

Evaluation of Operator Interventions in Autonomous Off-road Driving

Jean Scholtz, Brian Antonishek, and Jeff Young
National Institute of Standards and Technology
Gaithersburg, MD 20899

ABSTRACT

We participated in field trials of a semi-autonomous vehicle. This gave us an opportunity to collect data on operator interventions. In this paper we present an analysis of why and how operators intervene and examine the efficiency of these interventions.

KEYWORDS: *Human-robot interaction, operator interventions, autonomous vehicles, field studies*

1. INTRODUCTION

One metric for autonomous systems is the frequency and type of operator interventions that is needed. While completely autonomous systems are the ultimate goal of some robotics developers, we believe that operator interventions will always be needed for effective teaming of humans and robots. Therefore, robotic architectures and human-robot interfaces should be designed to support effective and efficient operator interactions.

We participated in a data collection effort to measure operator interventions and understand the current state of autonomous off-road driving. While the number and frequency of the interventions are of interest to us, we are also interested in the types of interventions. These include what actions the operator performed during various types of interventions, length of time of classes of interventions and the time operators needed to gain situational awareness to effectively handle the intervention.

2. DATA CAPTURE

We were able to capture a number of different types of data during these experiments. We captured two log files from the robot platform; a second by second status log and an event log. We also captured the graphical user interface (GUI) of the operator control unit (OCU) and the video from the robot's camera during real-time tele-operation. A video of the operator was used to determine the actions of the operator: monitoring the behavior of the vehicle, -operating, or interacting via the GUI. We also captured a video from the safety vehicle allowing us to see the terrain that the robot was traversing.

Capturing the video and audio data for off-road driving involved designing and building a robust, self-contained, self-powered data capture system. First the vehicles were outfitted and wired for the audio and video capture: a bullet-style

camera was mounted in the front center of the vehicle and angled to view the operator in the front passenger's seat. From this vantage point it is possible to see the operator's interactions with the keyboard and -op joystick. A video cable was added to the operator's -op monitor's LCD video-out port. In addition, two condenser microphones were installed, one was located directly above the operator and the other was located above the observer's seat which is directly behind the operator. The output wires from the camera, LCD monitor, and microphones were routed to the rear seats of the vehicle to a centrally located area. Connections were also made to be able to tap into the VGA signal (computer screen) coming from the rear-mounted computer to the operator's monitor.

Once these connections were available, then the capture devices (mini-DV camcorders) had to be set up for recording. Because of the rugged terrain that the vehicles were traversing, the data capture equipment needed to be able to sustain frequent and heavy vibrations and jolting. With that requirement in mind, a ruggedized video box capture case was created to hold three camcorders and a scan converter (used to convert the VGA signal to a recordable video stream). We wanted the video box to protect and group together the capture equipment in one easy to access container. Three camcorder remote controls were added to the top of the video box. The remote controls allowed the data collector to start, stop, power off, and display the current tape counter of the camcorders. Also, a feedback monitor along with switches was added to be able to select and the video that was being recorded by each of the camcorders. These arrangements assured that the only time the data collector needed to open the video box was between runs when he had to change videotapes.

In order to power the camera and microphones for several hours at a time, a car jump-starting battery was installed and the camera and microphone power cables were modified to plug into the battery's cigarette lighter sockets.

Figure 1 shows the outside of the video box. Figure 2 shows the inside of the video box complete with protective foam. Figure 3 shows how the video box was connected to the OCU, the tele -op monitor, the bullet-style camera, and the operator's audio input. We used remote controls on the outside of the box to start and stop the recording. These are not shown on the schematic drawing.



Figure 1: The outside of the video box



Figure 2: The inside of the video box

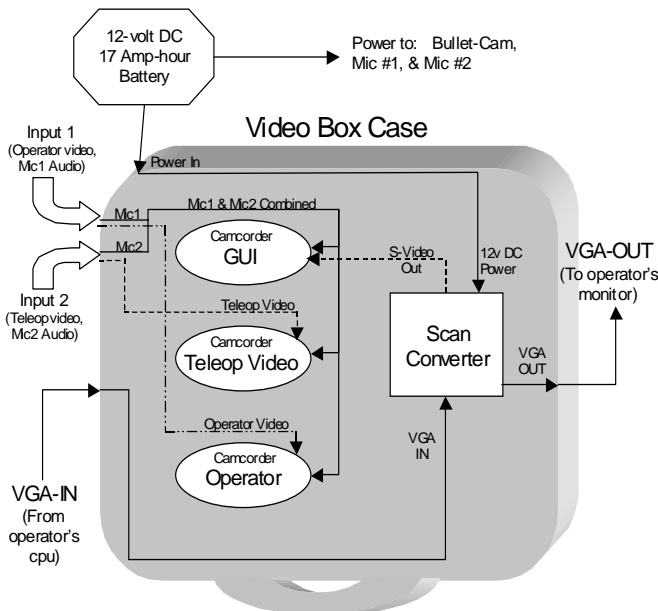


Figure 3: Schematic diagram of video box connections

3. EXPERIMENTS

The experiments were conducted in three different locations, with each experiment lasting approximately ten days. Two courses were setup at each location, a gold course and a black course with the latter covering more difficult terrain. Run distances of 500, 1000, and 2000 meters were conducted on each course. There were also two 7000 meter runs per course. Operators rode in a vehicle (OCU vehicle) behind the robot but were prevented from viewing the robot. There were two conditions on the runs: line-of-sight and non-line-of-sight. While the operator could never directly see the robot, in the line of site conditions the driver of the OCU vehicle could describe terrain conditions to the operator if requested.

Table 1 lists the trials in the one experiment analyzed in this paper and shows the breakdown of difficulty, length, and type. There were a total of 177 runs, including the 7000 meter runs. These trials were conducted between December 4th and December 13th, 2002 at Toelle Army Depot in Toelle, Utah.

Number of trials	500 Meters	1000 Meters	2000 Meters	7000 Meters	Totals
Gold/ LOS	12	21	12	1	91
Gold/ NLOS	12	21	12		
Black/ LOS	12	18	12	1	86
Black/ NLOS	12	19	12		
Total	48	79	48		177

Table 1: Number of trials by difficulty level and length

4. ANALYSIS

The operator control unit (OCU) in the study consisted of a graphical user interface (GUI), a keyboard, a joystick and a separate monitor for viewing real-time video from the vehicle. This is shown in Figure 4. The OCU used for the experiment is temporary and does not represent the interface that will eventually be implemented for the system. We were not concerned with the actual representation in our analysis but focused on what the operator had to do for different classifications of interventions. This analysis should be helpful in the redesign of the user interface.



Figure 4: The Operator Control Unit

The operator was able to view the status and event logs from the robot, the path of the robot, and the assigned way points on a map based display. The operator used the graphical user interface to switch from autonomous to tele-operation mode or to camera mode. In addition the operator could use a semi-autonomous mode and issue commands to the robot. The joystick was used to tele-operate the robot and to manually move the cameras. Switching between different cameras was accomplished in the graphical user interface. There was also a mode in which the camera automatically pans the area. The GUI is shown in Figure 5.

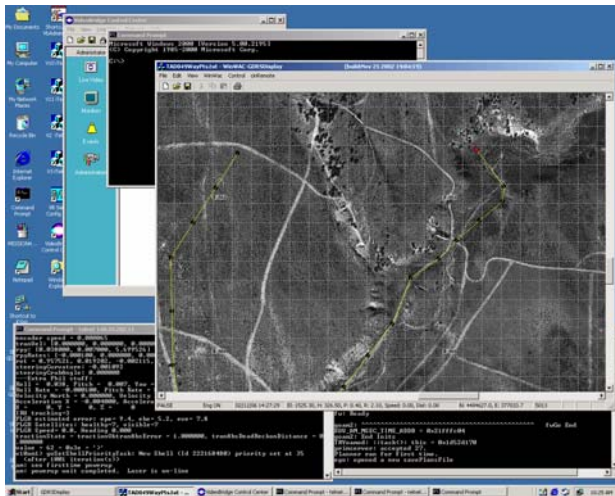


Figure 5: The Graphical User Interface

As the primary purpose of the experiment was to evaluate the autonomous performance of the robot, it was necessary to insure that all operators intervened consistently. The operator was only allowed to intervene for pre-defined situations. Table 2 lists these conditions. The operator could view these conditions from the GUI and was asked to note the reason for intervention each time. The reasons were broadcast to the test

administrator riding in the safety vehicle and captured in the trial log sheet.

Operator needed	OCU map display says it needs help Cannot back up Max backup attempts exceeded Other as displayed in the GUI window
Motion	The vehicle has not moved on the OCU map for more than 30 sec The vehicle loops back onto its path The vehicle stays in the same general area (20m) for more than 60 seconds
Speed	Vehicle speed stays below 0.2 m/s for more than 20 sec Speed is larger than assigned or less than -5m/s (negative)
Terrain	Vehicle pitch or roll is more than 20 degrees Traction is slipping and the vehicle is stuck Too rough
Communications	COMMS are lost with either of the two boards Loss of GPS fix
Obstacle	Bumper hit- starting to backtrack Bumper is stuck Bumper hit – cannot backup Stopping because the SIC says so Navigation confused Water is too deep
Path	The vehicle is more than 50 m off assigned path The vehicle is pointed in the wrong direction at start of mission The vehicle is about to cross into a restricted area
Plan failure	No good plan for a while Planner died
Mechanical	Engine too hot Ladar went down

Table 2: Reasons Operator was Allowed to Intervene

5. INTERVENTIONS

Our analysis is based on 177 trials. There were 35 trials in which interventions occurred ($\approx 1/5$ of the trials). And there were 44 interventions in total. Table 3 shows the types of interventions that occurred in the 35 trials.

Table 4 shows the occurrence of the interventions by course and length of trial. More interventions occurred on the Black course as compared to the Gold course. The 2000 meter trials had slightly more interventions than the 1000 meter trials even though there were more 1000 meter trials. Hence there were multiple interventions per trial on the 2000 meter runs.

Type of Intervention	Description	Number of Interventions of this type
Terrain	Interventions due to terrain obstacles (Traction and Tippage)	24
Movement	Interventions that affect movement (loopback, snailing, nomotion)	9
Plan Failure	Interventions due to a failure in the planner (noplan, deadplanner)	1
NeedOperator	Interventions where the Bot asked for operator assistance. This is a level by itself because the cause is not specified and requires operator interpretation to determine correct response.	10

Table 3: Types of Interventions that Occurred

Of Trials with Interventions [35]	%
Line-of-Sight missions	56.7
Non-line-of-sight missions	43.3
Black course	94.3
Gold Course	5.7
E-stops and Stops	14.3
500m trials	23.5
1000m trials	35.3
2000m trials	38.2
7000m trial	2.9

Table 4: Interventions percentages per course/ length of trial

Table 5 shows the number of interventions occurring by course and length of trial. There were more interventions on the 2000 meter trials even though there were significantly more 1000 meter trials.

Of # of Interventions [44]	%
Line-of-Sight missions	36.4
Non-line-of-sight missions	63.6
Black course	93.2
Gold Course	6.8
500m trials	15.9
1000m trials	36.4
2000m trials	45.5
7000m trial	2.3

Table 5: Interventions per course/ length of trial.

Table 6 shows the percentage of terrain, movement, and need operator classes of interventions that occurred on the course type and length of trial. All terrain interventions occurred on the Black course, which is to be expected. However, 46% of the terrain interventions occurred on the 2000 meter trails. All movement interventions also occurred on the Black course. The 'Needoperator' interventions were evenly dispersed between the distances.

	Terrain	Movement	Need Operator
Number	24	9	10
% in Black course	100	100	80
% on Gold course	0	0	20
% on 500 m trials	12.5	11.1	30
% on 1000 m trials	37.5	11.1	40
% on 2000 m trials	45.8	66.7	30
% on 7000 m trial	4.2	11.1	0

Table 6: Types of interventions occurring per course/length of trial

These tables give us a good picture of what types of interventions we can currently expect based on terrain difficulty. For purposes of looking at user interface requirements we are more concerned with what operators needed to do to respond to the various interventions. Because the operator was never able to directly see the robot or the terrain the robot was navigating, the first thing that operators did was to use the various cameras on the robot to gain situational awareness. Tables 7 and 8 show the breakdown of the interventions into the average time it took the operator to respond and the average time the operator needed to gain situational awareness. In these interventions, the operator could gain situational awareness by manipulating cameras and other sensors located on the robot or he could request information from the OCU driver if the trial was a line of sight trial.

Type of Intervention	Num of actual interventions	Average Time (sec)
Terrain	24	154
Movement	9	190
Plan Failure	1	29
NeedOperator	10	162
Totals	44	161

Table 7: Analysis of Interventions

Type of Intervention	Average SA time (sec)	Percent of intervention time
Terrain	25.1	16.3
Movement	28.4	15.0
Plan Failure	0	0
NeedOperator	31.8	19.5
Totals	26.7	16.6

Table 8: Further Analysis of Interventions

The highest average percentage of time required for situational awareness occurred for operator needed interventions. As previously described, these interventions required the operator to determine the cause and then to find a solution that would put the robot back on the correct path. Interpretation is costly so any error detection that can be done and reported to the operator will certainly help to lessen the time needed to identify the problem and gain the necessary situational awareness needed. The percentage of intervention time needed for situational awareness for the other interventions ranged from just over 10% to just over 17%. Given an average intervention time of 161 seconds, this amounts to almost 30 seconds per intervention. It might also be possible to perform the intervention more efficiently if situational awareness is improved.

6. INDIVIDUAL DIFFERENCES

In this experimental design the operator was also an independent variable. There were two operators and the design balanced the course type and length of trial that was assigned to them. In table 9 we break the tele-operation time down further into “Operator control” time. Tele -operation time is defined as the time from when the operator put the vehicle in tele-operation mode until he put it back into autonomous mode. The Operator control time refers to the actual time that the operator used the joystick to drive the vehicle. The remaining tele-operation time could be used to talk to the driver to gain situational awareness, to move the cameras around, or just to think about what to do.

Team	Number of Interventions	Average Op	Average SA	Average Operator Cntrl
Alpha	21	147	26	83
Bravo	24	268	45	117

Alpha: $N = 16$; Bravo: $N = 15$

**Note: only trials where both Op and SA times were recorded were used*

Table 9: Times for interventions by operator

We found significant differences between the two operators for operation time and situational awareness time. Tables 10 and 11 show the ANOVAs for these calculations. These differences could be due to skill level as

well as personality differences. More cautious operators may need more situational awareness before starting to operate and may move more slowly, stopping to check situational awareness as they drive.

	df	SS	MS	F
Team	1	294776.7	294776.66	9.51*
Error	31	960999.6	30999.98	
Total	32	1255776		

* $p < .01$

Table 10: Tele-operation time in relation to operator

	df	SS	MS	F
Team	1	5463.03	5463.03	9.44*
Error	27	15625.73	578.73	
Total	28	21088.76		

* $p < .01$

Table 11: Situation Awareness time in relation to operator

Another possible reason for the difference in times occurs from the operators’ different use of operating the cameras. When the robot vehicle reported a legal reason for operator intervention, the operator had two basic options: conduct a tele-operation (tele-op) or issue a semi-autonomous command. In the majority of cases, with only three exceptions, the OCU in charge decided to tele-op the robot. Once in tele-op mode the OCU had another three choices: issue a semi-autonomous command, take manual control of the cameras, or take manual control of the robot. In many cases the operator would use the cameras to gain situational awareness (see Table 4 for average times). It is interesting to note that camera use varied between the two operators used in the experiment. One operator always conducted semi-autonomous look-arounds (Bravo) while the other invariably used manual camera control (Alpha). Using the semi-autonomous mode may save on the operator’s cognitive load but it adds a fixed amount of time (≈ 18 seconds). This is an issue we intent to investigate further in subsequent analysis.

While we cannot draw any conclusions from a sample of two, this is certainly an issue to be further investigated. User interface designs must accommodate a wide variety of operational styles and we need to understand the range of styles.

7. OBSERVATIONS

We also were able to ride in the OCU vehicle during some of the trials and to observe the operator. Due to time constraints, we were only able to observe on several days of the experiment. However we were able to observe 26 trials in all. All of our observations were made during trials on the Black course.

In the majority of cases, the operator intervened only after the robot tried to extract itself from difficulties. Therefore,

there was some warning that the vehicle was having difficulty. Even though the operator couldn't see the actual terrain, he could gain some information about it from messages. For example, messages about traction and slipping told him that there was either a steep grade or snow or mud in the area.

In one instance the robot stopped and requested help because it was stuck. The operator couldn't see what the obstacle was using the video feed but when he intervened and issued a command to back up, he saw a tree root that may have caused the problem. In this instance there was no warning – the robot simply stopped moving.

In several trials the robot had difficulty getting up a hill or out of a ravine. When the operator took control he also had to make several attempts, trying different paths and different angles. In several of these instances the operator had to stop driving and use different camera view to gain a better understanding of an alternative path. In some cases, he had to switch to a back view prior to backing up. The operator also consulted the map displayed on the OCU for a better understanding of the terrain.

8. CONCLUSIONS AND FUTURE DIRECTIONS

From a numbers perspective, terrain caused the most operator interventions. However, the plan of action was straightforward. The operator just needed to find a more suitable path. This was accomplished by a series of actions of gaining situational awareness using camera views, backing up, more situational awareness, and moving ahead. These actions were repeated until a path was found for autonomous navigation. 'Operator Needed' interventions were the most

costly in terms of time as the operator had to decide what the problem was and what actions to take. A difficulty is that the operator has to stop moving in order to manually operate the cameras as the same control is used for both navigation and camera manipulation. Coupling actions such as backing up with automatic camera panning might be a possibility.

We did see that automation of subtasks does not necessarily produce the most efficient design. This was clear in looking at the automated panning of the camera that took a fixed 18 seconds with manual operation that often took less time.

Operators in these trials were free to watch the OCU and to monitor the robot at all times. In a real operational environment, operators may have multiple robots to monitor or other tasks to accomplish. This may increase the amount of time that is needed to gain situational awareness and to respond to a robot's request for help.

We also see the need to support individual differences in the user interface. In this early GUI there was little choice in configuring the display or in the ways the operator could interact.

We have two other sets of trials that we intend to analyze in a more focused fashion. By looking at data collected from realistic trials we have gained a much deeper understanding of what issues need to be addressed in the HRI. We plan to design some controlled experiments to further investigate these issues.

9. Acknowledgements

This work was supported by the DARPA MARS program and the Army Research Laboratory.