2020 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY SYMPOSIUM
MODELING, SIMULATION & SOFTWARE (MS2) TECHNICAL SESSION
AUGUST 11-13, 2020 - NOVI, MICHIGAN

PREDICTIVE DIPLAYS FOR HIGH LATENCY TELEOPERATION: EXTENSIONS AND EXPERIMENTS

Mark Brudnak, PhD[1]

[1]CCDC-Ground Vehicle Systems Center, Warren, MI

ABSTRACT

This paper presents developmental and experimental work beyond the initial presentation of the predictive display technology. Developmental work consisted of the addition of features to the predictive display such as image subsampling, camera stabilization, void filling and image overlay graphics. The paper then describes two experiments consisting of twelve subjects each in which the predictive displays were compared to both the zero latency case (baseline) and the unmitigated high-latency cases (worst case). The predictive display was compared using four objective performance and activity measures of mean speed, lateral deviation, heading deviation and steering activity. The predictive display was also assessed using subjective measures of workload and usability.

Citation: M.J. Brudnak, "Predictive Displays for High Latency Teleoperation: Extensions and Experiments", In *Proceedings of the Ground Vehicle Systems Engineering and Technology Symposium* (GVSETS), NDIA, Novi, MI, Aug. 11-13, 2020.

1. INTRODUCTION

One of GVSC's top objectives is to lead the Army and DoD Ground Vehicle Community in the research, development, engineering, demonstration and fielding of Unmanned Ground Vehicle (UGV) systems. Teleoperation is a near-term technology which has the potential to be a quick-win for UGVs. It is and has already been employed in the cases of small UGVs performing Counter Improvised Explosive Device (C-IED) and Explosive Ordinance Disposal (EOD) missions. In these cases the speeds are low and the operator is usually in relatively close proximity to the S-UGV yielding low latency and responsive control. Long-distance teleoperation on the other hand, introduces significant latency which degrades the operator's ability to drive/control the vehicle. As such, the mitigation of this latency is the most fundamental barrier to achieving teleoperation under high latency. In preceding work [1], the author presented the foundation and implementation of predictive displays as a proposed technology to mitigate this latency as illustrated in Figure 1. In that work the author developed the overall approach, developed the algorithms, implemented the method and presented a preliminary set of pilot data.

This paper builds on the results of the foundational work which has since been awarded a patent [2]. In this work the author will present the

results from a formal experiment across 12 subjects and three different latency conditions. This experimental result shows statistically significant improvement realized by the predictive display in four different measures of performance to include average speed, path deviation, heading deviation and steer effort. Subjective results were also collected using the NASA TLX [3] workload instrument as well as the system usability scale (SUS) [4]. Both these instruments show significant benefit gained by the predictive display. The experiment also collected open comments from the subjects which were used to improve the predictive display technology in several important ways: (1) Image subsampling was added to mitigate areas in the display for which there were no source pixels, (2) void filling was added to fill image voids when they do occur, (3) camera/image stabilization was added to mitigate the negative effects of vehicle pitch, (4) graphical overlays were added to give the operator visual cues regarding the predicted trajectory of the vehicle. The technical implementation of these additions will be presented in this paper.

With these additions, a second experiment was conducted to formally assess the benefits. This experiment also included 12 subjects and evaluated conditions under one latency condition (1 second) with and without the overlay and with and without the predictive display present. (The other features were present for all trials.) The experiment was conducted on the same terrain but introduced the additional feature of gates which the subject was directed to pass through. These gates were situated in all areas of the course. Again the presence of the predictive display showed significant improvement in all of the measures. Additionally, the overlay further improved heading error significantly. As with the first experiment subjective measures of workload and usability were collected, which indicated a significant improvement was observed by use of the predictive display and overlay.

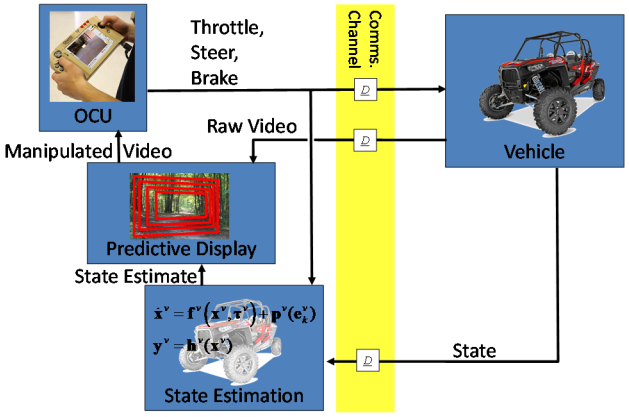

Figure 1 Block diagram of predictive display and state estimation approach.

2. OVERVIEW OF THE PREDICTIVE DISPLAY

The top-level approach to the predictive display system is illustrated in Figure 1. It is presented in detail in the prior paper [1], but this section briefly describes it at a high level. It is composed of a State Estimator (SE) and the Predictive Display (PD).

The SE functions in two modes simultaneously: feedforward and feedback. In feedforward mode the SE accepts the driver commands in the form of throttle, brake and steer (T, B, S) and (using the current state) predicts an immediate response. Since the SE runs at 100Hz, it has a very responsive reaction to the driver commands, on the order of 10 ms. This feedforward mode is based on a very simplified form of the vehicle dynamics with a prediction horizon on the order of the round trip delay, $2D$. If it were only operating in feedforward mode the states would drift from the actual values, so there is a correction term which keeps the states (mostly vehicle speed and yaw rate) roughly in-line with the states of the vehicle. Finally, the SE maintains a record of states so that it can look back to a past state to obtain the relative motion between the two.

The PD consumes the predicted state information from the SE as well as the raw video stream from the vehicle. The PD operates under the assumption that the scene observed at time t will be very similar to that observed at $t - 2D$. It therefore asks the SE

to give the difference in position from time $t - 2D$ to time t. It then uses this information to manipulate the latest video frame (which is $2D$ seconds old) to give a best estimate of what the operator would see as if there were no delay. This then is passed to the OCU for display to the operator.

3. EXTENSIONS

In the work subsequent to the initial development, the author added four features to the predictive display to address feedback from the first user experiment. These features are (1) Image subsampling, (2) void filling, (3) camera/image stabilization, (4) graphical overlays. These features are described in the following subsections.

3.1. Image Subsampling

In the first user experiment, comments and feedback centered on the presence of a void when turning (see Figure 2). This void is created when there are no source pixels to render the camera view from the updated perspective (see Figure 3). The void, although not present for slight turns, is present for tight turns. Image subsampling changes the camera aspect ratio to a wider aspect ratio to provide source pixels which can fill the void on turns. This concept is illustrated in Figure 4. In this particular implementation, the source image was captured at 16:9 ratio and transformed using the methods developed in [1]. After transformation, the image was cropped down to 4:3 ratio prior to presentation to the operator. This image manipulation was performed in OpenCV. This eliminated the void on most normal turns, but still was a problem on some of the sharpest turns. Prior to the implementation of image subsampling, the void could fill more than 50% of the operator's view. After image subsampling the void only fills approximately 20% in the worst case. Nonetheless, the presence of the void is highly distracting and calls attention to the fact that the perspective is being changed. To mitigate the occasions when the voids did occur a method called *void filling* was used to make them less visible to the operator. This is described in the next section.

Figure 2 Void problem when no source pixels.

3.2. Void Filling

The OpenCV perspective transform defaults to black when there are no image pixels in the source for the destination. Void filling attempts to fill this area with content that is *less* visible. There are multiple strategies to doing this. The easiest (perhaps) is to replace the black with another color which is more natural i.e. green or blue for the

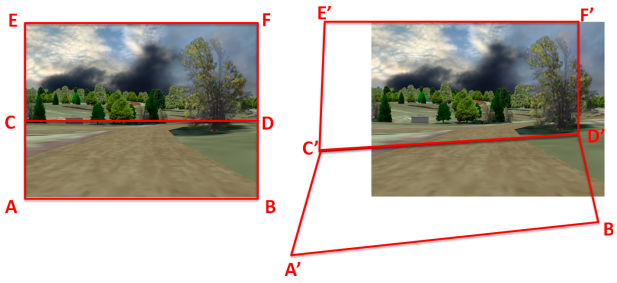

Figure 3 Illustration of original image sampling.

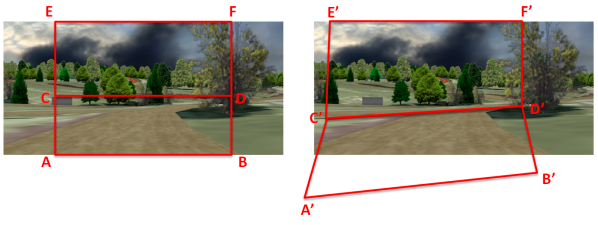

Figure 4 Illustration of image subsampling.

ground or sky, but still a monolithic color would likely be equally distracting. The method chosen was to use the source image pixels to determine the fill color. The method executes the fill on a pixel-line by pixel-line basis. The pixel-lines are filled using a linear interpolation between two colors. Color 1, c_1, is the color of the pixel immediately adjacent to the black void in that row and color 2, c_2, is the average color of all non-black pixels in the line. Each black pixel is then replaced by a linear mixture of these two colors. The mixture is 100% c_1 adjacent to the valid pixels and 100% c_2 at the edge of the image. This interpolation creates gentle transitions which are barely noticeable for small sections of void and even for substantial sections of void are not distracting if outside of the operator's primary vision. An illustration of the void filling feature is shown in Figure 5.

3.3. Image/Camera Stabilization

In the first experiment it was evident that momentary vehicle pitch affected the way that the predictive display worked. This was especially true for the pitching that occurs when the vehicle brakes aggressively. This pitching with a fixed camera changes the horizon in the image which the algorithm assumes is fixed. This causes the predictive display to give the illusion of going slower during a downward pitch and faster during an upward pitch. This was particularly noticeable when slowing down for a turn from a relatively high rate of speed. To address this phenomenon, the camera was stabilized on the robotic platform. Because camera stabilization is a solved problem, we elected to model perfect stabilization by removing the pitch from the camera. We were able to do this because experimentation was conducted on flat terrain. For hilly terrain, a camera stabilization control system would have to be implemented.

3.4. Graphical Overlays

The predictive display component compensates for the round trip delay by manipulating the image. This is intended to bring the operator from a 'past' perspective up to a 'present' perspective. The overlay presents to the operator what the vehicle will do in the near-term future, thus giving him a 'future' estimate from his/her 'present' perspective. The overlay is illustrated in Figure 6. In that figure, there are two elements, (1) the future path predictor and (2) the lateral acceleration limit lines. The future path predictor is designed as a semi-transparent overlay which presents the following key information to the operator. First it presents the estimate of the future path of the vehicle as represented by the green center line. Second it presents the swept path of the vehicle by the width of the green transparent area. Third, it presents the

Figure 5 Void filling using linear interpolation between adjacent color and average color on pixel-line basis.

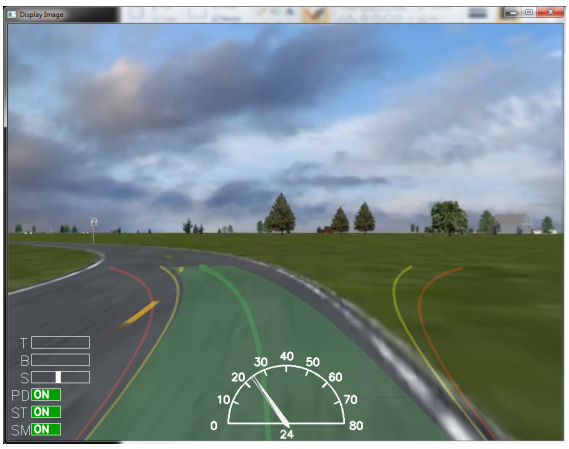

Figure 6 Illustration of the overlay for the predictive display.

expected distance traveled in a given time horizon. This is represented by the amount that the overlay is stretched forward. As illustrated, the green area stretches forward by 2 seconds into the future. The lateral acceleration limit lines represent paths of maximum curvature to stay below particular lateral acceleration limits given the current speed. In Figure 6 the yellow line represents a 0.2 g limit and the red line represents a 0.3 g limit.

The overlays are developed from the perspective of the present frame. They are defined by coordinates drawn at the Z=0 plane (see Figure 7) and then transformed to the camera frame "2" and then projected onto the image plane and then drawn as transparent shapes in OpenCV. The future path is defined in a differential sense as follows:

$$x(0) = y(0) = 0$$
$$\Delta x(0) = 0, \Delta y(0) = v\Delta t, \Delta\theta(0) = v\Delta t$$

Then the positions are updated based on the following update rules.

$$\begin{bmatrix}\Delta x(i+1)\\ \Delta y(i+1)\end{bmatrix} = R(\Delta\theta)\begin{bmatrix}\Delta x(i)\\ \Delta y(i)\end{bmatrix}$$
$$\begin{bmatrix}x(i+1)\\ y(i+1)\end{bmatrix} = \begin{bmatrix}x(i)\\ y(i)\end{bmatrix} + \begin{bmatrix}\Delta x(i+1)\\ \Delta y(i+1)\end{bmatrix}$$

where $R(\Delta\theta)$ is the 2-D rotation matrix associated with $\Delta\theta$. The values are the translated to the camera frame using

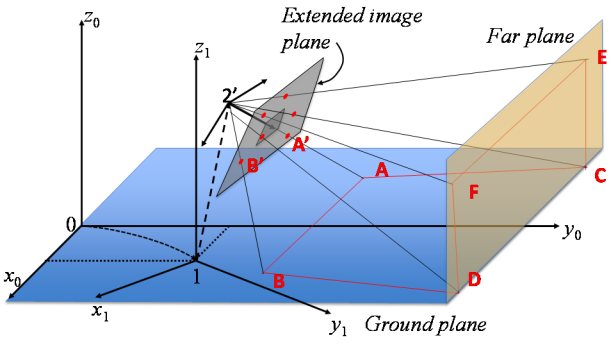

Figure 7 Illustration of frames used in the predictive display.

$$\begin{bmatrix}x_c\\ y_c\\ z_c\end{bmatrix} = \mathbf{H}_1^{2^{-1}}\begin{bmatrix}x(i)\\ y(i)\\ 0\end{bmatrix}$$

where $\mathbf{H}_1^2$ is the homogeneous transformation from the present frame to the camera frame [1]. And then converted to raster coordinates. Lines are then drawn on the screen using these coordinates.

The left and right extents of the swept path are likewise developed using the same procedure except that they start with $x(0) = \pm W_v/2$ where W_v is the width of the vehicle.

The lateral acceleration limit lines are drawn using the same procedure except that

$$\Delta\theta(0) = \pm\frac{a}{v^2}$$

where a is the limit acceleration and v is the vehicle velocity.

4. SIMULATION SETUP

The above design was implemented on two workstations running Microsoft Windows 7 Professional 64 bit as shown in Figure 8 and further described in the prior paper [1]. They were connected via a gigabit Ethernet LAN and passed all relevant information via the UDP/IP protocol. Although these computers have the power to run everything on one machine, they were run on two for a few reasons. First, because this is a simulation, information is very accessible. By running them on two different machines, it is assured that only valid/relevant state is being shared (i.e. no access to privileged information). The second is that because information is passed over a physical network, it provides well-defined "choke points" to monitor and/or control network

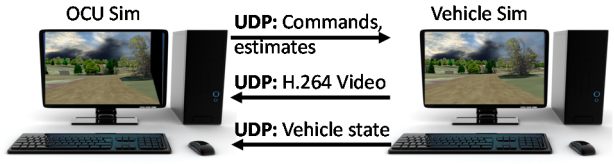

Figure 8 Setup of the computers.

behavior. This gives ample opportunity to control network performance via packet forwarding or the insertion of a network emulator.

5. EXPERIMENTAL RESULTS

The system which was described in the preceding section, was run experimentally in a within-subject experimental design and these results are discussed in this section. Each year (2016 & 2017) of the project conducted a separate experiment. Some aspects of the experiments were common and we will discuss the common aspects first.

5.1. The Terrain

The experiments were run on a flat terrain database which consists of four different types of tiles which are 200 m x 200 m each. These tiles consist of a straight section, a 75 m radius right turn, a 75 m radius left turn and an 'S' turn (which has six curves of 20 m radius each). A wire frame diagram of the terrain is shown in Figure 9. The course was negotiated in a counter-clock-wise direction, starting on the far right (east) portion facing up (north). The operator's goal is to stay in the center or right lane of the course which has two lanes and a width of 8 m. Speed on the course is regulated by means of speed limit signs which are as follows. Preceding each right or left turn with 75 m radius, there is a 30 mph (48 kph) speed limit sign, preceding each 'S' turn there is a 15 mph (24 kph) speed limit sign and preceding each straight section of length of two tiles there is a 40 mph (65 kph) speed limit sign. The course is approximately 5.8 km (3.6 miles) long and takes approximately 10 minutes to complete one circuit around the course. The advantage of a tile-based terrain is that a participant is subject to the same circumstance at multiple times through the course which multiplies the statistical sample size if one wishes to examine particular events. In particular this terrain contains 14 straight sections, 9 left turns, 5 right turns, and 2 'S' turns.

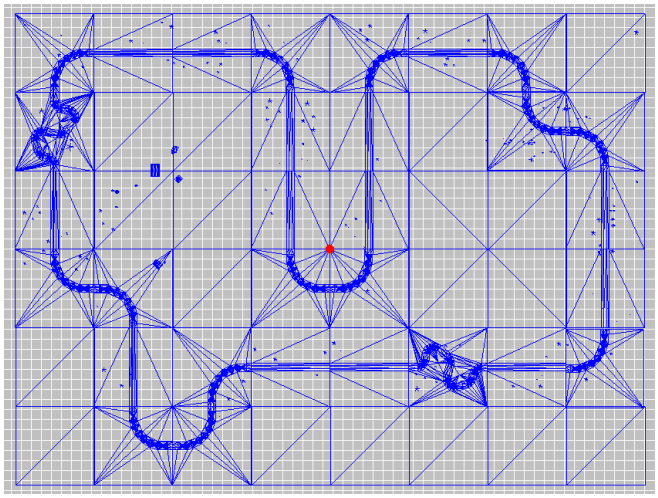

Figure 9. A wire-frame top-down view of the terrain database.

5.2. The Operator Station

The operator station consisted of a typical driving simulation setup and is illustrated in Figure 10. The apparatus consists of a seat, control and a display all setup in a static environmentally controlled laboratory. The controls consist of a driving game controller consisting of a steering wheel with force feedback, throttle, and a brake. The particular controller used was a Thrustmaster T300 RS. The display is a typical flat panel monitor (Dell U3014t) set at eye level. The control and display were connected to the OCU computer which reads the

Figure 10 Operator station consisting of a seat, display and controller.

controller and displays the view. In the second experiment, the buttons were also used to enable/disable the overlay, void filling, and stabilization.

6. EXPERIMENT 1

The first experiment was conducted from January to March 2017. Its objective was to evaluate effectiveness of the baseline predictive display under three different latency values of 0.5 sec, 1 sec and 2 sec against a baseline case of zero additional latency. Each case with non-zero latency was run with the predictive display both on and off. The experiments consisted of seven configurations as shown in Table 1. Configuration 0000_OFF with no additional latency and no predictive display is intended to represent the best possible scenario and should produce the best case. Configurations 0500_OFF, 1000_OFF and 2000_OFF on the other hand are intended to represent the baseline cases where latency is present but it is not actively mitigated. This may be thought of as the worst case baseline performance from which to improve for three different delay conditions. Configurations 0500_ON, 1000_ON and 2000_ON maintain the same latency as the "OFF" configurations however, the predictive display is added. Each configuration was run in the order prescribed in Table 2 and within this ordering two repetitions were executed one with the PD turned on and one with it off. Of the three latency conditions, the on/off sequencing was determined by Table 3. During the in-briefing, the subject was asked to choose A-F for the latency ordering and A or B for the PD ordering. As subjects chose these values, they were excluded as options from subsequent subjects to assure that all sequences were run. The zero latency condition was always run first and was only run without the PD. Subjects were allowed to practice all seven conditions prior to the start of experimental runs. The experiments were conducted on 12 subjects ranging in age from 18 to 55 years. The distribution of ages are illustrated in Figure 11.

Table 1 Design for experiment 1.

	Latency [sec]			
PD	0	0.5	1.0	2.0
OFF	0000_OFF	0500_OFF	1000_OFF	2000_OFF
ON		0500_ON	1000_ON	2000_ON

Table 2. Experiment 1 Latency Ordering

Latency Ordering				
A	None	Low	Medium	High
B	None	Low	High	Medium
C	None	Medium	Low	High
D	None	Medium	High	Low
E	None	High	Medium	Low
F	None	High	Low	Medium

Table 3. Experiment 1 PD ordering

PD Ordering			
A	OFF/ON	ON/OFF	OFF/ON
B	ON/OFF	OFF/ON	ON/OFF

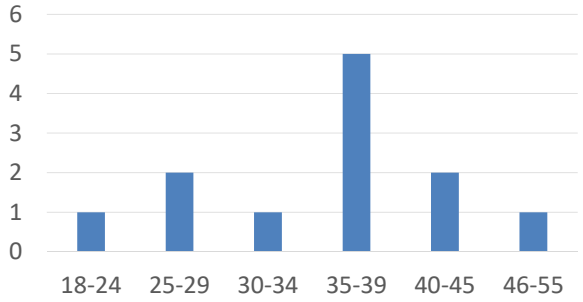

Figure 11. Experiment 1 Distribution of subjects' age.

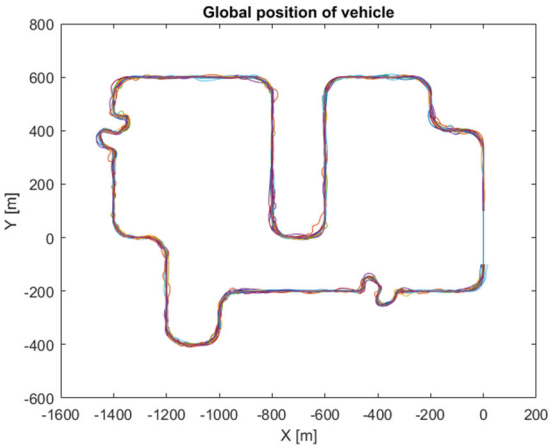

Figure 12 Experiment 1, Global position of the vehicle for all runs.

Data collected during the experiments consisted of four data logs associated with the OcuSim (see Figure 8) and the three processes run on the VehicleSim. On the OCU the data were logged at 100 Hz and on the VehicleSim the data were logged at 1,000 Hz. Most of the results hereafter presented speak to the effectiveness of the predictive display vs. the unmitigated case. The traces of the x-y position of these runs are illustrated in Figure 12. Since the section of the terrain where the simulation starts and ends has inconsistencies such as start time, stop location, etc., the data analysis omits the first portion of the run (i.e. analysis starts with the second tile), likewise, the same tile is omitted from the end of the analysis, so when the vehicle enters this tile data analysis stops. The key metrics of interest in teleoperation are speed and accuracy

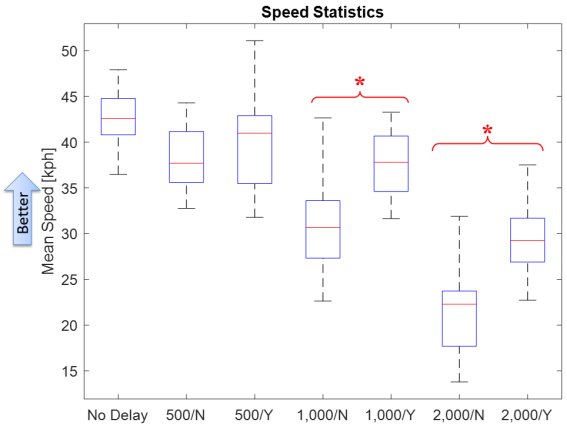

Figure 14 Experiment 1, Statistical distribution of mean speeds.

where both should be maximized. (These are normally mutually opposed objectives.)

In this analysis, accuracy is tracked using two metrics, namely path deviation and heading deviation (these are error metrics, so lower is better). The target path is not marked on the road but is regarded as the center of the 4 m wide right lane. The road way is defined by points along the center of the road, and the desired path is 2m to the right of the center of the road. Path heading error is measured as the angle difference between the path tangent and the vehicle heading direction. Speed is measured as the average speed along the course.

The mean speed for each of the configurations is shown in Figure 13 and the scatter plots of speed per condition are shown in Figure 14. As can be observed in Figure 13 some of the data are omitted from the analysis, namely subject 8 for the 1,000

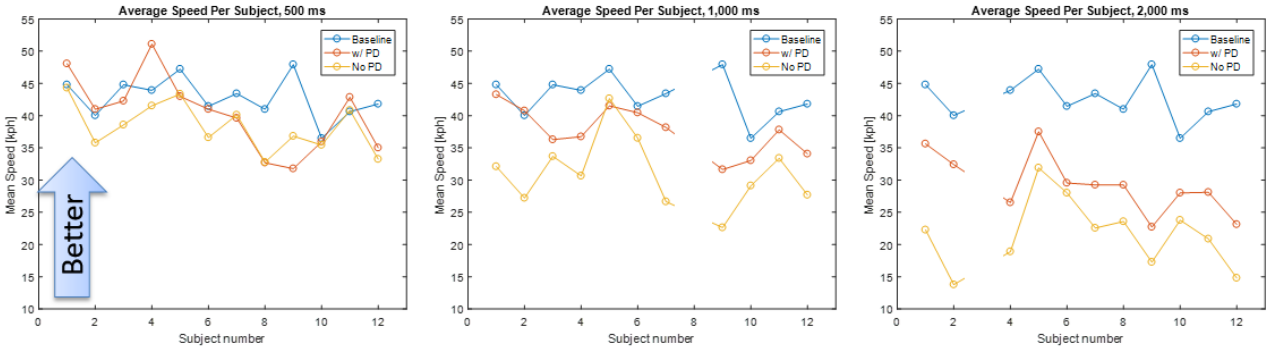

Figure 13 Experiment 1, Improvement in speed.

ms delay case and subject 3 for the 2,000 ms delay case. In both cases one of the respective subject's runs was lost so the respective PD and no-PD cases were removed from the analysis. In this case the analysis uses 12 sets of runs for the 500 ms case and 11 sets of runs for the 1,000 ms and the 2,000 ms cases. With respect to speed, it is clearly seen that the speed degrades as the latency increases. It can also be seen that the predictive display has significant impact on the speed achieved for the 1,000 ms and 2,000 ms cases but not for the 500 ms case.

The path deviation metric is computed as follows:

$$\int_0^L |e_p(s)| ds$$

where L is the length of the course under analysis (approximately 5,600 m) and $e_p(s)$ is the path deviation as a function of path length s. Figure 15 and Figure 16 address path accuracy. Figure 15 shows path deviation (which is the vehicle's distance from the center of the lane) for each subject under each condition. This metric obviously directly relates to accuracy and gives an intuitive sense as to how well the operator is keeping to his lane. Given that the vehicle is narrower than the lane, it is possible that the operator has a non-zero deviation but is still in the lane. Figure 16 illustrates the statistical differences between the experimental conditions as it relates to path deviation. There it is observed that there is no

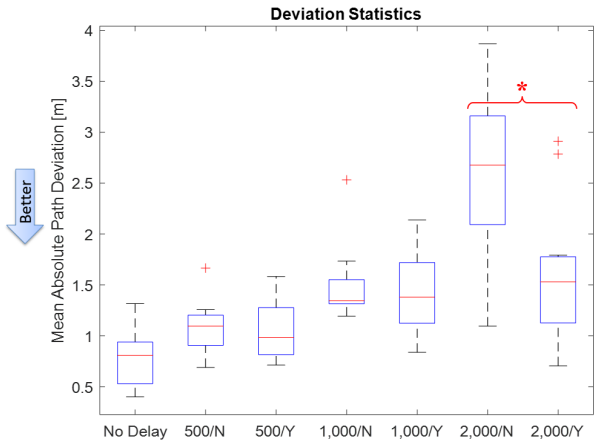

Figure 16 Experiment 1, Statistical distribution of path deviation.

significant difference observed as a result of the predictive display for the 500 ms and 1,000 ms latency cases, however, the 2,000 ms latency case did show significant differences.

The metric for heading error is computed as

$$\int_0^L |e_h(s)| ds$$

where $e_h(s)$ is the heading error. Figure 17 shows the heading error which indicates how well the operator is maintaining vehicle direction along the path. It is also an indicator of the ability (or inability) of the operator to attain and maintain a desired directional heading. Figure 18 shows the statistical distribution of heading error for all experimental conditions. There it can be seen that unmitigated latency adversely affects the heading

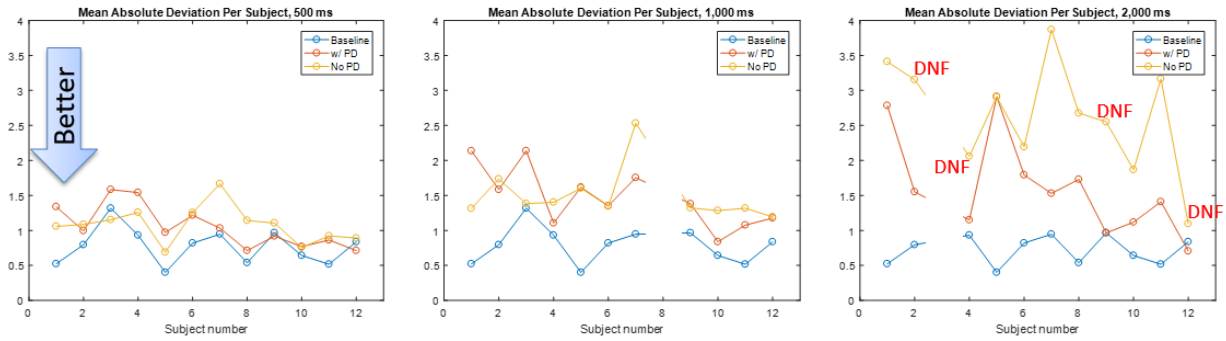

Figure 15 Experiment 1, Improvement in path deviation. (DNF = Did not finish)

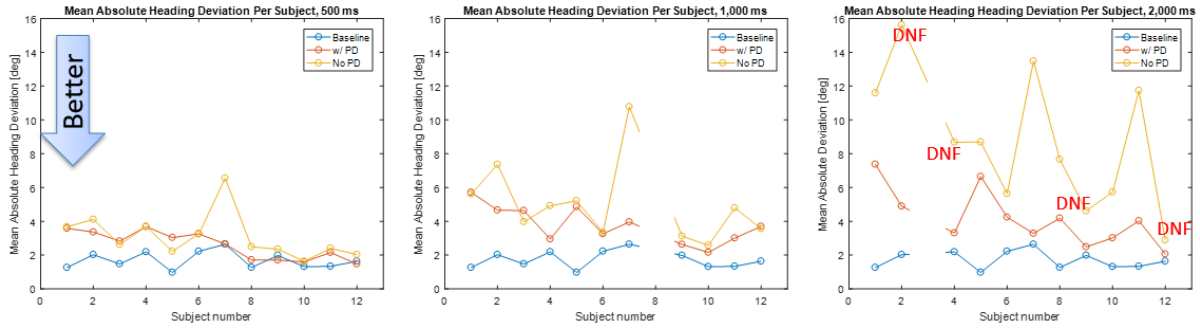

Figure 17 Experiment 1, Improvement in heading deviation.

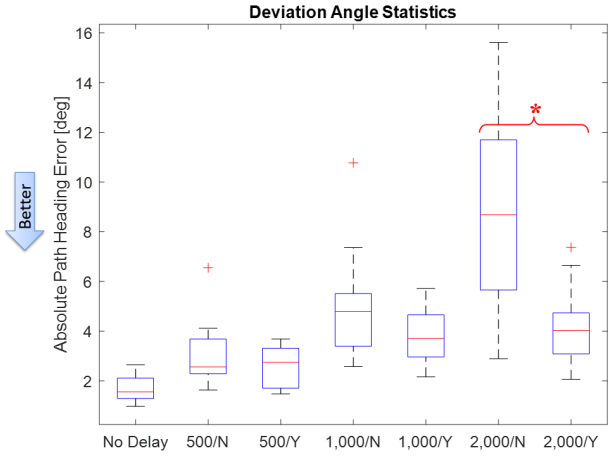

Figure 18 Experiment 1, Statistical distribution of heading deviation.

error. This correlates well with observations of subject operating under moderate to high latency, namely that under latency, overcorrection of heading and lateral errors is the primary challenge. Given that, it can be seen that the introduction of the predictive display substantially curbs this over-correction behavior as seen by the relatively slow growth of heading error as latency increases. Of the three latency conditions, the 2,000 ms case is the only one which shows significant improvement of heading error provided by the predictive display.

Finally a fourth metric to assess the quality of teleoperation performance is driver activity. Since the fundamental challenge is directional control, the steering command provides a suitable metric. Figure 19 shows the mean absolute steering command for all of the latency and predictive display conditions. Figure 20 shows the statistical distribution of steering angle activity. There we can see that increasing unmitigated latency correlates with increased steering effort while steering effort with the predictive display on remains relatively constant. This highlights another primary benefit of the predictive display in that it reduces the amount of effort required to operate the vehicle. Performing an ANOVA analysis reveals that the predictive display significantly improves the steering effort in the 1,000 ms (p=5.5%) and the 2,000 ms latency conditions.

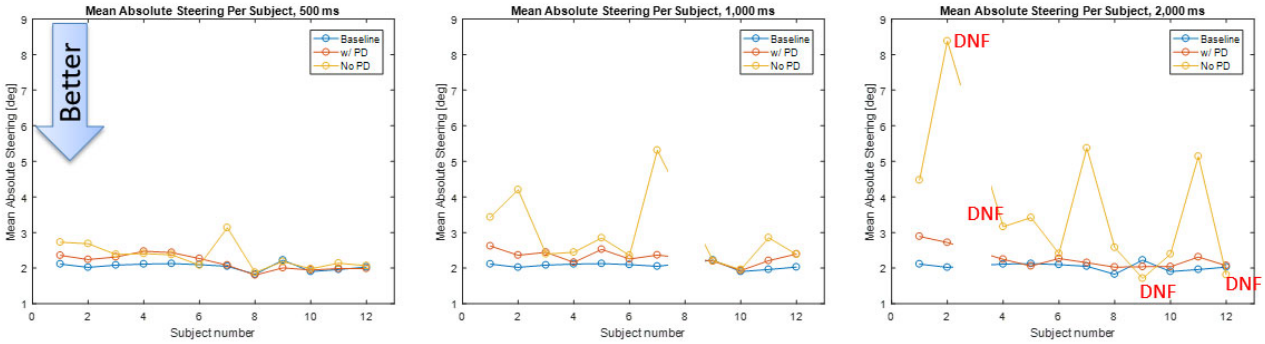

Figure 19 Experiment 1, Improvement in steering effort.

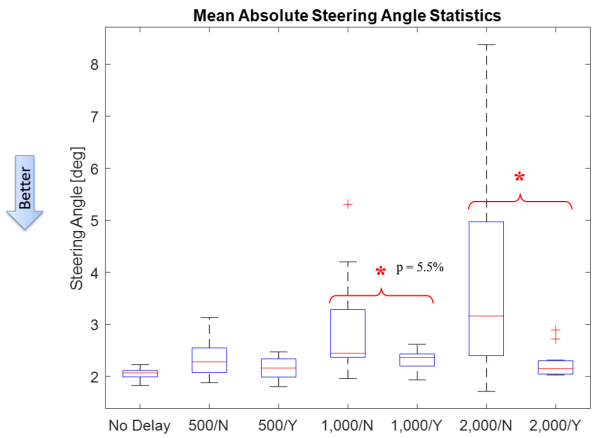

Figure 20 Experiment 1, Statistical distribution of steering angle.

The numerical means of the four performance parameters are shown in Table 4. Additionally the table includes a percent improvement which is attributable to the predictive display. This is computed as the relative distance from the unmitigated case normalized to the difference between the unmitigated (i.e. worst case) and the baseline (i.e. best case). This computation is as follows:

$$\%\text{Improvement} = \frac{|\text{PD} - \text{noPD}|}{|\text{Baseline} - \text{noPD}|} \times 100\%$$

which will be 0% if the PD performs like the no PD case (i.e. with latency and without PD) and will be 100% if the PD performs like the baseline (i.e. no delay, best case). So the closer to 100% the better. In Table 4 we can see improvement in the mean performance as influenced by the predictive display in all of the metrics for each latency case. Speed drops significantly in the 1,000 ms and 2,000 ms cases and the predictive display helps recover 55.7% and 36.2% of the lost speed due to latency. Path deviation also degrades from 0.77 m in the base line case to 2.63 in the worst case. The predictive display does not appear to improve the lane deviation in the 500 ms and the 1,000 ms cases, but for the 2,000 ms case path deviation is improved from 2.63 m to 1.60 m, an improvement of over 1 m or 55.2%. The path accuracy metric directly counterbalances the speed metric in that the faster one attempts to drive the more likely they are to deviate from the path. Although there is substantial scatter in the data, this trend is most evident in the higher latency cases as can be seen in Figure 21. Table 4 shows heading deviation immediately degrades even with the smallest amount of latency (i.e. 500 ms) which in the unmitigated case rises from 1.7 deg to 3.07 deg and for the highest latency mean heading deviation rises to 8.76 deg for the unmitigated case. This heading deviation is the clearest indication that in the unmitigated case, latency causes the operator to overshoot his/her directional corrections, producing a slalom or back and forth driving behavior. The introduction of the predictive

Table 4 Experiment 1, Mean performance of predictive display.

	Baseline	500 ms		1,000 ms		2,000 ms	
	w/o PD	w/o PD	w/ PD (% impr.)	w/o PD	w/ PD (% impr.)	w/o PD	w/ PD (% impr.)
Speed [kph]	42.8	38.3	40.4 (46.1%)	31.1	37.6 (55.7%)	21.6	29.3 (36.2%)
Path Dev [m]	0.77	1.08	1.05 (8.7%)	1.49	1.47 (3.4%)	2.63	1.60 (55.2%)
Heading Dev [deg]	1.70	3.09	2.59 (35.8%)	5.02	3.78 (37.5%)	8.76	4.14 (65.4%)
Steering [deg]	2.05	2.34	2.16 (61.3%)	2.95	2.32 (70.0%)	3.71	2.26 (87.3%)

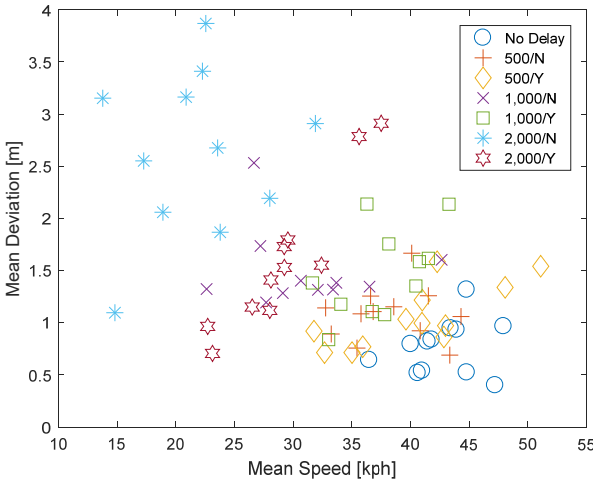

Figure 21 Experiment 1, Scatter plot of mean deviation vs. mean speed.

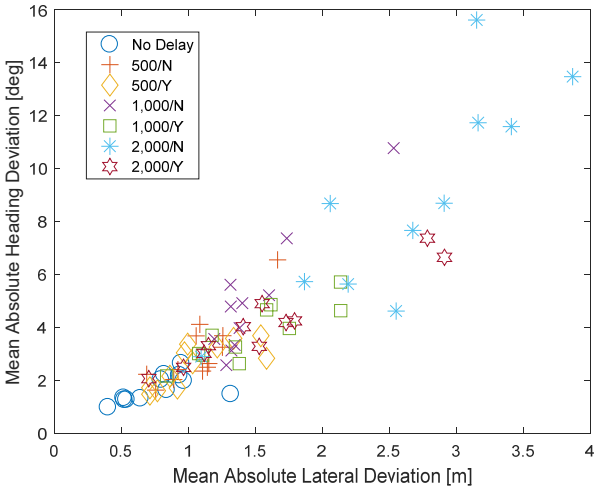

Figure 22 Experiment 1, Scatter plot of heading deviation vs. lateral deviation.

display improves this heading error by 35.8% and 37.5% for the 500ms and 1,000 ms latency cases respectively, but for the 2,000 ms case, heading error is improved by 65.4%. This provides evidence that fundamental challenge with teleoperation under latency is directional control and that the predictive display helps mitigate directional control issues exhibited by the vehicle operator. Figure 22 illustrates the strong correlation of heading deviation and lateral deviation. The fourth metric further highlights this phenomenon. The most marked improvements are seen in the steering effort metric since all of the

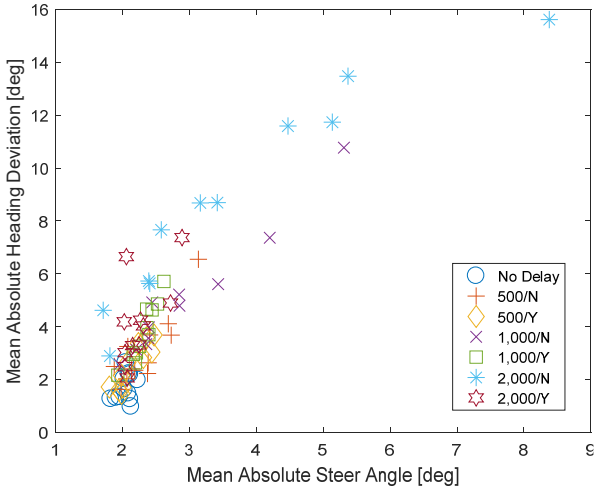

Figure 23 Experiment 1, Scatter plot of angular deviation vs. steer effort.

unmitigated latencies require a notable increase in steering effort to maintain directional control. By examining driver behavior, we get a measure of how much the operator is working to control the vehicle. It is observed that steering effort is greatly improved by the introduction of the predictive display by 61.3%, 70% and 87.3% for the three latency conditions respectively. We therefore believe that not only does the predictive display improve performance, but it also improves workload. Figure 23 shows the strong correlation of angular deviation and steer effort.

In addition to the performance measurements, the study also measured operator subjective workload using the NASA TLX instrument. The measured levels of workload for each of the conditions is as shown in Figure 24. There we see that the introduction of the predictive display lowers workload under all latency conditions.

Finally, we measured system usability [4]. The results are shown in Figure 25.

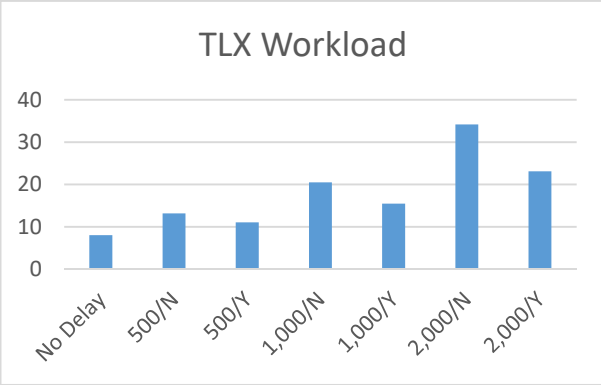

Figure 24 Experiment 1 work load.

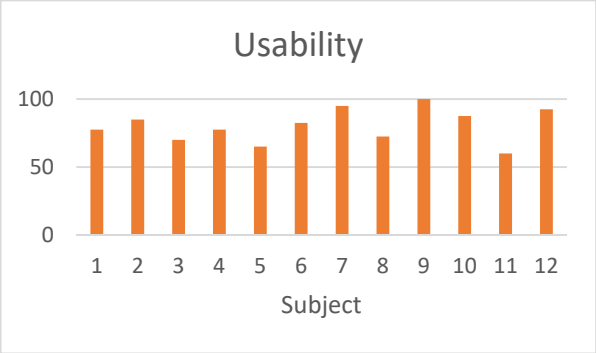

Figure 25 Experiment 1 usability.

7. EXPERIMENT 2

The second experiment was held in February and March of 2018. It was a within-subject study consisting of 12 participants. Each participant was asked to run four different configurations as outlined in Table 5. Unlike the first experiment which varied latency under three different values, this experiment chose to focus on a single value of latency, 1 second and vary the configuration of the mitigation technologies to include the predictive display, the camera stabilization, void filling and the overlay. The conditions consist of a Baseline condition which has no latency and no mitigation factors. The second has 1 second of latency with no mitigating technologies (This should be the worst performing condition.) The third also has one second of latency with the predictive display, stabilization and void filling turned on. The fourth is like the third with the addition of the overlay. All conditions employ the image subsampling. The four conditions were sequenced using a Latin Square design for the 12 subjects as shown in Table 6. The four sequences {A, B, C, D} were run for three subjects each. Assignments of sequences to subjects was random based on availability. The ages of the subjects ranged from 25 to 39 years. The distribution of ages is show in Figure 26.

The experiment used the same terrain as was used in experiment 1. Experiment 2 added 12 gates to the terrain consisting of Jersey Barriers set with a certain gap distance between them. The gates were configured with six of them being "wide" which allowed both lanes to pass through and six "narrow" which only allowed one lane to pass through. The subjects were told to safely pass through the gates without contacting them. The gates were setup on the turns, one wide and one narrow on the 75m radius turns and one wide and one narrow on the 20m turns. This was designed to test the precision of turning performance on road

Table 5 Experiment 2 experimental configurations.

Conditions	
Baseline	No Latency, no PD, no Overlay, no Stab
No-PD	1 sec latency, no PD, no Overlay, no Stab
PD	1 sec latency, w/ PD, Stab, & Void Fill
PD+OL	1 sec latency, w/ PD, Stab, Void Fill, & Overlay

Table 6 Experiment 2 Condition sequencing

A	Baseline	No-PD	PD+OL	PD
B	No-PD	PD	Baseline	PD+OL
C	PD	PD+OL	No-PD	Baseline
D	PD+OL	Baseline	PD	No-PD

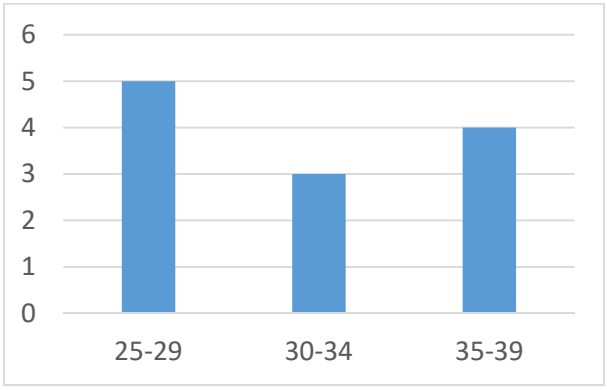

Figure 26 Experiment 2 Age distributions.

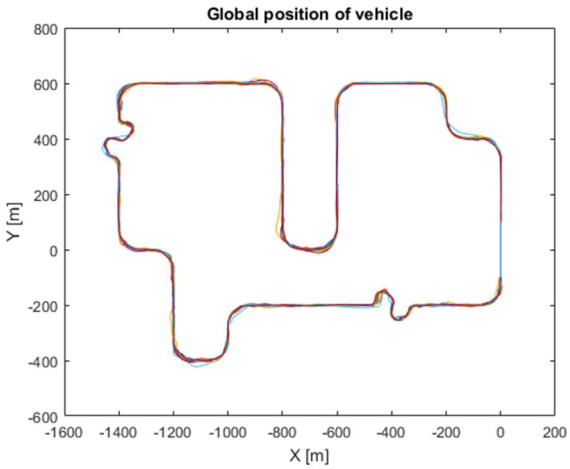

Figure 27 Experiment 2 paths driven.

features of significant curvature. Performance in this case was a measure of whether the participant pass through the gate or not. There was no consequence for not passing through the gate.

The paths driven by all participants is illustrated in Figure 27. Comparing the results to Figure 12, we can see that large deviations observed in Experiment 1 were not present. This is largely attributable to the omission of the 2 second latency condition in Experiment 1.

7.1. Performance Analysis

As in experiment 1, the four metrics of performance were used as: Average vehicle speed, path deviation, heading deviation and steering effort. In addition to these we measured if the operator passed through the gates.

Speed performance is shown in Figure 28. There we see that subjects performed best in the baseline condition and worst in the unmitigated condition in all cases. The addition of the predictive display and the overlay helped improve the unmitigated case by 55% and 56% respectively. An ANOVA analysis indicates a statistically significant difference in performance as provided by the predictive display and the overlay cases. The addition of the overlay did not significantly improve the speed beyond the predictive display.

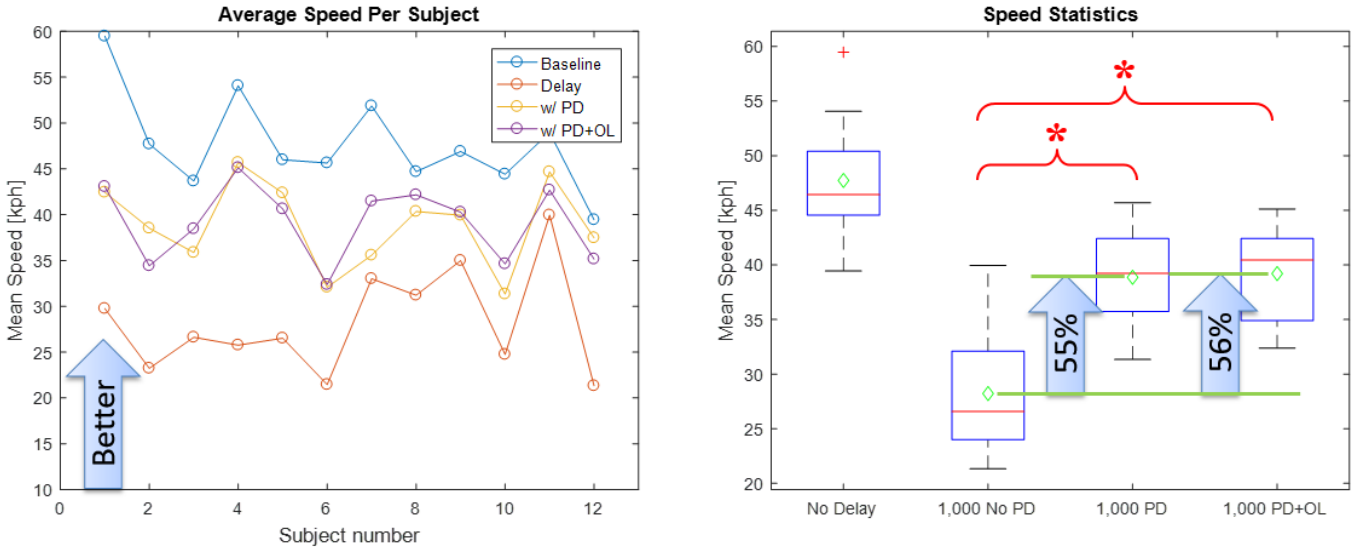

Figure 28 Experiment 2 Speed performance by subject (left) and by condition (right). (* Represents statistically significant difference.)

Path deviation performance is shown in Figure 29. There we see several large deviations in performance for the unmitigated case. In the left side of the figure, for subject 7 the #1 indicates that the unmitigated case was driven first in the sequence of the four conditions. Subjects 4 and 11 also had large unmitigated path deviation but they ran this case 4th in the sequence. This may be an instance of negative training in that the three easier cases were run first and the subject adapted to the lower amount of effort required to perform the task. These large values may also be attributable to the speed vs. accuracy trade because in two cases subjects 7 and 11 were among the fastest to complete the course in the unmitigated case (see Figure 28). In the right side of Figure 29 box plots of the 12 subjects' performance is show in all four test conditions. Here we find that the base line was the best and most consistent and the unmitigated case was the worst. We also find that the addition of the predictive display and the overlay improve performance by 73% and 93% respectively. Additionally an ANOVA analysis indicated that both distributions are significantly different than the unmitigated case. Although on average the overlay seemed to improve performance from 73% to 93%, these two samples were not found to be statistically significant.

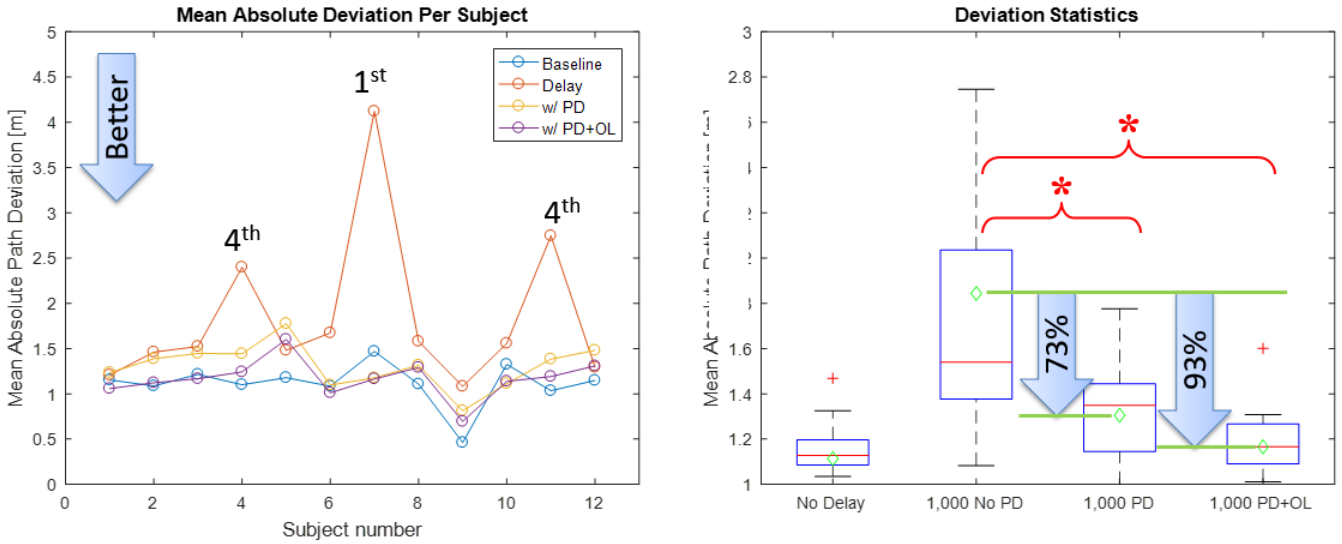

Figure 29 Experiment 2 Path Deviation performance by subject (left) and by condition (right). (* Represents statistically significant difference.) (1st and 4th in the left plot indicate which position in the sequence these particular runs were placed).

The third performance measure was heading error and the results are shown in Figure 30. There we see again that subjects 4, 7 and 11 had significantly poor performance as was the case in the path deviation metric. Indeed, it makes sense that the two would be/should be correlated. On the right side of the figure, we see box plots of the performance for all four conditions. There we see similar results to the path deviation metric. The introduction of the predictive display significantly improves the performance by 66% on average and then the addition of the overlay further improves by another 22% for a net performance improvement of 88%. In this case the addition of the overlay also significantly improved performance beyond the predictive display alone. This may be attributed to the projection of the overlay 2 seconds into the future. This notionally helps the operator see and

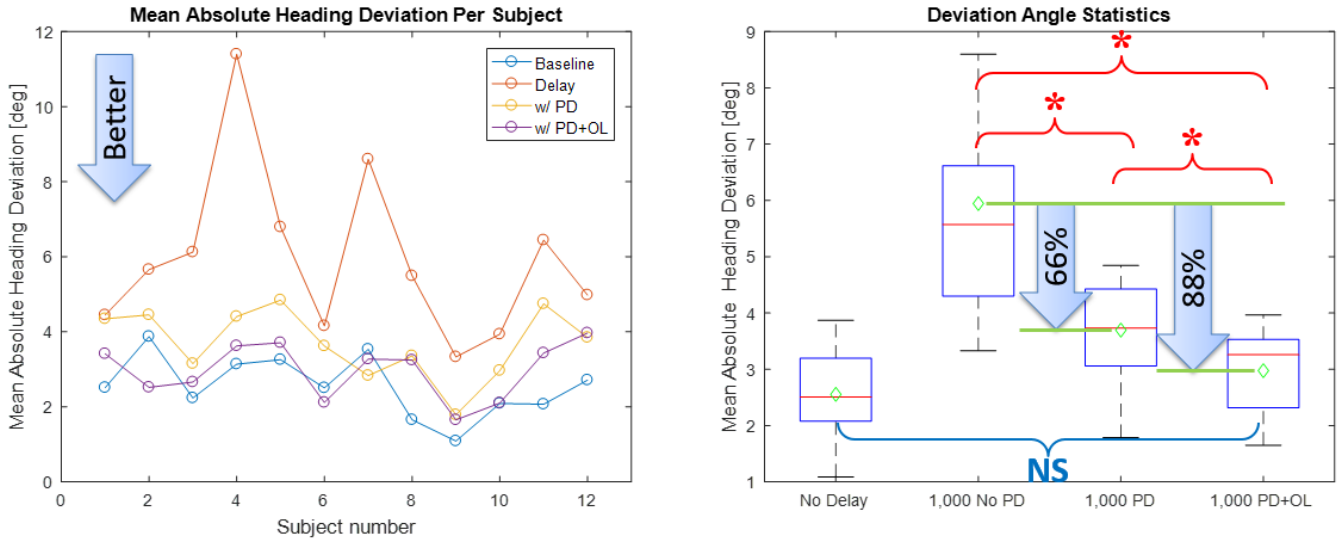

Figure 30 Experiment 2 Heading Error performance by subject (left) and by condition (right). (* Represents statistically significant difference.)

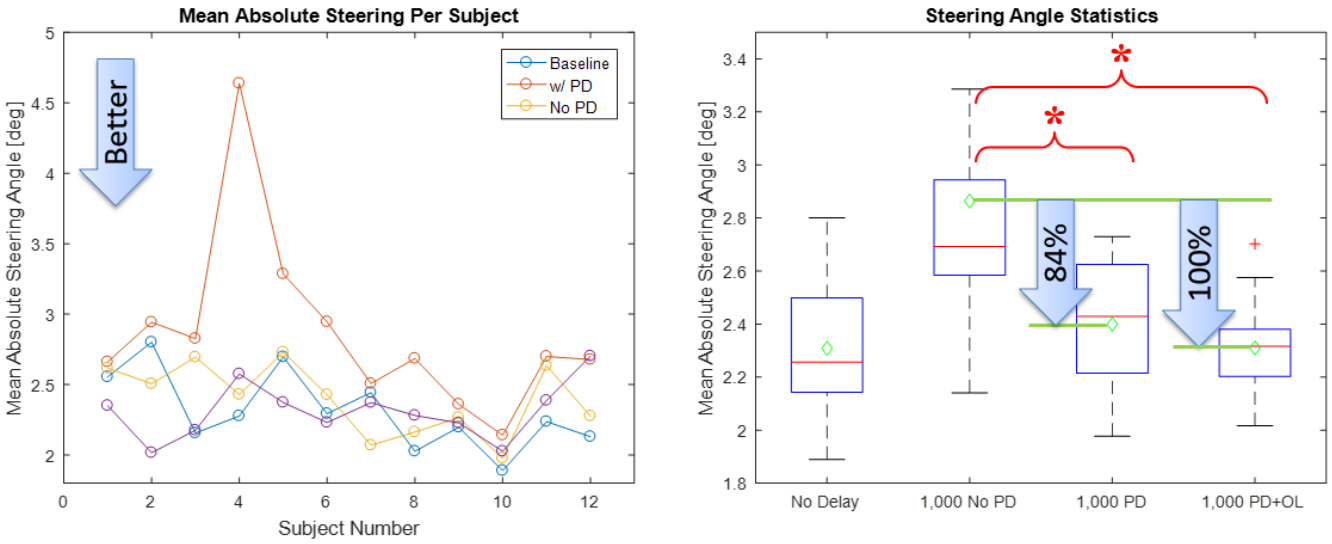

Figure 31 Experiment 2 Steering effort performance by subject (left) and by condition (right). (* Represents statistically significant difference.)

correct for heading errors. Finally, comparing the predictive display and overlay condition to the baseline, it was found that the performance between the two were not statistically significant.

The fourth performance measure of steering effort is shown in Figure 31. The left plot shows the performance of the individual subjects. Like in the prior two error conditions, subject 4 struggles with the unmitigated case. In the right side of the figure, the four conditions are illustrated with box plots. There we see that the addition of the predictive display and the overlay improve the performance by 84% and 100% respectively. Both are statistically different than the unmitigated case. Steering effort is an indicator regarding how much the operator is working to keep the vehicle on the path. The indication that the mean is the same for both the baseline and PD+OL cases meaning that at least in terms of physical effort, the predictive display and overlay do not require additional physical effort. We may speculate that the reduced physical effort corresponds to less cognitive effort. These will be expanded on in the discussion of the NASA TLX workload analysis.

Finally for experiment 2 we recorded the binary events of passing through the 12 gates. Results of this analysis are shown in Figure 32 and Figure 33. Figure 32 shows performance for the wide gates (2 lanes, 8.12 m) and shows performance for the narrow gates (1 lane, 4.12 m). Like the prior analysis the left plots indicate the per-subject performance and the right show the per condition box plots. The top two plots show the percentage of gates passed (note the y-axis scale does not go

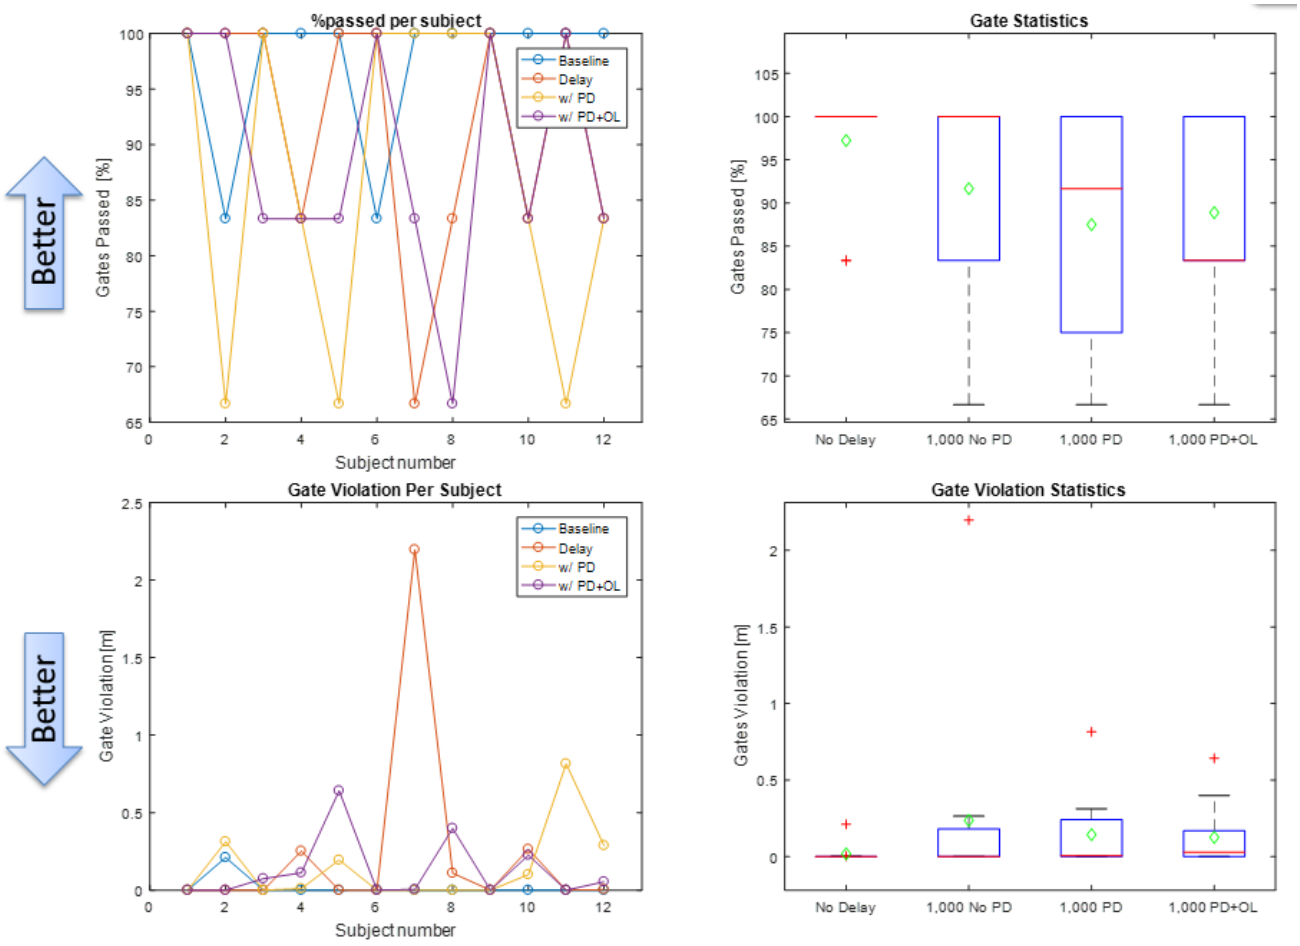

Figure 32 Experiment 2 Wide gate performance.

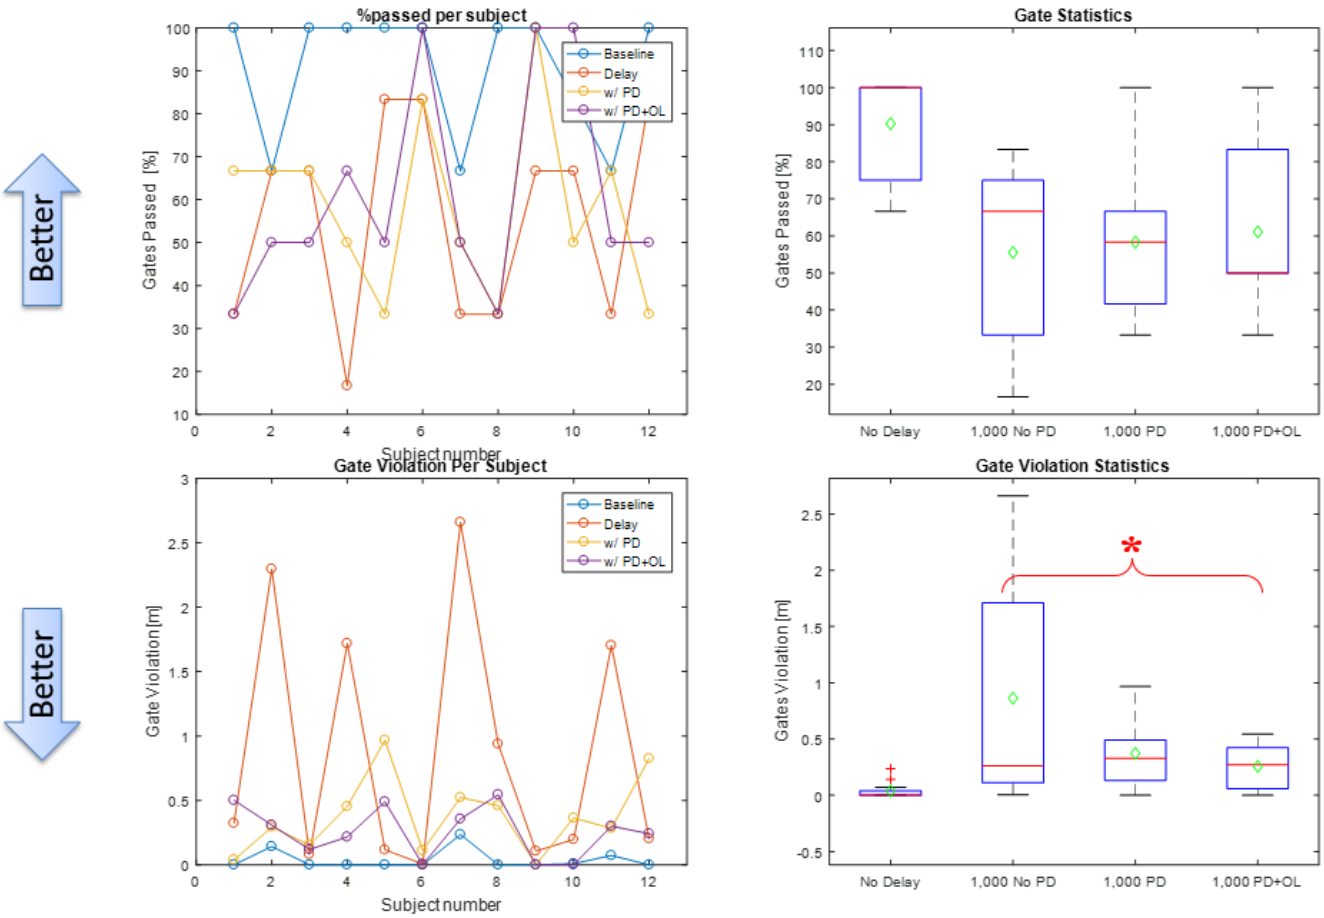

Figure 33 Experiment 2 Narrow gate performance.

down to 0). The bottom two plots show the gate violation meaning how much did the vehicle overlap the gate (i.e. how severe was the gate violation). Unlike the prior four analyses, this analysis was not as indicative of clear results. First, note that even in the baseline case, there were gate violations. This potentially is an artifact of teleoperation with a single camera in that peripheral vision is limited. Second, also note that even in the worst case of 1 second of delay and no mitigation, the subjects were able to negotiate the wide gates better than 90% of the time. The narrow gates seemed to be more discriminatory of precise control through the gates. Third, note that the percentage of gates passed does not improve with the addition of the predictive display. The degree of gate violation does however improve slightly. One possible explanation of this phenomenon is the way that the predictive display processes vertical objects. At a distance, vertical objects are included in the far-plane scene, but as the operator gets closer, the object gets increasingly included in the ground-plane which distorts the object and may potentially interfere with localization of the object. Since the gates were made with Jersey barriers, the gates were affected by this phenomenon.

With the addition of the overlay, the percent gate violations improved slightly as well as the degree of violation. This may be attributed to the predictive nature of the overlay in that it helps the operator to see where the vehicle is likely to go up to 2 seconds into the future. This helps the operator to steer the overlay through the gate and allow the vehicle to subsequently pass through the gate as well. For the narrow gates in Figure 33 the lower right plot, shows significant improvement for the

narrow gate case in terms of degree of violation. This is likely directly attributable to the control that the predictive display and the overlay have on reducing lateral error (see Figure 29).

7.2. Workload

After each experimental run, each subject was administered a NASA TLX questionnaire to assess workload. The data collected did not include the weighting aspects of the questionnaire. The results of this analysis are shown in Figure 34. There we can see the expected results that the baseline was the easiest and the unmitigated delay case was the most strenuous. The introduction of the predictive display alone improved workload over the unmitigated case by 61% on average and the addition of the overlay reduced workload by an additional 29% for a total reduction of 90%. This corroborates the conclusion observed in the steering activity reduction observed in Figure 31 that the predictive display and the overlay substantially reduce steering activity and perceived workload.

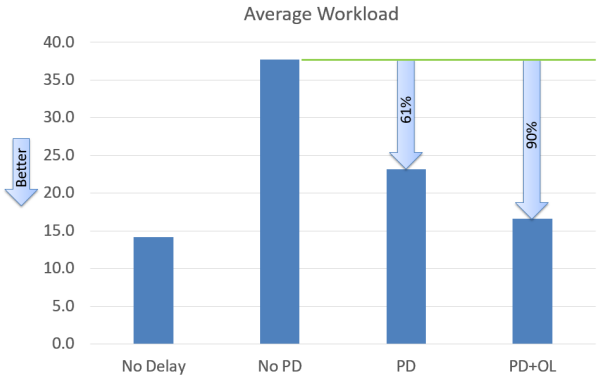

Figure 34 Experiment 2 Workload.

7.3. Usability

Each subject was administered the System Usability Scale [4] as part of their exit process. The subjects were not administered separate questionnaires for the predictive display and the overlay, but were asked to include the whole system in their assessment. The results of the survey are shown in Figure 35. There we see the usability ranged from 55 to 95 with an average of 78.1. Generally a score in excess of 68 is considered usable and on average this technology exceeds that threshold and 9 of the 12 subjects rated usability higher than 68. The system as evaluated is considered to be usable.

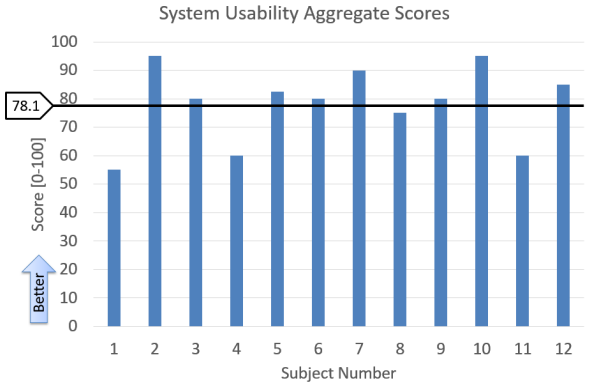

Figure 35 Experiment 2 System usability.

8. CONCLUSIONS

This paper is meant to complement and build on the prior paper [1] by adding capability and features which were the result of feedback from experiment 1. This paper also includes the result of two experimental studies from 2017 and 2018. The first evaluates the predictive display under three different latency conditions of 0.5, 1, and 2 seconds. The second evaluates two different technology levels under a fixed latency of 1 second. The two technology options are the predictive display (with stability, image subsampling and void filling) and the addition of a graphical overlay. Experimental results in both cases demonstrated significant improvement in performance and workload provided by the predictive display. The predictive display also showed acceptable usability.

9. REFERENCES

[1] M. J. Brudnak, "Predictive Displays for High Latency Teleoperation," in *Proceeding of The*

Ground Vehicle Systems Engineering Symposium (GVSETS), Novi, MI, 2016.

[2] M. J. Brudnak, "Method of generating a predictive display for tele-operation of a remotely-operated ground vehicle". USA Patent US10425622B2, 24 9 2019.

[3] NASA, "NASA TLX: Task Load Index," 15 9 2019. [Online]. Available: https://humansystems.arc.nasa.gov/groups/TLX/. [Accessed 11 5 2020].

[4] J. Brooke, "SUS: a "quick and dirty" usability scale," *Usability Evaluation in Industry,* 1986.