**Evaluating the Effects of a Programming Error on a Virtual Environment
Measure of Spatial Navigation Behavior**

Steven M. Weisberg¹, Victor R. Schinazi^{2,3}, Andrea Ferrario^{4,5}, and Nora S. Newcombe⁶

¹Department of Psychology, University of Florida

²Department of Psychology, Bond University

³Future Health Technologies, Singapore-ETH Centre, Campus for Research Excellence and
Technological Enterprise (CREATE)

⁴Department of Management, Technology, and Economics, ETH Zurich

⁵Mobiliar Lab for Analytics at ETH, ETH Zurich

⁶Department of Psychology, Temple University

Author Note

Steven M. Weisberg <https://orcid.org/0000-0001-7003-6082>

Victor R. Schinazi <https://orcid.org/0000-0002-2345-2806>

Andrea Ferrario <https://orcid.org/0000-0001-9968-9474>

Nora S. Newcombe <https://orcid.org/0000-0002-7004-6046>

Correspondence concerning this article should be addressed to Steven M. Weisberg, Department of Psychology, University of Florida, 945 Center Dr., Gainesville, FL 32608. E-mail: stevenweisberg@ufl.edu.

This work was supported by the National Institute of Aging (K01AG070333-01) to SMW and by the National Science Foundation (EHR 1660996) to NSN. The authors thank Derek J. Huffman for alerting us to the abnormalities resulting in bug discovery. We also thank Youmna Hassan, Brandi Reynolds, Rachel Geller, Amanda Myers, Sabrina Karjack, Madeleine Frazier, Sarah Hendricks, Elisabeth Boyce-Jacino, and Mitchell Decasere for help with data collection and Katherine Collier for performing a final code review.

All data, materials, and code are available at the links provided in the manuscript. This experiment was pre-registered on the Open Science Framework: https://osf.io/8ydzj/?view_only=e0a9b80d2cd347f4891e404939c9259b

Abstract

Relying on shared tasks and stimuli to conduct research can enhance the replicability of findings and allow a community of researchers to collect large data sets across multiple experiments. This approach is particularly relevant for experiments in spatial navigation, which often require the development of unfamiliar large-scale virtual environments to test participants. One challenge with shared platforms is that undetected technical errors, rather than being restricted to individual studies, become pervasive across many studies. Here, we discuss the discovery of a software bug in a virtual environment platform used to investigate individual differences in spatial navigation: Virtual Silcton. The bug, which was difficult to detect for several reasons, resulted in storing the absolute value of a direction in a pointing task rather than the signed direction and rendered the original sign of the direction unrecoverable. To assess the impact of the bug on published findings, we collected a new data set for comparison. Results revealed that although the bug caused suppression in pointing errors and had different effects across people (less accurate navigators had more suppression), the effect of the bug on published data is small, partially explaining the difficulty in detecting the bug. We also used the new data set to develop a tool that allows researchers who have previously used Virtual Silcton to evaluate the impact of the bug on their findings. We summarize the ways that shared open materials, shared data, and collaboration can pave the way for better science to prevent errors in the future.

Keywords: Open data, reproducible science, spatial navigation, replication

Evaluating the Effects of a Programming Error on a Virtual Environment Measure of Spatial Navigation Behavior

The 'replication crisis' in psychology has given rise to a new model of research. Data are now shared at higher rates, code used in the generation or analysis of results are stored transparently on the web, and experimental protocols are designed and shared on web-based platforms (e.g., Gorilla, Pavlovia) using a growing variety of tools (e.g., Psychopy, jsPsych). Psychological research has never been more open, transparent, and reproducible. However, open data, open code and shared materials are only useful insofar as they are accurate. Although mistakes may be easier to catch with open materials, one under-appreciated repercussion is that the more frequently materials are shared, the higher the chance that errors will proliferate across experiments. This illustrates a key difference between reproducibility and reliability. A particular finding may be reproducible using the same code and data, but may not be error-free. Without critical evaluation of data and code, we cannot be certain of the type and extent of errors that may lurk under the surface of publicly available information. Until recently, such errors were limited to individual papers and may have been unlikely to be found. The challenge of using open data and code is that while there is a higher chance that errors may be more discoverable, these same errors may also have a wider impact on the scientific literature until they are discovered.

In a new scientific reality of data-sharing, the detection and prevention of errors has increasing importance, and deserves thorough discussion. The motivation behind the current paper was the discovery of a >10 year-old programming error (a bug) in a virtual environment and data collection tool used to investigate individual differences in navigation, known as Virtual Silton (Weisberg & Newcombe, 2016; Weisberg et al., 2014). The bug introduced errors in data recording and affected multiple research groups, including ours. We have three goals in writing this paper. First, we provide an overview of the Virtual Silton bug and estimate its impact on the published literature. Second, we

develop a tool that allows researchers who have previously used Virtual Silcton to evaluate the impact of the bug on their findings. Third, we offer a model for reporting on such bugs along with our thoughts on the challenges and opportunities of dealing with sharing applications and code.

In service of the first goal, we describe the bug that was present in Virtual Silcton, the reasons we believe it went undetected until recently, and the impact that the bug may have had on published Virtual Silcton results. We conducted a large-scale replication study to evaluate the impact, using the original experiment method (Weisberg et al., 2014). We found that data compiled with the buggy code is highly correlated with correct data. We further exploit this high (but not perfect) correlation to estimate the effects of the bug on several datasets in published manuscripts. Here again, the results show that the difference in findings between buggy and correct data are minimal. In service of the second goal, we present efforts to reconstruct the correct data using the buggy data, which resulted in using a simple one-variable regression. These efforts culminated in a script that applies the correction to existing Virtual Silcton data. The tool is now provided in the Virtual Silcton Analysis Github repository (https://github.com/smweis/Virtual_Silcton_Analysis). In service of the third goal, we reflect on our own process of error discovery as it relates to broader efforts in programming error detection and prevention, the collection of replication data, and the potential impact of different types of programming errors.

Virtual Silcton

The Virtual Silcton platform includes a series of tasks designed to measure individual differences in spatial navigation in virtual environments (map view in Figure 1, screenshots in Figure 2). Current documentation is available at http://stevenmweisberg.com/Virtual_Silcton_Documentation. Virtual Silcton is based on a previous real world navigation experiment by Schinazi et al. (2013). In Virtual Silcton, participants are asked to explore and learn the names and positions of eight target

buildings along two main routes and two connecting routes. Specifically, participants are asked to walk along demarcated paths in the virtual environment, and then are tested using several different measures to see how accurately they learned the directions and distances between the target buildings. The two main tasks are a pointing task and a map building task. Virtual Silcton has been widely used by our group (Weisberg, Schinazi, Newcombe) and others (Blacker et al., 2017; Galati et al., 2017; Nazareth et al., 2018; Pagkratidou et al., 2020; Ruginski et al., 2019; Sutton et al., 2014; Weisberg & Newcombe, 2016; Weisberg et al., 2019; Weisberg et al., 2014; Youngson et al., 2019; Zhao et al., 2020). In general, results from these experiment have shown large variations in the accuracy of cognitive maps and specific correlations with other aspects of cognition (e.g., working memory, navigation strategy, mental rotation, and others). Consistently across experiments, people varied on both the pointing task (ranging from about 10° to 80°, with 90° as pure chance) and the map building task (ranging from about 0 to 1) as measured by bidimensional regression, which can be interpreted as the correlation between the participant's configuration and the actual configuration Friedman and Kohler (2003) and Tobler (1994). In other words, participants ranged from near perfect performance (10° of error and a perfect configuration of buildings on their map) to near chance (80° of error and no correlation between their configuration and the actual configuration). Importantly, the two tasks also correlated relatively highly (.5-.7). The consistency of these properties raised our confidence in the findings described in the papers we published.

The programming error occurred on two of the pointing tasks (i.e., onsite/egocentric pointing and offsite/judgments of relative direction) in the platform. However, we focus on the onsite pointing task because it has been more commonly used.¹ In the onsite pointing task, participants are placed next to one of the eight buildings that they have learned (directly beneath the gem in the virtual world, see Figure 2) and are prompted to point to

¹ In addition to being less commonly used, the offsite pointing task is also entirely correctable from old data.

Figure 1. Map view of the Virtual Silton environment (never seen by participants). The blue and yellow squares indicate the front doors of target buildings along each of the two main routes. Red arrows, which begin at the gem of each building and extend in the direction of the next building along the main route, depict the reference directions used to calculate pointing error. See Figure 4 for further explanation of the pointing bug.

one of the other buildings. Participants point to all of the other buildings from this first building and are then automatically teleported to the next building and asked to point to all the buildings from there. This process is repeated until participants complete all possible pointing combinations.

The dependent variable of the pointing task is the unsigned error, computed as the shortest angular distance between the correct direction and the participant's response direction. Canonical results from the Virtual Silton experiments tabulate error on the pointing task separately for within and between route trials. Within-route trials consist of pointing judgments on the same main route as the building the participant is pointing from. Between-route trials consist of pointing judgments to buildings from different routes.

Figure 2. *First-person view of the Virtual Silcton environment.*

The original Virtual Silcton experiment showed individual variability along both trial types. That is, participants varied both in their ability to point to buildings along one of the main routes (within-route trials) and their ability to point to buildings on one route from the other route (between-route trials). Figure 3 presents this pattern replicated across 4 separate experiments (Weisberg & Newcombe, 2016). However, the data used to create these figures were affected by the programming error that recorded the absolute value of any pointing judgment that should have been negative (i.e., to remap a pointing judgment that fell between -180° and 0° to be from 0° to 180°).

Figure 4 demonstrates the effect of the bug on one hypothetical pointing trial. The participant is located at Building 1 and is instructed to point to Building 3. The participant can rotate and point in any direction (recorded as the response direction in Figure 4 and Figure 5). The response direction and correct direction were recorded in

Figure 3. *Between-route and within-route pointing error (buggy data) revealed individual differences in both experiments.*

relation to a reference direction, defined on each trial as the angle between Building 1 and a reference building (e.g., Building 2). In Figure 4 and the bottom panel of Figure 5, the correct direction is 37° clockwise from the reference direction. If the participant points 90° counterclockwise from the reference direction, the actual error is calculated as

$|-90^\circ - 37^\circ| = 127^\circ$. The effect of the bug was to take the absolute value of the response direction (and of the correct direction) resulting in calculating buggy error as

$$||-90^\circ| - |37^\circ|| = |90^\circ - 37^\circ| = 53^\circ.$$

Figure 4. An illustration of one trial and the effects of the bug. At the top left, a schematic of how the reference directions are calculated. Each set of onsite pointing trials occurred at one of the eight learned buildings (eight unique pointing locations). From each pointing location, participants pointed to the other seven buildings. To calculate the pointing error on each trial, a reference direction (establishing 0°) was calculated as the vector between the pointing location and the next pointing location along that main route (see Figure 1). If the pointing location was the last pointing location, the previous pointing location was used to establish the reference direction. The circle illustrates one trial. The participant is free to point in any direction (from -180 to 180). In the trial illustrated, the correct direction (orange arrow) is 37° clockwise from the reference direction. The participant points 90° counterclockwise from the reference direction (stored as -90° ; blue arrow). Without the bug, the error would have been recorded as 127° (blue-orange arc). The bug instead records the response direction as $|-90^\circ| = 90^\circ$ (effect of bug; blue-red arc), which results in a buggy error of 53° (orange-red arc).

Figure 5 demonstrates the effects of the bug on two values for the correct direction, illustrating a large and small effect depending on features of the trials and participant response directions. Critically, some trials were completely unaffected because all response directions had the same sign as the correct direction whereas other trials were greatly affected because the majority of response directions had the opposite sign as the correct direction and the correct direction was not near 0° or 180° . Actual distributions of errors are described in the Results. For a more in depth discussion of the root cause of the bug, how it was discovered, and how it was fixed, see the Jupyter notebook:

https://github.com/smweis/silcton_2.0/blob/main/notebook/Silcton_2.0.ipynb.

In the following section we evaluate the effects of the bug on previously collected data. We performed a replication of Weisberg et al. (2014) that allowed us to evaluate the data collected in Virtual Silcton using the buggy and the corrected versions of the code. We proceed in three steps. First, we show that the replication effort resulted in data patterns that are comparable to the original experiment. Second, we evaluate the effects of the bug on the main results using the new data. Third, we calculate a correction to the data by predicting correct data based on the buggy data. We demonstrate that the corrected data results in identical correlations between the pointing task and other measures, and a highly similar pattern of results overall for two of our previous studies (Weisberg & Newcombe, 2016; Weisberg et al., 2014).

Method

Preregistration

The experimental protocol, sample size, and analysis plan were preregistered on the Open Science Framework:

https://osf.io/8ydzj/?view_only=e0a9b80d2cd347f4891e404939c9259b. Our preregistration specified how data would be collected and a sample size of approximately 50 participants per site (150 total).

Figure 5. An illustration of the effects of the bug on two different correct directions. Top Panel: The correct direction (orange line) is the same as the reference direction (0°). In this case, taking the absolute value of the response direction does not change the error (red and blue lines overlap). Bottom Panel: The correct direction (orange line) is misaligned with the reference direction. The bug will affect any pointing judgment if the correct direction and response direction have opposite signs. The potential magnitude of the effect of the bug increased as the correction direction approached 90° or -90° . The exact result of the bug depends on the response direction. Values between 0° - 180° will be unaffected by the bug because taking their absolute value results in the same value. Values between -180° - 0° show effects of the bug that increase as the response direction approaches the negative correct direction. In the case where the correct direction is negative, response directions on the positive side will be affected by the bug because the correct direction was also stored as its absolute value. Note, the bottom panel replicates the trial shown in Figure 4.

Participants

We recruited 147 undergraduate participants (96 women, 51 men) across three sites (University of Florida (UF): 65, Temple University (Temple): 26, and Bond University (Bond): 56). Fifteen participants recruited by UF and Temple were discarded because they did not have complete data due to technical issues (e.g., internet cut out, website was slow or non-responsive). An additional set of 43 participants from UF and TU who signed up did not complete the pointing task, meaning their data could not be used for this study. The majority of the reasons for dropping out were technical difficulties in loading the virtual environments. A small minority of participants' data could not be merged due to duplicate study IDs. S.M.W. merged the participant level data and removed participants for whom data were unusable.²

The average participant age was 20.85 years ($SD = 4.24$). Of the 91 participants in the sample collected in the United States, 23 participants were Hispanic, 67 were not, and one participant did not report. Self-reported ethnic category membership was 17 Asian, 11 Black/African-American, 55 White, three reported more than one race, one reported Other, and four did not wish to say. These data were not collected for the Australian sample. Participants were recruited via SONA systems in the three universities. These data collection procedures were reviewed by the following ethics committees: IRB of the University of Florida #201902385 - Cognitive and Neural Underpinnings of Spatial Navigation Behavior; IRB of Temple University #13394 - Computer-Based Spatial Abilities; Bond University Human Research Ethics Committee (BUHREC). BUHREC Reference Number: YH00014.

We compare this new data set ($n = 147$) to a set of previously collected studies ($n = 294$). For details of the methods and participant characteristics, see (Weisberg &

² Detailed experimenter notes are available in the online data set. We elected to be inclusive in our sample despite some participants experiencing mild to moderate technical issues. As we describe in the Results, the data are consistent with previously collected Virtual Silcton experiments and so we had no additional reasons to exclude subjects.

Newcombe, 2016; Weisberg et al., 2014).

Materials

Demographics. Participants were asked to report their biological sex, gender, ethnicity, age, education level, and handedness.

Santa Barbara Sense of Direction Test (SBSOD; (Hegarty et al., 2002)). This self-report measure of navigation ability consists of fifteen 7-point Likert-scale items such as 'I am very good at giving directions,' and 'I very easily get lost in a new city.'

Mental Rotation Test (MRT; (Peters et al., 1995; Vandenberg & Kuse, 1978)). This computerized version of the MRT consists of two 10-item sections of multiple-choice questions. Participants are given 3 minutes per section. Each item consists of one target two-dimensional image of a three-dimensional shape made up of connected cubes, and four answer choices that are also made up of connected cubes. Two of the answer choices have the same configuration of cubes, but rotated in 3D space. The other two answer choices are of different configurations. Participants received two points per correct choice, and lost two points per incorrect choice. Zero points were awarded for each omission. This measure was not collected for the Australian dataset.

Virtual Silcton (Schinazi et al., 2013; Weisberg & Newcombe, 2016; Weisberg et al., 2014). Virtual Silcton is a behavioral navigation test administered via desktop computer, mouse, and keyboard. Modeled after the route integration paradigm (Hanley & Levine, 1983; Holding & Holding, 1989; Ishikawa & Montello, 2006; Schinazi et al., 2013), participants learn two routes in separate areas of the same virtual environment by virtually traveling along a road indicated by arrows Figure 2. During navigation, participants are explicitly told to remember the names and locations of four buildings along each of these routes (8 buildings in total). Next, participants travel along two routes that connect the two areas of the virtual environment. Virtual travel consists of pressing arrow keys (or the W, A, S, and D keys) on a standard keyboard to move in the

environment, and moving the mouse (or trackpad) to look around. For more information and materials, see the Open Science Framework repository <https://osf.io/fykr7/>.³

During navigation, participants were constrained to travel only along the routes indicated with arrows. This was achieved by employing invisible walls in the virtual environment that prevent participant from navigating off route. Participants were told that they could move and look at whatever pace they chose but that they would have only one opportunity to learn each of the routes. At a minimum, we required participants to travel from the beginning to the end and back to the beginning of each route. Buildings were indicated by blue gems, which hovered over the path, and named with signs in front of the building.

Participants were tested on their knowledge of the direction among the buildings within each of the main routes, and among buildings between the main routes. Testing consisted of two tasks. For the onsite pointing task, participants performed 56 pointing judgments. While standing next to each of the eight buildings they were previously told to remember, participants pointed to each of the other seven buildings. On each trial, participants were placed alongside a building in the environment and prompted to point to each of the other seven buildings. The pointing judgment was conducted by moving the mouse to rotate the view and position a crosshair toward the target building and by clicking the left mouse button to record the direction. The dependent variable was calculated as the absolute angle error between the participant's pointing judgment and the actual direction of the building. If the angle difference was greater than 180°, it was corrected to measure the shorter of the two possible arcs. We calculated pointing error for within-route trials and between-route trials separately. This resulted in 32 between-route trials and 24 within-route trials.

Participants also completed a model-building task in which they viewed a rectangular

³ The names of the Virtual Silton buildings (e.g., Batty House, Lynch Station, etc.) refer to famous geographers.

box on a computer screen and birds-eye view images of the eight buildings. Scrolling over the buildings with the mouse revealed a picture of the front view of the building and its name. Participants were instructed to drag and drop buildings inside the box in the position they believed the building would be located in the virtual environment (as if they were creating a map), without regard to the orientation of the buildings or to the map. The model-building task was scored using bidimensional regression analyses (Friedman & Kohler, 2003; Tobler, 1994).

Procedure

Data were collected during the summer and fall of 2020 in the middle of lockdown as a result of the COVID-19 pandemic. Thus, all participants were recruited, screened, and run in a live Zoom session with the experimenter.

Participants were contacted via email by the experimenter and given a Zoom link along with instructions. Participants were told to plan for 1 to 2 hours of testing, have a stable internet connection, and be in a quiet, distraction-free environment. After logging into the Zoom session, participants were instructed to navigate to a Qualtrics survey via a link and afterwards share their screen with the experimenter. The participant began the experiment by reading over and e-signing the informed consent and completing the demographics questionnaire. Next, participants went to the Virtual Silcton website and completed the SBSOD and MRT. Finally, participants completed the Virtual Silcton test. Participants received instructions on using the keyboard and mouse/trackpad to navigate and the experimenters monitored their progress and provided feedback via Zoom. Experimenters could also view participants' progress via screen share, though for some participants this had to be disabled to preserve bandwidth.

One additional difference in collecting data online compared to previous Virtual Silcton experiments is that the majority of participants in this study used a trackpad (70 UF/TU participants used a trackpad compared to 21 who used a mouse; all 56 participants

at Bond used a mouse). All participants in the original studies used a mouse.

Results

Our first step was to assess whether our new sample was similar in performance levels to our original sample. We did these analyses using the buggy errors, simply to provide a common yardstick. Our second (and central step) was to use the new sample to assess the impact of the bug, by comparing buggy and actual errors. Having found minimal impact, in the third step we checked whether correlations with gender and with other measures changed with actual errors. We went on, in the fourth step, to propose a way to retrospectively correct buggy scoring, and in the fifth and final section, we evaluated whether the correction method resulted in changes to previously published findings.

Are the New and the Old Samples Comparable?

We began by comparing the similarity of performance between the original Silton participants and the new participants by performing two analyses using buggy scoring.

First, we compared performance on the pointing task. New participants ($N = 147$, $M = 40.46$, $SD = 16.06$) were significantly more accurate on between route pointing compared to original participants ($N = 294$, $M = 45.70$, $SD = 13.62$), $t(439) = 3.57$, $p = 0.0004$, $d = 0.35$. New participants ($M = 17.95$, $SD = 11.36$) were also significantly more accurate on within route pointing compared to original participants ($M = 23.66$, $SD = 11.85$), $t(439) = 4.82$, $p < 0.00001$, $d = 0.49$.

Second, despite these differences, the k-means clustering solutions were similar between the original and new data. Figure 6 illustrates the overlay of the three-group cluster solution implemented in scikit-learn (Abraham et al., 2014) using the original buggy data set and new buggy data. In the newly collected sample, 43% ($n=63$) of participants were classified as Integrators, 37% ($n=55$) as Non-Integrators, and 20% ($n=29$) as Imprecise Navigators. In the original data, there were 28% ($n=81$) Integrators, 43% ($n=128$)

Non-Integrators, and 29% (n=85) Imprecise Navigators. The distribution of the number of participants in each cluster was also significantly different, $\chi^2(2) = 11.11$, $p = .004$.

Figure 6. *K-means clustering solution for newly collected and originally collected datasets (buggy scoring). Overlap in the group solution is substantial (even excluding one outlier on the within-pointing trials). Convex hulls of all clusters are shown to illustrate overlap of the group solution.*

Within clusters, performance was slightly better overall in the new experiment. Participants classified as Integrators or Non-Integrators in the new experiment performed significantly better on within-route pointing compared to their Integrator or Non-Integrator counterparts in the original data sets ($t_{Integrators}(142) = 3.01$, $p = .003$, $d = 0.51$; $t_{Non-Integrators}(181) = 4.25$, $p = .00003$, $d = 0.70$; $t_{Imprecise}(112) = 1.00$, $p = .32$, $d = 0.22$). For between-route pointing, only the Integrators from the newly collected sample performed significantly better ($t_{Integrators}(142) = 2.19$, $p = .03$, $d = 0.37$). Neither of the other two groups differed significantly on between-route pointing ($t_{Non-Integrators}(181) = 0.78$, $p = .43$, $d = -0.13$; $t_{Imprecise}(112) = 1.07$, $p = .29$, $d = 0.22$).

These data suggesting more accurate spatial representations were corroborated by

the model building task. Compared to participants in the original data sets, new participants performed significantly better on the model building task both overall ($M_{New}=0.59$, $SD_{New}=0.29$; $M_{Original}=0.47$, $SD_{Original}=0.26$; $t(434)=4.12$, $p=.00005$, $d=0.41$) and using the configurations of each route separately ($M_{New}=0.71$, $SD_{New}=0.21$; $M_{Original}=0.63$, $SD_{Original}=0.21$; $t(432)=3.65$, $p=.0003$, $d=0.37$). Oddly, however, new participants self-reported significantly lower navigation ability on the SBSOD ($M_{New}=4.10$, $SD_{New}=1.05$; $M_{Original}=4.35$, $SD_{Original}=0.98$; $t(434)=2.49$, $p=.01$, $d=0.25$).

Although the samples are not identical, the newly collected data should suit our purpose of evaluating the effect of the bug on existing Silcton data. We reached this conclusion based on the fact that the distributions of pointing scores are qualitatively similar (e.g., over similar ranges), the cluster solution is highly aligned, and each group of navigators is well-represented. We now turn to the issue of how replacing the buggy pointing data with the correct pointing data affects the results.

Differences between the Actual and Buggy Pointing Data

To restate, the pointing bug sometimes resulted in pointing judgments being recorded as more accurate than they actually were. But different trials and different navigators may have been affected differently. We had two primary questions. First, how much worse were pointing judgments overall? Second, how were different types of navigators and trials affected by the bug?

The correlations between the buggy and actual pointing scores were nearly at ceiling. Using parametric and rank-based correlations (the latter of which are useful to determine whether the relative ordering of participants changed as a result of the correction) all combinations of pointing trials resulted in correlations $> .96$ and were significant at $p < 1 * 10^{-30}$ (within-route: Pearson's $r = 0.97$; Spearman's $r = 0.98$; between-route: Pearson's $r = 0.97$; Spearman's $r = 0.96$; total pointing: Pearson's $r = 0.98$; Spearman's $r = 0.97$). In addition to high correlation coefficients, cluster membership (i.e., whether participants

were defined as Integrators, Non-Integrators, or Imprecise Navigators on the basis of their between-route and within-route pointing errors) was also highly similar. Performing the k-means cluster analysis using the actual pointing errors results in only 16 out of 147 participants being reassigned compared to the buggy pointing errors. There was no significant difference overall in category membership, $\chi^2(2) = 0.69$, $p = .71$. Critically, different category assignments occurred at category boundaries, as can be seen in Figure 7.

Figure 7. Comparing cluster assignments for buggy and actual data reveals group membership changed for participant who fell along the group boundaries (white dots) and not a qualitatively different cluster solution.

Next, we quantified changes to the pointing task data as a result of the bug. The buggy pointing task data overall had significantly lower error compared to the actual data ($M_{Actual} = 40.07$, $SD_{Actual} = 17.44$; $M_{Buggy} = 30.81$, $SD_{Buggy} = 12.55$; $t(147) = 19.24$, $p < .00001$, $d_z = 1.12$). The bug impacted between-route trials ($M_{Actual} = 51.54$, $SD_{Actual} = 21.26$; $M_{Buggy} = 40.46$, $SD_{Buggy} = 16.11$; $t(147) = 19.15$, $p < .00001$, $d_z = 1.12$) and within-route trials ($M_{Actual} = 24.79$, $SD_{Actual} = 17.09$; $M_{Buggy} = 17.95$, $SD_{Buggy} = 11.40$; $t(147) = 12.23$, $p < .00001$, $d_z = 0.71$). The interaction between coding (actual vs. buggy)

and trial type (between vs. within) did not reach significance, $F(1,584) = 2.33$, $p = .13$.

Figure 8. Participant-level data show a greater impact of the pointing bug based on the accuracy of a participants' pointing judgments.

Our results also revealed that the bug impacted worse navigators more than better navigators. To quantify this, we subtracted each participants' actual pointing error from their buggy pointing error and correlated that difference with their actual pointing error. In other words, does the distance from the diagonal line in Figure 8 significantly correlate with how far down the y-axis a participant falls? As can be seen in Figure 8, these correlations for between-route $r(147) = -.81$, $p < .00001$ and within-route $r(147) = -.90$, $p < .00001$ are significant.

Another way to visualize this relationship is to plot an arrow from each participant's buggy pointing error to their actual pointing error. In Figure 9, the length (and color) of each arrow represents the distance between each participants' pointing errors. Consistent with the results described above, participants who performed better overall had pointing

errors less affected by the bug. This pattern can be shown in two ways. First, as in the previous correlations, we calculated the difference between the buggy and actual pointing errors for between- and within-route pointing separately. Correlating the difference in within-route corrections with between-route performance resulted in a significantly lower correlation than using between-route corrections, $r(147) = -.50, p < .00001$. Similarly, correlating the difference in between-route corrections with within-route performance was also lower, $r(147) = -.48, p < .00001$. Alternatively, we can calculate the overall distance (depicted by Figure 9 arrows) and correlate that with pointing errors. Correlations with distance are significantly lower (p 's $< .02$) for both between-route pointing $r(147) = .69, p < .00001$ and within-route pointing $r(147) = .69, p < .00001$. Using total pointing error is a better predictor of distance, and is not statistically lower than the between-route pointing correlation, $r(147) = .77, p < .00001$.

Figure 9. Arrows depict the amount of distance participants' pointing errors changed as a result of correcting the bug in the task.

Trial-specific Data

Overall, 29% of trials were affected by the bug. On average, if a trial was impacted by the bug, the actual error was 31.8° ($SD = 33.4^\circ$) higher than the buggy error. However, there was high variability across trials. Figure 10 displays the variability in proportion of participant responses affected by the bug and the magnitude of the effect across trials. One reason for variability is that the amount a trial could be affected by the bug is determined by the correct direction. The maximum amount a trial could change is double the correct direction.

As shown in Figure 10, the percent of responses affected and the difference between the actual error and buggy error varied widely across trials (min = 0%; max = 80%). Whereas some trials showed normal distributions around the correct direction with responses equally likely to fall to the left or right. Other trials had distributions closer to uniform with participants appearing to guess or showed bias, with responses more likely occur to the left or right. At the participant level, trials affected ranged from 3% to 59%. Beyond that, however, the degree to which a particular trial was affected was highly correlated with its difficulty (i.e., average error across all participants, $r(56) = .57$, $p < .00001$), with more difficult trials having a larger difference between buggy and actual responses (see the Jupyter Notebook for additional figures related to trial-level data including a version of Figure 10 with all trials).

Given the relatively high percentages of affected trials, what explains the high correlations across actual and buggy between and within pointing errors at the participant level? Pointing errors were highly correlated across trials. On average, if a participant makes a higher error response, this response is more likely to be affected by the bug; and this is true across trials. At the participant level, we observed high correlations between percent of trials affected with overall actual error, $r(56) = .88$, $p < .00001$, and with the change between a participant's actual and buggy error, $r(56) = .91$, $p < .00001$.

These data indicate that although raw scores are different, comparing the actual data

Figure 10. A comparison of pointing trials to and from a subset of four buildings (of eight total) from the new dataset. Two of the buildings (Batty House and Harris Hall) appear on one main route, whereas the other two (Golledge Hall and Sauer Center) appear on the other main route. Each row consists of pointing trials from the same location. Each column consists of pointing trials to the same location. Individual panels depict the distribution of actual participant responses for that trial. The correct direction for each trial is indicated with a vertical dashed line and each panel is annotated with the percent of responses affected by the bug ($X\%$ trials) and the mean difference between the actual error and buggy error for trials affected. Bars in grey, which share their sign with the correct direction, would have been unaffected by the bug. Bars in red, which have the opposite sign, would have been recorded as their absolute values. Some trials were unaffected by the bug because no actual participant responses occurred with the opposite sign as the correct direction (e.g., Sauer Center to Golledge Hall). Other trials were highly affected by the bug, but would have had little impact because the correct direction was near 0° or 180° (e.g., Sauer Center to Batty House).

with the buggy data did not result in substantial impact on the original findings. We observe similar individual differences, and a clear distinction in the between-route and within-route pointing judgments. The major difference is that the worse navigators seem to be yet worse than they had appeared formerly. We now turn to evaluating the impact of the pointing bug in relation to non-pointing measures.

Did the Pointing Bug Affect Other Results?

Gender differences

Men outperformed women in the new data on between pointing trials (for means and standard deviations, view the Supplementary Jupyter notebook). This result was not changed between the buggy and actual data. Here, a 2 (gender) x 2 (between/within pointing) x 2 (buggy / actual) ANOVA revealed only a significant main effect of gender, $F(1,580) = 52.64$, $p < .00001$. Gender did not interact with either variable ($p > .16$).

Model Building, SBSOD, MRT

The pointing correction did not alter the correlations for within-pointing, between-pointing, or total-pointing (all p 's $< .08$.) It is notable that all correlations between pointing error and Model Building and SBSOD became (non-significantly) stronger whereas the correlation between pointing error and MRT moved in both directions. Scatterplots and statistics available in the Jupyter Notebook.

Correcting the Buggy Pointing Data

Given the high correlations and similarity in findings between the actual and buggy pointing data, we suspected that the bug resulted in a low impact error for previously published findings. We wanted to test this empirically by correcting previously-collected Virtual Silton data. Recall that data affected by the bug cannot be recovered fully (i.e., we will never know the correct sign of the pointing angle). Applying a correction, however,

allows us to estimate and evaluate the impact of the bug. The purpose of this is two-fold. First, we can empirically evaluate the differences in published findings for actual and buggy data and determine what, if anything, changes. Second, we can provide a tool for other researchers that have used Virtual Silcton to apply to their data and evaluate differences in their own findings. The correction tool is accessible on GitHub (<https://mybinder.org/v2/gh/smweis/SilctonBugReport/HEAD>). Following the instructions in the notebook, the interested reader can 1) upload buggy data for both within and between pointing cases, 2) generate the corresponding corrected instances, for both types of pointing trials, and 3) save and download the corrected results. Here, we provide the results of the models and demonstrate the effects of the corrected results on two previously published manuscripts (Weisberg & Newcombe, 2016; Weisberg et al., 2014).

To generate the corrected pointing values we ran a linear regression model on the 147 newly collected participants, for within and between pointing trials separately. The regression model specifies the buggy pointing errors as the independent variable, and the corresponding corrected errors as the dependent variable. In other words, the model predicts actual pointing values (per participant) from the buggy pointing values. The resulting models show that the corrected pointing values could be predicted from the corresponding buggy pointing instances by the following formulae:

Within route: $pointing_{predicted} = 1.45 * pointing_{erroneous} - 1.18$. $r^2 = 0.93, p < .00001$.
 Between route: $pointing_{predicted} = 1.28 * pointing_{erroneous} - 0.09$. $r^2 = 0.94, p < .00001$.

Because the r^2 coefficients of the linear models are evaluated in-sample (i.e., they are computed on the same data of 147 participants on which the corresponding models are fit), we performed a robustness check on the r^2 coefficients of the linear models by means of repeated (N=100 repeats) 5-fold cross-validation (Krstajic et al., 2014). To achieve this, we fit $5*100=500$ linear models for both between and within trials, as above. At each of the 100 repeats, we fitted five linear models using a 5-fold cross-validation and computed the mean of the r^2 coefficient on the five validation folds. The goal of the robustness check is to

compute the mean and the standard deviation of the distribution of the 100 means of the r^2 coefficients and compare them with the in-sample estimates, as above.

For within trials, the distribution of the coefficients of determination has a mean equal to 0.92 and a standard deviation equal to 0.03. In the case of between trials, the mean of the distribution is 0.93 and the standard deviation equal to 0.02. As a result, for both pointing exercises the estimate of the amount of variance predicted by the erroneous pointing variables is robust with respect to the repeated random partition of the dataset of 147 trials by means of repeated cross-validation routines.

Additional Predictors

Although the variance explained was high, we attempted to include additional predictors in the model to see whether predictions could be improved and whether changes in the predictions resulted in a different impact on previously published findings. Detailed results can be found in the Silcton Correction Analyses Jupyter notebook.

First, we observed that between and within-route averages are better predicted separately compared to predicting total pointing error from the average of between-route and within-route. The high proportion of variance explained in these models suggests there is very little room for improvement by including additional factors. In particular, the model fit is better when predicting between and within-pointing averages separately than when predicting either from total pointing ($r^2 = 0.62$ when predicting within-route pointing from total pointing error; $r^2 = 0.87$ when predicting between-route pointing from total pointing error). In addition to providing further evidence that these are separable behavioral factors, these analyses show the importance of predicting each type of pointing trials separately.

Second, we observed that introducing additional random predictors did not improve model fit. To assess whether we were at ceiling in explainable variance, we simulated a new predictor by creating random vectors of the same length as our training data set (147) and only selected random vectors that improved the model fit. After 1000 simulations, the

maximum improvement in model fit for both within and between we observed was 0.94, suggesting there is very little variance that can be captured using a simple linear model.

Third, we found that more complex models did not improve model fit. We used a 5-fold cross-validated random forest algorithm for both within and between pointing trials using the erroneous pointing, SBSOD, and MRT. During the cross-validation routine we tuned the number and the maximum depth of the trees of the ensemble. For within pointing trials, the best performing random forest ensemble comprises 35 trees of maximum depth equal to three. This ensemble reaches a mean r^2 on validation folds equal to 0.92 (and standard deviation equal to 0.03). For between pointing trials, the best performing random forest ensemble comprises 40 trees of maximum depth equal to three. Similarly to the within case, the best performing ensemble reaches a mean r^2 on validation folds equal to 0.92 (and standard deviation equal to 0.03).

In summary, the use of machine learning algorithms (i.e., random forests), and additional variables beyond the buggy pointing data do not deliver a substantial performance improvement with respect to the linear models, for both within and between pointing trials.

Reanalyzing Previous Findings Using Corrected Data

Finally, we discuss whether the use of linear regression models from the previous section impacts prior research involving Virtual Silcton. Previously published findings using Virtual Silcton typically either correlate task scores with each other or with some other measure (e.g., SBSOD, MRT, performance on other virtual environment tasks), or calculate group membership (Integrators, Non-integrators, and Imprecise navigators) and investigate the differences between groups on another measure.

To address the correlational approach, due to the sign of the slope coefficient and intercept of the linear regression correction, any correlations calculated using the pointing errors should be numerically identical. To verify this, we replicated correlation tables from

(Weisberg & Newcombe, 2016) and found numerically identical results. Correlation coefficients will not change as a result of applying the linear regression correction.

To address the group-based approach, performing the cluster analysis using the corrected data results in four participants out of 294 changing cluster assignments compared to the erroneous data, all of whom went from Non-Integrators to Imprecise Navigators. On this basis, we conclude that the impact on previously published analyses that used one of these two approaches would be either none or minimal. For detailed analyses, please see the Jupyter notebook on Silcton Correction Analyses (https://github.com/smweis/silcton_2.0/blob/main/notebook/Silcton_Correction_Analyses.ipynb).

More complicated analyses or analyses depending on individual pointing trials may, however, be more impacted. For instance, we re-analyzed data in Ruginski et al. (2019), and observed numerical differences in the structural equation modeling results, though how meaningful those differences are remain an open question.

Discussion

A programming error discovered in a widely-used and highly-cited virtual test of spatial navigation behavior resulted in unrecoverable mistakes in data recording but minimal impact on published findings. To reach this conclusion, we collected data for a high-powered (three times as many subjects as the original study) multi-site replication, applied machine learning approaches to correct buggy data, and reanalyzed published findings. For Virtual Silcton and the researchers who use it, this result was nearly the best case scenario. With an eye toward the future, we discuss reasons the impact of the programming error was likely not worse, some lessons learned for the broader practice of science, silver linings of the replication results, and proposals for best practices in preventing, detecting, and reporting programming errors.

Before getting into the details of the current situation, it is worth contextualizing the kind of error we are reporting and provide the reasons we believe it went undetected for so

long. We propose two dimensions to describe programming errors. First, programming errors vary in their impact - the extent to which the error affects downstream analyses or data validity. Low impact errors result in trivial changes to the interpretation of the data whereas high impact errors result in drastic changes. Second, programming errors vary in their detectability. Here, we refer to detectability as the ease in which an error can be discovered by researchers scrutinizing the data themselves rather than inspecting the code that produced the data. To illustrate how impact and detectability relate, consider a hypothetical case in which a researcher is interested in the correlation between height (measured in feet and inches) and weight (measured in pounds). In the case where an error is easily detectable, impact matters less because (by definition) these errors should be caught. For instance, if a researcher systematically adds 100 to the height variable (no impact on the correlation between height and weight) or multiplies the height variable by -1 (high impact - complete reversal - of the correlation between height and weight), simple visualization of the data (histograms or boxplots) or inspection of the data descriptives (means) will reveal the mistake.

When errors are difficult to detect, discovering low impact errors is of questionable value. For instance, if the researcher systematically adds .001 inches to everyone's height, the correlation between height and weight will not change (low impact). However, this error would be difficult to detect because inspection of the means or standard deviations directly would not yield any obvious abnormalities. For low impact errors, although any analyses using the processed data may be impacted by the mistake, the error in the data is, by definition, very small and unlikely to lead to appreciable changes in conclusions based on the data set.⁴

The most problematic kinds of programming errors occur when a high impact error is difficult to detect. Using our height example, imagine the researcher had converted the

⁴ Note that this 'error' (at least for our height/weight example) could and likely does occur for any irrational number due to float-point precision on computers and may not be detectable in human-readable code. 'Bugs' with no measurable impact may remain silent forever.

variable height from feet and inches to feet. But rather than changing height using the following formula: $feet = feet + inches/12$ they rather append the inches onto feet as a decimal such that if someone is 6'11", their height is recorded as 6.11 feet. This error would result in incorrect data. People who are 6'11" will be recorded as shorter than people who are 6'5". The correlation between weight and height would likely still exist, but it would also certainly be affected by the error. Moreover, the effects of the error would be unpredictable for more complex analyses due to the nonlinearities that are introduced. Of course, such a mistake would be easy to detect in human-readable code (in practice, using inches after the decimal point is not a very easy mistake to make in code). However, for illustrative purposes, we emphasize the point that even for relatively simple data processing steps and analyses, effects of computer programming mistakes may introduce non-obvious errors with effects that can be difficult to predict.

Upon discovering the programming error in the Virtual Silcton code (which, we believe, falls under the difficult to detect category), our primary task was to estimate the impact of the error on previously published results. Although we contend that the impact of the error was low, we realize that the estimation of impact is subjective in the same way that estimating if a result replicated is subjective. For instance, along the impact continuum, if an error results in no change to the data (no impact) this is a better outcome than a small numerical change (low impact), which is better than a numerical change resulting in a change in significance (medium impact), which is better than a change in effect valence (high impact).

Amplifying this difficulty is the extent to which materials and code are shared among researchers. While we believe that sharing practices are beneficial to the scientific community, what happens when the original materials and code contain bugs? One back-of-the-envelope calculation estimates that high impact undetected bugs may impact as many as 5% of research findings (Soergel, 2015). Occasionally, these bugs are discovered and reported as corrections, altering the scientific record (Hutson, 2010; Society for

Neuroscience, 2015). As one real-world example, consider a bug in the neuroimaging software, AFNI, which went undetected for almost 15 years (Eklund et al., 2016). Detection occurred only through robust comparisons and simulations of data with known parameters. Although media coverage of the bug discovery indicated that the impact of the AFNI bug was high (one *Wired* article was titled 'Bug in fMRI software calls 15 years of research into question'), the bug author and discoverer agree that the impact was likely marginal (Cox et al., 2017). Needless to say, such comparisons and simulations were not completed on the Virtual Silcton paradigm until very recently and after many publications.

We believe there are three main reasons that made it likely for the Virtual Silcton bug to turn out to be low impact. First, if the impact of the error were larger, it would have been more detectable in early sanity checks; pointing errors could have been spuriously low or high, or all pointing trials could have been affected rather than some. Second, built-in redundancy in dependent measures made Virtual Silcton more robust to an error on one measure. The results from the model-building task provided convergent validity with the pointing task. Correlations between the buggy pointing task and the unaffected model-building task were high and in the expected direction. In that sense, the model-building task provided a fail safe - results relying on the model-building task were consistent with the pointing task results and remain unchanged. This redundancy is particularly important in this case because the bug was systematically related to the magnitude of pointing error. That is, the further the response direction was from the correct direction the larger the potential impact of the bug. As such, the best performers were affected less than the worst performers, whose pointing errors were recorded as lower than they actually were. This resulted in a reduction of the range of the overall performance and the introduction of bias in the opposite direction of individual differences. Third, if materials were not made public and widely shared, independent verification and sanity checking would have meant the error may never have been caught. Independent data verification led to detection of the abnormality (that recorded pointing directions only

ranged between 0° and 180° rather than a 360° range), which allowed us to diagnose the programming error. Although this error appears obvious in retrospect, confirmation bias and results that made sense and were mutually reinforcing (e.g., model-building and pointing correlating) blinded us to lower impact but pernicious errors. Nevertheless, if the Virtual Silcton measures were less widely-used the programming error may never have been detected. Ironically, if Virtual Silcton were not shared or widely-used, the error also would have been confined to fewer datasets. Still, although this is an underappreciated pitfall of open materials, science has the opportunity to self-correct only if the same measures and code are used - otherwise mistakes from scientific history are doomed to repeat themselves undetected.

On the positive side, discovering the Virtual Silcton bug resulted in the collection of a large-scale multi-site direct replication of Weisberg et al. (2014). Findings in the original experiment were supported in the replication, including broad variability in navigation ability and differences in within-route and between-route learning. Although we did observe differences in the overall accuracy and distributions of participants across pointing groups, we note that data was collected entirely online (with experimenter supervision), that a considerable number of participants used a trackpad (rather than a computer mouse) and that part of our replication sample was from a different country (Australia). Despite these inconsistencies, the replication sample builds confidence in the use of Virtual Silcton to measure individual differences in navigation ability and provides even more data for future data mining. These results serve to enhance the robustness of the Virtual Silcton results.

Many scientists are not professional computer programmers, although many do spend much of their time writing computer code. Catching errors in computer code, particular errors that do still allow the code to compile and run, is paramount to ensuring mistake-free analyses. Yet, there is also an issue of diminishing returns - catching low impact errors is likely of little value and these are likely the only ones remaining after high impact errors are discovered and eliminated. There are four suggestions we offer to any

researcher wishing to reduce the chance of a high impact error. We are not alone in making these suggestions and refer the reader to other more in depth resources (Hutson, 2010; J. Strand, 2020; Wilson et al., 2014).

1. Enhance reproducibility and legibility. Write clean, well-documented, reproducible code. Code that is easier to read and check will be easier to verify. In general, code that has clear comments and is legible will be easier to debug. To this end, one of the current authors has used Jupyter Notebooks as the containers for all data processing, analyses, and visualizations. Using Binder, any reader or researcher can reproduce the results using the exact same code (and package versions) reported in the original manuscripts.

2. Use validated code or validate custom code. Packaged code that has been reused and tested many times will more likely be robust than custom code. If an analysis or experiment does call for custom code, try to use unit testing wherever possible so it is clear, documented, and validated how the code is supposed to work. In the case of Virtual Silcton, the code was difficult to debug in part because it required demonstrating code functionality with Unity3d integrated with a website. Providing unit tested validity for the code snippet with the bug may have alerted us to the mistakes, though would still be difficult to implement. Code review is another technique used to reduce the chance of programming errors. In a code review, a second user (not the person who wrote the code), reads through the new code line by line to provide independent interpretation and verification. Code review can enhance code legibility and validate unit tests. A third way to enhance validation is to reuse code snippets. Rather than writing new code for each application, reusing validated code snippets will reduce the need to check every line of code.

3. Sanity check outputs. Similar to validation, sanity checking is low-level verification that the data output from a script or analysis is as expected. For example, if a line of code is supposed to randomize presentation of images in an experiment, a researcher might verify this by generating a large set (e.g., 1000) of sequences and visualizing the

number corresponding to when each image is being presented. The average presentation number should be the total number of images presented divided by 2. In the case of Virtual Silcton, the outputs of the pointing error were sanity checked, but not using data that could verify that the code was functioning properly. Original sanity checking occurred when one researcher (SMW) visualized the pointing angles of all participants for each trial. However, real participant data would not necessarily reveal the programming error because participants may not have pointed completely randomly. That is, perhaps participants were biased in their pointing judgments for one particular trial or never pointed very inaccurately (although this latter was not true, it could have been true in the correct pointing data). Catching the bug required a systematic sanity check on the output of the pointing task - a researcher pointed at each 45° increment and verified that the output was in the expected range.

4. Share code and materials. Independent verification and reproduction of findings is the cornerstone of science. Even incorporating the best possible code reviews, bugs are still the rule, rather than the exception. By sharing the materials and making the data and code as transparent and accessible as possible, we can maximize the chances that high impact errors are mitigated.

Once bugs are discovered, which they inevitably will be, how should researchers respond? Although relatively rare, reporting errors in scientific projects is becoming more common. In looking back through the scientific record, we identified several cases of programming errors, each of which resulted in a different result and none of which were quite the same as the case here. In the simplest cases, a programming error resulted in a high impact mistake that fundamentally changed the interpretation of the data. Such an example occurred in the Journal of Neuroscience (Society for Neuroscience, 2015) when researchers discovered a matrix transposition error in the analyses upstream of their final results. The error changed the results and the authors voluntarily retracted the paper and were working on replicating the findings. In a somewhat similar situation, J. F. Strand

et al. (2020) discovered a programming error in which the experiment recorded the timing for a control condition before the stimulus appeared on screen causing it to appear that reaction times were slower for control but not experimental conditions. In this case, although the data were recoverable, the impact of the bug was large - there was no effect of condition. *Psychonomic Bulletin and Review* allowed Strand and colleagues to amend the (significant) results in the original manuscript, replacing them with the new (null) results. Finally, the AFNI bug (Cox et al., 2017) described earlier probably affected many published scientific findings, although the impact of the bug is estimated to be low. Perhaps for that reason, very little has been done to correct the scientific record. Affected publications were, to our knowledge, not retracted nor very often reanalyzed. The buggy version of the software has been replaced with updates, but it is possible that older versions of AFNI are still being used in neuroimaging labs.

The nature of reporting errors in the scientific record is constantly evolving and, in our opinion, would benefit from nuance and case-by-case decisions. In all but the most straightforward of errors, replicating the original results and reproducing the analyses using independently written or verified software may be the best way to move forward. Here, we have also leveraged machine learning to introduce the concept of applying an empirically-generated correction to erroneous data. Using replication data, we found that original findings were likely to be unchanged, suggesting little need to for those findings to be called into question.

We would also note the great costs undertaken by researchers who attempt these corrections. Collecting and analyzing new data, sanity checking experimental paradigms and analysis scripts, and implementing version controlling, clean code, and documentation are largely thankless tasks, which do not earn academic awards, promotions, or grants. As scientists, we require evidence to support theories and do not estimate the likelihood of theories on their face value. Why do we not apply the same skepticism and invest the same intellectual resources and training to the validity of the evidence itself? Although the

Virtual Silcton bug had a negligible impact on published findings, we hope it does have an impact on scientific practice more broadly, encouraging open science, reproducible methods, and replication.

References

- Abraham, A., Pedregosa, F., Eickenberg, M., Gervais, P., Mueller, A., Kossaifi, J., Gramfort, A., Thirion, B., & Varoquaux, G. (2014). Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics*, 8. <https://doi.org/10.3389/fninf.2014.00014>
- Blacker, K., Weisberg, S., Newcombe, N., & Courtney, S. (2017). Keeping track of where we are: Spatial working memory in navigation. *Visual Cognition*. <https://doi.org/10.1080/13506285.2017.1322652>
- Cox, R. W., Chen, G., Glen, D. R., Reynolds, R. C., & Taylor, P. A. (2017). FMRI Clustering in AFNI: False-Positive Rates Redux. *Brain Connectivity*, 7(3), 152–171. <https://doi.org/10.1089/brain.2016.0475>
- Eklund, A., Nichols, T. E., & Knutsson, H. (2016). Cluster failure: Why fMRI inferences for spatial extent have inflated false-positive rates [Publisher: National Academy of Sciences tex.eprint: <https://www.pnas.org/content/113/28/7900.full.pdf>]. *Proceedings of the National Academy of Sciences*, 113(28), 7900–7905. <https://doi.org/10.1073/pnas.1602413113>
- Friedman, A., & Kohler, B. (2003). Bidimensional regression: Assessing the configural similarity and accuracy of cognitive maps and other two-dimensional data sets. *Psychological methods*, 8(4), 468–491.
- Galati, A., Weisberg, S., Newcombe, N., & Avraamides, M. (2017). When gestures show us the way: Co-thought gestures selectively facilitate navigation and spatial memory. *Spatial Cognition and Computation*. <https://doi.org/10.1080/13875868.2017.1332064>
- Hanley, G. L., & Levine, M. (1983). Spatial problem solving: The integration of independently learned cognitive maps [Publisher: Springer-Verlag]. *Memory & Cognition*, 11(4), 415–422. <https://doi.org/10.3758/BF03202457>

- Hegarty, M., Richardson, A. E., Montello, D. R., Lovelace, K., & Subbiah, I. (2002). Development of a self-report measure of environmental spatial ability. *Intelligence*, 30(5), 425–447. [https://doi.org/doi:DOI:10.1016/S0160-2896\(02\)00116-2](https://doi.org/doi:DOI:10.1016/S0160-2896(02)00116-2)
- Holding, C., & Holding, D. (1989). Acquisition of route network knowledge by males and females. *Journal of General Psychology*, 116, 29–41.
- Hutson, S. (2010). Data handling errors spur debate over clinical trial. *Nature Medicine*, 16(6), 618–618. <https://doi.org/10.1038/nm0610-618a>
- Ishikawa, T., & Montello, D. R. (2006). Spatial knowledge acquisition from direct experience in the environment: Individual differences in the development of metric knowledge and the integration of separately learned places. *Cognitive Psychology*, 52(2), 93–129. <https://doi.org/10.1016/j.cogpsych.2005.08.003>
- Krstajic, D., Buturovic, L. J., Leahy, D. E., & Thomas, S. (2014). Cross-validation pitfalls when selecting and assessing regression and classification models. *Journal of Cheminformatics*, 6(1), 10. <https://doi.org/10.1186/1758-2946-6-10>
- Nazareth, A., Weisberg, S. M., Margulis, K., & Newcombe, N. S. (2018). Charting the development of cognitive mapping. *Journal of Experimental Child Psychology*, 170, 86–106. <https://doi.org/10.1016/j.jecp.2018.01.009>
- Pagkratidou, M., Galati, A., & Avraamides, M. (2020). Do environmental characteristics predict spatial memory about unfamiliar environments? *Spatial Cognition & Computation*, 20(1), 1–32. <https://doi.org/10.1080/13875868.2019.1676248>
- Peters, M., Laeng, B., Latham, K., Jackson, M., Zaiyouna, R., & Richardson, C. (1995). A Redrawn Vandenberg and Kuse Mental Rotations Test - Different Versions and Factors That Affect Performance. *Brain and Cognition*, 28(1), 39–58. <https://doi.org/10.1006/brcg.1995.1032>
- Ruginski, I. T., Creem-Regehr, S. H., Stefanucci, J. K., & Cashdan, E. (2019). GPS use negatively affects environmental learning through spatial transformation abilities.

- Journal of Environmental Psychology*, 64, 12–20.
<https://doi.org/10.1016/j.jenvp.2019.05.001>
- Schinazi, V. R., Nardi, D., Newcombe, N. S., Shipley, T. F., & Epstein, R. A. (2013). Hippocampal size predicts rapid learning of a cognitive map in humans. *Hippocampus*, 23(6), 515–528. <https://doi.org/10.1002/hipo.22111>
- Society for Neuroscience. (2015). Author-Initiated Retraction: Anderson et al, Induced Alpha Rhythms Track the Content and Quality of Visual Working Memory Representations with High Temporal Precision. *Journal of Neuroscience*, 35(6), 2838–2838. <https://doi.org/10.1523/JNEUROSCI.0074-15.2015>
- Soergel, D. A. W. (2015). Rampant software errors may undermine scientific results. *F1000Research*, 3, 303. <https://doi.org/10.12688/f1000research.5930.2>
- Strand, J. (2020). Error Tight: Exercises for Lab Groups to Prevent Research Mistakes. Retrieved July 26, 2021, from [10.31234/osf.io/rsn5y](https://doi.org/10.31234/osf.io/rsn5y)
- Strand, J. F., Brown, V. A., & Barbour, D. L. (2020). Talking Points: A Modulating Circle Increases Listening Effort Without Improving Speech Recognition in Young Adults. *Psychonomic Bulletin & Review*, 27(3), 536–543.
<https://doi.org/10.3758/s13423-020-01713-y>
- Sutton, J. E., Buset, M., & Keller, M. (2014). Navigation Experience and Mental Representations of the Environment: Do Pilots Build Better Cognitive Maps? *PLOS ONE*, 9(3), e90058. <https://doi.org/10.1371/journal.pone.0090058>
- Tobler, W. R. (1994). Bidimensional Regression. *Geographical Analysis*, 26(3), 187–212.
<https://doi.org/10.1111/j.1538-4632.1994.tb00320.x>
- Vandenberg, S. G., & Kuse, A. R. (1978). Mental rotations, a group test of three-dimensional spatial visualization. *Perceptual and Motor Skills*, 47(2), 599–604.
- Weisberg, S. M., & Newcombe, N. S. (2016). How do (some) people make a cognitive map? Routes, places, and working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(5), 768–785. <https://doi.org/10.1037/xlm0000200>

- Weisberg, S. M., Newcombe, N. S., & Chatterjee, A. (2019). Everyday taxi drivers: Do better navigators have larger hippocampi? *Cortex*, *115*, 280–293.
<https://doi.org/10.1016/j.cortex.2018.12.024>
- Weisberg, S. M., Schinazi, V. R., Newcombe, N. S., Shipley, T. F., & Epstein, R. A. (2014). Variations in cognitive maps: Understanding individual differences in navigation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(3), 669–682. <https://doi.org/10.1037/a0035261>
- Wilson, G., Aruliah, D. A., Brown, C. T., Chue Hong, N. P., Davis, M., Guy, R. T., Haddock, S. H. D., Huff, K. D., Mitchell, I. M., Plumbley, M. D., Waugh, B., White, E. P., & Wilson, P. (2014). Best practices for scientific computing. *PLoS biology*, *12*(1), e1001745. <https://doi.org/10.1371/journal.pbio.1001745>
- Youngson, N. L., Vollebregt, M., & Sutton, J. E. (2019). Individual differences in cognitive map accuracy: Investigating the role of landmark familiarity [Place: US Publisher: Educational Publishing Foundation]. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, *73*(1), 37–46.
<https://doi.org/10.1037/cep0000165>
- Zhao, J., Sensibaugh, T., Bodenheimer, B., McNamara, T. P., Nazareth, A., Newcombe, N., Minear, M., & Klippel, A. (2020). Desktop versus immersive virtual environments: Effects on spatial learning. *Spatial Cognition & Computation*, *20*(4), 328–363. <https://doi.org/10.1080/13875868.2020.1817925>