Modeling and Influencing Human Attentiveness in Autonomy-To-Human Perception Hand-Offs

Yash Pant[1], Balasaravanan Thoravi Kumaravel[2], Ameesh Shah[2], Erin Kraemer[2], Marcell Vazquez-Chanlatte[2], Kshitij Kulkarni[2], Bjoern Hartmann[2], Sanjit A. Seshia[2]
[1]Univ. of Waterloo, [2]UC Berkeley
Award ID: CNS-1545126

Introduction

Perception-based errors are the cause of many failures in Autonomous Vehicles (AVs) - how can we rely on a human operator to solve perception errors that the AV is uncertain about?

We seek to formalize the process of when an AV cannot confidently perceive an object in its environment and therefore must ask the human operator a Perception Query to identify it for them - called The Perception Handoff Problem.

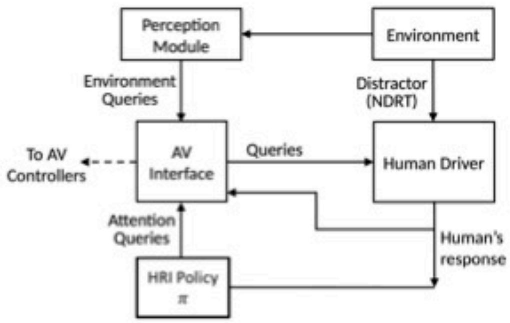

Figure 1: An overview of the human-AV interaction in the hand-off process. The AV, through a human-machine interface (HMI), can query the human driver when its perception module requires help in decision making, or to gauge/influence the human's state of attentiveness. Further influencing the human state is a Non-Driving Related Task (NDRT.)

Through our formalization, we aim to estimate human attentiveness in the Perception Handoff Problem, and we design a query-based Human-Machine Interface (HMI) that can influence human attentiveness to ensure that Perception Queries are answered correctly and efficiently.

Experimental Design: Human-Study Web Experiment

We seek to answer the following questions:

- Does the presence of a *distractor task* (e.g. texting on a cell phone) negatively impact the performance of a human answering a perception query?
- If we ask *non-necessary* 'attention' perception queries to the human during AV operation, will the performance on *necessary* perception queries improve, if the user is distracted?

We design a web-based user study to evaluate how distractor tasks and attention queries affect users.

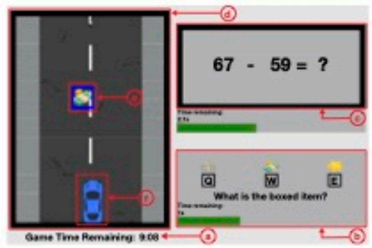

Figure 2: The human subject experiment design for studying the hand-off process. UI Text emphasized for clarity.

Experimental Details:

1. The AV operates with the user observing passively. Objects occassionally appear on the road - some dangerous, some safe.
2. On occassion, we ask a true perception query or an attention perception query about the objects on the road, and measure how quickly and correctly the user answers the query.
3. We use basic arithmetic questions to serve as our distractor task. The user is told explicitly to answer such questions to the best of their ability to ensure they are properly distracted.

Experimental Results and Takeaways

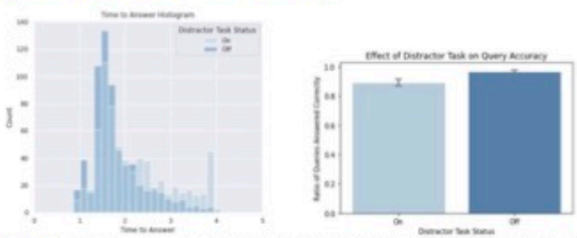

Figure 3: Response Times (RT) and Fraction of queries correctly answered (f) for the different conditions

Takeaways:

1. The presence of the distractor task negatively impacts the speed and accuracy of the primary task.
2. The presence of attention queries improves the speed and accuracy of the primary task.

Finding a Policy for when to ask Attention Queries

When is it most ideal to ask attention queries to maximize the accuracy and efficiency of true perception query responses?

We choose to model the human operator's attentiveness as a Partially Observable Markov Decision Process (POMDP):

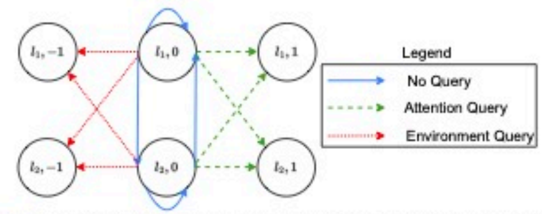

Figure 4: The state space for the POMDP formalizing the perception hand-off experiment. Shown here are the possible 1-step state transitions when starting in states $l_1, 0$ or $l_2, 0$ and under the different possible actions.

1. The POMDP is in state zero until either an attention query is asked by our policy or a perception query occurs.
2. Once a query is asked, states will transition per timestep, and into either an attentive or inattentive state at each timestep.
3. We cannot see these true states, we must infer them based on observations (the responses from the human operator.)
4. We *learn* the transition probabilities from our human study data.

From the learned POMDP, we infer an optimal policy for when to ask attention queries to maximize efficient and accurate responses.

Policy	Reward (R)	T_{resp}	$f * 100$	$\#a^{AIGA} : \#a^{PER}$
Learned	**15.52** $\pm$ 5.27	1.56 $\pm$ 0.05	**98.2** $\pm$ 2.8	**0.83** $\pm$ 0.12
No AIGA	11.29 $\pm$ 5.55	1.57 $\pm$ 0.07	92.8 $\pm$ 3.9	0
Random	11.78 $\pm$ 6.42	1.55 $\pm$ 0.04	95.4 $\pm$ 3.1	1.48 $\pm$ 0.14
Belief	13.83 $\pm$ 4.39	**1.54** $\pm$ 0.03	95.9 $\pm$ 2.8	2.9 $\pm$ 0.11

Our policy asks fewer queries and gets faster and more accurate responses from the human than baselines.