

PORTOBELLO: Extending Driving Simulation from the Lab to the Road

FANJUN BU, Cornell University, U.S.

STACEY LI, Cornell University, U.S.

DAVID GOEDICKE, Cornell University, U.S.

MARK COLLEY, Institute of Media Informatics, Ulm University, Germany and Cornell Tech, U.S.

GYANENDRA SHARMA, Woven Planet, USA

WENDY JU, Cornell University, U.S.



Fig. 1. Portobello enables *platform portability*, so that the same study can be run on in-lab (left) and on-road (right) driving simulation platforms.

In automotive user interface design, testing often starts with lab-based driving simulators and migrates toward on-road studies to mitigate risks. Mixed reality (XR) helps translate virtual study designs to the real road to increase ecological validity. However, researchers rarely run the same study in both in-lab and on-road simulators due to the challenges of replicating studies in both physical and virtual worlds. To provide a common infrastructure to port in-lab study designs on-road, we built a platform-portable infrastructure, Portobello, to enable us to run twinned physical-virtual studies. As a proof-of-concept, we extended the on-road simulator XR-OOM with Portobello. We ran a within-subjects, autonomous-vehicle crosswalk cooperation study ($N=32$) both in-lab and on-road to investigate study design portability and platform-driven influences on study outcomes. To our knowledge, this is the first system that enables the twinning of studies originally designed for in-lab simulators to be carried out in an on-road platform.

CCS Concepts: • **Human-centered computing** → **Systems and tools for interaction design**; *Interaction design process and methods*; • **Computing methodologies** → **Simulation tools**; **Simulation environments**.

Additional Key Words and Phrases: Human-Autonomous Vehicle Interaction, Driving Simulations

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made for profit or for advertising or promotional purposes, for resale or distribution. This notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 ACM.
Manuscript submitted to ACM

ACM Reference Format:

Fanjun Bu, Stacey Li, David Goedicke, Mark Colley, Gyanendra Sharma, and Wendy Ju. 2024. PORTOBELLO: Extending Driving Simulation from the Lab to the Road. In *CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11-May 16, 2024, Honolulu, USA. ACM, New York, NY, USA, 21 pages. <https://doi.org/XXXXXXX.XXXXXXX>

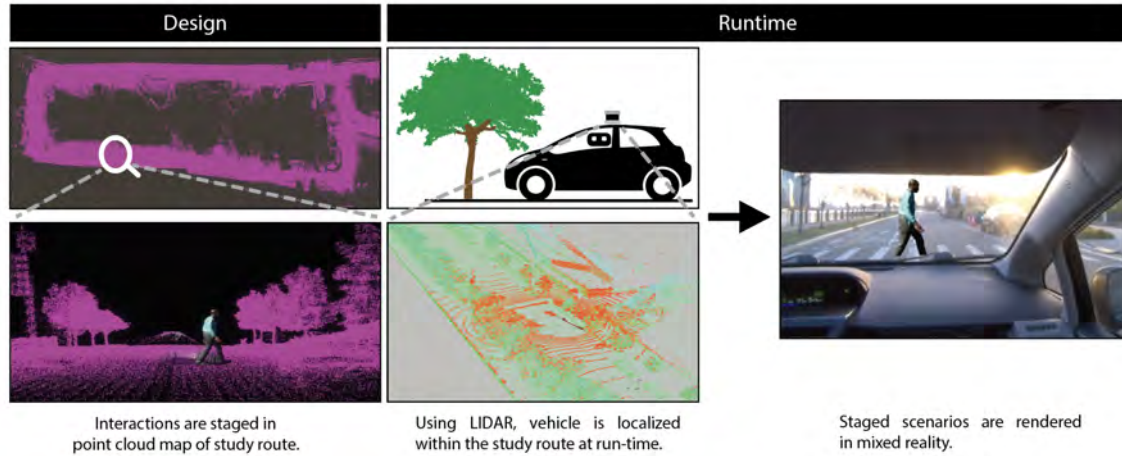


Fig. 2. Complete execution pipeline using the Portobello infrastructure. During the design phase, researchers can drag and drop virtual objects on the point cloud map. At run time, the LiDAR-based navigation system locates the vehicle's position within the same map and sends the relative vehicle position to Unity. As a result, passengers wearing the video-see-through headset can see virtual objects at their corresponding real-world locations.

1 INTRODUCTION

Driving simulators play a critical role in human-centered automotive research applications because they allow people to experience different driving scenarios in a safe and repeatable fashion. Researchers have full control over the simulation setup and can program complex events in virtual environments. However, one of the major challenges for driving simulation has always been the difficulty of replicating the inertial forces and vehicle dynamics present in on-road driving [43]—even the most high-end motion platform driving simulators only replicate a fraction of the forces felt in real-world driving [19]. These forces are more critical in testing automated vehicle (AV) scenarios, where participants are often given non-driving-related tasks that keep them from looking at the screens of the simulations (e.g., [56]). Without the physical sensation and the visual engagement of the simulation environment, the immersiveness of the environment drops, making the evaluation of novel user experience and interaction techniques such as VEmotion [5] or SoundsRide [30] difficult. One way to address this problem is to incorporate driving simulation into a real vehicle driving on real streets (on-road simulators) [2].

On-road simulators are possible due to the maturity of XR, where digital displays blend reality and virtuality to increase the level of a user's immersion. The use of XR-in-the-car using simulated or actual dash-mounted heads-up displays have been explored by prior researchers such as Tonnis et al. [53], Kim and Dey [32], Schall Jr et al. [49], Ghiurău et al. [15], McGill et al. [39], Colley et al. [9], von Sawitzky et al. [54], Narzt et al. [41], and Bark et al. [3].

105 Despite this influx of on-road simulators under development, it remains challenging to migrate studies from in-lab
106 simulators to on-road simulators. From a study design perspective, the key difference between in-lab and on-road
107 simulators is that with in-lab simulators, designers have full access not only to the virtual vehicle but also to the virtual
108 environment. The position and orientation of every brick are available to the designer with high precision, which makes
109 event staging as simple as dragging-and-dropping modules into a map. However, researchers do not have easy access to
110 objects outside the vehicle in on-road simulators. As a result, no on-road driving simulation system to date considers
111 the surrounding context outside the vehicle for event staging, which limits the range of applicable studies. To replicate
112 in-lab simulation on on-road platforms, access to the out-of-vehicle surrounding context is crucial.
113
114

115 In this paper, we describe a novel driving simulation infrastructure called Portobello, which enables *platform*
116 *portability* in virtual driving simulation by incorporating localization technology and software from robotics. We
117 define *platform portability* as the ability to run the same study on different (in-lab vs. on-road) platforms, an approach
118 which we refer to as the twinning of studies. For this demonstration, we extended XR-OOM, a state-of-art XR driving
119 simulation [16], to support on-road, location-based event staging (see Figure 2.) To validate Portobello's *platform*
120 *portability*, we developed a within-subjects crosswalk-cooperation study ($N=32$) to be run on both an in-lab fixed-based
121 vehicle chassis driving simulator and the on-road driving simulator built on top of the Portobello system [55]. As
122 part of this work, we investigated how the different simulation platforms may affect the design process and results of
123 user studies. The primary contribution of this work is the technical infrastructure system of Portobello, as validated
124 by the proof-of-concept study. In addition, we provide a definition of platform-portability in virtual reality driving
125 simulation and contribute insights into the process needed to develop twinned studies whose deployment is intended
126 across multiple platforms. Finally, we demonstrate the relative strengths of different study platforms in the course of
127 automotive research.
128
129
130
131

132 2 RELATED WORK

133 Driving simulation platforms are intended to be proxy environments that enable researchers to conduct studies where
134 real-world experiments are dangerous or impossible. The standard for such platforms is face validity [14]: when
135 participants take a simulation seriously, researchers can have greater faith that the study's results will be applicable to
136 the real world. It is more important that the simulation allows participants to behave as if they are in a realistic setting
137 than it is for the simulation itself to replicate reality in fine detail. In-lab and on-road simulators provide different
138 approximations of driving scenarios.
139
140
141

142 2.1 In-Lab Driving Simulators

143 In-lab driving simulators are used to test interactions between drivers and the vehicle, the driving environment, and
144 other in-world agents [27]; simulation allows researchers to observe in-vehicle behaviors safely, without physical risk
145 to study participants [52]. Simulations can be implemented with low fidelity but still provide insights into how drivers
146 will react sitting in a real vehicle [10]. When used to simulate outdoor environments, in-lab driving simulators can
147 be used to test the usage of AR on roads [53]. In-lab simulators can also be used to test how other road users, such as
148 pedestrians and cyclists, interact with autonomous vehicles [23, 38]. However, simulator sickness remains a large risk
149 for in-lab driving simulators because a user's vestibular senses do not align with their visual senses when taking part
150 in a simulation [4]. Researchers have tried to address this with methods such as aligning vehicular motion with VR
151 content [7] or simulating movement [11, 22], but even the highest-end simulators replicate a fraction of the forces felt
152 in normal on-road driving. [19]
153
154
155

2.2 XR On-Road Driving Simulators

XR technology enables simulating AV driving either by allowing users to move through completely virtual spaces with real vehicle dynamics or by overlaying virtual objects on top of real-time video footage of the surroundings to increase the immersion of pre-programmed interfaces or interactions [6, 16, 17, 21, 45, 46, 48, 58]. Recently, XR systems have been deployed on-road to take advantage of the realistic road environment and vehicle dynamics. The XR-OOM system developed by Goedicke et al. [16] employed an XR headset for drivers to drive through virtual, external obstacles in a parking lot. The MAXIM system developed by Yeo et al. utilized a virtual reality headset coupled with 360° cameras for subjects to experience an autonomous virtual vehicle situated in a live environment created from live streamed 360° videos [58]. Ghiurău et al. [15] showed a proof-of-concept headset-based XR driving experience revealing that original equipment manufacturers (OEMs) such as Volvo use such technology. Finally, McGill et al. [39] presented PassengXR, an open-source toolkit to create passenger XR experiences. While providing XR experiences, they did not compare their system to an in-lab simulator.

Although all the systems mentioned above track the vehicles' dynamics to accurately position virtual objects related to the participant, they do not natively support high-precision interaction staging based on the surroundings outside of the vehicle, which we refer to as *surrounding context* in this paper. In essence, previous approaches are mostly concerned with aligning virtual and physical motions, such as CarVR, but not the worlds themselves [21]. The authors of PassengXR introduced a hypothetical application that requires high-precision alignment between virtual and physical worlds, where passengers on a bus tour could view AR-style information overlaid on historic buildings outside the vehicle, but their demonstration was still carried out indoors without discussing how feasible it was to implement such application on-road [39]. Potential issues for designing such applications with existing platforms are two-fold. First, current systems make it challenging for designers to stage interactions, where they need to manually locate event trigger positions (in the bus tour example, buildings' coordinates either in GPS coordinates or relative coordinates in the vehicle frame) and program corresponding AR information boards to be at those precise coordinates. Second, the GPS-based tracking system may not provide enough accuracy for small-scale interaction, especially in cities where buildings can disrupt GPS signals. The authors for PassengXR commented that their "approach prioritizes perception of motion over location accuracy," which is not best suited for location-based XR experience [39]. These hardships limit the capability of these on-road systems to act as participant testing platforms compared to traditional in-lab simulators. Most autonomous driving studies require sufficient staging and surrounding context [8]. For example, driver-to-driver communications mostly happen at intersections, and pedestrian-vehicle interactions usually occur at crosswalks [24, 35, 37, 50]. Our work makes high prevision interaction staging possible while keeping the design process intuitive.

2.3 Platform Portability

Previous research has focused on how to replicate on-road scenarios in in-lab simulators, which is crucial when studying problems that are dangerous to experiment on the road, such as near-collision scenarios and passive rail level crossing [13, 36]. However, compromises are necessary to compensate for the lack of motion and sensory cues in in-lab simulators, and little research has been conducted to reduce the performance gap between on-the-road and in-simulators [18]. As such, merely pursuing statistical significance with in-lab simulators may result in overlooking issues of practical relevance in real-world contexts. [18].

209 Because in-lab and on-road driving simulation environments offer different strengths when it comes to control and
210 realism of driving scenarios, it can be desirable to run the same study in both when possible— an approach we call
211 twinning of studies — to understand how study results from different environments relate to each other. Hammel et al.
212 [20] found that, when they replicated an on-road study in a fixed-based simulator, participants' eye-scanning behavioral
213 patterns were similar, which demonstrates fix-based simulators' ecological validity. In a systematic review of validation
214 studies featuring comparisons of driving simulation and on-road driving between 1977 and 2017, authors Wynne et al.
215 [57] found only 44 validation studies comparing simulation to real driving. This is out of the 21,312 found by the same
216 researchers to be English-language publications of original research having to do with driving simulation. Such studies
217 are so rare that the 44 represent less than 0.25% of the published driving simulation research Wynne et al. [57]. They
218 report that "There was little consistency in the dependent measures used to assess differences between the simulator
219 and on-road drive...Of particular concern is the fact that only half of the driving simulators were found to be valid and
220 some were valid for one measure but not others." They note that since policy, legislation, and training are built off of
221 simulator studies, a better understanding of which aspects of simulated studies are likely to carry over to real road
222 conditions, and which are not is critical [57].
223

224 Frequently, we believe, the lack of validation studies is due to the significant challenge of creating "twinning" studies
225 in both environments. The advent of on-road mixed reality simulation [2, 16, 39, 58, 59] makes it possible to bridge the
226 divide using software events in the real world. However, no system has yet ported the same study course, code, and
227 event design from one environment to the other. By making it possible to port studies developed for in-lab simulators
228 to be run on-road—what we call *platform portability*—we improve the ability for automotive researchers to extend their
229 in-lab studies to the real road, and thereby improve the validity of simulation research. This would improve the ability
230 of researchers to validate their simulation studies in on-road environments, as recommended by Wynne et al. [57]:
231 "Ideally this would see authors report empirical validation evidence for their own simulator, and not relying on other
232 simulators as support for validity. Even if modeled on a previously validated simulator, each set-up is unique and should
233 be validated for those specifications."
234

240 3 SYSTEMS

241 In this work, we present a study we developed meant to run on two driving simulator environments: a lab-based driving
242 simulator and an on-road XR driving simulator equipped with our Portobello system. Here, we recap the features of
243 both environments and introduce the key adjustments made to accommodate the Portobello system.
244

247 3.1 In-Lab Driving Simulator

248 Our in-lab fixed-base driving simulator features a modified Fiat 500 in front of three projector screens (see Figure 1).
249 The projector screens cover participants' visual field when they sit in either the driver's seat or the passenger's. The
250 three projectors are DLP-based and can produce an image with low latency on the projector screens. The projectors
251 are connected to the computer over HDMI and use the TripleHead2Go to split one DisplayPort signal into the three
252 outputs.
253

254 The vehicle is coated with non-reflective material to reduce the backscatter onto the projector screens, increasing the
255 contrast of the projector screens. ButtKicker haptic transducers are installed under the front seats to provide realistic
256 tactile feedback from road noise and the engine. The simulation software is run by an Alienware Area-51 R4 computer
257 with two NVIDIA GTX1080 in SLI. The vehicle's side mirrors are small digital displays rendered by the same computer.
258




Fig. 3. The Portobello system uses a LiDAR-based navigation stack to localize the vehicle's runtime position within a given map. The location information is transmitted to the Unity Desktop through ROS. The black dotted line indicates a virtual divider appearing only in the headset.

In our simulation software, our vehicle model was imported from the open-source Genivi Vehicle Simulator, which has been widely used in driving simulation studies [1].

3.2 On-Road XR System

To enable cross-environment study deployment, we used Portobello with the XR-OOM system designed by Goedicke et al. [16]. In the original XR-OOM system, tracking and positioning of virtual objects are managed by a ZED 2 camera (for visual-SLAM) and the ART SMARTTRACK3¹ (for headset tracking within the vehicle) [16]. An onboard desktop running the Unity 3D game engine in version 2020.3.26f1 overlays virtual objects on top of the "passed-through" video of surroundings in the XR headset.

3.3 Portobello System

In the Portobello system, we used a LiDAR-based navigation system on the car roof driven by the Robot Operating System (ROS 1 Noetic) [44]; this replaces the XR-OOM's ZED 2 camera in front of the vehicle. The Portobello navigation system features an Ouster OS-1 3D LiDAR with a built-in Inertial Measurement Unit (IMU) and a ZED-F9R u-Blox GPS module. The communication between the navigation system and the Unity Desktop is managed through the ROS-TCP-Connector² provided by Unity [29].

Platform portability also drove augmentations to the XR-OOM software structure. Whereas XR-OOM used real-time visual SLAM to compute short-term vehicle odometry, the Portobello LiDAR system enables global vehicle localization

¹<https://ar-tracking.com/products/tracking-systems/smartrack3/>, accessed Jan 20, 2023

²<https://github.com/Unity-Technologies/ROS-TCP-Connector>, accessed Oct 20, 2022

313 within a given map. As a result, virtual objects' positions are no longer associated with the vehicle position directly as
314 its children. Rather, the vehicle and virtual objects share a common parent — the world frame — which is introduced by
315 a map of the environment. Instead of staging virtual events around the vehicle, as was done in XR-OOM (and in all the
316 other XR-based driving simulation systems mentioned in [15]), Portobello can stage virtual events in a static shared
317 map through which the vehicle drives. As Portobello replaces the car-centered reference frame with a map-based global
318 frame, out-of-vehicle virtual objects remain fixed with respect to the map instead of to the arbitrary starting position of
319 the vehicle. We detail the map generation process and staging process in Section 4.
320
321

322 At runtime, Portobello's LiDAR-based navigation system updates the position of a virtual vehicle in Unity. The relative
323 position and orientation between the virtual vehicle and the participant's headset are managed by SMARTTRACK3.
324 The virtual vehicle has the same shape as our research vehicle and is aligned with the research vehicle throughout the
325 ride. To simulate proper depth ordering, the virtual vehicle is transparent with the alpha clipping option enabled. By
326 acting as a cutout shader, the virtual vehicle occludes virtual objects outside. From the participants' view, the virtual
327 objects are occluded by the real car they sit in. As we are rendering virtual events in XR over a large area (half of an
328 island), we capped the headset's maximum rendering distance to 45 meters in order to improve motion parallax.
329

330 From a system design perspective, the computer running Portobello with ROS is a separate computer from the
331 computer running XR-OOM with Unity. In essence, the communication between the ROS localization algorithm and
332 Unity is achieved through the ROS transform package (TF), which is a hierarchical tree structure that tracks the relative
333 position of multiple coordinate frames (map, LiDAR, etc.). These coordinate frames can be accessed from Unity as game
334 objects. Isolating the system on a hardware level allows researchers to develop ROS and Unity in their own environments
335 and for one driving simulator to be swapped out for another. Designers can focus on designing studies by placing
336 objects in the course map in Unity, rather than being concerned with low-level ROS localization of objects. Another
337 benefit of this practice is that the Portobello system consumes no computational resources in the original on-road
338 platform at run time. The LiDAR-generated point-cloud map is rendered into Unity in the design phase. At runtime,
339 the on-road platform computer imports the Unity map at start-up, and does not require the additional computational
340 resources that would be needed to manage the point cloud map data. (This does mean that the Unity map might be
341 missing physical features that change between the design phase and runtime.)
342
343
344
345
346
347

348 4 ENABLING PORTABILITY OF STUDY DESIGN USING COMMON MODELS

349 To run twinned studies across platforms, we must keep portability in mind during the design of the study. We outline
350 the necessary components to guarantee equivalent performance in cross-environment twinned studies and detail our
351 system pipeline to showcase the components' connections using our system.
352
353
354

355 4.1 Course Design

356 On-road simulators are limited by real-world road infrastructures. Based on the study focus, researchers should carefully
357 consider test routes to ensure efficiency and reproducibility.
358

359 As our study focused on interactions at crosswalks, we chose the southern loop of Roosevelt Island as our study
360 course. Most of the 0.9-mile long route is single-track and has a high density of crosswalks, 15 in total. The drive takes
361 about 8 minutes. The route has no traffic lights and overlaps with two bus lines.
362
363
364

4.2 Map

A map (model of the study area) is the starting point of event design, and its precision and resolution greatly affect design complexity and quality. In cross-environment twinned studies, a map is a bridge between simulation and the real world, and it is also the shared common ground on every simulator platform.

To generate a high-quality map, we used the LiDAR-based navigation system to scan the entire test area. Specifically, we ran the real-time LiDAR-inertial odometry package (LIO-SAM) developed by Shan et al. [51] to create a true-to-scale point cloud mapping of the study area. We drove through the testing route multiple times to ensure loop closure. The resulting map is a monochromatic digital twin that includes over one million points and captures the test area in fine detail. Researchers can manually contextualize the point cloud map in Unity with colored assets and use the map as the background in in-lab simulators.

4.3 Event Design

Staging events along the planned course requires researchers to consider two major questions: where and when events occur. Staging is relatively easy for in-lab simulations, where agents' movement and speed profiles can be carefully controlled to guarantee timing and location. In this section, we discuss how we stage events using on-road simulators.

4.3.1 Planned Events.

Where? A one-to-one scaled map is necessary to plan events for on-road simulations because the vehicle will drive through the real world during the study. Any scaling or shifting on the map will cause significant errors at runtime. With a loaded digital twin in Unity, researchers can drag and drop virtual agents and objects just like they would for in-lab simulators.

When? Timing of events for on-road simulators can be controlled through the placement of collision-based triggers. Triggers in Unity are colliders that trigger events upon external contact. For example, invisible triggers can be placed at some distance x in front of a virtual traffic light. Once the vehicle collides with the trigger, the traffic light starts changing colors, and participants should react accordingly. The distance x governs the start of the interaction, which essentially affects the maximum response time for participants.

Of course, a vehicle on-road cannot "collide" with the virtual collider in simulation. In our system, the LiDAR-based navigation system synchronized the position of the real vehicle with a virtual vehicle in the digital twin through the *hdl_localization* and *ROS-TCP-Connector* packages in real-time [33]. As the real vehicle drives through the world, the virtual vehicle simultaneously moves through the digital twin to trigger planned events. Note that this is the same virtual vehicle mentioned in section 3.2 for proper depth ordering.

4.3.2 Unplanned Events. We define events outside the simulation, which researchers have no control over, as *unplanned events*. In in-lab simulations, unplanned events are rare and typically caused by system failures or external interruptions. However, in on-road simulators, unplanned events are common and can even be valuable for ecological validity; they let researchers know if their findings are robust to real-world variation. Findings that are only true in the tightly controlled environment of a study have little practical application.

Nevertheless, researchers must factor in potential unplanned events during the design phase to ensure safety and preserve meaningful study results. Unplanned events come in different forms, from unexpected appearances of



Fig. 4. Our study area for on-road simulator on Roosevelt Island. The pre-determined test route is highlighted in blue. The crosswalks with staged interactions are highlighted in the red bounding boxes. The start and end locations are denoted by the green and purple dots, respectively.

pedestrians to weather changes. For instance, in our study, we encountered the following unplanned events: real pedestrians and geese crossing the street, other vehicles passing the research vehicle from the bicycle lanes, and rain.

4.4 Platform Measures

Another aspect of *platform portability* is whether researchers can obtain the same set of measures from twinned studies. While measurements should be equally attainable across all simulation platforms, the characteristics of each simulator naturally encourage and discourage different sets of measures.

4.4.1 Behavioral Response. Behavioral responses refer to the participants' elicited behavior during the study. Examining behavioral responses is crucial when studying interactions between drivers, vehicles, and infrastructure [47]. For example, Jansen et al. [26] was interested in differences in participants' responses to different stimuli in automotive user interfaces. With appropriate sensors, collecting behavioral responses in in-lab and on-road simulators is possible.

4.4.2 Performance Response. Paas and Van Merriënboer [42] define performance as efficiency in completing tasks. We distinguish performance from behavioral responses based on the availability of ground truth. Researchers can collect performance responses during the study when participants are assigned tasks with general guidelines and standards. One example of performance response is the lateral vehicle position when the driver is distracted [34]. While extra sensors might be needed for on-road simulators to obtain vehicle-related performance measures (e.g., vehicle speed, acceleration, or trajectories are not easily attainable in on-road simulators as they are in in-lab simulators), we do not anticipate significant challenges in obtaining performance responses in both in-lab and on-road simulators.

4.4.3 Survey Response. Surveys can be conducted through different devices (pen and paper, tablets) in various formats (interviews, multiple choice, open-ended questions). In portable study design, it is important to consider the timing of the survey. One natural advantage of in-lab simulators over on-road simulators is the ability to pause at any point of the study and prompt participants with questions in situ [12]. On-road simulators cannot be paused easily, so surveys need to be planned so that participants can take them when it is safe to do so.

4.5 Additional Instrumentation

Detailed runtime recording of the environment is crucial for post-facto data analysis, particularly of unplanned events, for both on-road and in-lab simulators. Some measures need to be recorded differently in the different platforms and translated. For example, geo-location data from the on-road vehicle GPS must be correlated with the virtual world coordinates in the lab simulator. Head orientation and gaze direction obtained from the XR headset in the on-road simulator can be correlated with camera-tracked head-pose in the in-lab driving simulator.

5 TWINNING OF STUDIES

As proof of concept that we can run the same study design in the lab and on-road (twinning of studies), we conducted a within-subjects experiment with $N=32$ participants (under IRB#1806008105). We describe the cross-platform deployment of twinned studies and compare the differences in between. We counterbalanced the experiment conditions, where half of the participants experienced the study in the indoor simulator first and in the on-road simulator second, and the other 16 participants experienced the simulators in the reverse order.

We employed the in-lab and on-road driving simulators to run twinned Crosswalk Cooperation studies, which we adapted from Walch et al. [55]. In this previous study, Walch et al. used an in-lab driving simulator to evaluate the usability of a novel car UI and staged interactions around a driving loop. In this current work, we are not seeking to validate the results of the previously published study; we are not expecting or arguing that our study results would be the same. Rather, we are merely using this study design to evaluate the capacities and key influences of both systems.

5.1 Study Setup

5.1.1 Protocol. The experiment is a within-participants experiment design; each participant experiences two study sessions in counterbalanced order. In one session, participants experience the crosswalk cooperation study in the in-lab driving simulator. Afterwards, they fill out a post-session questionnaire, which collects information on their experience with the simulator. In the other session, the participants are escorted to the curbside and experience the crosswalk cooperation study in the on-road driving simulator. Afterwards, they fill out the same post-session questionnaire. Finally, they fill out a post-study questionnaire, which collects information on their perceived differences between the two simulators.

During each simulation session, participants are seated in the front passenger seat and informed that they will experience automated driving: The vehicle will stop at all crosswalks automatically. The vehicle will proceed autonomously when road conditions are clear (e.g. crosswalks without virtual pedestrians) and will ask for input from the participant via a smartphone interface, on how to proceed in unclear situations (e.g. virtual pedestrians walking towards crosswalks). Specifically, the vehicle will ask, "Is now safe to proceed?" on the phone interface while waiting at the crosswalk; when the participant feels it is safe to proceed, they press the "Proceed" button, and the car resumes its predefined route. If the participants decide it is not safe to proceed, they should wait until it is safe to press the button. The researcher driving the vehicle monitors when the participant presses the button and manually proceeds with the course if it is safe.

5.1.2 Differences in Simulator Setup. As much as possible, we maintained identical setups for the twinned studies. The key difference was that during the on-road simulation session, the researcher driving the car was mindful of the actual road conditions before proceeding with driving.

To maintain the narrative that the vehicle was driving autonomously, the researcher driving during the on-road simulation was masked by a black divider in the video pass-through headset (shown in Figure 3). A similar black divider

521 was also installed in the in-lab simulator to maintain setup parity. During pilot sessions, we found it difficult to disguise
522 the on-road vehicle as an AV due to differences in sound profile. In complex on-road conditions, the sound of pressing
523 the pedal and rotating the steering wheel broke the illusion quickly. Thus, we decided to inform the participants of the
524 divider's purpose and that there was a real researcher in the car with them operating the vehicle.
525

526 While a virtual map environment is required in in-lab simulation, the point cloud map is not required for on-road
527 simulation due to the benefit of the video-pass-through headset. Therefore, after the event staging phase, we disabled
528 the point cloud rendering in Unity to save computation power.
529

530 *5.1.3 Scenarios.* We recreated four scenarios from the original Crosswalk Cooperation study by Walch et al. [55]. Each
531 scenario was engineered so that pedestrian interactions would only happen at crosswalks. In each scenario, virtual
532 pedestrians interacted with each other on the sidewalks near the crosswalks. In half of the scenarios, one pedestrian
533 walked to the stop sign and crossed the street after giving clear body language signals that they intended to cross
534 (looking left and right). In the other half of the scenarios, the pedestrian stopped at or walked past the crosswalk. In our
535 version of the Crosswalk Cooperation study, we constrained our study area to the southern loop of Roosevelt Island.
536
537
538

539 5.2 Participants

540 Out of our 32 participants ranging from 20 to 47 years old (age: $M = 27.38 \pm 5.92$), 19 participants identified themselves
541 as male, 11 as female, and two as non-binary. Six participants had experience with AV simulations and/or AV research,
542 and five had experience with commercialized AVs (Tesla, demos at car shows). The others had little experience with AVs.
543
544

545 5.3 Study Measures

546 After each session, participants filled out a post-session questionnaire, which collected information on their experience
547 with the simulator. After completing both sessions and corresponding questionnaires, participants filled out a post-study
548 questionnaire, which collected information on their perceived differences between the two simulators.
549

550 We also recorded video and audio for all sessions run in both in-lab and on-road driving simulations to investigate
551 participants' behavioral responses. For the in-lab driving simulator, a go-pro camera was pointed towards the participant
552 to record their upper body. The simulated virtual environment was recorded using screen recording software. For the
553 on-road driving simulator, similar to the setup in the in-lab simulator, a camera is mounted in the glove compartment
554 to record the participant's upper body. An additional camera is mounted near the rear mirror facing forward to record
555 the road condition ahead. The participant's XR view (video pass-through with overlay) is also recorded using the Varjo
556 Base software.
557
558
559

560 6 RESULTS AND ANALYSIS

561 We analyzed the video footage and the questionnaires from both sessions with the goal of investigating the influences
562 of both driving simulators on the user study experience and understanding what results from one platform predict for
563 results on the other. Our evaluation of Portobello is based on being able to run and gather comparable results from the
564 studies on both the in-lab and on-road platforms and not on the originality, validity, or significance of the study itself.
565
566
567

568 6.1 Study Results

569 *6.1.1 Measures.* To compare the overall participant measures in the study on both simulators, we asked the participants
570 ($N = 32$) to rate their feelings of anxiety, safety, and trust on a 5-point Likert scale for each of the platforms. We designed
571




Fig. 5. Participants' ratings of each simulator on a five-point Likert scale across seven distinct dimensions where 1=low and 5=high.

the questions based on the questionnaire used in Walch et al. [55]'s original Crosswalk Cooperation study. We ran a Bayesian factor analysis on the captured measures with the null hypothesis that there is no difference between the platforms [31] with R version 4.3.2 and the BayesFactor package [40] using Jeffreys-Zellner-Siow (JZS) priors, i.e., the default non-informative Jeffreys prior. Interpretations were made according to Jeffreys [28]. All packages were up to date in November 2023. Since we have limited data, and we are not interested in how different variables (ratings) interact with each other, we chose to report single variate analysis over multivariate analysis. Nonetheless, our claims hold under multivariate analysis as well. We will provide R script for both analyses.

One participant left multiple answers empty, so we dropped their results in the following analysis.

Anxiety. More participants reported reduced anxiety with the in-lab simulator ($M = 1.32 \pm 0.60$) than with the outdoor simulator ($M = 1.45 \pm 0.72$). We found *moderate* evidence ($BF = 0.32$) in favor of the null model, suggesting that there is no significant difference in the anxiety generated by the simulators.

Safety. Participants consider the in-lab simulator ($M = 4.58 \pm 0.76$) safer than the on-road simulator ($M = 4.16 \pm 0.86$). We found *moderate* evidence ($BF = 3.53$) against the null model, suggesting a moderate difference in favor of the in-lab simulator.

This may have been because the in-lab simulator's roads did not involve any real vehicles or people, e.g., P6 explained that "... there were more actual obstacles...to take into account [in the outdoor simulator] whereas the in-lab [simulator] had a preset number."

Trust. Participants reported their trust in the simulated autonomous driving to be higher in the on-road simulator ($M = 3.71 \pm 0.86$) than in the in-lab simulator ($M = 3.52 \pm 1.06$). We found *moderate* evidence ($BF = 0.31$) in favor of the null model, suggesting there is no significant difference in trust.

The preference for the on-road simulator may have been influenced by people's perception of each vehicle's performance. At the same time, this may have been a breakdown in face validity. For example, P14 said, "One of the survey questions asked 'how much do you trust this car' and I think I forgot to pretend that this was an AI driving the car when answering that..." Participants were informed during the on-road simulator session that a researcher would be

625 driving the car; their trust rating might have been an indicator that they trusted the driver to obey local traffic laws,
626 rather than an indicator of their trust in the simulated autonomous driving, as we had intended.
627

628
629 **6.1.2 Cooperation Behavior.** From recorded videos, we analyzed 32 participants' cooperative behaviors with the
630 vehicle at crosswalks. Video recordings from the same session were synchronized before the analysis. One researcher
631 watched recordings for each participant and labeled their behaviors in terms of the timing of cooperation behaviors for
632 each scenario. The researcher then noted down the behavioral changes, if any, for each participant between different
633 simulators.
634

635 15 participants cooperated with the vehicle perfectly in both simulation platforms in all scenarios, where they waited
636 until the pedestrians fully crossed the street or waited for clear non-crossing signals before instructing the vehicle to
637 proceed. During the first crossing scenario, 11 participants instructed the vehicle to proceed while the pedestrian was
638 about to cross; this led to six virtual collisions. One of these participants made the vehicle run over the same virtual
639 pedestrian in both in-lab and on-road simulators. Three participants did not wait for any virtual pedestrian to cross in
640 both in-lab and on-road simulators. Six participants who had cooperated perfectly in the first session made different
641 decisions in the second session; they chose not to wait for the crossing participants when they believed it was safe, and
642 one of them ended up running over the virtual pedestrian. One participant who ignored both crossing pedestrians in
643 in-lab simulators waited for one of the crossing pedestrians in the on-road simulator. Overall, 14 participants made
644 different decisions in the second session from the first session.
645

646 We feel compelled to point out that the fact that participants made poor crossing decisions is not a sign that the
647 study or the driving simulation platforms were designed poorly; instead, it is precisely these sorts of outcomes that
648 indicate the necessity for simulation platforms that enable studies with virtual pedestrians to be conducted prior to
649 putting real pedestrians in harm's way. We do not expect this means that participants would run over real people in
650 subsequent tests with real cars, but this does point out that participants are aware that they are not exposed to real
651 danger in driving simulators [25]; impatience and lack of conscientiousness amongst some portion of the population
652 are factors that any cooperative autonomous driving system would have to account for.
653
654
655
656

657
658 **6.1.3 Ordering Effect.** While we counterbalanced our study, we noticed some differences in cooperation behavior that
659 may be attributed to the ordering of simulators. Of the six participants who made mistakes in the first crossing scenario,
660 four participants were experiencing the on-road driving simulator. We hypothesize that the on-road driving simulator
661 is more overwhelming than the in-lab simulator to familiarize the participants with the study setup. P28 mentioned in
662 their post-study questionnaire that "Visual noise in outdoor sim [made] task completion more difficult but [was] more
663 realistic in that regard." It is worth mentioning that the three participants who made the mistake in the first crossing
664 scenario had limited (e.g., they had only their learner's permits) to no driving experience.
665
666

667 668 **6.2 Simulation Evaluation Results**

669
670 **6.2.1 Simulator Measures.** To compare the overall experience between the two simulators, we also asked participants
671 to rate their feelings about car performance, system enjoyment, discomfort, and immersion on a 5-point Likert scale.
672 These questions were based on questionnaires used during the validation process of Goedicke et al. [16]'s XR-OOM
673 system. We again ran a Bayesian factor analysis on the captured measures with the null hypothesis that there is no
674 difference between the platforms [31].
675

677 *Car Performance.* Participants considered the autonomous vehicle in the on-road simulator ($M = 4.35 \pm 0.71$) to
678 perform better than in the in-lab simulator ($M = 3.42 \pm 0.96$). We found *extreme* evidence ($BF = 1.67e+04$) against the
679 null model, suggesting a significant difference in favor of the on-road simulator. Participants felt that the in-lab car
680 simulator did not appear to drive smoothly and stopped rather abruptly at times and thus thought that the driving felt
681 more natural in the on-road simulator.
682

683
684 *System Enjoyment.* More participants reported increased levels of system enjoyment with the in-lab simulator ($M =$
685 3.84 ± 0.97) than with the outdoor simulator ($M = 3.03 \pm 0.98$). We found *extreme* evidence ($BF = 225.00$) against the
686 null model, suggesting a significant difference in favor of the in-lab simulator. Participants seemed to have preferred
687 the in-lab simulator because of the graphics quality and comfort. For example, when asked to describe their experience
688 with both simulators in the post-study questionnaire, P14 said, "The pedestrians seemed to "appear out of nowhere"
689 in the [on-road simulator], whereas it seemed like they were always part of the scenery in the [in-lab] simulator (i.e.,
690 came into view naturally in the simulator). ... [The] turns in the [in-lab simulator's] road felt unnatural/like the scenery
691 was clicked and dragged in front of my eyes, instead of me moving through the scenery."
692
693

694
695 *Discomfort.* Participants reported less discomfort with the in-lab simulator ($M = 1.71 \pm 0.86$) than the on-road simulator
696 ($M = 2.35 \pm 1.14$). We found *anecdotal* evidence ($BF = 2.24$) against the null model, suggesting a mild difference in favor
697 of the in-lab simulator. For the in-lab simulator, discomfort mainly arose from the unrealistic vehicle dynamics. For the
698 on-road simulator, many comments were related specifically to the XR headset ($n = 11$), which people found heavy and
699 uncomfortable. The occasional misalignment of virtual objects caused by bumps in the road also induced a considerable
700 amount of motion discomfort. For example, P13 said, "...both [simulators] cause some discomfort, but I think the outdoor
701 one is more uncomfortable due to the [pass-through] being very [shaky] and more motion-sickness-inducing."
702
703

704
705 *Immersion.* As shown in Figure 5, participants considered the in-lab simulator ($M = 4.00 \pm 0.77$) to be more immersive
706 than the on-road simulator ($M = 3.74 \pm 1.06$). We found *anecdotal* evidence ($BF = 0.040$) in favor of the null model.
707 However, in the post-study questionnaire, when both simulators were presented on the same Likert scale, the numbers of
708 participants in favor of either simulator were the same. Figure 6 shows that 14 participants thought the in-lab simulator
709 was more immersive, and 14 thought the on-road simulator was more immersive. Three participants considered both
710 simulators equally immersive.
711

712 Two participants reported verbally during the on-road session that it was difficult to distinguish virtual from real
713 pedestrians.
714

715 Many participants reported difficulty with the weight and technical maturity of the XR headset in the outdoor
716 simulator, which impeded their attention and may have contributed to its lower immersion rating. P8 said, "headset
717 jitter made visuals blurry, which impaired decision-making/attention." P2 said, "[Putting] something on my head is so
718 uncomfortable. I couldn't [focus on] the view or be relaxed. The [in-lab] simulator is not so real[,] but I could be so
719 relaxed."
720

721 **6.2.2 In simulator behavior.** We weighed the observed and reported behavior of participants to assess differences in
722 behavior introduced by each platform.
723

724 *Natural Head Motions.* We noticed that some participants' head pose motion patterns were different when experi-
725 encing the two simulation platforms. When sitting in the in-lab static simulator, participants tilted their heads mostly
726 during staged events to track the motion of virtual pedestrians. For the rest of the ride, they faced forward. However,
727

728

for the on-road simulator, the participants' head motions were more varied, and participants naturally looked to view the surroundings more often.

Decision-making. For the crosswalk cooperation study, the key measure is the participant's decision-making around whether the vehicle should go or not go. Participants reflected that the level of complexity and severity in decision-making was greater in the on-road simulator. The staged interactions are identical in both sessions, but researchers have less control over surrounding factors during the on-road sessions. Other road users, including real pedestrians, other vehicles, and geese native to Roosevelt Island complicated the staged scenarios.

P5 said, "I think that in the on-road simulator, I was a little more nervous because real people were on the street. The in-lab simulator does not deal with real people, so any mistake I make does not have as much weight." P28 said, "Outdoor provides [a] more generally immersive feeling and is the only one in which I can realistically feel unsafe, which is a positive in terms of validity. Visual noise in outdoor sim makes task completion more difficult but is more realistic in that regard." P31 said, "I felt anxious after saying to proceed and wondering how the pedestrians would move afterwards." In contrast, P11 reflected: "Perhaps due to its nature of being [an in-lab] simulator, I felt at ease at all times." P20 reflected: "the [polished nature] of the [in-lab] simulator makes the experience [feel] more entertaining/performative rather than realistic. ... The [on-road] simulator feels more like a functional approximation of an actual autonomous driving car, while the indoor simulator feels more like a fun, polished experience." P21 reflected: "I felt much safer/less anxious in the indoor simulator which also probably means it was less realistic." P29 reflected: "[the] indoor [simulator] seems to have lower stakes, even though it was the same virtual people.




Fig. 6. After experiencing both simulators, participants also directly compared the in-lab simulator and the on-road simulator on the same Likert scale. Participants' responses are shown as histograms.

7 TECHNICAL VALIDATION

Since the Portobello system is an infrastructure meant to be used in conjunction with existing platforms, the technical performance largely depends on the system on top of Portobello. Therefore, we investigated the change in performance of the on-road platform after adapting Portobello. Rendering a total of 12 virtual pedestrians, the headset runs at 60 FPS consistently with a display latency of around 35ms, which is on par with the original XR-OOM system. The localization frequency is 10 Hz, limited by the 10Hz LiDAR. The fact that there is no change in performance is expected since the Portobello system operates on a computer separated from the original XR-OOM system.

8 DISCUSSION

8.1 LiDAR-based vs. GPS/IMU-based systems

Our LiDAR-based Portobello system resolves the two challenges posed by traditional GPS/IMU-based systems in surrounding context-based interaction staging: localization accuracy and design hardships. LiDAR-based localization systems work reliably in cities where buildings serve as landmarks instead of GPS signal blockers. A point cloud 3D map of the environment generated from LiDAR-based SLAM algorithm saves designers from staging using hardcoded coordinates, and simplifies the design process to drag-and-drop within the map. Lastly, we want to point out that the sensors are not mutually exclusive. We can fuse in GPS data as an additional data source into the LiDAR-based algorithm if necessary.

8.2 Platform Portability Challenges

The primary goal of this research effort was to establish a proof-of-concept demonstration of *platform portability* through our Portobello system. *Platform portability* is necessary to run twinned studies across different platforms, which is desirable because the in-lab simulator can help establish causal differences across experimental conditions using well-controlled studies, and the on-road simulator can help validate the ecological validity of such study results when the same study conditions are moved into the less-controlled environment of the real world. By incorporating robotics mapping, sensing, and localization capability in the Portobello system, we can set up twinned studies to run in different environments. We believe that this is of relevance to any of the simulators discussed in Section 2.2, which could be deployed on top of Portobello.

8.2.1 Randomness in the Wild. For on-road simulators, randomness persists throughout the entire study. For example, during mapping, we generated a snapshot of the test area. While major landmarks such as buildings and land topology will not change significantly over time, the map also captures transient objects (e.g., parked cars). Such randomness may cause inaccuracy in real-time localization.

During the study, unplanned events were the most salient form of randomness. Unplanned occurrences and interruptions from the real world may increase immersion. P19 reported that "the extra people and cars in the outdoor study made the experience feel more immersive and interesting." However, randomness also brings concerns regarding study reproducibility. We tried to eliminate the co-occurrence of planned and unplanned events by staging events in less populated areas. In general, we encourage researchers to plan for all possible unplanned events during the study design phase.

8.2.2 Event Timing. One major challenge we faced with the on-road simulator is the trigger and timing of staged events, and we foresee such a problem persisting in future similar studies. For in-lab simulators, vehicle speed and travel distances can be coded in detail. For on-road simulators, it can be difficult to maintain the same speed curve as in-lab simulators due to obstacles and unplanned events. In our study, we expect the vehicle to stop at the crosswalk simultaneously as a virtual pedestrian reaches the stop sign on the sidewalk. The researcher who operates the car has access to a mini-display monitoring the location of the virtual pedestrians and adjusting vehicle speed accordingly. However, we have noticed that participants made different decisions across platforms due to event timing differences.

The timing misalignment between the simulators is the natural consequence of the intentional difference between running studies in a controlled environment (the lab) and an uncontrolled environment (the real world). We are arguing that it is desirable to run both kinds of studies and that it is easier to do this if *platform portability* exists. The in-lab

833 simulator is more suitable for quantitative analysis, and the on-road simulator is more suitable for qualitative analysis
834 of the factors that complicate the outcomes learned from the more controlled in-lab simulator.
835

836 8.3 Platform Effects 837

838 Our study results were intended to help us understand *whether and how* our twinned studies were the same across
839 the two platforms and to help us understand the differences across the platforms, a comparison made possible by the
840 Portobello system. For research purposes, it would be best if the study results between the two platforms were similar
841 (i.e., that the platform effects were negligible), or at least that the results were biased consistently across the platforms
842 (i.e., that the platform effects were predictable).
843

844 From the study results, we can see that around half of the participants made similar decisions in both the in-lab
845 simulator and the on-road system. Notably, one participant made the same mistake in each simulator. The fact that
846 some participants made different decisions in different simulators indicates that participants' decision-making was not
847 affected by their existing knowledge about the study. Even if they were aware of the pedestrians' behaviors from the
848 first session, they took into account the event timing difference and made the most appropriate decision at the moment.
849

850 Some notable differences between the platforms were centered around the participant's decision-making behavior
851 and their resulting trust in the autonomous driving system. In some ways, there were indications that the on-road
852 simulator might have failed to maintain face validity for at least one of the participants; their trust rating might have
853 shown their trust in the research driver rather than their trust in the simulated autonomous driving, as we had intended.
854 On the other hand, many participants reported greater weight in the decision-making around whether the vehicle
855 should go or not go in the on-road simulator, a sign that face validity is higher in the real-world environment.
856
857

858 While our participants favored the experience of the in-lab simulation platform on a whole, most of their complaints
859 pertained to aspects of the XR system—the weight of the headset, the jitter in the display—which are likely to improve
860 with advancing technology. It seems like the naturalism of the on-road environment is more likely to yield naturalistic
861 behaviors than the in-vehicle environment, and hence this platform can help industrial and academic researchers better
862 understand how people will engage with autonomous vehicle technologies in the real world.
863
864

865 8.4 Limitations and Future Work 866

867 Some limitations in the study results are inherent to a driving simulation study. In this section, we focus on limitations
868 in the design and execution of studies using the XR driving simulation system augmented by Portobello that should be
869 accounted for and discuss future developments that could improve such systems.
870

871 *8.4.1 Real-time adjustment of Depth Ordering.* Although we render the occlusion of virtual objects caused by the
872 research vehicle (e.g. we do not render the pedestrians over the front pillar of the car), the current system does not
873 provide the same occlusion for runtime dynamic objects. If a bus drove between the research vehicle and the location
874 where the virtual pedestrians were supposed to appear, for example, participants would see the virtual pedestrians in
875 front of the bus. To correct the depth order, future systems could use real-time LiDAR scans of the environment.
876
877

878 *8.4.2 Pedestrian Appearance.* We capped the maximum rendering distance for on-road virtual objects for technical
879 reasons; distant virtual objects are less salient and require better alignment between the in-vehicle and out-of-vehicle
880 reference frames to be placed believably in the mixed-reality view. The artifice of having pedestrians suddenly appear,
881 however, may affect study results. One participant (P24) said, "I think being able to see the passengers from further away
882 in the in-lab simulator made a big difference because, by the time I got to the intersection, it was easier to anticipate
883

885 their movements." Future technology could improve the motion parallax issues, enabling longer rendering distances
886 and smoother transitions when virtual objects approach the rendering threshold.
887

888 **8.4.3 Headset Discomfort.** Many participants complained about the bulkiness and narrow field-of-view of the headset.
889 This platform-level discomfort was pronounced enough that it drowned out our ability to measure experiential aspects
890 (system enjoyment, discomfort) of the autonomous driving scenario. While the weight and limitations of the XR headsets
891 were beyond our control, we believe that anticipated advancements in XR headset technology are necessary to use
892 these systems in experiments wherein the experiential aspects of automated driving are critical.
893
894

895 **8.4.4 Simulating Autonomous Driving on Road.** In our current study, we informed the participants that there was an
896 actual driver behind the scenes in the on-road simulator. The driver's maneuver sound easily breaks the AV illusion for
897 participants who have previous experience with AV. Future research can benefit from disguising the driver by playing
898 the recorded AV sound profile during acceleration and deceleration to cover the driver's maneuver sound.
899
900

901 **9 CONCLUSION**

902
903 Driving simulations can be used to create scenarios for driving interactions, which enable researchers to better
904 understand how people will behave and respond to future driving scenarios. In this work, we presented the Portobello
905 system, an on-road driving simulation infrastructure that enables *platform portability*. By advancing the capabilities of
906 driving simulators, we can better anticipate what aspects of driving interaction will work well or poorly.
907

908 This paper outlines the first-ever deployment of twinned studies across in-lab and on-road simulators. We found that
909 participants preferred the experience of the in-lab simulator but displayed more natural head movements in the on-road
910 simulator; they also reported that the decisions made in the on-road system carried more weight. Based on our findings,
911 we suggest researchers working in driving simulations also take the twinning of studies approach: they should first run
912 studies within a controlled, in-lab environment to collect statistical measures and form hypotheses and then port their
913 studies to a less-controlled, on-road simulator and test their hypotheses in a more complex, realistic environment.
914

915 This experiment looking at the platform-driven influences on study outcomes demonstrates the utility of *platform*
916 *portability*, as the same study design was able to be run both in-lab and on-road. This was made possible by the
917 Portobello system's common model and vehicle localization; using robotics mapping and localization technology, we
918 were able to capture surrounding environments for study and event staging for our on-road simulator. We anticipate
919 that Portobello will advance the state of open-source and accessible driving simulation by extending the reach and
920 translational capabilities of VR and XR driving simulation systems and thus, in turn, enable wider-scale development
921 and testing of safe driving systems.
922
923
924

925 **OPEN SCIENCE**

926
927 The source code for Portobello has been made publicly available. It can be accessed via the following link:
928 <https://anonymous.4open.science/r/Portobello-588C/>.
929

930 **ACKNOWLEDGMENTS**

931
932 We thank all study participants. This research was made possible by sponsorship by Toyota Research Institute. Studies
933 were conducted under Cornell Tech's IRB Protocol #IRB0008105. This work was supported by a fellowship of the
934 German Academic Exchange Service (DAAD).
935
936

REFERENCES

- [1] GENIVI Alliance. 2021. Genivi Vehicle Simulator. <https://github.com/GENIVI/genivi-vehicle-simulator>.
- [2] Sonia Baltodano, Srinath Sibi, Nikolas Martelaro, Nikhil Gowda, and Wendy Ju. 2015. The RRADS platform: a real road autonomous driving simulator. In *Proceedings of the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. ACM, New York, NY, USA, 281–288.
- [3] Karlin Bark, Cuong Tran, Kikuo Fujimura, and Victor Ng-Thow-Hing. 2014. Personal navi: Benefits of an augmented reality navigational aid using a see-thru 3d volumetric hud. In *Proceedings of the 6th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. ACM, New York, NY, USA, 1–8.
- [4] Giovanni Bertolini and Dominik Straumann. 2016. Moving in a moving world: a review on vestibular motion sickness. *Frontiers in neurology* 7 (2016), 14.
- [5] David Bethge, Thomas Kosch, Tobias Grosse-Puppenthal, Lewis L. Chuang, Mohamed Kari, Alexander Jagaciak, and Albrecht Schmidt. 2021. VEmotion: Using Driving Context for Indirect Emotion Prediction in Real-Time. In *The 34th Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (*UIST '21*). Association for Computing Machinery, New York, NY, USA, 638–651. <https://doi.org/10.1145/3472749.3474775>
- [6] Björn Blissing, Fredrik Bruzelius, and Johan Ölvander. October 23-25, 2013. Augmented and mixed reality as a tool for evaluation of vehicle active safety systems. In *RSS2013 Road Safety and Simulation-International Conference*. Aracne, Rome, Italy, 7 pages.
- [7] Hyung-jun Cho and Gerard J. Kim. 2020. RoadVR: Mitigating the Effect of Vection and Sickness by Distortion of Pathways for In-Car Virtual Reality. In *Proceedings of the 26th ACM Symposium on Virtual Reality Software and Technology (VRST '20)*. Association for Computing Machinery, New York, NY, USA, 1–3. <https://doi.org/10.1145/3385956.3422115>
- [8] Mark Colley, Julian Britten, Simon Demharter, Tolga Hisir, and Enrico Rukzio. 2022. Feedback Strategies for Crowded Intersections in Automated Traffic – A Desirable Future?. In *Proceedings of the 14th International Conference on Automotive User Interfaces and Interactive Vehicular Applications* (Seoul, Republic of Korea) (*AutomotiveUI '22*). Association for Computing Machinery, New York, NY, USA, 243–252. <https://doi.org/10.1145/3543174.3545255>
- [9] Mark Colley, Benjamin Eder, Jan Ole Rixen, and Enrico Rukzio. 2021. Effects of Semantic Segmentation Visualization on Trust, Situation Awareness, and Cognitive Load in Highly Automated Vehicles. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) (*CHI '21*). Association for Computing Machinery, New York, NY, USA, Article 155, 11 pages. <https://doi.org/10.1145/3411764.3445351>
- [10] Mark Colley, Pascal Jansen, Enrico Rukzio, and Jan Gugenheimer. 2021. Swivr-car-seat: Exploring vehicle motion effects on interaction quality in virtual reality automated driving using a motorized swivel seat. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 4 (2021), 1–26.
- [11] Mark Colley, Pascal Jansen, Enrico Rukzio, and Jan Gugenheimer. 2022. SwiVR-Car-Seat: Exploring Vehicle Motion Effects on Interaction Quality in Virtual Reality Automated Driving Using a Motorized Swivel Seat. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 4, Article 150 (dec 2022), 26 pages. <https://doi.org/10.1145/3494968>
- [12] Joost De Winter, Peter M van Leeuwen, Riender Happee, et al. 2012. Advantages and disadvantages of driving simulators: A discussion. In *Proceedings of measuring behavior*, Vol. 2012. 8th.
- [13] Sebastien Demmel, Gregoire S Larue, Andry Rakotonirainy, and Dominique Gruyer. 2011. Using cooperative vehicle systems to collect near collision incidents events: a simulation study. In *2nd International Workshop on Traffic Data Collection & its Standardisation*.
- [14] Lorin Dole and Wendy Ju. 2019. Face and Ecological Validity in Simulations: Lessons from Search-and-Rescue HRL. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (*CHI '19*). Association for Computing Machinery, New York, NY, USA, 1–8. <https://doi.org/10.1145/3290605.3300681>
- [15] Florin-Timotei Ghiurău, Mehmet Aydın Baytaş, and Casper Wickman. 2020. ARCAR: On-Road Driving in Mixed Reality by Volvo Cars. In *Adjunct Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology* (Virtual Event, USA) (*UIST '20 Adjunct*). Association for Computing Machinery, New York, NY, USA, 62–64. <https://doi.org/10.1145/3379350.3416186>
- [16] David Goedicke, Alexandra W.D. Bremers, Sam Lee, Fanjun Bu, Hiroshi Yasuda, and Wendy Ju. 2022. XR-OOM: MiXed Reality Driving Simulation with Real Cars for Research and Design. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (*CHI '22*). Association for Computing Machinery, New York, NY, USA, Article 107, 13 pages. <https://doi.org/10.1145/3491102.3517704>
- [17] David Goedicke, Jany Li, Vanessa Evers, and Wendy Ju. 2018. Vr-oom: Virtual reality on-road driving simulation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–11.
- [18] Paul Allan Green. 2005. *How driving simulator data quality can be improved*. University of Michigan Transportation Research Institute, Human Factors Division.
- [19] Jeff Greenberg, Bruce Artz, and Larry Cathey. 2003. The effect of lateral motion cues during simulated driving. *Proceedings of DSC North America* (2003), 8–10.
- [20] Kim R Hammel, Donald L Fisher, and Anuj K Pradhan. 2002. Verbal and spatial loading effects on eye movements in driving simulators: A comparison to real world driving. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 46. SAGE Publications, Los Angeles, CA, USA, 2174–2178.

- 989 [21] Philipp Hock, Sebastian Benedikter, Jan Gugenheimer, and Enrico Rukzio. 2017. CarVR: Enabling In-Car Virtual Reality Entertainment. In *Proceedings*
990 *of the 2017 CHI Conference on Human Factors in Computing Systems* (Denver, Colorado, USA) (CHI '17). Association for Computing Machinery, New
991 York, NY, USA, 4034–4044. <https://doi.org/10.1145/3025453.3025665>
- 992 [22] Philipp Hock, Mark Colley, Ali Askari, Tobias Wagner, Martin Baumann, and Enrico Rukzio. 2022. Introducing VAMPIRE – Using Kinaesthetic
993 Feedback in Virtual Reality for Automated Driving Experiments. In *Proceedings of the 14th International Conference on Automotive User Interfaces and*
994 *Interactive Vehicular Applications* (Seoul, Republic of Korea) (*AutomotiveUI '22*). Association for Computing Machinery, New York, NY, USA, 204–214.
995 <https://doi.org/10.1145/3543174.3545252>
- 996 [23] Ming Hou, Karthik Mahadevan, Sowmya Somanath, Ehud Sharlin, and Lora Oehlberg. 2020. Autonomous Vehicle-Cyclist Interaction: Peril and
997 Promise. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). Association for Computing
998 Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376884>
- 999 [24] Maura Houtenbos. 2009. Interacting drivers at intersections: what can make them more safe and more efficient. In *Proceedings of the 21st International*
1000 *Technical Conference on the Enhance Safety Vehicles*. Citeseer, Citeseer.
- 1001 [25] Samantha L Jamson and A Hamish Jamson. 2010. The validity of a low-cost simulator for the assessment of the effects of in-vehicle information
1002 systems. *Safety Science* 48, 10 (2010), 1477–1483.
- 1003 [26] Pascal Jansen, Julian Britten, Thilo Segschneider, Mark Colley, and Enrico Rukzio. 2023. AutoVis: Enabling Mixed-Immersive Analysis of Automotive
1004 User Interface Interaction Studies. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23).
1005 Association for Computing Machinery, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3544548.3580760> conditionally accepted.
- 1006 [27] Pascal Jansen, Mark Colley, and Enrico Rukzio. 2022. A Design Space for Human Sensor and Actuator Focused In-Vehicle Interaction Based on a
1007 Systematic Literature Review. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6, 2, Article 56 (jul 2022), 51 pages. <https://doi.org/10.1145/3534617>
- 1008 [28] Harold Jeffreys. 1998. *The theory of probability*. OUP Oxford.
- 1009 [29] Arthur Juliani, Vincent-Pierre Berges, Ervin Teng, Andrew Cohen, Jonathan Harper, Chris Elion, Chris Goy, Yuan Gao, Hunter Henry, Marwan
1010 Mattar, et al. 2018. Unity: A general platform for intelligent agents. *arXiv preprint arXiv:1809.02627* (2018).
- 1011 [30] Mohamed Kari, Tobias Grosse-Puppendahl, Alexander Jagaciak, David Bethge, Reinhard Schütte, and Christian Holz. 2021. SoundsRide: Affordance-
1012 Synchronized Music Mixing for In-Car Audio Augmented Reality. In *The 34th Annual ACM Symposium on User Interface Software and Technology*
1013 (Virtual Event, USA) (*UIST '21*). Association for Computing Machinery, New York, NY, USA, 118–133. <https://doi.org/10.1145/3472749.3474739>
- 1014 [31] Robert E Kass and Adrian E Raftery. 1995. Bayes factors. *Journal of the American statistical association* 90, 430 (1995), 773–795.
- 1015 [32] SeungJun Kim and Anind K Dey. 2009. Simulated augmented reality windshield display as a cognitive mapping aid for elder driver navigation. In
1016 *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 133–142.
- 1017 [33] Kenji Koide, Jun Miura, and Emanuele Menegatti. 2019. A portable three-dimensional LIDAR-based system for long-term and wide-area people
1018 behavior measurement. *International Journal of Advanced Robotic Systems* 16, 2 (2019), 1729881419841532.
- 1019 [34] Georgios K Kountouriotis and Natasha Merat. 2016. Leading to distraction: Driver distraction, lead car, and road environment. *Accident Analysis &*
1020 *Prevention* 89 (2016), 22–30.
- 1021 [35] Sven Krome, David Goedicke, Thomas J Matarazzo, Zimeng Zhu, Zhenwei Zhang, JD Zamfirescu-Pereira, and Wendy Ju. 2019. How people
1022 experience autonomous intersections: taking a first-person perspective. In *Proceedings of the 11th International Conference on Automotive User*
1023 *Interfaces and Interactive Vehicular Applications*. ACM, New York, NY, USA, 275–283.
- 1024 [36] Grégoire S Larue, Christian Wullems, Michelle Sheldrake, and Andry Rakotonirainy. 2018. Validation of a driving simulator study on driver behavior
1025 at passive rail level crossings. *Human factors* 60, 6 (2018), 743–754.
- 1026 [37] Markus Lemmer, Stefan Schwab, and Sören Hohmann. 2020. Driver interaction at intersections: A hybrid dynamic game based model. In *2020 IEEE*
1027 *International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, IEEE, New York, NY, USA, 2269–2276.
- 1028 [38] Karthik Mahadevan, Elaheh Sanoubari, Sowmya Somanath, James E. Young, and Ehud Sharlin. 2019. AV-Pedestrian Interaction Design Using a
1029 Pedestrian Mixed Traffic Simulator. In *Proceedings of the 2019 on Designing Interactive Systems Conference* (San Diego, CA, USA) (*DIS '19*). Association
1030 for Computing Machinery, New York, NY, USA, 475–486. <https://doi.org/10.1145/3322276.3322328>
- 1031 [39] Mark McGill, Graham Wilson, Daniel Medeiros, and Stephen Anthony Brewster. 2022. PassengXR: A Low Cost Platform for Any-Car, Multi-User,
1032 Motion-Based Passenger XR Experiences. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology* (Bend, OR,
1033 USA) (*UIST '22*). Association for Computing Machinery, New York, NY, USA, Article 2, 15 pages. <https://doi.org/10.1145/3526113.3545657>
- 1034 [40] Richard Morey. 2022. *Richardmorey/BayesFactor*. Retrieved 6 September, 2022 from <https://github.com/richardmorey/BayesFactor>
- 1035 [41] Wolfgang Narzt, Gustav Pomberger, Alois Ferscha, Dieter Kolb, Reiner Müller, Jan Wieghardt, Horst Hörtnner, and Christopher Lindinger. 2006.
1036 Augmented reality navigation systems. *Universal Access in the Information Society* 4 (2006), 177–187.
- 1037 [42] Fred GWC Paas and Jeroen JG Van Merriënboer. 1993. The efficiency of instructional conditions: An approach to combine mental effort and
1038 performance measures. *Human factors* 35, 4 (1993), 737–743.
- 1039 [43] Maria Pinto, Viola Cavallo, and Théophile Ohlmann. 2008. The development of driving simulators: Toward a multisensory solution. *Le travail*
1040 *humain* 71, 1 (2008), 62–95.
- 1041 [44] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, Andrew Y Ng, et al. 2009. ROS: an open-source
1042 Robot Operating System. In *ICRA workshop on open source software*, Vol. 3. Kobe, Japan, Japan, 5.
- 1043 [45] Michael Quinlan, Tsz-Chiu Au, Jesse Zhu, Nicolae Sturca, and Peter Stone. 2010. Bringing simulation to life: A mixed reality autonomous intersection.
1044 In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, New York, NY, USA, 6083–6088. <https://doi.org/10.1109/IROS.2010>

- 1041 5651993
- 1042 [46] Andreas Riegler, Andreas Riener, and Clemens Holzmann. 2020. A research agenda for mixed reality in automated vehicles. In *19th International*
- 1043 *Conference on Mobile and Ubiquitous Multimedia*. ACM, New York, NY, USA, 119–131.
- 1044 [47] Daniele Ruscio, Luca Bascetta, Alessandro Gabrielli, Matteo Matteucci, Dedy Ariansyah, Monica Bordegoni, Giandomenico Caruso, and Lorenzo
- 1045 Mussone. 2017. Collection and comparison of driver/passenger physiologic and behavioural data in simulation and on-road driving. In *2017 5th IEEE*
- 1046 *International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*. IEEE, New York, NY, USA, 403–408.
- 1047 [48] Hiroaki Sawano and Minoru Okada. 2005. A car-navigation system based on augmented reality. In *ACM SIGGRAPH 2005 Sketches*. ACM, New York,
- 1048 NY, USA, 119–es.
- 1049 [49] Mark C Schall Jr, Michelle L Rusch, John D Lee, Jeffrey D Dawson, Geb Thomas, Nazan Aksan, and Matthew Rizzo. 2013. Augmented reality cues
- 1050 and elderly driver hazard perception. *Human factors* 55, 3 (2013), 643–658.
- 1051 [50] Friederike Schneemann and Irene Gohl. 2016. Analyzing driver-pedestrian interaction at crosswalks: A contribution to autonomous driving in
- 1052 urban environments. In *2016 IEEE intelligent vehicles symposium (IV)*. IEEE, IEEE, New York, NY, USA, 38–43.
- 1053 [51] Tixiao Shan, Brendan Englot, Drew Meyers, Wei Wang, Carlo Ratti, and Rus Daniela. 2020. LIO-SAM: Tightly-coupled Lidar Inertial Odometry via
- 1054 Smoothing and Mapping. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, IEEE, New York, NY, USA, 5135–5142.
- 1055 [52] Yoshisuke Tateyama, Hiroki Yamada, Junpei Noyori, Yukihiro Mori, Keiichi Yamamoto, Tetsuro Ogi, Hidekazu Nishimura, Noriyasu Kitamura, and
- 1056 Harumi Yashiro. 2010. Observation of drivers' behavior at narrow roads using immersive car driving simulator. In *Proceedings of the 9th ACM*
- 1057 *SIGGRAPH Conference on Virtual-Reality Continuum and its Applications in Industry (VRCAI '10)*. Association for Computing Machinery, New York,
- 1058 NY, USA, 391–396. <https://doi.org/10.1145/1900179.1900263>
- 1059 [53] Marcus Tonnis, Christian Sandor, Gudrun Klinker, Christian Lange, and Heiner Bubb. 2005. Experimental evaluation of an augmented reality
- 1060 visualization for directing a car driver's attention. In *Fourth IEEE and ACM International Symposium on Mixed and Augmented Reality (ISMAR'05)*.
- 1061 IEEE, New York, NY, USA, 56–59.
- 1062 [54] Tamara von Sawitzky, Philipp Wintersberger, Andreas Riener, and Joseph L. Gabbard. 2019. Increasing Trust in Fully Automated Driving: Route
- 1063 Indication on an Augmented Reality Head-up Display. In *Proceedings of the 8th ACM International Symposium on Pervasive Displays (Palermo, Italy)*
- 1064 *(PerDis '19)*. Association for Computing Machinery, New York, NY, USA, Article 6, 7 pages. <https://doi.org/10.1145/3321335.3324947>
- 1065 [55] Marcel Walch, Stacey Li, Ilan Mandel, David Goedicke, Natalie Friedman, and Wendy Ju. 2020. Crosswalk cooperation: A phone-integrated
- 1066 driver-vehicle cooperation approach to predict the crossing intentions of pedestrians in automated driving. In *12th International Conference on*
- 1067 *Automotive User Interfaces and Interactive Vehicular Applications*. ACM, New York, NY, USA, 74–77.
- 1068 [56] Marcel Walch, Marcel Woide, Kristin Mühl, Martin Baumann, and Michael Weber. 2019. Cooperative Overtaking: Overcoming Automated
- 1069 Vehicles' Obstructed Sensor Range via Driver Help. In *Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive*
- 1070 *Vehicular Applications (Utrecht, Netherlands) (AutomotiveUI '19)*. Association for Computing Machinery, New York, NY, USA, 144–155. <https://doi.org/10.1145/3342197.3344531>
- 1071 [57] Rachael A. Wynne, Vanessa Beanland, and Paul M. Salmon. 2019. Systematic review of driving simulator validation studies. *Safety Science* 117
- 1072 (2019), 138–151. <https://doi.org/10.1016/j.ssci.2019.04.004>
- 1073 [58] Dohyeon Yeo, Gwangbin Kim, and SeungJun Kim. 2019. MAXIM: Mixed-reality Automotive Driving XIMulation. In *2019 IEEE International*
- 1074 *Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. IEEE, IEEE, New York, NY, USA, 460–464.
- 1075 [59] Dohyeon Yeo, Gwangbin Kim, and Seungjun Kim. 2020. Toward immersive self-driving simulations: Reports from a user study across six platforms.
- 1076 In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–12.