  $A_p \leftarrow \emptyset$
  for all $\alpha \in A$ **do**
    if $\lambda(\alpha, \pi, A_w) = True$ **then**
      $A_p = A_p \cup \{\alpha\}$
    end if
  end for
  if $A_p = \emptyset$ **then**
    continue
  end if
  $\alpha = argmax_{\alpha \in A_p} \sum_{\xi} w_{\xi} * \mu(\pi, \alpha, \xi, \rho_{\tau})$
  Recommend α
  $A_w.enqueue(\alpha)$
  if $A_w.size() > 10$ **then**
    $A_w.dequeue()$
  end if
  if α was performed **then**
    update ρ_{τ}
  end if
end for

3.4 Computing and updating the personalized weight parameter

The calculation of a personalized weighting parameter is necessary for the prediction of the expected time for the execution of an action. This, in turn, is necessary for predicting how much time the action is expected to take. Knowing the expected time for an action, one can estimate when it is reasonable to perform an action. The initial parameter was determined by the expected reaction time, which was measured for each user individually using the short online version of the PVT test [21]. At each iteration, the parameter is updated using the exponential moving average update method introduced by Roberts [49]. The parameters chosen are designed to give significant weight to history relative to the current sample.

3.5 The advice interface

The recommendations that are generated for the overtaking scenario are represented using visual overlays on the surroundings. In particular, we use a green arrow for lane change recommendation (see Figure 2) and a red triangle for slowing down or braking scenarios (see Figure 4). The visual overlays on the surroundings and the specific notations were studied by Eriksson et al. [19] for supporting a human driver who is located inside an autonomous car and needs

Scenario (type of π)	relevant objects
Overtaking in a highway	front, rear left, front left
Traffic Light	front, back, traffic light

Table 1. Relevant objects per scenario, in each direction only the closest vehicle is considered.

Object (ω)	Relevant Parameters $B_{\pi,\omega}$
vehicles	type, direction, distance and speed
pedestrians	direction, distance and speed
road signs	type, distance and position
traffic signals	signals color, distance and position

Table 2. Relevant parameters per object for problems where the object is relevant (for all problems π, including overtaking and traffic light).

to take control over the car. Eriksson et al. compared this approach to other methods and notations and showed that in overtaking scenarios the proposed approach effectively improved the number of correct decisions to brake or overtake in an overtaking scenario.

4 IMPLEMENTATION OF THE TELEOPERATOR ADVISOR IN SPECIFIC SCENARIOS

In this section, we show how the above methodology is implemented in two scenarios using CARLA simulation. The first scenario is overtaking a vehicle on a highway and the second is running a traffic light. In both scenarios, we decided on 30 ms between time ticks through trial and error. This decision was made by trying multiple gaps - first at 10 ms and then in jumps of 10. For each gap, we ran the advisory system in multiple scenarios. The experts analyzed the videos and decided whether the quality of the advice was appropriate. To save computational resources, we chose the largest value that resulted in appropriate advice.

The source code of our implementation, together with some example videos are available in a public repository [1].

4.1 Overtaking in a Highway

First, we defined the relevant properties of the ego vehicle, $B_{\phi,\epsilon}$, to be the velocity, acceleration, and the vehicle's braking and speed capabilities. In our case, they were obtained from the CARLA simulation. We then decided which objects in the environments were relevant (see Table 1), and what parameters were relevant to them (see Table 2).

Then the set of possible actions A was set to "slow down", "start an overtake" and "no action". As mentioned above, for each operator o, ρ_o was initiated by the PVT test results. The selected goals for the overtaking scenario are collision avoidance (0.6), avoidance of a dangerous distance to other vehicles or obstacles (0.1), avoidance of emergency braking (0.1), and avoidance of excessive cognitive load (0.2). The weights of the goals are given in parentheses.

In realtime, for each time tick in the "for" loop in Algo 1[Line 3], the distances to the leading slow vehicle and to the other vehicles nearby are calculated, as defined in Table 1. These distances are first used to determine whether A_p (Algo 1[Lines 6-7]) includes an overtaking action. Overtaking is not possible in situations where the ego-vehicle is in the leftmost lane, cases where a recommendation to slow down was made in the last 0.3 seconds, and cases where the distance to the lead vehicle is more than 12 seconds.

[1] https://github.com/yohayt/Advice_Provision_in_Teleoperation_of_Autonomous_Vehicles

Next, we defined a set of rules that, if met, we know with certainty that the start of an overtaking maneuver leads to the highest-scoring action (Algo 1[Line 13]). We used these rules rather than computing the expected benefit directly because lack of data. The first rule states that an overtaking recommendation will not be executed if the ego vehicle is too close to the leading vehicle (time to collision less than 2 seconds), since the safety of such an overtaking operation is problematic. The second rule states that an overtaking attempt will not be started if the speed in the current lane is higher than the speed in the left lane. The next rules refer to the left lane. These rules determine whether the left lane is free for an overtaking attempt. This is done by estimating the expected time for the overtaking maneuver and, assuming a constant speed, checking whether the next vehicle on the left behind and the next vehicle in front of the vehicle in the left lane are expected to be too close to the ego-vehicle during the overtaking maneuver (less than 0.25 second). The estimation of the overtaking performance time is based on a multiplicative combination of the personal weight parameter (ρ_o) and a constant time that is needed to perform an overtaking attempt based on our dataset.

If all the rules are met, a recommendation to overtake is given. The actual result of the advice system is a green arrow to the left as seen in Figure 2.

The slow-down recommendation is simpler. It is recommended if and only if the time to collision with the leading vehicle is less than 2 seconds and the speed of the ego-vehicle is at least 1.2 meters per second. If both the "overtake" and the "slow down" recommendations were not given, then no recommendation is given. Finally, when an overtake has been performed and has lasted $t$seconds, the personal weight parameter ρ_o is updated according to the formula: $\rho = 7/8 * \rho + 1/8 * t$. The formula and its parameters (7/8 and 1/8) were taken from the TCP Round Trip Time (RTT) estimation procedure and typical parameters (See Kurose and Ross [37] for more details). The flow of choosing an action is available in Figure 1.

4.2 Passing through a traffic light

First, the relevant properties of the ego vehicle were defined similar to the overtaking scenario. Similarly, the relevant objects and what parameters were relevant to them are defined in Tables 1 and 2. The possible actions A was set to "stop" and "no action". As mentioned above, for each operator o, ρ_o was initiated by the PVT test. The selected goals for the traffic light scenario (and their weights) are collision avoidance (0.4), avoidance of a dangerous distance to other vehicles or obstacles (0.1), avoidance of emergency braking (0.15), avoidance of excessive cognitive load (0.1) and avoidance of red light violations(0.25). The weights of the goals are given in parentheses.

In real-time, in the "for" loop in Algo 1[Line 3], the distances to the traffic light and to the vehicles nearby are calculated for each time tick, as defined in Table 1. Then, situations in which stopping is not relevant are not included in A_p (Algo 1[Lines 6-7]). These included situations where the ego-vehicle is too close to a yellow traffic light and cases where the distance to the traffic light is more than 60 meters.

Next, we defined a set of rules that, if at least one of them is met, we know with certainty that stopping leads to the highest-scoring action. The rules and their values were inspired by Rittger et al [48] and the values were finalized by some trial and error. The rules and the whole flow are presented in Figure 3. The actual outcome of the advice system is a stopping sign in the front, see Figure 4.

5 A USER STUDY

We assigned participants to play the role of remote operators in driving scenarios. Our goal was to show that the TOA and its advice are helpful in two scenarios: overtaking and approaching traffic lights.

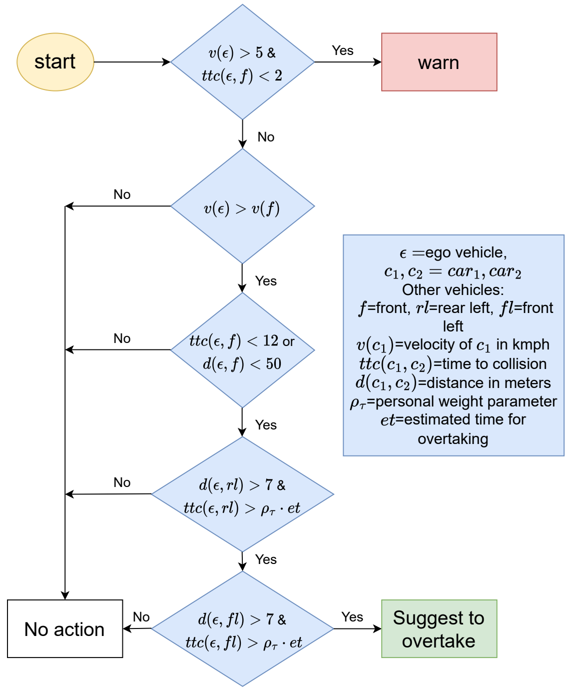

Fig. 1. Rules for choosing an action in overtaking in a highway.

Fig. 2. An example for advising to start an overtake attempt in CARLA simulation.

5.1 Participants

Participants were invited to a university laboratory to take part in the experiment. Most participants were Computer Science students who were encouraged to participate in order to receive a small bonus on their university course grades. Participants had the option of participating in another experiment if they felt uncomfortable with this one and receiving the same bonus. Additional participants were recruited through a Facebook group of students from other disciplines. These participants were paid an amount of $15 for their participation.

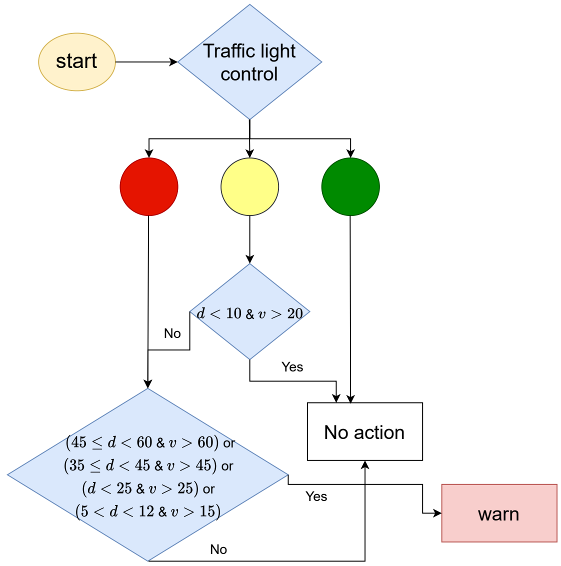

Fig. 3. Rules for choosing an action in passing through a traffic light. Speeds (v) are in kilometers per hour and distances (d) are in meters.

Fig. 4. An example for advice to stop before a red light in CARLA simulation.

There were a total of 22 participants, of which 7 were females and 15 were males. Their mean age was 27 with a standard deviation of 4.2. Most of the participants (19) were in their third year of study or were masters students in the Department of Computer Science and had good computer skills. 18 of the participants had at least 4 years of driving experience, 2 had $2-4$ years, 1 had $1-2$ years, and 1 had less than one year. 17 of the participants drive daily, 2 drive at least once a week, and 2 drive less frequently.

5.2 Experimental Settings

We performed the experiment using a university laboratory computer. For the simulation, we used the CARLA environment [16]. The CARLA simulation includes a few city maps and the environment consists of buildings, traffic signals, traffic lights, intersections, moving cars, stationary cars, and pedestrians. In our experiment, we used two of the maps: a city with a highway for the passing scenario and another city with many traffic lights for the traffic light scenario. As in some previous works (e.g., [62] and [17]), driving is done with the help of the arrow keys. The right and left arrow keys are used to turn right and left, up to accelerate, and down to slow down. The 1 and 2 keys are used to switch between forward and reverse. For the ego-vehicle, we used 3 cameras. One for the front view and two side rear cameras. There was also a navigation map that was used to navigate to the desired location.

The simulation consists of short driving tasks with a specific location to which the participant must drive a vehicle. Participants were instructed to drive the vehicle safely and efficiently to a location marked on the map. Each task begins with a few seconds of an autonomous driver controlling the vehicle. After some time, a light in the upper left corner of the screen changes from green to red. From then on, the driver can take control of the vehicle by pressing the 1 key. In this way, the teleoperator can observe the environment of the ego-vehicle before taking control.

When the vehicle arrives at the destination, the participants are instructed to return the control to the autonomous agent by pressing a dedicated key. All simulation tasks were performed in good weather conditions and in daylight.

A supervisor was present throughout the experiment. The supervisor's role was to give the participants some explanations about the simulation and the experiment, to check that the participants were concentrating on performing the experiment, and to decide when it was appropriate to stop the training tasks and start the main tasks of the experiment (more details below).

5.2.1 Training Tasks. The first training session took place in a city environment with no advice given. The goal of this session was to practice staying in the lane, turning, maintaining a constant speed, recognizing traffic lights, and stopping when needed. Side mirrors were also introduced. A navigation map with a destination was also suggested and explained (see Figures 2 and 4). The second training session took place on a highway where participants practiced overtaking. In this scenario, the TOA suggested to the participants whether they should slow down or try to overtake. In this scenario, the notations of the possible advice were explained. The third training task took place in a city setting. In this scenario, the TOA suggested to the participants when to stop in front of a traffic light. Again, the notation of the stopping advice was explained. At the end of this scenario, the supervisor evaluated the participant's driving performance and, depending on the participant's performance, decided whether to continue with the main experiment or repeat one or more of the training sessions to improve the participant's performance. The decision to repeat a training session is made when there has been a red light violation and/or a collision and/or at least two lane violations.

5.2.2 Main Tasks. After the training sessions, the participants were given a total of 12 simulation tasks, which were divided into two groups of 6. All overtaking tasks took place on a highway with 4 lanes. Traffic light scenarios took place in a city and involved several traffic lights turning from green to red as the vehicle approached.

5.3 Experimental Procedure

The experiment was conducted at the lab one participant at a time. We first explained the procedure of the experiment to the participant. The participant was then asked to complete an informed consent confirming that he/she was aware of the procedure and agreed to participate. The participant was then asked to fill out demographic information, information about driving experience and habits, and information about experience playing simulation games.

After completing the pre-experiment questionnaires, participants were asked to complete the PVT test (we used the online version in [21]), which measures reaction speed to visual stimuli. The results of the PVT test are related to driving performance ([2]). Participants were then given instructions on how to use the driving simulation and then practiced driving in three simulation training tasks as listed above. When they felt ready, they took a driving test in the simulation to verify that they could operate the car properly. During the test, the supervisor checked the following: The participant is able to turn right and left, keep in lane while driving and turning, obey traffic signs and avoid collisions with other vehicles and pedestrians, and reach the location on the navigation map. If the supervisor determined that the participant could not operate the car properly, he or she was given more time to practice in the driving simulation before proceeding to the next steps. Each participant completed 12 simulation tasks: 6 with the TOA and 6 without the TOA. The order of the tasks was counterbalanced so that half of the participants (chosen randomly) completed the first 6 tasks without the TOA followed by 6 tasks with the TOA, while half of the participants completed the tasks the other way around. The first 3 tasks were always the overtaking scenario, while the next 3 tasks were the traffic light scenarios. Once participants have completed the first set of 6 tasks, they are asked to complete the NASA-TLX questionnaire ([30]), which is a standard way to measure perceived workload. They then began the second set of 6 tasks. After completing the second set of tasks, participants were asked to fill out the NASA-TLX questionnaire again.

Finally, participants were asked to fill out a post-experiment questionnaire with three parts. The first part is about the usefulness of the TOA in the overtaking scenario, the second part is about its usefulness in the red light scenario and the last part contains some questions about both scenarios.

5.4 Data Analysis

We collected the logs provided by the CARLA simulation. Logs record the status of the system every 45 milliseconds. We analyzed logs both for brake usage and collisions. We considered multiple collisions as one if they occurred simultaneously, and split them into multiple collisions otherwise. We also wrote code to decide whether a red light violation occurred in a frame based on the locations of the ego vehicles and the status of the traffic light. Finally, we used the location of the ego-vehicle, its direction of travel, and the location of the traffic light to distinguish between different red light violations.

We performed a $t - test$ or similar statistical procedure for each of the measurements to test whether the difference between each measurement with and without the TOA was significant.

5.4.1 Evaluate Overtaking Scenarios. To evaluate the participants' performance in the overtaking scenarios, we first summed up the number of collisions with and without the TOA. In addition, we measured how long (in seconds) the brake was used and formed averages for the cases with and without the TOA.

To evaluate the safety and efficiency of each overtake, we assigned three human evaluators to watch the videos of the scenarios and rate each overtake. Each evaluator watched the overtaking videos in random order and was asked to rate each overtake on a scale of $1 - 7$, indicating how safe and how efficient the overtaking was. The videos presented to the taggers were presented without the advice notations. For each video and question, we averaged the ratings of the three raters to obtain an average rating of safety and efficiency. Figure 5 shows the taggers' interface.

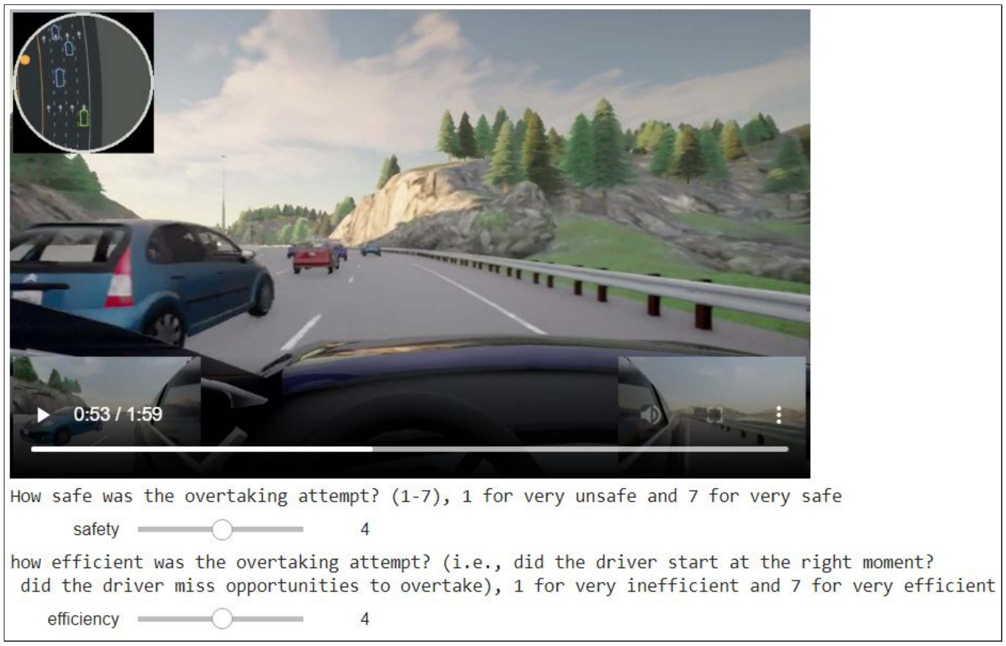

Fig. 5. Manual tagging interface as presented to taggers.

5.4.2 Evaluate Traffic Light Scenarios. The evaluation of the traffic light scenarios was completely automatic. We counted the number of collisions and measured how long (in seconds) the brake was used and formed averages for the cases with and without TOA. The detailed results are available in the next section.

6 USER STUDY RESULTS

We first present the performance measures both in the overtaking and the traffic light scenarios, followed by the cognitive load and subjective results.

6.1 Performance measures

6.1.1 Overtaking Scenario. No difference was found when comparing the number of collisions. Overall, there was a low number of collisions during the tasks: 15 collisions out of the 66 tasks with the TOA and 15 collisions in the 66 tasks without the TOA.

Two measures were proposed to examine the usefulness of the TOA in the overtaking scenario. The first is the use of the vehicle brake. In this case, participants using the TOA used the brake for shorter periods of time (M=3.23 seconds, SD=2.04 with the TOA and M=3.75 SD=1.98 without). A one-tailed t-test showed that this difference was significant (p = 0.034). A lower number of brakes could indicate better speed planning and perhaps more efficient driving.

The second measure was the evaluators' scores of safety and efficiency of the overtake. Results of our analysis showed that the raters rated both safety (M=4.95, SD=1.64 with the TOA and M=4.73, SD=1.85 without the TOA) and efficiency (M=5.12, SD=1.49 with the TOA and M=4.79, SD=1.57 without) higher with the TOA compared to without it.

subscale	Score with TOA	Score without TOA	$P(T <= t)$
Effort	136.82 (122.34)	172.82 (126.61)	0.032
Frustration	81.73 (107.53)	134.5 (125.46)	0.005
Mental Demand	95.91 (74.09)	154.14 (137.79)	0.024
Performance	160.82 (115.06)	141.91 (138.00)	0.248
Temporal Demand	103.95 (102.95)	98.14 (95.35)	0.358

Table 3. TLX mean results (and STDEV) for the different subscales.

Cronbach's alpha score is 0.78 for safety and 0.71 for efficiency. The difference in safety is not significant ($p = 0.19$) while the difference in efficiency is significant ($p = 0.036$).

6.1.2 Traffic Light Scenarios. Analysis of the results revealed that significantly fewer red light violations were performed when TOA was involved (M=0.18, SD=0.40) compared to when the TOA was not involved(M=0.82, SD=1.22). A one tailed t-test was performed and verified that this difference is significant (p=0.005). In addition, participants used the brake more often with the TOA (M=4.42 seconds, SD= 1.95 vs. M=3.53 seconds, SD=1.63), indicating that participants were more cautious when an advice was given. However, the difference in braking behavior could be explained by poorer planning when advice is present. A t-test on these results showed that the difference is significant (p=0.014).

6.2 Cognitive load

The mean value of NASA-TLX questionnaire data was significantly lower when TOA was involved (M=40.77,SD=22.1) compared to without the TOA(M=47.55,SD=24.45), p=0.014.

The results for the subscales are presented in Table 3. The values when TOA is involved are lower for the effort, frustration, and mental demand subscale and this difference is significant for these three subscales. For the performance and temporal demand subscales, results were not significant. Physical demand was not measured because it was not relevant in our simulation, which only required participants to use a computer with a keyboard and mouse. The difference in TLX results indicate that the TOA does indeed help reduce operator workload.

6.3 Subjective Results

On average, subjective ratings indicate that participants thought the advice given by the TOA was good, helpful and safe. As can be seen in Table 6, participants indicated that the TOA helped them understand the situation (3.91 on average for the overtake scenario and 4.59 for the traffic light scenario on a 5-point Likert scale) and made a correct decision (4.36, 4.68). In addition, most participants indicated that the advice was safe enough (3.5 and 3.82 on average). As can be seen in Table 7, most participants also indicated that the system helped them react faster (average rating of 4.05) and that they would like to have such advice when driving (average 4.14).

7 DISCUSSION

We describe a method for developing an advisory agent for teledriving (TOA) and implement it in two scenarios: overtaking a vehicle and stopping at a traffic light. The results of an experiment evaluating the advice in these two scenarios show that the TOA is beneficial in reducing the operators' workload and improving their performance. Furthermore, most participants thought that the advice was useful and helpful for the teleoperation task.

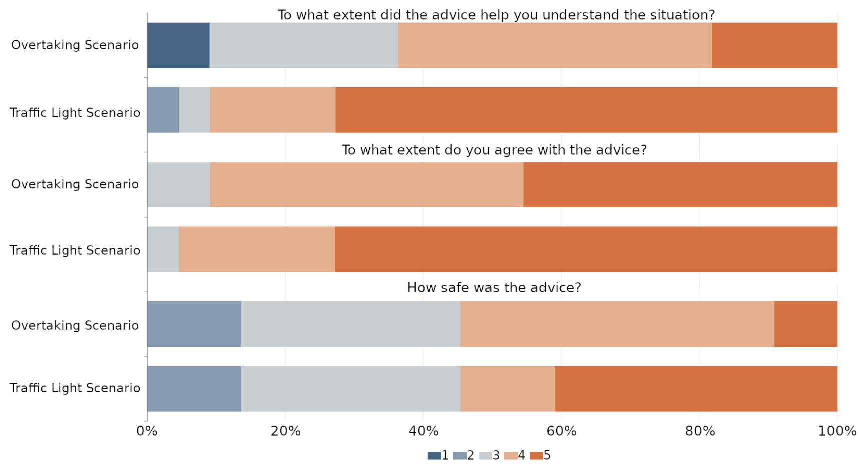

Fig. 6. A survey on the overtaking scenario and the traffic light scenario.

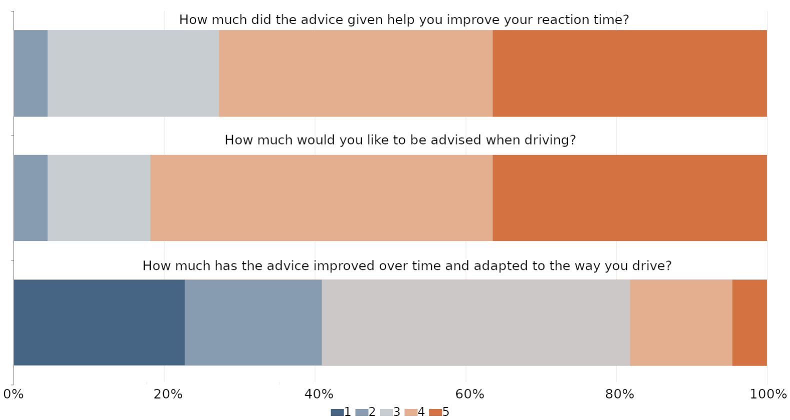

Fig. 7. A general survey on the advisory system.

As we mentioned earlier, many advanced driver assistance systems (ADAS) warn drivers of certain dangerous events (e.g., [58], [61]). However, these works focus on warning drivers of various hazards. In contrast, our system focuses on helping drivers in decision-making which is much more challenging. The main reason is that it is difficult to model people's behavior and decisions (e.g., [36, 51, 52]) that are needed for advice provision. What can be helpful when developing advice for teleoperators is the fact that it runs in centralized teleoperation centers. This can facilitate the collection of teleoperators' behavior data in a much easier way than when trying to collect data on actual human drivers. Teleoperation centers can also use strong computational powers for the computations needed for providing real-time advice that may not be available in regular cars.

Eriksson et al. [19] have shown that displaying advice to drivers while driving can help drivers' performance. However, they did not provide a general methodology for how to come up with the advice. The subjects in our study indicated that they would be happy to obtain such advice when driving regular cars. Thus, future work would examine whether it is possible to use our method to provide such real-time in-vehicle advice while driving.

What could be helpful to solve the challenge of advising drivers located in their vehicles is the important outcome of our advice provision methodology that allows considering only a small subset of the objects in the environment and some of their properties, and there is no need to consider all of them.

Because centralized teleoperation centers have yet to be active, for this study, we collected data specifically for our two implementation scenarios. Thus, the data we collected was insufficient to predict an appropriate score for achieving each goal in a given action. Therefore, we had to define a set of rules that implicitly aimed at maximizing the score. As more data on teleoperator behavior becomes available, we will be able to make more accurate and more realistic predictions about the expected outcome of an action given a problem. We also hope to be able to build appropriate machine-learning models and make more accurate predictions. In addition, with more data per teleoperator, we will be able to better tailor and personalize recommendations to each teleoperator. As mentioned above, once teleoperation centers are available, it will be possible to collect large amounts of data for this purpose.

The results from our user study show that the given advice reduced the overall perceived workload of the participants, specifically reducing their effort, frustration, and mental demand. This is encouraging as a high cognitive workload is one of the challenging factors in teledriving [59]. Supporting cognitive-demanding tasks such as overtaking [45] with visual decision support aids may take some of the burdens off the driver. However, this should be done with caution. Too many warnings and recommendations, as well as incorrect advice, can backfire and actually increase the driver's workload [47]. Thus both the accuracy of the advice and the number of recommendations should be carefully examined.

It is interesting to compare the two scenarios studied in terms of trust. One possible explanation for the fact that the advice in the traffic light scenario was somewhat more successful than in the overtaking scenario is that in the traffic light scenario, the reason for the advice is straightforward. It is just a matter of noticing the traffic light and stopping accordingly. With overtaking, on the other hand, the goal of the advice is to tell the driver what to do in a complex situation. This is less intuitive and more complex, and the reasons and validity of the advice (i.e., when to perform an overtake) are not entir; new to the participant. Therefore, it is more difficult for the participant to simply follow the advice. In addition, it is more difficult for the agent to model the overtaking scenario and therefore give correct advice to the participant compared to the traffic light scenario. On the other hand, the need for advice in such difficult situations might be greater. It might be interesting to investigate how to increase participants' trust in the agent's advice and whether following the advice improves performance if the agent is trustworthy. One way to increase trust might be to provide a confidence level in addition to the advice itself (see, e.g., Zhang et al. [66]). However, even if this leads to an improvement, it could be problematic if it increases the teleoperator's workload.

An advice agent faces similar problems to AV algorithms and many methods have been developed for tasks such as overtaking and crossing intersections. However, it is somewhat more difficult to give advice to a teleoperator compared to actually controlling an AV because the capabilities of the human operator are not always known and her actions can be unexpected, whereas AVs are fully aware of their capabilities and performance. Moreover, there is a time lag between giving advice and actually performing the advice, while an AV fully controls the start time of the activity. On the other hand, the teleoperator is not obliged to accept the advice, thus providing an additional level of validation for the decision. In teleoperation, the human teleoperator is accountable for the final decision, making the task of the advising agent somewhat easier than that of the AVs.

Finally, the timing of the advice seems to play an important role in its effectiveness. One direction for future research is to examine exactly when advice should be given. To do so we may need to modify our model to account for more than one unit of time at a time, as we do now.

7.1 Limitations

Our work has several limitations. First, in the current work, we have considered only two scenarios to examine general provision of driving advice. Other scenarios which may present further challenges should be examined in the future. Second, to evaluate the recommendations, we used a simple simulator setting rather than a real vehicle control stting in a teleoperation center. Also, we used arrow keys for steering, which is different from other work that uses a steering wheel (e.g., [59]). Since drivers are used to using a steering wheel for driving, we assume that teleoperator centers use a steering wheel, at least initially. In addition, a teleoperator center might dedicate even more resources to emulate the vehicle's setting. For example, in a simulation, the participant does not feel the forces acting on the car and does not hear external sounds. Better simulators (e.g., imposing forces) can possibly better resemble the remote driving scenario and the cognitive efforts needed for such a task.

Third, our personalization approach and its effect is limited. This is because each participant went through very few simulation tasks (6 overtaking tasks and 6 traffic light tasks). A higher number of tasks per participant could help us to better tailor the system to each participant and provide better personalized advice.

Forth, as mentioned earlier, all simulation tasks were performed under good weather conditions and in daylight. Driving in poor weather conditions and dark environments may require appropriate adjustment of the teleoperation itself as well as the advising system parameters (See for example Graf and Hussmann [27]).

Finally, the current status of the teleoperator is not considered in our TOA. Some work in other areas of teleoperation such as Jia et al. [31] and Singh et al. [54] consider questions about a teleoperator's mental state, such as whether he/she is tired and how much stress is present in the teleoperation center at the time of advising. Investigating the effect of such questions on the teleoperator's performance is relevant in our setting and may require adjusting the durations and safety distances calculated in our setting.

8 CONCLUSIONS

Since AVs cannot be expected to handle all driving situations, at least not in the near future, teleoperation centers are expected to remotely support AVs using teledriving. In this paper, we assume a model in which an AV sends a request for intervention to a teleoperation center in a situation it cannot handle, after which a human operator remotely controls the vehicle until the problematic situation is resolved. A major challenge with this model is that the teleoperator must quickly take control of the vehicle, understand the situation, and remotely control the vehicle while keeping an eye on the environment and the changing situation of the vehicle. To support the teleoperator in this challenging task and reduce his or her cognitive load, we have developed a teleoperator assistance agent (TOA). The TOA uses parameters from the remote environment, including parameters derived from common knowledge, as well as personalized and adaptive parameters. It focuses only on relevant objects and parameters to work efficiently. The TOA was implemented in two challenging driving scenarios: Overtaking a slow vehicle on a highway and crossing a signalized intersection when the vehicle arrives just before the traffic light turns red. We conducted an experiment with human participants to test the TOA in a simulation. Results indicate that the TOA significantly reduced the cognitive load of the human participants and improved their performance in some areas: it reduced the use of brakes and improved the rated efficiency in the overtaking scenario and reduced the number of red light violations in the traffic light scenario. In addition participants reported that they were satisfied with the teleoperation agent's advice.

ACKNOWLEDGMENTS

The research was supported in part by the Israeli Innovation Authority through the Andromeda consortium.

REFERENCES

[1] Melissa Acevedo and Joachim I Krueger. 2004. Two egocentric sources of the decision to vote: The voter's illusion and the belief in personal relevance. *Political Psychology* 25, 1 (2004), 115–134.

[2] Al-Baraa Abdulrahman Al-Mekhlafi, Ahmad Shahrul Nizam Isha, and Gehad Mohammed Ahmed Naji. 2020. The relationship between fatigue and driving performance: A review and directions for future research. *J. Crit. Rev* 7 (2020), 134–141.

[3] Ersin Armağan and Tufan Kumbasar. 2020. An intelligent overtaking assistant system for autonomous vehicles. In *International Conference on Intelligent and Fuzzy Systems*. Springer, 1068–1076.

[4] Aviram Aviv, Yaniv Oshrat, Samuel A Assefa, Tobi Mustapha, Daniel Borrajo, Manuela Veloso, and Sarit Kraus. 2021. Advising Agent for Service-Providing Live-Chat Operators. *arXiv preprint arXiv:2105.03986* (2021).

[5] Amos Azaria, Ya'akov Gal, Sarit Kraus, and Claudia V Goldman. 2016. Strategic advice provision in repeated human-agent interactions. *Autonomous Agents and Multi-Agent Systems* 30, 1 (2016), 4–29.

[6] Amos Azaria, Sarit Kraus, Claudia V Goldman, and Omer Tsimhoni. 2014. Advice provision for energy saving in automobile climate control systems. In *Twenty-Sixth IAAI Conference*.

[7] Christopher R Baker and John M Dolan. 2008. Traffic interaction in the urban challenge: Putting boss on its best behavior. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 1752–1758.

[8] Hillel Bar-Gera, Tal Oron-Gilad, and Oren Musicant. 2013. In-vehicle stopping decision advisory system for drivers approaching a traffic signal. *Transportation research record* 2365, 1 (2013), 22–30.

[9] Daniel Bogdoll, Stefan Orf, Lars Töttel, and J Marius Zöllner. 2022. Taxonomy and Survey on Remote Human Input Systems for Driving Automation Systems. In *Future of Information and Communication Conference*. Springer, 94–108.

[10] Wändi Bruine de Bruin, Andrew M Parker, and Baruch Fischhoff. 2007. Individual differences in adult decision-making competence. *Journal of personality and social psychology* 92, 5 (2007), 938.

[11] Yilun Chen, Chiyu Dong, Praveen Palanisamy, Priyantha Mudalige, Katharina Muelling, and John M Dolan. 2019. Attention-based hierarchical deep reinforcement learning for lane change behaviors in autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 0–0.

[12] JD Chovan, L Tijerina, G Alexander, and DL Hendricks. 1994. *Examination of lane change crashes and potential IVHS countermeasures. Final Report.* Technical Report.

[13] Nancy J Cooke. 2006. Human factors of remotely operated vehicles. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 50. SAGE Publications Sage CA: Los Angeles, CA, 166–169.

[14] Cindy Dietrich. 2010. Decision making: Factors that influence decision making, heuristics used, and decision outcomes. *Inquiries Journal* 2, 02 (2010).

[15] David F Dinges and John W Powell. 1985. Microcomputer analyses of performance on a portable, simple visual RT task during sustained operations. *Behavior research methods, instruments, & computers* 17, 6 (1985), 652–655.

[16] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An open urban driving simulator. In *Conference on robot learning*. PMLR, Mountain View, United States, 1–16.

[17] Henrikke Dybvik, Martin Løland, Achim Gerstenberg, Kristoffer Bjørnerud Slåttsveen, and Martin Steinert. 2021. A low-cost predictive display for teleoperation: Investigating effects on human performance and workload. *International Journal of Human-Computer Studies* 145 (2021), 102536.

[18] Avshalom Elmalech, David Sarne, Avi Rosenfeld, and Eden Erez. 2015. When suboptimal rules. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29.

[19] Alexander Eriksson, Sebastiaan M Petermeijer, Markus Zimmermann, Joost CF De Winter, Klaus J Bengler, and Neville A Stanton. 2018. Rolling out the red (and green) carpet: supporting driver decision making in automation-to-manual transitions. *IEEE Transactions on Human-Machine Systems* 49, 1 (2018), 20–31.

[20] Luke Fletcher, Seth Teller, Edwin Olson, David Moore, Yoshiaki Kuwata, Jonathan How, John Leonard, Isaac Miller, Mark Campbell, Dan Huttenlocher, et al. 2008. The MIT–Cornell collision and why it happened. *Journal of Field Robotics* 25, 10 (2008), 775–807.

[21] Sleep Disorders Center Florida. 2022. PVT test. http://www.sleepdisordersflorida.com/pvt1.html. Accessed: 2022-09-18.

[22] Terrence Fong and Charles Thorpe. 2001. Vehicle teleoperation interfaces. *Autonomous robots* 11, 1 (2001), 9–18.

[23] Andrei Furda and Ljubo Vlacic. 2011. Enabling safe autonomous driving in real-world city traffic using multiple criteria decision making. *IEEE Intelligent Transportation Systems Magazine* 3, 1 (2011), 4–17.

[24] Jean-Michael Georg, Johannes Feiler, Frank Diermeyer, and Markus Lienkamp. 2018. Teleoperated driving, a key technology for automated driving? comparison of actual test drives with a head mounted display and conventional monitors. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 3403–3408.

[25] Sebastian Gnatzig, Frederic Chucholowski, Tito Tang, and Markus Lienkamp. 2013. A System Design for Teleoperated Road Vehicles.. In *ICINCO (2)*. 231–238.

[26] Noah Goodall. 2020. Non-technological challenges for the remote operation of automated vehicles. *Transportation research part A: policy and practice* 142 (2020), 14–26.

[27] Gaetano Graf, Hao Xu, Dmitrij Schitz, and Xiao Xu. 2020. Improving the prediction accuracy of predictive displays for teleoperated autonomous vehicles. In *2020 6th International Conference on Control, Automation and Robotics (ICCAR)*. IEEE, 440–445.

[28] GreyB. 2022. Top 30 Self Driving Technology and Car Companies. https://www.greyb.com/autonomous-vehicle-companies. Accessed: 2022-09-28.

[29] Robert C Hampshire, Shan Bao, Walter S Lasecki, Andrew Daw, and Jamol Pender. 2020. Beyond safety drivers: Applying air traffic control principles to support the deployment of driverless vehicles. *PLoS one* 15, 5 (2020), e0232837.

[30] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*. Vol. 52. Elsevier, 139–183.

[31] Yunyi Jia, Ning Xi, Shuang Liu, Yunxia Wang, Xin Li, and Sheng Bi. 2014. Quality of teleoperator adaptive control for telerobotic operations. *The International Journal of Robotics Research* 33, 14 (2014), 1765–1781.

[32] Hossein Jula, Elias B Kosmatopoulos, and Petros A Ioannou. 2000. Collision avoidance analysis for lane changing and merging. *IEEE Transactions on vehicular technology* 49, 6 (2000), 2295–2308.

[33] E Ásgeir Juliusson, Niklas Karlsson, and Tommy Gärling. 2005. Weighing the past and the future in decision making. *European Journal of Cognitive Psychology* 17, 4 (2005), 561–575.

[34] Nidhi Kalra and Susan M Paddock. 2016. Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? *Transportation Research Part A: Policy and Practice* 94 (2016), 182–193.

[35] Prabhjot Kaur, Samira Taghavi, Zhaofeng Tian, and Weisong Shi. 2021. A survey on simulators for testing self-driving cars. In *2021 Fourth International Conference on Connected and Autonomous Driving (MetroCAD)*. IEEE, 62–70.

[36] Sarit Kraus. 2016. Human-agent decision-making: Combining theory and practice. *arXiv preprint arXiv:1606.07514* (2016).

[37] Jim Kurose and Keith Ross. 2017. Computer networking: A top-down approach, global edition.

[38] Qingqing Li, Jorge Peña Queralta, Tuan Nguyen Gia, Zhuo Zou, and Tomi Westerlund. 2020. Multi-sensor fusion for navigation and mapping in autonomous vehicles: Accurate localization in urban environments. *Unmanned Systems* 8, 03 (2020), 229–237.

[39] Xin Li, Xin Xu, and Lei Zuo. 2015. Reinforcement learning based overtaking decision-making for highway autonomous driving. In *2015 Sixth International Conference on Intelligent Control and Information Processing (ICICIP)*. IEEE, 336–342.

[40] Mobileye. 2022. Driving the autonomous vehicle evolution. https://www.mobileye.com/. Accessed: 2022-08-23.

[41] Abhishek Mukhopadhyay, Vinay Krishna Sharma, Prashant Tatyarao Gaikwad, Ajay Kumar Sandula, and Pradipta Biswas. 2022. Exploring the Use of XR Interfaces for Driver Assistance in Take Over Request. In *Adjunct Proceedings of the 14th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. 58–61.

[42] Clare Mutzenich, Szonya Durant, Shaun Helman, and Polly Dalton. 2021. Updating our understanding of situation awareness in relation to remote operators of autonomous vehicles. *Cognitive research: principles and implications* 6, 1 (2021), 1–17.

[43] Stefan Neumeier, Nicolas Gay, Clemens Dannheim, and Christian Facchi. 2018. On the way to autonomous vehicles teleoperated driving. In *AmE 2018-Automotive meets Electronics; 9th GMM-Symposium*. VDE, 1–6.

[44] Julia Nilsson and Jonas Sjöberg. 2013. Strategic decision making for automated driving on two-lane, one way roads using model predictive control. In *2013 IEEE Intelligent Vehicles Symposium (IV)*. 1253–1258. https://doi.org/10.1109/IVS.2013.6629638

[45] Nishant Mukund Pawar and Nagendra R Velaga. 2021. Investigating the influence of time pressure on overtaking maneuvers and crash risk. *Transportation research part F: traffic psychology and behaviour* 82 (2021), 268–284.

[46] Daniele Pinotti, Fabio Tango, Maria Giulia Losi, and Marco Beltrami. 2014. A model for an innovative Lane Change Assistant HMI. *Human Factors and Ergonomics Society Europe* (2014), 149–159.

[47] Omar Raddaoui and Mohamed M Ahmed. 2020. Evaluating the effects of connected vehicle weather and work zone warnings on truck drivers' workload and distraction using eye glance behavior. *Transportation research record* 2674, 3 (2020), 293–304.

[48] Lena Rittger, Gerald Schmidt, Christian Maag, and Andrea Kiesel. 2015. Driving behaviour at traffic light intersections. *Cognition, Technology & Work* 17, 4 (2015), 593–605.

[49] SW Roberts. 1959. Control Chart Tests Based on Geometric Moving Averages. *Technometrics* (1959), 239–250.

[50] Ariel Rosenfeld, Noa Agmon, Oleg Maksimov, and Sarit Kraus. 2017. Intelligent agent supporting human–multi-robot team collaboration. *Artificial Intelligence* 252 (2017), 211–231.

[51] Ariel Rosenfeld, Joseph Keshet, Claudia V Goldman, and Sarit Kraus. 2016. Online prediction of exponential decay time series with human-agent application. In *Proceedings of the Twenty-second European Conference on Artificial Intelligence*. 595–603.

[52] Ariel Rosenfeld and Sarit Kraus. 2018. Predicting human decision-making: From prediction to action. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 12, 1 (2018), 1–150.

[53] SAE. 2021. Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles. (2021).

[54] Gaganpreet Singh, Sergi Bermúdez i Badia, Rodrigo Ventura, and José Luís Silva. 2018. Physiologically attentive user interface for robot teleoperation: real time emotional state estimation and interface modification using physiology, facial expressions and eye movements. In *11th International Joint Conference on Biomedical Engineering Systems and Technologies*. SCITEPRESS-Science and Technology Publications, 294–302.

[55] Autonomous Solutions. 2022. Teleoperated Mining Vehicles. https://asirobots.com/teleoperated-mining-vehicles/. Accessed: 2022-08-23.

[56] Keith E Stanovich and Richard F West. 2008. On the relative independence of thinking biases and cognitive ability. *Journal of personality and social psychology* 94, 4 (2008), 672.

[57] Chengrui Su, Huanghe Li, and Xiaodong Wu. 2021. Artificial Steering Feel for Teleoperated Road Vehicle with Disturbance Observer. In *2021 5th CAA International Conference on Vehicular Control and Intelligence (CVCI)*. IEEE, 1–6.

[58] Qinyu Sun, Hongjia Zhang, Zhen Li, Chang Wang, and Kang Du. 2019. ADAS acceptability improvement based on self-learning of individual driving characteristics: A case study of lane change warning system. *IEEE Access* 7 (2019), 81370–81381.

[59] Felix Tener and Joel Lanir. 2022. Driving from a Distance: Challenges and Guidelines for Autonomous Vehicle Teleoperation Interfaces. In *CHI Conference on Human Factors in Computing Systems*. 1–13.

[60] Shara Tibken. 2018. Waymo CEO: Autonomous cars won't ever be able to drive in all conditions. *CNET. Retrieved from Nov* 13 (2018).

[61] Jianqiang Wang, Lei Zhang, Dezhao Zhang, and Keqiang Li. 2012. An adaptive longitudinal driving assistance system based on driver characteristics. *IEEE Transactions on Intelligent Transportation Systems* 14, 1 (2012), 1–12.

[62] Xiaohua Wang, Teng Zheng, and Zhiqiang Zhang. 2018. Design and Implementation of Intelligent Vehicle Control System Based on ROS. In *2018 Chinese Automation Congress (CAC)*. IEEE, 1661–1665.

[63] Waymo. 2022. Waymo Driver. https://https://waymo.com/. Accessed: 2022-08-23.

[64] Junqing Wei, John M Dolan, and Bakhtiar Litkouhi. 2010. A prediction-and cost function-based algorithm for robust autonomous freeway driving. In *2010 IEEE Intelligent Vehicles Symposium*. IEEE, 512–517.

[65] Changxi You, Jianbo Lu, Dimitar Filev, and Panagiotis Tsiotras. 2018. Highway traffic modeling and decision making for autonomous vehicle using reinforcement learning. In *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 1227–1232.

[66] Yunfeng Zhang, Q Vera Liao, and Rachel KE Bellamy. 2020. Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 295–305.

[67] Yuting Zhang, Xuedong Yan, Xiaomeng Li, Jiawei Wu, and Vinayak V Dixit. 2018. Red-light-running crashes' classification, comparison, and risk analysis based on General Estimates System (GES) crash database. *International journal of environmental research and public health* 15, 6 (2018), 1290.