Precise display time measurement in JavaScript for web-based experiments

Gáspár Lukács ^{a,b}, Andreas Gartus ^a

^a Department of Cognition, Emotion, and Methods in Psychology, University of Vienna

^b JSPS International Research Fellow at the Aoyama Gakuin University

Authors' Note

Gáspár Lukács <https://orcid.org/0000-0001-9401-4830>

Andreas Gartus <https://orcid.org/0000-0002-9066-1114>

Corresponding author: Gáspár Lukács, Aoyama Gakuin University, 4-4-25 Shibuya, Shibuya-ku, Tokyo, 150-8366, Japan. E-mail: lkcs gaspar@gmail.com

This preprint is the final version (prior to typesetting) of the following paper (unfortunately not open-access, but viewable via <https://rdcu.be/cNMzm>):

Lukács, G., & Gartus, A. (2022). Precise display time measurement in JavaScript for web-based experiments. *Behavior Research Methods*. <https://doi.org/10.3758/s13428-022-01835-2>

Abstract

Conducting research via the internet is a formidable and ever-increasingly popular option for behavioral scientists. However, it is widely acknowledged that web-browsers are not optimized for research: In particular, the timing of display changes (e.g., a stimulus appearing on the screen), still leaves room for improvement. So far, the typically recommended best (or least worst) timing method has been a single `requestAnimationFrame` (RAF) JavaScript function call within which one would give the display command and obtain the time of that display change. In our Study 1, we assessed two alternatives: Calling the RAF twice consecutively, or calling the RAF during a continually ongoing independent loop of recursive RAF calls. While the former has shown little or no improvement as compared to single RAF calls, with the latter we significantly and substantially improved overall precision, and achieved practically faultless precision in most practical cases. Our two basic methods for effecting display changes, plain text change and `canvas` color filling, proved equally efficient. In Study 2, we reassessed the “RAF loop” timing method with image elements in combination with three different display methods: We found that the precision remained high when using either `visibility` or `opacity` changes – while drawing on a `canvas` element consistently led to comparatively lower precision. We recommend the “RAF loop” display timing method for improved precision in future studies, and `visibility` or `opacity` changes when using image stimuli. We publicly share the easy-to-use code for this method, exactly as employed in our studies.

Keywords: computer, display, timing, measurement, JavaScript, precision, online research

Introduction

The dramatic increase and impact of behavioral research conducted via the internet in the past decades, and in particular in the past years, have been reviewed in numerous papers and handbooks (e.g., Anwyl-Irvine et al., 2020; Birnbaum, 2004; Krantz & Reips, 2017; Peer et al., 2021; Ratcliff & Hendrickson, 2021; Wolfe, 2017). In brief, such reviews usually point out a number of great advantages but also some disadvantages and potential pitfalls. One limitation of online experiments relates to the software (Bridges et al., 2020; Pronk et al., 2020). The simplest and by far the most popular way to conduct online experiments is via web-browsers, using the JavaScript (JS) programming language to manipulate and monitor the HTML/CSS graphical interface for (interactive) stimulus display and recording responses (e.g., Grootswagers, 2020; technical mechanisms are extensively detailed by, e.g., Garaizar & Reips, 2019). However, as often noted (e.g., Garaizar & Reips, 2019; Pronk et al., 2020), browsers are not optimized for such specific purposes, and, in particular, they do not necessarily provide the timing precision that is desired in behavioral experiments. Plenty of studies have been devoted to comparing results of online studies and technologies to laboratory-based ones (e.g., Bridges et al., 2020; Crump et al., 2013; De Leeuw & Motz, 2016; Miller et al., 2018; Reimers & Stewart, 2015): By and large, the consensus is that the online alternatives work very well and produce very similar results as laboratory-based ones, but they are still to some extent inferior, and timing precision still leaves room for improvement (Anwyl-Irvine et al., 2020; Bridges et al., 2020; Grootswagers, 2020; Pronk et al., 2020; Reimers & Stewart, 2015). In the present paper, we demonstrate how the timing of any specific display change, measured in a browser via JS, can be improved, so that in most practical cases it would provide consistent millisecond precision similarly to conventional laboratory-based software.

There are two essential qualifiers of display change timing, typically termed, by convention, *accuracy* and *precision*. In this context, *accuracy* refers to the average deviation of the recorded time as compared to the real-world time: For example, for an eye tracking experiment, it may be important to know how many milliseconds passed from the start of a given eye movement to the display of a given stimulus. However, the need for the reliable timing of this kind of accuracy is relatively rare in behavioral research, and, importantly, typically relates to psychophysiological measurements (e.g., eye tracking, EEG) that are not generally used in online research in any case (due to their need for dedicated equipment and expert supervision). On the other hand, *precision* refers to the variability among the recorded times, that is, how consistent they are relative to each other: For example, response time

experiments typically measure the time passed between the display (i.e., appearance) of a given stimulus and the subsequent keypress response of the participant. To experimentally measure the average time taken for the response for different stimulus types, it is important for the measured time intervals to be consistent relative to each other. The more random variability is introduced to the data due to lack of precision in the measurement, the more the real differences are masked and difficult to ascertain, leading to reduced statistical power (e.g., Lord & Novick, 2008).¹ The level of precision is the main question regarding browser-based display change timing, and it is hence also the focus of the present study.

An often repeated technical recommendation has been that the best precision is achieved by using the `requestAnimationFrame` (RAF) native JS function (Barnhoorn et al., 2015; Gao et al., 2020; Garaizar et al., 2014; Pronk et al., 2020; Yung et al., 2015). Computer monitors actualize changes on their screen by periodically “repainting” it (having given pixels change color), at fixed intervals. The interval, the duration of one “frame,” is defined by the given monitor’s refresh rate, which is usually 60 Hz, meaning potential repaints (if any change is needed) at a constant interval of 16.7 ms. According to web standard specifications,² the RAF function makes it possible to execute commands in a callback³ immediately before the upcoming repaint (for the upcoming next frame), and thereby also to capture the time of that repaint, either via the dedicated internal timestamp of the RAF or by using a JS timing function, which simply returns the current time, within that callback (e.g., Pronk et al., 2020).

However, empirical testing is crucial, because the results can contradict what is expected (Garaizar & Reips, 2019), and indeed the RAF’s behavior too can differ unexpectedly depending on the circumstances (as demonstrated in the present study, but also as indicated by repeated sanity checks throughout the past years; Lukács, 2018). Relatively few studies have been devoted to empirically comparing different JS methods for online experiments – and these studies focused not on display change timing precision, but primarily on ensuring the correctness of the duration of a given stimulus being displayed on the screen (Gao et al., 2020; Garaizar & Reips, 2019; Garaizar et al., 2014; Pronk et al., 2020). For

¹ In some cases, researchers may even be interested in the magnitude of variability itself. In that case, an unknown extent of added variability due to measurement error is even more of a problem.

² <https://developer.mozilla.org/en-US/docs/Web/API/window/requestAnimationFrame>

³ A “callback” or “callback function” is any executable code passed into a main function (here: RAF), to be subsequently executed within that main function at a time specified in that main function.

example, in a given online behavioral response time experiment (and in any given trial), a stimulus may be intended to be displayed on the screen for 50 ms. In that case, counting from the instance the stimulus is drawn on the screen, it should stay there for the duration of three frames, assuming a 16.7 ms duration for each frame. Depending on the JS method employed, there will be a certain ratio of trials in which the display duration is incorrect: The stimulus may be displayed either in too few frames (e.g., two frames, for 33.3 ms), or in too many frames (e.g., five frames, for 83.3 ms). The repeated pivotal finding is that correct display duration is most frequently achieved when using consecutive RAF calls (Gao et al., 2020; Garaizar & Reips, 2019; Garaizar et al., 2014; Pronk et al., 2020): Within the callback of an initial RAF function, the display command (for the stimulus to appear) is given together with another RAF call, which is to be executed in the next frame. This other RAF call then also includes yet another RAF call in its callback, to be executed in the frame coming afterwards, and so forth. Thereby, the number of frames can be counted via the number of RAF calls, and the command to make the stimulus disappear can be given when the desired number of frames is reached (or, similarly, when the time measured via the timestamps of the RAF functions reaches the desired time).

The precision of display change timing in JS is however a different question. As explained above at the beginning, the more precisely the display changes are timed, the more precisely (i.e., with less noise) one can record its time distances from other events, such as a keypress response (or any other input or output preceding or succeeding the given display change).⁴ Such precision has been reported in various manners and in various methodological contexts, typically showing that the precision is acceptable, though not perfect (e.g., Anwyl-Irvine et al., 2020; Bridges et al., 2020; Pronk et al., 2020; Reimers & Stewart, 2015). However, we know of no paper that reported any comparisons of different JS methods in view of display timing precision.⁵ The current best practice seems to be a single RAF call (regardless of whether or not it is followed or preceded by other RAF calls), in which the

⁴ The precision of display change timing can nonetheless be relevant for correct display duration too, in that the latter can be assessed by the former (e.g., with sufficiently precise timing, the display duration can be measured in JS in each trial to, e.g., verify the durations' correctness generally, or to exclude trials with incorrect duration).

⁵ De Leeuw and Gilbert (2019) informally reported some tests using a "double RAF" timing method (see below), but the results for this method seem inconclusive (not unlike in our study).

display command is executed, and whose timestamp is used for timing the corresponding display change.⁶

In our study, we used automatized keypresses as a baseline, and the JS command to effect the display change (make a stimulus appear on the screen) was given as soon as the keypress was detected. In our analysis, we compared the reaction time distances, from the keypress to the display change, as measured (a) in JS and (b) via external devices. Since the external measurements as well as the keypress detection in JS are all extremely precise (e.g., unlikely to be off by more than 2 ms; as also separately verified in our study), any substantial variance in the difference between the JS-timed and externally-timed distances (e.g., more than 2 ms deviations from the mean difference) indicates the level of imprecision of the JS display change timing, for which we used a variety of JS methods to be assessed and compared to each other, in order to find the most precise one(s).

In Figure 1, via a preview of our findings, we illustrate how the results of display timing via a single RAF call can vary depending on browser and Operating System (OS): In Google Chrome on Microsoft Windows, the returned RAF timestamp generally indicates the last repaint (before the RAF command is executed), in Safari on Apple’s macOS it indicates the next repaint (right after the RAF command, hence probably close to when the element is actually painted), and in Mozilla Firefox on Microsoft Windows it simply returns the current time (whenever the RAF command is executed; hence not indicating any repaint time, unless by chance). Whether the measurement captures the last or the next repaint (or any other), as long as it is consistent, should not matter for precision (although the difference could bias comparison between participants who use different browsers/OS combinations) – however, not capturing any repaint time introduces a random error, each time ranging, if we assume a 60 Hz refresh rate, from zero to 16.7 ms. Furthermore, even when the timestamp indicates a repaint time, it may not always capture the correct one (i.e., it may be inconsistent relative to the true display time, see Figure 1, right panel).

⁶ The only exception we know of is the JS framework Labvanced (Finger et al., 2017), in which, to improve precision, the “RAF loop” method (see below) is already implemented at least when Chrome browser is detected (based on the recommendation by Fraipont, 2020). However, at least one other JS framework (PsychoJS) as well as one previous study (Pronk et al., 2020) also included a continual RAF loop, although not with any explicitly noted intention of improving the precision of display change time measurements, but primarily as a corollary of their technical implementation. Nonetheless, this might at least in part explain why PsychoJS was found, by Bridges et al. (2020), to provide higher precision than other JS frameworks.

Figure 1*Error Distribution Examples*

Note. The figure depicts first of all, in two shades of dark blue, three typical examples of JS timing measurement error distribution, in histograms with values calculated as “*JS-timed reaction time*”: (time from JS keypress time detection to JS display time detection, using RAF timestamp in all three examples) minus “*real reaction time*” (time from external keypress time detection to external display time detection). Secondly, in gray and positioned higher (beginning at count 50; so as not to overlap with the timing counts), the relative JS-timed keypress is depicted (the *JS-timed reaction time* subtracted from the difference described above); so that one can see, in gray, when the keypresses were detected (in JS), and, in comparison, in dark blue, when the consequent display changes (invoked by the given corresponding keypresses) were detected (in JS). The macOS/Safari combination (left panel) seems to capture the desired time: The detected display change is extremely consistent with the external measurement, and clearly indicates either an exact repaint time or a consistent time distance from a repaint time (with the practically negligible few ms deviations likely attributable to hardware limitation). Furthermore, it is preceded by keypress detections within a range of ca. 17 ms (in addition to a brief, ca. 5 ms gap), indicating that the browser probably captured the change that was actualized in the next available repaint following each keypress detection. The Windows/Firefox combination (middle panel) seems to completely miss repaints, it simply returns the time when the RAF was called immediately following the keypress detection. Finally, the Windows/Chrome combination (right panel) seems to capture repaint times, but not the desired ones: The timestamp indicates repaints that preceded the actual display, and, more importantly regarding precision, the relative distance is not consistent, the detected time is at sometimes one, sometimes two frames distance from the real one. Less importantly, the figure also differentiates black-to-white display change (background: black) and white-to-black display change (background: white), showing general equivalence (or very minute difference in case of Windows/Chrome).

Part of the explanation is that the RAF was implemented in browsers not with the aim of obtaining the precise time of a single physical display change (for which there is hardly any use in conventional web pages and applications), but rather, primarily, to achieve smooth animation by continuously calling it frame by frame to effect a new display change in each new frame – which may, incidentally, also explain why consecutive RAF calls work well for ensuring correct display duration. In the present paper, we assess two alternatives that address this issue in view of display change timing.⁷ The first is the “double-RAF”: calling the RAF a second time within the callback of the first one, and using the timestamp of the latter (De Leeuw & Gilbert, 2019). The rationale is that, even if the first RAF’s timestamp is unrelated to repaint times, the callback with the second RAF is executed near the repaint time, and hence the latter’s timestamp may yield a correct value (i.e., a repaint time either at, or at a consistent distance from, the real repaint time in question). The second alternative is the “RAF loop”: calling an instance of the RAF during an independently ongoing RAF loop where the function keeps calling itself recursively in each new frame (Fraipont, 2020).⁸ To increase reactivity, some browsers, when there is no concurrent ongoing animation (i.e., anything to be rendered for display), force an immediate display change that disregards the RAF callback’s expected time of execution.⁹ However, calling an independent RAF beforehand initiates an animation event (as its very name indicates, a “requested” new frame including rendering update, even despite no actual display change),¹⁰ which prevents this forced reaction, and the display change will be in line with the RAF callback and its

⁷ We have in fact examined altogether eight different alternatives, but all of these fall into the category of one of the two main alternatives, and yield practically identical results. For example, to make sure it makes no difference, we made separate measurements for placing the command for display change before the RAF instead of within its callback. The details and the results are available in the online Appendix A (<https://osf.io/7h5tw/>).

⁸ This method is not the same as the consecutive RAF calls intended for correct display duration. The latter method too creates a loop at least while the stimulus is being displayed, but this was not specifically intended as a method for precise display change timing as here. The initial RAF call when the stimulus is to appear is not necessarily preceded by other RAF calls, and none of the related papers, to our knowledge, have brought up the relevance of preceding RAF calls in view of timing precision. For our RAF loop method, it is essential that the loop be initiated and already ongoing prior to (and independently of) any display changes and their timings.

⁹ See https://bugzilla.mozilla.org/show_bug.cgi?id=1675614.

¹⁰ See <https://html.spec.whatwg.org/multipage/webappapis.html#event-loop-processing-model>.

timestamp. How well these methods actually perform for our purpose of precise display timing is an empirical question.

As an extension, to more thoroughly address the question of ideal display timing in various cases, we also assess, in a second study, how different methods for displaying images may affect precision.

Study 1: Timing Methods

To ensure technical integrity, extensive piloting with various sanity checks and trial-level visual inspection preceded the main measurements below (using different computers and earlier versions of the script). The main pilot findings are briefly presented in online Appendix B (<https://osf.io/5wc9r/>; all raw data is available via the OSF repository), and all correspond to the main findings presented below.

The computer code examples throughout the paper are simplified (e.g., with shortened variable names) for better illustration and, while having an identical mechanism, they are not format-wise identical to the ones we actually used in the study. (All complete original scripts are available via <https://osf.io/wyebt/>.)

Method

Apparatus

Computers and Browsers

We performed measurements on four different computers: two computers with Windows 10, designated as “Win (A)” (CPU: Intel Core i3-550; GPU: NVIDIA GeForce GT 710), and as “Win (B)” (CPU: Intel Core i5-8500; GPU: Intel UHD Graphics 630); an iMac18,3 “macOS” computer (CPU: Intel Core i7-7700K; GPU: AMD Radeon Pro 580); and one computer with Linux Mint 20 (“Linux”; CPU: Intel Core i5-7200U; GPU: Intel HD Graphics 620). The monitors were always set to maximal brightness and contrast, and always had a 60 Hz refresh rate. Detailed system and hardware information for each computer is available via the OSF repository. To note, we thought it worthwhile to include Linux (e.g., to compare our results on this OS to those in previous studies), but it should be kept in mind that Linux-based OSes are very rarely used by research participants (e.g., Anwyl-Irvine et al. [2020] estimated their prevalence at ca. 3%, with Windows at 74%, macOS at 22%; well in line with general OS market share estimates).

We measured all major browsers available for the given OSes: Mozilla Firefox, Google Chrome, Opera, Microsoft Edge (for Windows), and Safari (for macOS). We used recent browser versions, in most cases the latest one available at the time of the data collection (September-October, 2021). For each measurement, the specific version of each

browser is recorded in the corresponding data files. Here too, we would like to point out that browsers differ in popularity, and in particular Opera is only very rarely used. For example, the browser distribution in recent representative samples from the US and China were, respectively (regardless of OS): 71.8% and 81.9% Chrome, 16.0% and 8.9% Firefox, 4.4% and 7.9% Edge, 7.1% and 0.5% Safari, and only 0.8% and 0.5% Opera users (Kawai et al., 2021; see also Anwyl-Irvine et al., 2020; again well in line with general browser market share estimates).

Keypress simulation and external timing

For keypress simulation we used an Arduino Leonardo microcontroller (ATmega32u4, <https://store.arduino.cc/products/arduino-leonardo-with-headers>) as a virtual keyboard. For the detection of brightness changes, we used a Texas Instruments monolithic photodiode and single-supply transimpedance amplifier (OPT101P; powered by a 9V block battery). Both devices were connected to a virtual oscilloscope and logic analyzer (LHT00SU1), which in turn was connected to a regular computer on which the keypress and brightness data was recorded using the software Sigrok and fx2lafw with PulseView frontend (<https://sigrok.org/>).

Procedure

Having set up all the devices, the Arduino was connected via USB to the computer (A) on which the stimuli were displayed in the browser. Following a starting key (“x”), keypresses of “q” were simulated at an interval that each time randomly varied between 1,000 and 1,100 ms. (The keypresses stopped only with the manual removal of the USB, after all trials had been completed.) At the same time as the keypress start (*keydown* command), a trigger was sent to the logic analyzer and recorded on the other computer (B). As soon as the keypress (*keydown*) was detected by the browser, a JavaScript command was given to start displaying the given stimulus on the screen of the former computer (A). The display (brightness) change was recorded on the other computer (B) via the photodiode and the logic analyzer, with 20 kHz resolution.

In each full recording, the display change was either always a fully black square appearing (immediately after the keypress detection) and disappearing (within 500 ms, see below) on a fully white background, or always the reverse colors (i.e., white squares on black background). Figure 1 showcases that, despite general equivalence, in a few instances (here on Windows/Chrome), there may be minute (1-2 ms) differences between black-to-white display change (background: black) and white-to-black display change (background: white) – but this of course also depends on how we define and measure physical color (i.e., brightness) change. However, the relevant variance differences (between different timing methods) are

far larger than the largest discrepancy caused by the background difference (see supplementary figure at <https://osf.io/xedms/>). Hence, for all further analysis, these two types of display changes are merged.

We used two different display methods, both very basic and straightforward, in order to make sure that the results were not specific to either: plain text change using a square symbol (as `di.textContent = '■'`; where `di` is a *div* type graphical element in the HTML – with its text emptied before the end of each trial) and canvas filling (as `CanvasRenderingContext2D.fillRect(0, 0, canv.width, canv.height)`; where `canv` is a *canvas* type graphical element in the HTML – cleared before the end of each trial). This method difference did not affect the results either (see supplementary figure at <https://osf.io/7j8pc/>), and therefore, for all further analysis, these two types of display methods are merged.

The display duration setting varied between 16.7, 50, 150, 300, and 500 ms. Since the most straightforward measure of display timing variance is the time from the keypress to the initial display change (square stimulus appearance), this is our main dependent variable, and the duration (i.e., the time of disappearance) for this measure does not matter.¹¹ However, the precision for duration timing was also analyzed and presented in online Appendix A (<https://osf.io/7h5tw/>), with results generally closely corresponding to those with keypress-to-display timings. In a single recording, the stimulus was presented 10 times per each of the 5 durations, per each of the two presentation methods (plain text or canvas), for a total of 100 presentations per each timing method. Since our black and white background measurements are merged, we have altogether 200 measurements for each timing method.

To provide the perspective of a “gold standard” baseline, we also conducted an analogous recording with PsychoPy (version 2021.2.1; Bridges et al., 2020; Peirce, 2008), on the Win (A) computer, with 50 trials per duration and per background (white or black), hence altogether 500 trials.

Timing Methods

¹¹ The precision differences in estimating duration (appearance-to-disappearance time), presented in the appendix, are very similar to those in estimating keypress-to-display time. However, the former is confounded by the timing methods affecting durations (especially those lasting one or few frames): As it has been repeatedly shown, display command within RAF, and counting time via RAFs, leads to more consistent durations. However, to avoid interference with timing methods, we set all durations using the much less reliable `setTimeout` JS function.

We compare four different JS methods: (1) *direct*, (2) *RAF 1*, (3) *RAF 2*, and (4) *RAF loop*. In each trial, the same JS method was used to time the appearance as well as the disappearance of the stimulus.¹² The *direct* method means merely recording the current time immediately after the display command (i.e., the current time is called in the next line in the code), without the use of the RAF (and hence without any regard of the refresh rate and the actual repaint time); see Listing 1. (The current time was always recorded using the `performance.now()` native JS function, which returns, with at least one-millisecond resolution, the time passed since the loading of the page. The timestamp in the RAF is a conceptually equivalent JS timer.) For the RAF 1 method, the display command was given within a RAF callback, and timed via the RAF's timestamp¹³ (see `RAF1` function in Listing 1). For the RAF 2 method, the display command was given within a RAF callback, where immediately afterwards another RAF was called, and the timestamp of this latter RAF was used for timing (see `RAF2` function in Listing 1).

Listing 1

JavaScript code for the direct, RAF 1, and RAF 2 methods (display functions)

```
function direct() {
    di.textContent = '■'; // display command
    display_time = performance.now(); // time recording
}

function RAF1() {
    requestAnimationFrame(function(timestamp) {
        di.textContent = '■'; // display command
        display_time = timestamp; // time recording
    });
}
```

¹² The given method's display function (see Listing 1) was called immediately after the detection of the keypress, to make a stimulus appear, and, after the given display duration (specified via `setTimeout`), the same method was called again to make the stimulus disappear. In case of the RAF loop method, the loop was initiated well in advance.

¹³ In all cases of all methods, directly calling `performance.now()` in the RAF callback (in the next line after the display command) gave results either equivalent to the timestamp or otherwise entirely disregarded repaint times, same as if no RAFs were called at all (see histograms at <https://osf.io/c3eum/>). We therefore used the timestamp in all cases.

```
function RAF2() {
  requestAnimationFrame(function() {
    di.textContent = '■'; // display command
    requestAnimationFrame(function(timestamp) {
      display_time = timestamp; // time recording
    });
  });
}
```

For the RAF loop method, the display and timing is the same as for RAF 1, except that there is an independently ongoing RAF loop function (see Listing 2). For our measurements, the RAF loop function was initiated about 100 ms after the disappearance of the preceding stimulus, so that it started at least 400 ms before the appearance of the upcoming stimulus. The display was then timed within one RAF, using its timestamp (see RAF 1 in Listing 1). The loop was ended, in each trial, after the disappearance of the stimulus.

Listing 2

JavaScript code for the RAF loop

```
function loopFunction() {
  if (do_loop) {
    requestAnimationFrame(loopFunction);
  }
}

function startLoop() {
  do_loop = true;
  loopFunction(); // initiate RAF looping (before trial)
}

function endLoop() {
  do_loop = false; // end RAF looping (end of trial)
}
```

Data Analysis

In the case of white background, all brightness values were reversed (multiplied by minus one), so that the rest of the analysis would be identical regardless of background. The exact point of display change was measured as a value passing one-third of the distance

between minimum and maximum brightness values. Apart from extensive visual inspection of trial-level segments (that all verified the correctness of detections), a number of programmatic safety checks were also implemented to automatically verify the integrity of each given detection during the analysis (for details, see the R scripts at the OSF repository). Assuming that the measurement error is negligible for all practical purposes, we term as “real time” the times of these externally recorded and detected display changes, as well as the times of the externally recorded keypresses.

Given our focus on precision, of main interest is the variance of the JS-timed minus real time keypress-to-display differences. As for statistical reporting, previous papers that compared browser performance (typically: duration of a stimulus display) with externally measured “real” time have taken a variety of approaches. For one, almost all papers report at least the means (accuracy) and *SDs* (precision) for the difference between the browser performance and the real times (e.g., Bridges et al., 2020; Kuroki, 2021; Reimers & Stewart, 2015). We too report these as our main results, with a focus on *SD* (which reflects precision).

Specific to assessing correct display duration, which always primarily concerns the number of whole frames, some papers reported the numbers or ratios, in each given test set, of displays with incorrect numbers of frames (Garaizar & Reips, 2019; Garaizar et al., 2014; Pronk et al., 2020). Since display duration was not our focus, this is not applicable here.

Finally, some studies included histograms or similar graphic depiction (Anwyl-Irvine et al., 2020; Gao et al., 2020; Pronk et al., 2020). Such graphic depiction seems to us most informative for our results too. However, the unusually large number of our examined cases (60 OS/browser/method combinations for main results alone) prevents any sensible way of detailed presentation within the manuscript. Nonetheless, importantly, we did upload, to our OSF repository, all histograms per OS/browser combination, which the readers are encouraged to consult.

However, we actually have a way to characterize the distribution (together with variability) in a succinct and readable manner. This is mainly because each OS/browser combination that we tested resulted in an outcome that may be informally categorized as belonging to one of the three types of outcomes depicted in Figure 1, and each case shows a fairly similar distribution to the corresponding exemplar: (1) practically faultless capture of repaint time with consistent distance (or no distance) from the actual display (left panel); (2) no capture of any repaint time, hence random error between ca. 0-17 ms (middle panel); and (3) capture of certain repaint times, but at inconsistent distances from the actual display (right panel). These three types of outcomes can be quite easily distinguished by reporting, in

addition to SD, median absolute deviations (MADs; e.g., Arachchige & Prendergast, 2019). The SD well reflects the measurement errors in both of the inferior outcome types: no repaint time capture (2), and inconsistent repaint time capture (3). However, while the MAD will also be high in case of no repaint time capture at all (2), it will be low in case of inconsistent repaint time capture (3).

For illustration, in the examples of Figure 1, faultless repaint capture (1) gives SD = 0.56, 95% CI [0.51, 0.62] and MAD = 0.56, 95% CI [0.47, 0.64]; no repaint time capture (2) gives SD = 5.33, 95% CI [4.85, 5.91] and MAD = 6.15, 95% CI [5.42, 6.89], and inconsistent repaint time capture (3) gives SD = 7.20, 95% CI [6.55, 8.00] and MAD = 1.30, 95% CI [0.98, 1.63]. For the gold standard perspective, we may consider the PsychoPy test results: For the PsychoPy-timed minus real time keypress-to-display differences, the SD is 1.30, 95% CI [1.23, 1.39], and the MAD is 1.90, 95% CI [1.71, 2.10].

Hence, wherever relevant, in addition to SD (which is still the primary statistic, representing all types of timing errors), we always report MAD too as a complementary statistic (representing the case of no repaint time capture).

Previous related studies usually did not report formal tests of differences, such as statistics of null hypothesis testing. This seems reasonable given that the statistical power in such studies is typically evidently extremely large, with a very high number of observations and comparatively low variance (i.e., all in all, the findings are obviously reliable) – as it seems to be the case in our study too. Nonetheless, we decided to use significance testing, because, especially in case of such a large number of comparisons, such additional statistics show the strength of evidence in a more immediately interpretable manner (see also, e.g., Fricker et al., 2019). In a similar vein, we include the 95% CIs of all reported point estimates (means, medians, SDs, and MADs) to provide a tangible measure of certainty (e.g., Lakens et al., 2018; Simon, 1986).

For variance equality tests specifically, we always use both a Brown-Forsythe test (Levene's test using medians; LBF; Brown & Forsythe, 1974) and a Fligner-Killeen test (FK; Fligner & Killeen, 1976). These are both relatively robust, but the FK is in particular less sensitive to outliers (Conover et al., 1981; Fox & Weisberg, 2019). Hence, as apparent in our results, while both tests' significance values primarily reflect changes in SD, the FK tends to be more likely to detect changes in MAD as well.

For all reported *p* values, in each set of significance tests for any given independent variable, we use the Benjamini-Hochberg procedure to correct for multiple-testing (Benjamini & Hochberg, 1995). In most cases, we indicate statistical significance in figures,

with three asterisks (***) for $p < .0001$ (very strong evidence); two asterisks (**) for $p < .001$ (substantial evidence); one asterisk (*) for $p < .01$ (weak evidence); and hyphen (-) for $p \geq .01$ (given the exploratory nature of our analyses, we do not consider p values no less than .01 as noteworthy evidence). All exact p values per test, along with the exact values of all other reported or depicted statistics, are available in data tables via the OSF repository.

All analyses were conducted in R (R Core Team, 2020; with extension packages by Lukács, 2021; Signorell, 2021; Wickham, 2016).

Results

Keypress Detection

We performed a preliminary analysis to ensure that our main comparison of timing methods was not biased by measurement errors of keypress detection (or hardware limitations). These results, however, may be interesting in their own right: Although several studies have explored the precision of accuracy of detecting keypress responses in online experiments (e.g., Pinet et al., 2017; Reimers & Stewart, 2015), we do not know of any previous findings that pinpointed the extent of variability caused by given browsers alone.

For each separate recording (each OS/browser combination, per each background type), we calculated the difference between JS-timed keypress detection and the real keypress time, for each trial. Since the mean difference (which simply depended on when the browser was loaded relative to when the key simulation was started) is completely irrelevant here, for easier visual comparison, the mean of all values (differences) within each given recording was subtracted from each value within that recording, to obtain a mean of zero. Furthermore, a linear regression coefficient was calculated and the slope and intercept were also subtracted from each value. This was needed because there was a constant skew between the browser's timings and the external timing, though extremely small (with JS-time ca. 50 ms behind external time by the end of the ca. 17 min recording; see the supplementary figures via OSF), and the deviations due to the passing of time, without controlling for it (i.e., for the passing of time), would have introduced an irrelevant surplus variance far larger than the actual measurement error related to the relevant consistency (i.e., precision) of the time elapsing between the real keypresses and the corresponding keypress detections.¹⁴

¹⁴ Such adjustment was not necessary in previous studies or in the rest of the analyses in the present paper, because the relevant time distances were compared within-trial, using, within each trial, a certain dedicated mutual baseline (e.g., a keypress recorded externally and internally, or the internal command and the external detection of a stimulus appearing). In this specific case, however, we

Figure 2*Keypress Detection Precision per OS and Browser*

Note. Keypress detection precision indicated as the variance (SD and MAD, with 95% CI in error bars) of the standardized difference between the JS-timed and externally timed keypresses. The statistical significance for the comparison with PsychoPy is indicated in parentheses by asterisk(s) or hyphen (see Data Analysis section), with LBF on the left and FK on the right (separated by whitespace; e.g., $p = .0645$ for LBF and $p = .0003$ for FK would be indicated as “- **”).

The resulting variances are presented in Figure 2. To avoid excessive statistical testing (which is hardly relevant here in any case), each browser’s variance was compared only to that of PsychoPy. However, regardless of mostly statistically significant differences, the variance was extremely small in all cases (both SD and MAD always below 0.9 ms), and presumably negligible for all practical purposes – especially in view of the far larger variances caused by suboptimal timing methods (see below).

wanted to assess the consistency of distances between the internal and external keypress detections throughout the entire duration of each ca. 17-min testing, for which the mutual points of connection consisted of all the keypress detections themselves (internal paired with external, at each point) in the given testing. The skew became relevant only in this latter case, and only in this case did it need to be adjusted.

Timing method comparisons

Our main analysis concerns the difference between the (a) JS-timed and (b) real (externally timed) keypress-to-display time distances. The variance (SD, MAD) in the difference between JS-timed and real time distances reflects the precision of the given JS timing method. The results of comparing these different JS timing methods are presented in Figure 3. It can be seen that, in comparison to the direct method, the single RAF (RAF 1) and the double-RAF (RAF 2) almost never improve (decrease) variances as measured with SDs, but only when measured with MAD (and even so not in all cases). The RAF loop method in contrast shows substantial advantage: In each OS and browser combination, and in comparison with each other method, the RAF loop either provides similar variances, or, especially in Windows, significantly lower variances. Using RAF loop on Windows, in seven out of eight cases (the one exception is Firefox on the “Win (B)” computer), both SD and MAD drops to a near-zero level that is practically faultless, that is, indistinguishable from errors caused by hardware and keypress detection errors (see above). On macOS, it provides no advantage over single or double RAF, but the SD and MAD are very low in all cases except with Firefox. On Linux however, the variance remains rather high in all cases, with no improvement with single or double RAF, and little or no improvement using RAF loop.

The RAF loop, however, has a side effect: it causes a one-frame delay in display (in all cases except in Safari on macOS, where it has the reverse effect – which also means that, curiously, Safari introduces a one-frame delay by default). That is, when the display command is executed during a given frame, the physical display will not be actualized at the beginning of the next frame (i.e., next repaint), but only at the beginning of the one after that (i.e., the second repaint following the display command). This may be best illustrated by keypress-to-display (reaction) times, depicted in Figure 4, where there is a clear ca. 16.7 ms difference in all cases of using RAF loop (but see further related checks in online Appendix B; note also that this is independent of whether the display command is given within a RAF callback or immediately following the keypress detection, as shown in online Appendix A). The delay is presumably related to the prevention of forcing immediate display, as explained in the Introduction.¹⁵ Nonetheless, this small reaction delay to external input is not usually

¹⁵ We have no precise technical explanation for this phenomenon, but setting `CanvasRenderingContext2D.desynchronized` to `true` for canvas elements removed this delay on all Chromium-based browsers (Chrome, Edge, Opera) on Windows, which indicates that the delay is on the OS level (see <https://html.spec.whatwg.org/multipage/canvas.html#concept-canvas->

relevant to online research (see General Discussion).

Figure 3

Display Detection Precision in Study 1, per Timing Method, OS, and Browser

Note. Display detection precision indicated as the variance (SD and MAD, with 95% CI in error bars) of JS-timed minus real (externally timed) keypress-to-display time differences. The statistical significance for each comparison is indicated by asterisk(s) or hyphen (with LBF on the left and FK on the right; see Data Analysis section and the note under Figure 2).

desynchronized). This setting however is currently only available for these specific combinations (i.e., canvas elements in Chromium-based browsers on Windows), hence it does not provide a reasonable general solution.

Figure 4*Reaction Time in Study 1, per Timing Method, OS, and Browser*

Note. Display reaction time means (with 95% CI in error bars) and medians: The externally measured time from the keypress to the physical appearance of the corresponding stimulus.

Study 2: Image Display Methods

In Study 1, we used two different display methods, (a) plain text change and (b) canvas filling. Previous studies have mostly used the canvas filling method (e.g., Gao et al., 2020; Garaizar et al., 2014; Pronk et al., 2020), but, as noted in Study 1, we found that plain text change is equivalent to it in terms of display timing precision. This is unsurprising, given that both methods are computationally undemanding and their processing is in itself unlikely to influence precision. However, displaying HTML image elements, such as ones loaded from large photographic image files, is a more complicated process, and has rarely been studied – even though images (loaded from files) as stimuli are widely used in behavioral experiments. A notable exception is the study by Garaizar and Reips (2019), who also gave an extensive review of the technical background of the subject. However, they examined and reported only display duration errors (in terms of frame numbers), and did not address display

change timing. Most importantly, the RAF loop timing method itself has also never been examined with images before.

Therefore, we extend the findings of Study 1 by assessing its best-performing “RAF loop” timing method for images (image files loaded as image elements), displayed via three different JS methods that are to be compared with each other. First, changing `visibility` (`hidden` vs. `visible`), as a “default” method, since this is what is most generally used for having a graphical element appear or disappear in browsers. Second, changing `opacity` (0 vs. 1, with `willChange` style property preset to `opacity`), as recommended in at least two previous papers (Garaizar & Reips, 2019; Pronk et al., 2020). Third, drawing images to a `canvas` element (via the `CanvasRenderingContext2D.drawImage()` method): Unlike the other two methods, this is a synchronous procedure that should circumvent complete page reflow and all its side effects (for details, see Garaizar & Reips, 2019), and therefore has been suggested as a potentially optimal solution for precise timing (Fraipont, 2019; Garaizar et al., 2014).

Method

We used the same apparatus (computers, browsers, and keypress simulator and recording) as in Study 1, and the procedure was also the same except that the timing method was always “RAF loop” (as described in Study 1) and that the stimuli were images as described below (and displayed using changes in `visibility` or `opacity`, or by drawing on `canvas`).

The image stimuli were either fully black and fully white, always presented on a background of the reverse color (white or black, respectively). For a realistic sample, we took random sets of images from the Geneva affective picture database (GAPED; Dan-Glauser & Scherer, 2011; 20 BMP images), the Open Affective Standardized Image Set (OASIS; Kurdi et al., 2017; 10 JPG images), and the Bicolor Affective Silhouettes and Shapes (BASS; Kawai et al., 2021; 10 PNG images). Each of these images was converted into one fully black and one fully white image. While the large uncompressed GAPED images and the very small and already simple (bicolor, black-and-white) BASS images retained their file size (GAPED: $M \pm SD = 615 \pm 311$ kB, full range: 308-922 kB; BASS: $M \pm SD = 2.3 \pm 2.3$ kB, full range: 0.9-8.2 kB), the originally large and complex but compressible JPG type OASIS images lost most of their size during color conversion and became very small as well ($M \pm SD = 3.9 \pm 0.1$ kB, full range: 3.8-4.2 kB). Altogether, this provided a balanced and relatively wide range of commonly employed image types and sizes. In each recording, each of the 40 images was

presented one time per duration, for a total of 200 presentations per display method (and 400 with black and white merged). All images were preloaded (and the test starting key worked only after all preloading was completed).

Results

The results of our main analysis of comparing display methods is presented in Figure 5, and seem clear-cut: Drawing on `canvas` leads to larger variance than either of the other methods in all cases (all browsers) on Windows and in one case (though less robustly) in Linux. There are no other significant or substantial differences – `visibility` or `opacity` changes perform equally well, and they are in most OS/browser combinations once again practically faultless, same as in case of the simpler stimuli in Study 1. Once again the exception is Linux, where variance is very high in all three browsers regardless of display method.

Drawing on `canvas`, as compared to the other two methods, also delays keypress-to-display times in some cases, while, in this respect too, there are no substantial differences between `visibility` or `opacity` changes; see Figure 6. The reaction times when using `visibility` or `opacity` changes are generally similar to those with the non-image stimuli in Study 1 (using RAF loop), although somewhat (up to ca. one frame) slower in Linux, and, interestingly, faster (up to ca. one and a half frames) in several cases on Windows.

Supplementary analysis in online Appendix C (<https://osf.io/fs89t/>) shows that precision is consistently lower for the larger GAPED images as compared to the smaller OASIS and BASS images. With the latter small images, in case of the browsers in Windows with `visibility` and `opacity` changes, even the small remaining variance is reduced (to near-zero) in most cases. Display reaction times were not found to be substantially affected by image size.

Figure 5*Display Detection Precision in Study 2, per Timing Method, OS, and Browser*

Note. Display detection precision indicated as the variance (SD and MAD, with 95% CI in error bars) of JS-timed minus real (externally timed) keypress-to-display time differences. Same as above, the statistical significance for each comparison is indicated by asterisk(s) or hyphen.

Figure 6*Reaction Time in Study 2, per Timing Method, OS, and Browser*

Note. Display reaction time means (with 95% CI in error bars) and medians: The externally measured time from the keypress to the physical appearance of the corresponding stimulus.

General Discussion

Generally, even without special timing methods, the variance introduced by timing errors is not too drastic, and it should normally be possible to compensate by certain amounts of increase in sample size or within-test observations (e.g., Bridges et al., 2020; Miller et al., 2018; Reimers & Stewart, 2015; see also the simulations by Pronk et al., 2020). Nonetheless, better precision, by increasing reliability and statistical power, helps achieve larger effects, improves diagnostics and estimates of true variance, and, in the end, saves time and resources.

Here, we introduced a simple and costless method to increase timing precision: one simply needs to run a “RAF loop” in the background, add the display command in a RAF, and use the latter RAF’s timestamp as the timer.¹⁶ In our study, in contrast to other timing methods, this led to practically negligible variance in most cases, in particular on Windows and macOS, with simple text changes and canvas filling, as well as when displaying images with `visibility` or `opacity` changes.

The exception is Linux, where, unfortunately, the variance remained rather high in all cases, regardless of timing method. This is in line with previous studies that also failed to achieve high precision in Linux (Gao et al., 2020; Garaizar & Reips, 2019). Nonetheless, it is crucial that (as detailed in the Methods sections) Linux users generally represent only a very small fraction of online research participants (ca. 3%). If timing precision is critical, one radical solution would be to simply exclude Linux users. Then again, if timing precision is really critical, laboratory-based studies with specialized software (and hardware) should be preferred: Notwithstanding the improved method we introduced, faultless precision can hardly ever be guaranteed without direct supervision and dedicated equipment (ideally pre-verified using external measurements, such as an oscilloscope; see also, e.g., Plant & Turner, 2009).

The RAF loop also results in the side-effect of a one-frame delay in displaying the desired element (except in Safari on macOS, where it actually eliminated the delay that apparently exists by default). Since the delay is consistent (e.g., it affects the appearance and disappearance of the stimuli equally), the only concern is the delayed reaction time to external input signals. However, online experiments do not generally require extremely fast display reactions to external input, and a 16.7 ms delay (assuming a monitor with the regular 60 Hz refresh rate) has no relevance in most cases. As similarly noted in our Introduction section in relation to response time accuracy, this would only be important for the type of experiments, typically psychophysiological, that are not usually used in online research. Moreover, even without the RAF loop, extremely fast reactions without specialized hardware cannot be expected anyway, since the response times range between at least ca. 40-60 ms in

¹⁶ A brief script with the RAF loop method (exactly as used in the present study; along with some other useful related functions and usage instructions) is available via <https://github.com/gasparl/dtjs> (permanently stored via <https://doi.org/10.5281/zenodo.5912068>). This script can be inserted or loaded in any JS-based experiment, and the loop can be started or ended anytime with the simple commands `DT.loopOn()` and `DT.loopOff()`.

the first place (Figure 4). Finally, if the timing of the upcoming display change is not (or less) important, the RAF loop may be easily switched off when waiting for input, and switched back on only afterwards. In such cases, the double-RAF may be used instead for the timing (as long as the RAF loop is disabled), which might provide better precision in a few cases than (and otherwise is not inferior to) direct or single RAF methods (Figure 3).

Relatedly, our findings also highlight that researchers (and future studies) should take great care to distinguish between calling RAF a single time, and calling it multiple times, or when another RAF has already been called during the same frame: As seen in our results, this can robustly affect the obtained timestamps.

As for image display, we found that our synchronous method (drawing on a `canvas` element) was never superior to asynchronous methods (as it would have been expected based on the theoretical workflow of browsers), but in fact in most cases it was clearly inferior. Hence, anecdotally recommended practices for image display with the aim of minimizing layout changes are once again contradicted by empirical findings (as in Garaizar & Reips, 2019). This finding also highlights that there can be important differences between simply filling a `canvas` element with a different color (using `fillRect()`) and drawing an image element on a `canvas` element (using `drawImage()`): While, same as plain text change, filling gave practically faultless precision in most cases (with the RAF loop method; Study 1, Figure 3), drawing image elements introduced, contrariwise, substantial variance in most cases (Study 2, Figure 5).

Finally, we found (online Appendix C) that the large GAPED images consistently led to lower precision than the smaller BASS and (compressed) OASIS images. Importantly, this is confounded by file type and other image properties (e.g., height and width), but, even so, we have at least demonstrated that, specifically, large BMP images are less optimal for online response time research, while small JPG and PNG images allow very high precision (when using RAF loop timing and `visibility` or `opacity` changes). More generally, this finding provides a first important proof that different kinds of images can lead to very different levels of precision in online studies – future research could address this subject with a more comprehensive study using systematic combinations of image file sizes, file types, and image properties (height, width, transparency, etc.).

Conclusions

We have introduced a new timing method that, unlike previously recommended methods, captures the time of displaying stimuli in browsers with essentially faultless

precision in most practical cases. The precision remains high even in the case of presenting images (when using `visibility` or `opacity` changes – while we found that image display by drawing on a `canvas` element lowers precision and should therefore be avoided). The method is easy to use and the necessary computer code is available via our OSF and GitHub repositories. Finally, we have also shown that certain different image types lead to different levels of precision (most likely due to differences in file size), inviting more research in this topic.

Declarations

Funding

Gáspár Lukács has been funded as an International Research Fellow of Japan Society for the Promotion of Science. This funder had no role or involvement related to the submitted work. The rest of the authors did not receive support from any organization for the submitted work.

Conflicts of interest/Competing interests

The authors have no relevant financial or non-financial interests to disclose.

Ethics approval

Not applicable.

Consent to participate

Not applicable.

Consent for publication

Not applicable.

Availability of data and materials

All data and material are available via <https://osf.io/wyebt/>.

Code availability

All computer codes used in our studies (for testing as well as for analysis) are available via <https://osf.io/wyebt/> (and https://github.com/gasparl/disp_time).

Authors' Contributions

Concept, software, design, measurements, and analysis by GL; custom hardware and firmware (keypress simulation) by AG; manuscript by GL and AG.

Acknowledgments

We most of all thank Tristan Fraipont (Kaiido), who provided invaluable help with JavaScript (including the suggestion for the RAF loop method), but declined co-authorship. We also thank Kiki Kawai for lending a laptop and helping out with some of the measurements; Giorgia Silani and Emilio Chiappini for lending a macOS computer; and all

those who shared their JS timing codes openly (such as the maintainers of jsPsych, Labvanced, PsychoJS, and OSWeb).

References

- Anwyl-Irvine, A., Dalmaijer, E. S., Hodges, N., & Evershed, J. K. (2020). Realistic precision and accuracy of online experiment platforms, web browsers, and devices. *Behavior Research Methods*. <https://doi.org/10.3758/s13428-020-01501-5>
- Arachchige, C. N. P. G., & Prendergast, L. A. (2019). Confidence intervals for median absolute deviations. *ArXiv:1910.00229 [Math, Stat]*. <http://arxiv.org/abs/1910.00229>
- Barnhoorn, J. S., Haasnoot, E., Bocanegra, B. R., & van Steenbergen, H. (2015). QRTEngine: An easy solution for running online reaction time experiments using Qualtrics. *Behavior Research Methods*, 47(4), 918–929. <https://doi.org/10.3758/s13428-014-0530-7>
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1), 289–300. <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
- Birnbaum, M. H. (2004). Human Research and Data Collection via the Internet. *Annual Review of Psychology*, 55(1), 803–832. <https://doi.org/10.1146/annurev.psych.55.090902.141601>
- Bridges, D., Pitiot, A., MacAskill, M. R., & Peirce, J. W. (2020). The timing mega-study: Comparing a range of experiment generators, both lab-based and online. *PeerJ*, 8, e9414. <https://doi.org/10.7717/peerj.9414>
- Brown, M. B., & Forsythe, A. B. (1974). Robust tests for the equality of variances. *Journal of the American Statistical Association*, 69(346), 364–367. <https://doi.org/10.1080/01621459.1974.10482955>
- Conover, W. J., Johnson, M. E., & Johnson, M. M. (1981). A comparative study of tests for homogeneity of variances, with applications to the outer continental shelf bidding data. *Technometrics*, 23(4), 351–361. <https://doi.org/10.1080/00401706.1981.10487680>
- Crump, M. J. C., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon’s Mechanical Turk as a Tool for Experimental Behavioral Research. *PLoS ONE*, 8(3), e57410. <https://doi.org/10.1371/journal.pone.0057410>
- Dan-Glauser, E. S., & Scherer, K. R. (2011). The Geneva affective picture database (GAPED): A new 730-picture database focusing on valence and normative significance. *Behavior Research Methods*, 43(2), 468–477.

- <https://doi.org/10.3758/s13428-011-0064-1>
- De Leeuw, J. R., & Gilbert, R. A. (2019). Testing different methods of displaying stimuli in JavaScript. *GitHub Repository*. <https://github.com/vassar-cogscilab/js-display-durations>
- De Leeuw, J. R., & Motz, B. A. (2016). Psychophysics in a Web browser? Comparing response times collected with JavaScript and Psychophysics Toolbox in a visual search task. *Behavior Research Methods*, 48(1), 1–12. <https://doi.org/10.3758/s13428-015-0567-2>
- Finger, H., Goeke, C., Diekamp, D., Standvoß, K., & König, P. (2017). LabVanced: a unified JavaScript framework for online studies. International Conference on Computational Social Science (Cologne).
- Fligner, M. A., & Killeen, T. J. (1976). Distribution-Free Two-Sample Tests for Scale. *Journal of the American Statistical Association*, 71(353), 210–213. <https://doi.org/10.1080/01621459.1976.10481517>
- Fox, J., & Weisberg, S. (2019). *An R companion to applied regression* (Third edition). SAGE.
- Fraipont, T. [Kaiido]. (2019). *One problem is that officially, onload only tells us about the network status*. Stack Overflow. <https://stackoverflow.com/a/59300348>
- Fraipont, T. [Kaiido]. (2020). *What you are experiencing is a Chrome bug*. Stack Overflow. <https://stackoverflow.com/a/57549862>
- Fricker, R. D., Burke, K., Han, X., & Woodall, W. H. (2019). Assessing the statistical analyses used in Basic and Applied Social Psychology after their p-value ban. *The American Statistician*, 73(sup1), 374–384. <https://doi.org/10.1080/00031305.2018.1537892>
- Gao, Z., Chen, B., Sun, T., Chen, H., Wang, K., Xuan, P., & Liang, Z. (2020). Implementation of stimuli with millisecond timing accuracy in online experiments. *PLOS ONE*, 15(7), e0235249. <https://doi.org/10.1371/journal.pone.0235249>
- Garaizar, P., & Reips, U.-D. (2019). Best practices: Two Web-browser-based methods for stimulus presentation in behavioral experiments with high-resolution timing requirements. *Behavior Research Methods*, 51(3), 1441–1453. <https://doi.org/10.3758/s13428-018-1126-4>
- Garaizar, P., Vadillo, M. A., & López-de-Ipiña, D. (2014). Presentation Accuracy of the Web Revisited: Animation Methods in the HTML5 Era. *PLoS ONE*, 9(10), e109812. <https://doi.org/10.1371/journal.pone.0109812>
- Grootswagers, T. (2020). A primer on running human behavioural experiments online.

- Behavior Research Methods*, 52(6), 2283–2286. <https://doi.org/10.3758/s13428-020-01395-3>
- Kawai, C., Lukács, G., & Ansorge, U. (2021). A new type of pictorial database: The Bicolor Affective Silhouettes and Shapes (BASS). *Behavior Research Methods*. <https://doi.org/10.3758/s13428-021-01569-7>
- Krantz, J. H., & Reips, U.-D. (2017). The state of web-based research: A survey and call for inclusion in curricula. *Behavior Research Methods*, 49(5), 1621–1629. <https://doi.org/10.3758/s13428-017-0882-x>
- Kurdi, B., Lozano, S., & Banaji, M. R. (2017). Introducing the Open Affective Standardized Image Set (OASIS). *Behavior Research Methods*, 49(2), 457–470. <https://doi.org/10.3758/s13428-016-0715-3>
- Kuroki, D. (2021). A new jsPsych plugin for psychophysics, providing accurate display duration and stimulus onset asynchrony. *Behavior Research Methods*, 53(1), 301–310. <https://doi.org/10.3758/s13428-020-01445-w>
- Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence Testing for Psychological Research: A Tutorial. *Advances in Methods and Practices in Psychological Science*, 1(2), 259–269. <https://doi.org/10.1177/2515245918770963>
- Lord, F. M., & Novick, M. R. (2008). *Statistical theories of mental test scores*. Information Age Publishing Inc.
- Lukács, G. [gaspar]. (2018). *Exact time of display: RequestAnimationFrame usage and timeline*. Stack Overflow. <https://stackoverflow.com/questions/50895206/>
- Lukács, G. (2021). neatStats: An R package for a neat pipeline from raw data to reportable statistics in psychological science. *The Quantitative Methods for Psychology*, 17(1), 7–23. <https://doi.org/10.20982/tqmp.17.1.p007>
- Miller, R., Schmidt, K., Kirschbaum, C., & Engle, S. (2018). Comparability, stability, and reliability of internet-based mental chronometry in domestic and laboratory settings. *Behavior Research Methods*, 50(4), 1345–1358. <https://doi.org/10.3758/s13428-018-1036-5>
- Peer, E., Rothschild, D., Gordon, A., Evernden, Z., & Damer, E. (2021). Data quality of platforms and panels for online behavioral research. *Behavior Research Methods*. <https://doi.org/10.3758/s13428-021-01694-3>
- Peirce, J. W. (2008). Generating stimuli for neuroscience using PsychoPy. *Frontiers in Neuroinformatics*, 2(10). <https://doi.org/10.3389/neuro.11.010.2008>
- Pinet, S., Zielinski, C., Mathôt, S., Dufau, S., Alario, F.-X., & Longcamp, M. (2017).

- Measuring sequences of keystrokes with jsPsych: Reliability of response times and interkeystroke intervals. *Behavior Research Methods*, 49(3), 1163–1176.
<https://doi.org/10.3758/s13428-016-0776-3>
- Plant, R. R., & Turner, G. (2009). Millisecond precision psychological research in a world of commodity computers: New hardware, new problems? *Behavior Research Methods*, 41(3), 598–614. <https://doi.org/10.3758/BRM.41.3.598>
- Pronk, T., Wiers, R. W., Molenkamp, B., & Murre, J. (2020). Mental chronometry in the pocket? Timing accuracy of web applications on touchscreen and keyboard devices. *Behavior Research Methods*, 52(3), 1371–1382. <https://doi.org/10.3758/s13428-019-01321-2>
- R Core Team. (2020). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
- Ratcliff, R., & Hendrickson, A. T. (2021). Do data from mechanical Turk subjects replicate accuracy, response time, and diffusion modeling results? *Behavior Research Methods*. <https://doi.org/10.3758/s13428-021-01573-x>
- Reimers, S., & Stewart, N. (2015). Presentation and response timing accuracy in Adobe Flash and HTML5/JavaScript Web experiments. *Behavior Research Methods*, 47(2), 309–327. <https://doi.org/10.3758/s13428-014-0471-1>
- Signorell, A. C. (2021). *DescTools: Tools for descriptive statistics*. <https://CRAN.R-project.org/package=DescTools>
- Simon, R. (1986). Confidence Intervals for Reporting Results of Clinical Trials. *Annals of Internal Medicine*, 105(3), 429. <https://doi.org/10.7326/0003-4819-105-3-429>
- Wickham, H. (2016). *ggplot2: Elegant graphics for data analysis* (Second edition). Springer.
- Wolfe, C. R. (2017). Twenty years of Internet-based research at SCiP: A discussion of surviving concepts and new methodologies. *Behavior Research Methods*, 49(5), 1615–1620. <https://doi.org/10.3758/s13428-017-0858-x>
- Yung, A., Cardoso-Leite, P., Dale, G., Bavelier, D., & Green, C. S. (2015). Methods to Test Visual Attention Online. *Journal of Visualized Experiments*, 96, 52470.
<https://doi.org/10.3791/52470>