

# Automotive Turing Test

Steven F. Kalik and Danil V. Prokhorov  
Toyota Technical Center-TEMA  
Boston, MA, Ann Arbor, MI, USA

**Abstract**— The Turing Test is often cited when the intelligence level of a presumably intelligent computer program is to be assessed. We propose to extend the scope of the test to the domain of intelligent automotive vehicles. We discuss possible formats for such a test, and consider different measures produced by these tests.

**Keywords:** *Turing test, imitation game, human-like driving behavior, human driving intelligence, intelligent vehicle, robot driver, vehicle intelligence, road test, driving simulator, vehicle interaction, safe driver.*

## I. INTRODUCTION

Alan Turing’s original paper [1] proposed two central ideas that revolutionized thinking about modern computing. Of the two, perhaps the more revolutionary idea was the conception of the Universal Computing machine. But, the one which in many ways did more to capture the imagination of society, is the Turing Test (TT), based on the “Imitation Game” as Turing originally coined it.

As originally proposed, the imitation game pits an interrogator (C) against a man (A) and a woman (B). The interrogator is separated from the man and woman in such a way as to limit his observations to only the interactions they have with him and with each other. As Turing originally conceived it, those interactions were all verbal, relayed to the interrogator via an intermediary or, more preferably, through a set of teletype terminals. Thus, the responses of each player to the questions of the interrogator, and to each other’s responses to those questions, define the only input the interrogator receives to decide which is the man and which the woman. What makes the game far from simple is that only one of the players is trying to be correctly identified. In this case, let’s describe it as Turing did, as player B. The remaining player, A in this case, seeks to deceive the interrogator by giving them answers and statements that they hope will convince the interrogator to apply the opposite labels to the two players. The interrogator may ask questions and observe the responses until such time as he feels he knows which labels to apply, but the ability of A to give false answers or to respond in ways that would deliberately mislead the interrogator ultimately leaves the interrogator with doubt about which player to trust, and therefore with only an estimate of which player is which. This makes the game challenging with the placement of labels decided in a sense statistically, based on a set of responses observed by the interrogator.

As is well known, Turing’s ultimate variant of this game replaced A with a computer, posing the critical question: “Will the interrogator decide wrongly as often when the game is played like this, as he would when the game is played between a man and a woman?” [1]. This form of the question creates a statistical measure of distinguishability out of the set of decisions made by the interrogator as they play the imitation game multiple times. A comparison of the distinguishability measures between the cases when humans exclusively play the roles of A and B in the game, and when machines enter the game and take the role of the deceiver (A), utilizes these measurements to quantify the original question Turing sought to improve upon, which was simply “Can machines think?” [1]. Thinking in this case carries the normal but poorly-defined meaning humans usually ascribe to it when they speak of human thought.

We propose a similar approach in this paper: use human decisions about their notions of human and machine behaviors to determine statistically over a set of decisions whether behavior observed appears intelligent enough to be called by the same name. Thus we propose the Automotive Turing Test (ATT), and explore some basic ideas and properties suggested by such a test.

We begin by identifying the central features of the original Turing test in Section II, and consider the implementation of those features in the automotive domain in Section III. Section IV presents several different ways the automotive Turing test may be constructed. We discuss different features of those implementations. Section V considers the goals and measures of the ATT with regard to human goals for intelligent vehicles, pointing to required future work. Finally, Section VI concludes the paper.

## II. KEY COMPONENTS OF THE TURING TEST

Turing’s original conceptualization of the TT provided the following benefits:

1. A way to explicitly quantify a vague idea (“thinking” in Turing’s case).
2. Use of a human as an active sensor to identify human-like behavior
3. Simplification of the test structure to probe the essence of the assessed behavior while maintaining impartiality.

It is this third point that attracts our attention here.

To legitimately retain the TT label in the work we propose for this paper, we must retain the key elements of the original TT :

1. Limited observations between the players and the

interrogator

2. The ability of the players and the interrogator to interact directly in the domain of the test and not beyond
3. The opportunity to try to deceive the interrogator, and
4. The ability of both players and interrogator to recognize and influence the context of the observations and the responses returned during the game.

While we feel that it is critically important to maintain the quantitative nature of the TT, with its central measure that compares ensembles of label applications between cases when (1) two humans act as players and (2) a machine and a human act as the players, (in both cases the interrogator being human), we recognize that this is not the only version of the game people think of today (see, e.g., [3]). But, to the extent possible we adhere to these key features of the test as originally conceived. Where we diverge from this, we note it explicitly.

### III. TURING TEST IN THE AUTOMOTIVE DOMAIN

Just as Alan Turing posed the question “Can machines think?”, we now pose the question “Can machines drive?” Just as Turing’s use of “think” carried with it human aspects of thinking, so too do we carry components of human driving into this test, seeking to capture aspects that we associate with a human driver into our use of the word “drive”. As this paper and perhaps future work will illustrate, the exact form of what we will call the **Automotive Turing Test (ATT)** is still evolving. The Turing test goal of being indistinguishable from a human may or may not be the target we as a society choose to approach, given the number of traffic accidents observed each year across the globe. On the other hand, as we point out, there may be some benefits to the limitations humans show, too.

To convert the Turing test into the automotive domain, we must transform the key elements from section II into a form appropriate for testing driving behavior, and into interactions appropriate to a driver controlling an automobile. To do so, we remove from consideration the verbal interactions central to the original TT, and replace them with behavioral and signaling interactions available to vehicles.

So far we conceive the ATT to be administered in two different environments, and in each environment, the test can be administered in two different ways. Each of these provides opportunities for the test to address different aspects of human driving behavior.

#### *Environments*

The environments in which we initially imagine the test being administered are:

- On real vehicles
- In a life-like simulation environment

In both cases, to stay with the spirit of the TT, the interrogator must be prevented from seeing anything inside any vehicles to which they will be expected to apply labels,

and by extension, from seeing any information other than the vehicles behavior.

In either testing environment we require that the space over which the vehicles will drive be large enough, in terms of road length, space covered, and potential driving situations that can be encountered, so that drivers will need to learn a map of the environment and patterns of dynamics in that environment to get around conveniently and effectively. This assures the exercise of driving skills at both tactical (on the order of one second) and strategic (seconds – minutes – hours) scales during the administration of the test.

As some situations in the real world occur on the road much less frequently than others, we recognize the importance of explicitly building the opportunity for these events to occur into the testing environment. For example, near-accident situations are much less common than normal driving situations [2]. Yet, such rare situations, if recreated either by surprising behavior from another vehicle, or through the inclusion of unexpected behavior from other agents in the testing environment, would present great opportunities to assess the limits of human and intelligent machine capabilities. Hence the ATT environment, whether in a real or a virtual car, should include the ability to implement driving surprises.

#### *Administration methods*

The administration methods we initially imagine are:

- Vehicle to Vehicle format: Each participant – the interrogator and both players A and B - drive separate vehicles in the environment.
- Road Test format: A single player is to be labeled as a human or machine driver by an interrogator who sits in the vehicle to be observed, similar to the way a human driving instructor would pass judgment on a driver’s license applicant during a student driving test.

In both of the administration methods, human drivers have access to maps or navigation systems to support their route choice decisions. We offer driver-readable maps or driver-understandable navigation systems to support preparation for and driving in the testing environment, so that a standard access to knowledge of the environment can be assured.

Since today’s vehicles communicate with each other in only limited ways (obviously excluding driver sticking his head out the window and yelling or gesturing) if the test were to be run today, only standard signaling devices such as turn signals and brake lights would be available for communication from the driver. Observations of the environment including all vehicles, agents, objects, etc., are also visible to all players and the interrogator to inform everyone involved of the actions happening around them.

However, if at a future time, additional technologies were added to standard vehicles such as message passing equipment or warning systems for vehicle to vehicle communication, or aircraft-like “black box recorders” to record all vehicle and environmental data, one could easily

imagine incorporating such tools into the ATT. We elaborate on the use of data recorders for the purpose of the ATT in Section IV.

*How the tests are run*

We distinguish the test vehicles of players from all others in the environment by clearly marking or labeling such vehicles. To run an instance of the imitation game allowing only vehicle to vehicle interaction for the case of two test vehicles, we place vehicles labeled X and Y respectively at their starting locations. The interrogator gives them goal locations to reach which may differ from each other's, but which will require them to traverse a common subset of streets and intersections, and which will lead them to encounter common driving situations, or to create driving situations for each other. The interrogator knows to which destinations these players will need to drive.

When the players reach their destinations, they may be given new destinations that will again take them over another common set of roads and through common intersections and traffic situations according to the requirements set in the previous paragraph. The interrogator always knows these destinations so that he can choose his own actions to allow observation of and/or desired behavioral interaction with the players over the course of their trip. We allow this process to iterate until the interrogator decides he knows which labels A and B to ascribe to the vehicles originally designated X and Y.

Once the labels have been applied, a round of the game is considered over, and the correctness of the decision recorded. New assignments, new starting positions and destinations are selected for the test vehicles.

Although Turing originally constructed the test to provide an elaborately quantitative measure of whether a human interrogator interacting with an intelligent machine could be convinced to mislabel the players in the game as frequently when the players were both human as when they were human and machine, a more common usage today asks a simpler variant of that question: can a human observer distinguish the behavior of the intelligent vehicle from the behavior of a vehicle driven exclusively by a human driver.

When this alternate version of the TT is used, as in the Road Test administrations, we modify the previous route granting procedure to allow the interrogator to explicitly modify the route as necessary at any time during the test to exercise the drivers' capacities and to explore their decisions and behaviors in the face of route selections or modifications, as will be discussed in more detail later in the paper.

In the next section, we explore the implications of the different testing environments and the opportunities raised by testing in each combination of administration and environment.

IV. VARIANTS OF THE AUTOMOTIVE TURING TEST

Given the two different administration methods described above, and the two different environments in which the test

could be given, there are already four different variations of the test under consideration. A fifth, passively interrogated form of the test will also be briefly considered later, wherein only databases are reviewed and compared to decide whether the driver is human or machine. We begin here by describing four variants with active interrogators, as shown in Table 1.

TABLE I  
AUTOMOTIVE TURING TEST VARIANTS

AUTOMOTIVE TURING TEST VARIANTS		ENVIRONMENT	
		Real World	Simulated Environment
ADMINISTRATION	Vehicle-to-Vehicle Format	Variant 1: Real World Environment with only Vehicle-to-Vehicle Interaction	Variant 2: Simulated Environment with only Vehicle-to-Vehicle Interaction
	Road Test Format	Variant 3: Real World Environment with Road Test Format	Variant 4: Simulated Environment with Road Test Format

*Variant 1: Real World environment with Vehicle-to-Vehicle interactions.*

In this first variant described, we imagine an interrogator and two separate players in their respectively marked vehicles, receiving destinations and route suggestions and driving to them as described earlier.

Remember that the interrogator in this vehicle-to-vehicle interaction variant can supply new destinations and route guidance only before a player sets out on their route to the destination. Along the way, the interrogator, when seeking to change a vehicle's route, must interact with the players' vehicles only in the same ways that other vehicles in the environment can interact with the players. That is, interrogators influence the players' vehicles by driving close enough to them to make the actions of the interrogators' vehicle relevant to the driving decisions of the players. The type of interaction described here is imagined to primarily influence a player's tactical behavior in the time surrounding the interaction.

The interrogator may also want to test strategic decisions by a driver, such as choosing when to take a new route to a destination. To do this, the interrogator might try something like trying to tie up traffic so the player's vehicle must re-evaluate, and perhaps revise, the route they planned to take. However, performing actions such as these influence a

potentially very large number of people in addition to the players in the game, and may be considered inappropriate for the Real World Vehicle-to-Vehicle variant of the ATT. The route change can be more safely addressed in the domain of simulated environment, and more explicitly addressed through the Road Test variants.

*Variants 2: Simulated environment with Vehicle-to-Vehicle interactions.*

Just as the Vehicle-to-Vehicle interaction variants can be administered in the real world, so too can they be administered in a high-fidelity simulation environment, or in a virtual reality world. In this variant we remove the risk inherent in real world experiments to provide a safer test environment that still permits direct but controlled interactions between the players and the interrogator. This permits the interrogator to consider a broader range of actions with which to challenge the players than they had available to them before, when lives, limbs, property, and in fact even simply the convenience of non-involved parties were at stake. Now the penalties for behaviors that would influence these concerns can be modulated and explored, and if desired, even ignored in the construction of the testing environment.

For example, imagine an interrogator wishing to test the sensitivity of players to loss of time resulting from perturbations of the traffic flow through the environment in which the test is taking place. The interrogator could influence and test what it takes to get a response from the players by moving to a position he knows both vehicles will need to pass through or inescapably close to. (In many older cities, such central transit points are easy to imagine.) By uncivilly blocking a high traffic intersection, or perhaps more simply, by just adjusting the gaps at any intersection that the interrogator will accept to enter or pass across a traffic flow, the interrogator can modify some gross behaviors of the traffic system that will ripple out from that part of the environment. This puts players in the position of having to decide whether to wait out the delay, or to select an alternate route. Patterns in this behavior (and the resulting interactions with vehicles and other agents in the environment) can then be used to help the interrogator distinguish between different players before they make their decisions which player is human and which a robot driver. Possibly, in the simulated environment, players themselves might also select some of these behaviors in an effort to lead the interrogator to the labeling conclusion that the player seeks to achieve.

However, this experimental freedom for the interrogator comes at a price. Now, many of the things that came for free in the Real World environment such as sensory stimulation, and the inherent costs humans associate with risk to their person or property must be artificially created. For example, to provide the same sights, sounds, smells, and other sensory phenomena to the driver, we require novel simulation platforms that can re-create those effects. While a few high fidelity simulators exist in the world today, in most cases

many variables are ignored if they are not critical to the experiment of interest. (Olfactory stimulation, for example, need not be included as part of following a truck in a simulator, but in the real world it might influence the decision a human driver would make about whether to pass a truck or remain behind it on the road.)

The key advantage to both Vehicle-to-Vehicle interaction variants of the ATT are that they adhere most closely to Turing's original test.

Whether the simulated reality in the ATT is a perfect recreation of any known location in our world is not important for this vehicle-to-vehicle variant of the test. But, since one typical purpose in measuring the "intelligence" of a vehicle is to help decide when such a vehicle would be safe to introduce to our roads in the real world, we would benefit by requiring all vehicles to treat the simulated environment the same way they would treat the environment in the real world. To achieve this, we need them to act in the following ways and have the following characteristics:

- They must obey traffic laws and traffic control devices and customs;
- They must avoid curbs, trees, vehicles and pedestrian figures;
- They must have similar fields of view and limitations on their knowledge of the environment as they would in the real world (limited but wide field of view dependant on modeled head orientation; olfactory, auditory and tactile feedback separated from the visual field of view; reasonable directionality to all sensory input to provide the normal input available for a human driving in a vehicle); and
- They must have limited but reasonable control capabilities over the virtual vehicles they are driving, similar to the capabilities they would have over vehicles they would drive in the real world.

Having made these demands of our test to help fit it to a human goal for the domain of intelligent vehicles, we observe that this strong requirement can actually be relaxed without upsetting the pure *intelligence* testing measure developed by the imitation game. This is based upon the notion that the imitation game versions of the TT and ATT are not really about *safety*, or obeying traffic laws *per se*, but about producing behaviors that are indistinguishable from human behaviors in the environment in which they are observed.

We claim here that in imitation game based ATT's, obeying traffic laws will only matter to the players of the game to the extent that humans also choose to obey those laws in the given environment. However, if we ultimately hope to convince ourselves and society that an intelligent vehicle is as safe as, or safer than, human driven vehicles on the road (and therefore worthy to share the road with humans), we will want to consider ways to create the test environment so the value of obeying traffic laws and civil behavior (and the cost of disobeying them) is clearly understood and similar for all players in the testing environment, and is similarly weighted

to what it is now in real world. The often raised issue about how similar is human behavior in a simulated environment is then revisited here, but this time with an eye to parameters that can be built into the simulation environment to help enforce the kind of behavior we might desire. One considered approach to enforce this is to offer something of value (perhaps points that can be inherently valuable to a machine, or exchanged for prizes or money by a human player after they finish taking part in the test), and to place “law enforcement” drones in the simulated environment (location-locked camera-like systems or mobile police-like vehicles in the environment) that either deterministically or probabilistically penalize the player for violations of the laws in that environment. Such additions to the environment might be a noted increase in the effort necessary to construct it, but might still be beneficial for future consideration.

However, a test that tests for differences from a standard obviously carries with it all of the strengths and weaknesses of that standard, should one achieve the goal of being indistinguishable from the standard. With this in mind, we now turn our attention to other variants of the ATT whose measures perhaps allow more freedom in what is being measured. The Road Test format, with its implication of Pass or Fail grading for driving at a level suitable to be considered worthy of being called human, is perhaps more easily adapted to address the societal goal of intelligent vehicles that are not just human-like, but as clever as a human and at least as safe.

#### *Variant 3: Real World environment with Road Test format.*

This variant is similar to the variant 1 except that it prevents the interrogator from driving his or her own vehicle. However, in contrast to the variant 1, the interrogator is always present in the vehicle, and he can observe its behavior constantly, rather than episodically, as in the first two variants.

The interrogator can provide instructions to the driver in at least two ways, verbally or electronically. The instructions should be relevant to driving, e.g., “make lane change”, “turn right over there”, etc. This variant of the ATT should not turn into a domain-nonspecific testing the driver for natural language understanding.

The instructions provided to the driver must be given with a reasonable lead time for the driver to execute the appropriate maneuver safely. Alternatively, the instructions could be delivered at any time so long as the driver is empowered to just ignore illegal or unsafe instructions from the interrogator.

#### *Variant 4: Simulated environment with Road Test format.*

This variant is similar to variant 3 except that the real world environment is replaced by the simulated environment. In contrast to variant 2, the interrogator cannot drive his own virtual vehicle, but may enjoy the benefit of constantly observing the player’s behavior.

#### *Variant 5: Passive interrogation via recorded behavior*

#### *comparison to database of human behavior.*

A fifth variant, alluded to earlier, is available if recorded data becomes available through black-box type recorders. We proceed here to describe the ATT in that form, and to bring out the interesting points it reveals.

It bears stating that this departs from the original TT in two ways. First, passive interrogation removes any aspect of en-route interaction with the driver to be assessed, reducing the interrogator’s role to that of mere observer. Second, this test could be implemented as a computerized classifier system, which offers interesting opportunities, but places the human even further out of the testing loop, so that now they only prescribe the statistical tests employed, with the rest of the test executable without human intervention.

To explore this version in some detail, in this variant of the ATT, the behavioral estimate of the player is made not on second-to-second unfolding observations, but upon the ensemble recording of their behaviors over a set of routes driven during the test run. This passively interrogated, after-the-fact analysis of the data recorded from a new driver is then compared to a larger set of recorded data from known human drivers driving under similar conditions. Such a comparison would help reveal the relationship of the newly recorded dataset to the pool of human data encapsulated by the larger multi-driver ensemble.

We note here that even if a newly recorded player’s behavior differed significantly from that of other vehicles in the database of vehicles driven by humans, it would not necessarily mean that that new player was a machine. But, such a finding of significant difference from the behavior of the comparison set would certainly be worthy of further review, as it identifies a driver whose behavior doesn’t fit the behaviors of the others in the ensemble to which it was compared. Obviously all proper statistical concerns must be observed to make sure that both the sample of new driver behavior and the ensemble of behavior to which it is compared are large enough and rich enough to provide statistical reliability in the estimates they provide. Interestingly, we see no reason why data recordings from the real world couldn’t be just as valid as those collected in simulation.

#### *Other variants*

During considerations of the idea of the ATT, a number of other variants have come up.

One as yet poorly explored variant considers the inclusion of real vehicles driven by a player or interrogator via remote control. In both the vehicle-to-vehicle, and Road Test administrations of the ATT, this option could be included. This is interesting in particular because it raises a basic question. “Would the driving behavior be distinguishable between a human driving in a car as is normal today, and a human driving a real car via remote control?” Such a test adheres closely to the imitation game form of the ATT. But, more importantly, it gets directly at the heart of what it means to have your “skin in the game”. We hypothesize that there

will be discernable differences in behavior and interactions with other vehicles when the remote driver is safely ensconced away from the vehicle that is at risk.

If it were of interest, perhaps in a future where automated machines were the safer way to travel, another variant could be imagined, in which the passive interrogator version of the ATT is reversed to compare a human driver's recorded behavior data to a database of super-safe machine drivers. In that environment, this method might decide when the human driver's behavior differed little from super-safe machine drivers to be indistinguishable from them, making it safe to allow them to share the roads. Such a situation could also arise as part of the integration of human driven vehicles into environments that were established as isolated roadways initially built for the exclusive use of intelligent autonomous vehicles, as has been suggested by the earlier automated vehicle projects like those demonstrated by the DOT in the late 1990's.

#### V. WHAT DOES THIS AUTOMOTIVE TURING TEST (ATT) ACTUALLY MEASURE, AND DO WE WANT THAT?

The measure of intelligence in the original TT [1] is actually quite interesting, as are the related measures coming out of the ATT variants described in this paper. In the original TT, the judgment made by the human interrogator is based upon their observations of the behaviors of the players. But, the behaviors of the players are actually based upon B's (likely) honest attempt to represent herself faithfully, and on the deceiver's (A's) ability to model, duplicate, and display B's behavior appropriately to deceive the interrogator into thinking that player A is in fact player B. The model that A builds could also be augmented to include a model of the expectations of the interrogator C, capturing what A expects C to look for in his attempt to recognize honest player B from deceiver A. This would allow deceiver A to exhibit that behavior first, or to prepare a counter behavior that either influences C's expectations of B or nullifies the effect of B's behavior (as in Turing's suggested example statement by B "I'm the woman", followed by A's comment "Don't listen to him, I'm the woman").

While human guile, and to some extent its duplication by a machine, are inherently part of the definition of human thinking in the original TT and imitation game, the other interpretations of the TT idea ease this requirement. Instead, they replace it with a demand for sufficient skill in a particular task or a set of tasks. This implies that outside of conversational games, guile is not necessarily the ultimate definition of intelligence. It is at this point that the Road Test version of the ATT comes into play, because for most people the goal of vehicle intelligence is not to take the risks that humans would, but rather to be as safe and effective on our roadways as humans are today, or more so. Thus, for the Road Test version of the ATT, we seek a vehicle that is *smart enough to be safe in situations where a human might not be*.

This differs substantially from the Vehicle-to-vehicle ATT variants, which would give us a measure of whether or not a

human observing the behavior of our intelligent vehicle could reasonably have the same expectations of this vehicle's behavior that they would have of a human driver's. Given the number of accidents on our roadways every year, it is reasonable to ask even then whether this would be an acceptable standard. To add to the Vehicle-to-Vehicle variants the element of safety, we would need to require that our set of human players be limited to humans known to be notably good or safe drivers. This is a standard which, while obvious when grossly violated, may still be somewhat loosely defined currently and harder to assure than we might like.

So, if not absolute safety from an intelligent vehicle, what do we gain by using a purely human standard of intelligence? We may gain two important things: (1) a behavioral estimate of the driver we observe as matching a model of behavioral expectations which we build in our own mind, and (2) an estimate of how far we should trust that driver to actually use that model.

The creation of a model of the observed driver behaviors allows us to pre-plan actions to take if the observed driver either maintains their currently estimated course of action, or changes to one or the other of the predicted actions. This pre-computing and caching of solutions to more probable future situations is a valuable skill that allows faster identification of situation changes and faster reactions to those changing situations. This pre-computation is lost or significantly curtailed when encountering vehicles that drive in ways very different from the way we expect, as can be observed when driving for the first time in an unfamiliar location. This can be most striking when traveling in a land where cars drive on the opposite side of the road from a driver's previous experience. But, over time, through observation and mental modeling we reconstruct a model of the behavior of others around us. We can then incorporate their actions into our own repertoire and expectations. The inherent flexibility associated with this skill of learning and adapting is central to the aspects of intelligence that come to mind when we describe humans as smarter than a brittle artificial system that is otherwise highly trained but limited to a single fixed set of already-solved problems. But, underlying this aspect of intelligence, is the action of observing others, modeling their behavior, and adapting our own to be indistinguishable from theirs. This is the key measure tested by the TT and the Vehicle-to-Vehicle ATT, and this is what makes us safer in new situations, which is something we would like to see for intelligent vehicles.

If predictability similar to a human's is valuable because it allows us to prepare for upcoming situations, knowing the limit of that predictability is also valuable, but in a different way. One value is that, when we remember this limitation and explicitly include it in our thinking, it keeps us watching for the unexpected. This reduces our chances of becoming overconfident, and limits the risk to which we are willing to expose ourselves.

Another, and perhaps more important value of the limitation of the predictability of other vehicles, is that it creates a *social*

*contract* between vehicles on the road to ensure their mutual safety. The expectation of some unpredictability in the behavior of others around us (the element of surprise mentioned in Section III), and a respect for the sovereign right of others to behave in ways that we might not have predicted, leads us to create buffer zones around vehicles that are larger than would be required if the behavior of other vehicles was perfectly predictable.

A good example illustrating the points above in everyday driving can be seen in severely foggy or snowy weather, when vehicles often close the gaps between themselves and the vehicle ahead and rely to a larger extent than usual on following the taillights of the vehicle in front of them. When a driver of the lead vehicle overestimates the predictability of the rest of the environment, that mistake can ripple back through the entire chain of vehicles, with each in turn overestimating the reliability of the actions of the vehicles and environment in front of them. These chains of unmet expectations in the reliability of other drivers may lead to much larger accidents than might occur if each driver adjusted their estimates of how far to trust the other drivers more appropriately to match the actual environmental conditions.

Concluding this section, we wish to touch upon the amount of intelligence required to pass Road Test variants of the ATT.

In general, it does not require a lot of conscious efforts for an experienced human driver to drive a vehicle safely, especially in a familiar environment, e.g., repetitive drives from home to work and back. Potentially significant mental efforts seem to be used only for complex navigational tasks in a busy traffic environment, or when a driver chooses to do several tasks simultaneously (e.g., talking on a cell phone while checking directions on the map and driving), likely at the expense of an elevated risk of an accident.

Sometimes a driver must exercise a quick and correct judgment and decision making to avoid an accident or minimize its severity. For example, if a child suddenly jumps out on the road in front of the vehicle, a driver must quickly execute a suitable avoidance maneuver (humans would do so instinctively, but unfortunately not always successfully). A driver might opt for driving in a ditch next to the road to avoid hitting the child.

When an animal suddenly appears on the road, drivers sometimes choose the “stay the course” behavior for their vehicles, as evident in many US states by the sight of dead animals lying on the road.

What if a ball suddenly appears in front of the moving vehicle? Does this mean that a child might follow the ball in the next second? This hints at the need for an intelligent vehicle to possess intelligence broader than what is immediately applicable to driving or simply distinguishing humans from animals. It is highly desirable that the intelligent vehicle have enough intelligence not only to react quickly but also to exceed any human driver in ability to avoid the collision with a human or an animal.

## VI. CONCLUSION

In this paper we proposed several variants of the Automotive Turing Test (ATT) based on both the original imitation game version of the TT, and on subsequent extensions of the TT into domains of expert performance. We sketched out several implementations of the ATT for consideration and discussion on the topic. The first implementation is administered with only vehicle-to-vehicle interactions, while the second is administered “Road Test” style in a way very similar to student driver testing. These two administration styles and methods were described for implementation on both cars in the real world, and in a high fidelity simulation environment where risk to life, limb, or property can be removed. We discussed the existing value of testing for human levels of intelligence, including its inherent limitations and perhaps diverging goals. We distinguish it from a more typical automotive definition of vehicle intelligence, which uses the safety of vehicle occupants as a proxy for vehicle intelligence. We also pointed out how modeling human weaknesses within a system can actually strengthen the robustness of the system to protect against system perturbations or failures.

We acknowledge readily that there is a great deal more work to be done to instantiate the ATT, and to deliver intelligent vehicle driver systems and intelligent driver support systems that can compete with human drivers in the ATT. While focusing initially on human-like driving qualities, we point out how the ATT discussion might prove useful to target development of intermediate steps in the process of realizing intelligent driving systems in which accidents no longer occur. We also foresee opportunities through this work to explore the system of interacting human behaviors that take place on our roads, and to explore the true nature of human intelligence as it is reflected in vehicle control decisions.

## ACKNOWLEDGEMENT

SFK also thanks his family (DK, RK, & VK) for their support and understanding during the preparation of this paper. We also want to thank Michael Samples for useful comments about a draft of this paper.

## REFERENCES

- [1] A.M. Turing, “Computational Intelligence and Machinery”, *Mind*, Vol LIX, No. 236, pp 433-460, 1950
- [2] Dingus, T.A., Klauer, S.G., Neale, V.L., Petersen, A., Lee, S.E., Sudweeks, J., Perez, M. A., Hankey, J., Ramsey, D., Gupta, S., Bucher, C., Doerzaph, Z.R., Jermeland, J., and Knipling, R.R., “The 100-Car Naturalistic Driving Study, Phase II – Results of the 100-Car Field Experiment”, DOT HS 810 593, April 2006
- [3] S.J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, p. 2, Prentice Hall, 2nd edition, Pearson Education Inc, Upper Saddle River NJ, 2003.