

Frustration: Still a Common User Experience

Hertzum, Morten; Hornbæk, Kasper

Published in:
ACM Transactions on Computer-Human Interaction

DOI:
[10.1145/3582432](https://doi.org/10.1145/3582432)

Publication date:
2023

Document Version
Peer reviewed version

Citation for published version (APA):
Hertzum, M., & Hornbæk, K. (2023). Frustration: Still a Common User Experience. *ACM Transactions on Computer-Human Interaction*, 30(3), Article 42. <https://doi.org/10.1145/3582432>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain.
- You may freely distribute the URL identifying the publication in the public portal.

Take down policy

If you believe that this document breaches copyright please contact rucforsk@kb.dk providing details, and we will remove access to the work immediately and investigate your claim.

Frustration: Still a Common User Experience

Morten Hertzum

Roskilde University, Denmark

Kasper Hornbæk

University of Copenhagen, Denmark

Abstract

When computers unexpectedly delay or thwart goal attainment, frustration ensues. The central studies of the extent, content, and impact of such frustration were done more than 15 years ago. We revisit this issue after computers have become more mature and computer use more extensive. To this end, we had 234 crowdsourced participants log the frustrating episodes they experienced with their computer during one hour of computer use. The average time lost due to frustrating episodes was between 11% and 20% of the one-hour period. Though this is less time lost than in the earlier studies, frustration remains a common user experience. While shorter, the median level of frustration during the episodes was high (7 on a 9-point scale). The frustration level correlated with task importance and time lost but was unaffected by computer experience and largely unaffected by computer self-efficacy. In addition, participants indicated that 84% of the episodes had happened before, that 87% could happen again, and that they were unable to resolve 26% of the episodes. This high rate of recurrence and lack of control likely added to the frustration level. The episodes spanned various issues pertaining to performance (49%), usability (36%), and utility (16%).

CCS concepts

Human-centered computing → User studies; HCI theory, concepts and models; Empirical studies in HCI

Additional keywords and phrases

frustration, time lost, computer self-efficacy, user experience

1 Introduction

Emotion is key to the user experience [12,26,46]. While positive emotions in relation to computer use have received considerable attention [e.g., 11], this study is about negative emotion. We have all experienced episodes where we cannot use computers the way we want. Or when something that we usually do with an interactive system is not working anymore. Or when we need to fix the same problem in a user interface over and over. Such episodes are frustrating because computers were supposed to help us attain our goals, but in these cases do not.

Frustration, understood in this way, has been discussed since the beginning of the field of HCI. Golden [23] pointed out that “it is perfectly possible to have a program which is structured, modular, readable, flexible, self-documenting, maintainable, which performs its specified function, and which is a source of constant frustration and irritation to its users”. Thus, even if computers are useful, they

might be (and apparently were) frustrating. This was also pointed out in early definitions of usability, for instance by Shackel [48]. His definition mentioned frustration as one type of negative attitude toward computers. Finally, early views of usability problems also suggested that problem severity might be rated according to the frustration it causes the user [20]. Thus, the second-severest type of usability problem—after being unable to do a task—is when the user faces significant frustration [20]. This interest in frustration has continued over the past 40 years [e.g., 3,6,24,46].

Notwithstanding the continued interest, empirical data on computer frustration are limited. The central studies of frustration were done more than fifteen years ago [9,10,15,36,37]. Their findings were striking. In particular, they found that 44 to 50 percent of the time spent on computers was lost to frustrations [15,37]. This alarming figure provides a sobering contrast to the many statements that computers improve productivity and yield stimulating user experiences. Moreover, the frustrations were often recurring and not one-time occurrences. Are computers really that substantial a stumbling block to attaining our goals?

Answering this question is difficult. First, we are aware of only one replication of the above results [27]. This replication had fewer participants but found only half the rate of frustration among users. Second, much has changed over the past fifteen years regarding technology and its use. Perhaps computer technologies have matured. User-centered design may have made computers easier to use. One could hope so, but it remains unsettled. Computer use has certainly expanded, for example online activity has skyrocketed. This, too, may influence the level of frustration. Third, the method used in the original studies included the requirement that the participants had to report at least three frustrations [15]. Skeptics might argue that this inflates the frequency of frustrations. In addition, the participants in one of the original studies were computer-science students, who may be more alert to computer shortcomings but also better equipped to resolve them.

The present study therefore revisits the question of the extent, content, and impact of frustrations with computers. We use a web-based logging tool to collect frustrations from each of 234 users during one hour of computer use. We analyze those frustrations quantitatively and qualitatively to show that substantial amounts of time are lost to a variety of frustrating incidents, which are often recurring rather than resolvable. These results suggest that frustrations with computers are still common. We therefore end the paper with discussing what to do to alleviate the frustrations experienced by computer users.

2 Background

Fifteen years ago, a series of five papers aimed to conceptualize and quantify users' frustrations with computers [9,10,15,36,37]. We will refer to these five papers as the computer user frustration (CUF) studies.

2.1 Conceptualizing user frustration

Bessière et al. [10] stated that “Frustration is almost universally accepted as the emotional outcome of a negative computing experience”. Three properties of this statement deserve note. First, frustration is an emotional outcome. Treating frustration as an emotional outcome is consistent with contemporary theories, such as Berkowitz’s [8] reformulation of the frustration-aggression hypothesis. While the original frustration-aggression hypothesis [18] defined frustration as an event (that caused aggression), the reformulated hypothesis shifts the focus from the event to the negative affect caused by the event. This negative affect (frustration) may, in turn, cause further responses, including increased effort or learned helplessness [19]. Second, negative experiences with computers cause frustration. While many kinds of events may cause frustration, the CUF studies are specifically

concerned with frustration caused by computer use. In general, events are held to be frustrating if they unexpectedly thwart goal attainment, that is, if they interrupt task completion [4,55]. In relation to computers, the events that thwart goal attainment are broadly known as usability problems. That is, frustration with computers is tightly associated with experiencing usability problems. Third, the association between the negative computing experience (task interruption) and the emotional response (frustration) is mediated by numerous factors. Bessièrè et al. [9] distinguished between incident factors and individual factors. In Bessièrè et al. [10] these factors were relabeled to situational mediators and dispositional mediators, see Figure 1. These mediators co-determine the frustration that users experience with computers.




Figure 1. The computer frustration model, based on Bessièrè et al. [10].

An important feature of the computer frustration model in Figure 1 is its distinction between immediate frustration and the escalation of frustration to a broader range of emotional outcomes. Immediate frustration is a pre-emotional appraisal. It is caused by an impediment to attain a goal and serves to “redirect limited attentional resources away from the central task or goal at hand to peripheral features of the information environment that may now have become obstacles” [10]. That is, immediate frustration is an adaptive response that aims to ensure task completion by, temporarily, shifting attention to an obstacle that must be dealt with to be able to progress on the task. In contrast, the escalation of frustration to a broader range of emotional outcomes is often described as maladaptive. This is for example apparent in the many studies that link frustration to aggression, anger, fear, helplessness, rage and so forth [13,55]. These emotional outcomes are maladaptive because they imply that “getting more and more frustrated can make the problem solving situation more difficult, rather than less so” [10]. However, thwarted goal attainment tends to produce an emotional response, even if maladaptive.

To assess how the mediators influence frustration, Bessière et al. [10] conducted a study similar to that used for quantifying user frustration, see Section 2.3 for a description of the method. They assessed the effect of the mediators at four different points in time: when the frustrating experience occurred, aggregated across the computer session in which the frustrating experience occurred, after the session, and in terms of the effect of the frustrating experience on the participant’s day. Hierarchical regression showed that demographics were not a significant predictor of the level of frustration at any of the four points in time, while computer experience, situational mediators, mood, and computer attitudes/self-efficacy significantly predicted the level of frustration at one or more points in time, see Table 1. Incident-level frustration was mainly predicted by situational mediators, session-level frustration mainly by computer attitudes/self-efficacy, post-session mood mainly by computer attitudes/self-efficacy and mood, and effect on day mainly by situational mediators. Other studies show that frustration with computers is also mediated by personality traits such as neuroticism and conscientiousness [21,24].

Table 1. Determinants of the level of frustration at four different points in time, $N = 144$ participants.

Factor	R^2 change			
	Incident level frustration	Session level frustration	Post-session mood	Effect on day
Demographics (age, sex, education)	.001	.03	.06	.04
Computer experience (hours per week, years of use)	.01	.02	.02	.04*
Situational mediators (task importance, time to fix, time lost, occurrence)	.16*	.07	.004	.18*
Mood (satisfaction with life, upset often, pre-mood)	.02*	.03	.15*	.06*
Computer attitudes/self-efficacy (computer anxiety, comfort, ability to fix...)	.05*	.14*	.12*	.05

* Significant R^2 change, $p < .05$ (hierarchical regression)

2.2 Quantifying user frustration

The most troubling finding from the CUF studies was their quantification of the amount of time lost to frustration. Ceaparu et al. [15] found that their 111 participants lost an average of 50% of the time they spent using computers to frustrating episodes. The time lost varied substantially across individual participants but only moderately across participant groups. For example, the data collected from University of Maryland students showed roughly the same amount of time lost (48%) as the data collected from Towson University students (53%). And the data collected by students who self-reported their frustrating experiences showed the same amount of time lost as the data collected by students who observed others’ computer use, see Table 2. Lazar et al. [37] replicated the study with 50 workplace professionals as participants, rather than students. The professionals self-reported that they lost on average 44% of the worktime they spent using computers to frustrating episodes. In both studies, the time lost to frustration was calculated by summing the time the participants reported that they spent solving the problem and the time they reported spending to recover from any work lost. The participants in Ceaparu et al. [15] spent more time solving the problems than the professionals

in Lazar et al. [37]; the two studies were similar with respect to the time spent recovering from lost work, see Table 2.

Table 2. Time lost to frustrating experiences.

Study	Group	Participants	Frustrating episodes	Percent of time lost to frustration		
				Solving problem	Recovering from any work lost	Total
Ceaparu et al. [15]	Self-report	59 students	199	28	22	50
Ceaparu et al. [15]	Observation	52 ^a	174	29	21	50
Lazar et al. [37]	Self-report	50 professionals	149	22	22	44
Hertzum [27] ^b	Self-report	21 students	50	16	11	27

^a It is not reported whether the 52 observed computer users were studying, working or engaged in leisure activities. ^b A small-scale replication of the CUF studies.

The median frequency of the frustrating episodes was that they happened several times a month, see Table 3. That is, most of the frustrating episodes had happened before. Only 26% [15] and 20% [37] of the frustrating episodes happened for the first time, and as much as 16% [15] and 17% [37] of them recurred as frequently as several times a day. In addition to being frequent, the frustrating episodes were also associated with high levels of frustration. The median level of frustration was 7 and 8 on a nine-point rating scale with 9 indicating *Very frustrating*, see Table 3. In both studies at least 25% of the frustrating episodes received the top rating of very frustrating, while the lower half of the scale (ratings 1 through 5) accounted for less than 25% of the frustrating episodes.

Table 3. Problem frequency and level of frustration.

Study	Participants	Frustrating episodes	Median problem frequency ^a	Median level of frustration ^b
Ceaparu et al. [15]	111	373	5: several times a month	7
Lazar et al. [37]	50	149	5: several times a month	8
Hertzum [27] ^c	21	50	4: once a week	7

^a On an eight-point rating scale with the categories: *More than once a day* (1), *One time a day* (2), *Several times a week* (3), *Once a week* (4), *Several times a month* (5), *Once a month* (6), *Several times a year* (7), and *First time it happened* (8). ^b On a nine-point rating scale with the end points: *Not very frustrating* (1) and *Very frustrating* (9). ^c A small-scale replication of the CUF studies.

Finally, the CUF studies documented that the frustrating episodes were caused by a range of applications. Ceaparu et al. [15] and Lazar et al. [37] found that the three applications causing the largest numbers of frustrating episodes were web browsing, email, and text processing. However, the participants were also frustrated by their chat and instant messaging applications, database programs, file browsers, graphic design programs, hardware, operating systems, presentation software, programming tools, spreadsheet programs, video/audio software, and “other internet use” [15]. No one application was the cause of a large fraction of the frustrating episodes.

2.3 Method for investigating user frustration

The above results were obtained by means of “modified time diaries” [15,37] in which the participants were to record each frustrating episode as it occurred during their computer use. This concurrent

recording was considered “an improvement over retrospective survey questions, because estimates from memory often lead to inflated or incorrect answers” [15]. Both studies [15,37] applied essentially the same method, detailed below.

The 59 self-report participants in Ceaparu et al. [15] were 33 computer science undergraduate student from University of Maryland and 26 computer information systems undergraduate students from Towson University. After these participants had completed their self-report session, they each recruited and observed one of the 52 participants in the observation session. The observed participants were an average of 26.1 years old ($SD = 13.1$), suggesting that many of them were students but also that some of them were probably not. The 50 participants in Lazar et al. [37] were recruited “through listservs and word of mouth in the Baltimore/Washington professional community”. As a requirement for participation, these professionals had to be college graduates and use computers for part of their work.

The participants were asked to “spend at least an hour using a computer and report their frustrating experiences” [15]. To make the data more representative of the participants’ usual tasks, the participants were “simply asked to carry on with their normal tasks” [15]; no specific tasks were assigned or expected. The participants – self-report as well as observed – in Ceaparu et al. [15] were however asked to “report at least three frustrating experiences that took place when they were performing their common tasks”. The participants in Lazar et al. [37] received similar instructions, except that they were apparently not required to report at least three frustrating experiences.

Each frustrating experience was documented by filling out a frustrating-experience report. The reporting form [15] stated that frustrating experiences included both major problems and minor issues, exemplified by “computer or application crashes” and “a program not responding the way that you need it to”, respectively. The reporting form further stated that “Anything that frustrates you should be recorded”. This statement left it entirely to the participants to define what constituted a frustrating experience and, thus, bypassed the authors’ conceptualization of user frustration (see Section 2.1). The study procedure was facilitated by a web site at which all study data were entered. For the observed participants, the observer filled out the frustrating-experience reports. To support the observers in determining when the observed participants experienced frustration, the observers were instructed to ask the observed participants to think aloud and to prompt them if it was unclear whether they experienced frustration.

3 Method

The present study aimed to investigate the extent, content, and impact of frustrations during computer use. We collected data on such frustrations with a self-report tool deployed through a crowdsourcing website. Subsequently, we analyzed that data with qualitative and quantitative methods.

3.1 Study design

Our study design resembles that in Ceaparu et al. [15]. Compared to methods such as critical incidents, this design allows for estimating the percentage of time that participants experience frustration. However, in contrast to the CUF studies, we recruited participants through crowdsourcing. The rationale for recruiting participants in this way was to sample frustrations (a) from a large group of participants, (b) during their normal computer use, and thereby (c) across a variety of hardware, software, tasks, and use contexts.

We relied on participants to log a frustration episode whenever they thought frustration occurred, so-called event-based monitoring [50]. This approach allowed participants to use their computer relatively uninterrupted during logging compared to, for instance, sampling at random times for

frustrations. However, earlier work suggests that participants in crowdsourced studies do not pay as much attention as people recruited in traditional ways (e.g., students recruited in a classroom [5]). Therefore, if participants had not reported anything for 10 minutes, we reminded them to either log or dismiss the reminder.

Participants were required to log frustrations for one hour. Contrary to Ceaparu et al. [15], we did not ask participants to report a particular number of frustrating episodes; we suspected that such a requirement might inflate the number of frustration episodes. An hour of logging seemed feasible for a crowdsourced study and corresponded to the logging period in the CUF studies.

Participants were free to choose *when* they would complete the one-hour logging. The intention was to ensure that the logging of frustrations would not happen immediately after participants had signed up for the study. If it did, the logging might happen when the participants would normally do crowdwork (e.g., during breaks or in the evening) or when the participants were simultaneously doing other crowdwork (which might not be representative of computer work in general). We therefore asked participants to choose a time of participation “where you will be using your computer as you would on any normal day. What you are using your computer for does not matter; it may be work or leisure, gaming or taxes, browsing or writing”. Thus, in contrast to Ceaparu et al. [15], we did not require participants to do the study when they “would be using a computer for their personal use as opposed to tasks assigned to them”. We simply aimed to get participants at a time that was typical of their work with computers.

We were broadly interested in any episode of computer interaction that participants experienced as frustrating. Therefore, the definition of frustration was left to participants, like in the CUF studies. Besides reporting on specific frustrations, participants described what they generally considered frustrating about computers. These descriptions will be used to understand differences in the participants’ general concept of computer frustration. We did not define computer beyond “the computer you normally use”.

3.2 Participants

A total of 306 participants signed up to participate in the study at the crowdwork website Prolific (<https://www.prolific.co/>). Participants were compensated GBP 5. This compensation matched the time spent signing up, answering questions, and so forth (based on an estimation that these took altogether 30 minutes), but not the time spent during the primary task.

3.3 Data collected

The central data were collected during the one hour of logging. In addition, participants answered questions before and after the logging.

3.3.1 Prelogging questions

Before beginning the logging, participants answered several questions on demographics, including age, gender, and occupation. Following earlier work on frustrations, these questions also concerned the traits that might influence the frequency and impact of frustrations. We asked participants about how many hours a day they used computers (with five response options: 0-2, 3-5, 6-8, 9-11, 12 or more). We also asked about perceived computer experience indicated by two questions, both taken from Ceaparu et al. [15]:

- “Computers make me feel”, answered between “very uncomfortable” (1) and “very comfortable” (9).

- “How experienced do you think you are when it comes to using a computer”, answered between “very inexperienced” (1) and “very experienced” (9).

To gauge computer self-efficacy, we asked four questions:

- “When you run into a problem on the computer or an application you are using, do you feel”, answered between “anxious” (1) and “relaxed/indifferent” (9).
- “When you encounter a problem on the computer or an application you are using, how do you feel about your ability to fix it?”, answered from “helpless” (1) to “confident that I can fix it” (9).
- “When there is a problem with a computer that I can't immediately solve, I would stick with it until I have the answer”, answered from “strongly disagree” (1) to “strongly agree” (9).
- “If a problem is left unresolved on a computer, I would continue to think about it afterward”, answered from “strongly disagree” (1) to “strongly agree” (9).

These four questions were taken from earlier work on computer self-efficacy and used by Ceaparu et al. [15]. Before beginning the logging, we also asked participants to describe what a frustrating episode was to them (“Please describe what you consider a frustrating episode with computers. Please also give a few examples of such episodes that have occurred to you in the past”).

3.3.2 Logging of frustrations

During the hour of logging, participants logged any number of frustration reports. For each frustration report, we collected the following:

- “What were you trying to do?”, answered with at least 50 characters. This question was taken from Ceaparu et al. [15].
- “Why were you trying to do it?”, answered with at least 50 characters. The intent of this question was to allow participants to explain the context of the frustration (i.e., the goal interruption).
- “How important was it for you to get it done?”, answered from “not at all important” (1) to “extremely important” (9). This question was taken from Ceaparu et al. [15].
- “Which program, app or webpage did you use?”
- “Please describe your frustration”, answered with at least 50 characters.
- “How did you resolve the frustrating episode, or were you not able to resolve it”, answered with at least 50 characters.
- “How much time did you lose”, answered in an hour:minute:second format.
- “How frustrating was this episode for you”, answered from “not very frustrating” (1) to “very frustrating” (9). This question was taken from Ceaparu et al. [15].

After the one hour of logging, a popup made the participants aware that the period of logging had ended. If needed, they could fill a final frustration report, and then they were taken to the postlogging questions.

3.3.3 Postlogging questions

At the end of the study, participants would review what they entered for each frustration report, edit their responses, and answer a few further questions. This was deemed feasible because participants

in the study by Ceaparu et al. [15] on average entered just 3.3 episodes. The participants were asked five further questions about each frustrating episode:

- “Will this episode, or how you felt during it, affect the way you subsequently use your computer?” (yes or no). The intention of this question was to understand the consequences of the frustration; it was adapted from Tuch and Hornbæk [54].
- “If yes, how? Can you give an example of how you used your computer differently?” This question was to gauge the broader impact of the frustration.
- “Did the frustration happen before?” The intention behind this question was to understand if frustrations were occurring regularly or merely isolated, non-recurring episodes.
- “Could the episode you described happen again?” (yes or no). The intention behind this question was to understand better the downstream effects of frustrations; it was adapted from Tuch and Hornbæk [54].
- “If yes, why? If no, why not?” This question allowed participants to expand on the above.
- “Add any other information you wish to share about this frustration”.

The intention of the above questions was to learn whether frustrations had effects beyond the individual episode. To avoid participant drop-out during this last part of the study, the post-logging questions were not mandatory.

3.4 Procedure

Participants signed up for the study at Prolific and immediately completed the prelogging questions. They participated in the study between April and September 2021. In the announcement of the study, we did not explicitly mention the focus on frustrations to avoid biasing the pool of participants. Participants then registered their email address and gave some timeslots within which they would be able to do the one-hour logging. At the beginning of the period indicated by participants, they received an email with a link to the self-report tool. The link opened a webpage that they were asked to keep open during the one-hour logging session; it could be minimized or kept in a separate tab in their browser. If they accidentally closed it, the session could be restarted by following the link. The self-report tool had two parts: a logging part and a postlogging part.

The logging part allowed participants to register any frustration (or lack thereof) they wanted, at any time. To register a frustration, participants would click the self-report tool and were then brought to a webpage that allowed them to answer the questions described in Section 3.3.2. If 10 minutes passed without any registration, participants were reminded that they needed to log their frustrations. The reminder consisted of a pop-up notification from the self-logging tool and an audio cue. Clicking the notification brought the self-report tool into focus. In the self-report tool, participants had the option to report a frustration (as just described) or to report that no frustration had occurred. Participants were asked to allow pop-up notifications to be shown and to keep the sound on during the logging session to ensure that they would see and hear notifications.

The postlogging part of the self-report tool allowed the participants to look over all their frustration reports; they were also requested to fill out the five postlogging questions for each frustration. If the participant had not reported any frustrations, they were prompted for an explanation.

3.5 Data analysis

As a precursor to the data analysis, we performed three quality checks of the answers to identify and exclude those that were clearly pasted, incomprehensible, or otherwise incoherent. The first quality

check consisted of looking through the data from each of the 306 participants to flag incomprehensible or missing answers and people who had misunderstood the study or experienced trouble with the self-report tool. This check led to the exclusion of 72 participants, leaving 234 (76%) for analysis. The second quality check consisted of looking through the 207 frustrating incidents reported by the included participants to flag the reportings that did not qualify as frustrations. This check led to the exclusion of 22 reportings, leaving 185 (89%) for analysis. The most common reason for exclusion was participants who responded to the reminder to log their frustrations with a reporting indicating that they had not experienced a frustrating incident during the past ten minutes. The third quality check consisted of looking through the 185 frustrating incidents to flag those reporting more than an hour of lost time. Within the one-hour reporting period, it was not possible to lose more than an hour so the reportings in excess of an hour were either wrong or they included time lost beyond the reporting period. Because some participants reported that on their browser and system configuration they had trouble setting the hours of the time-lost indicator to zero, we also considered whether the reported time lost would make more sense if the hours were reset to zero (but the minutes and seconds kept unchanged). We made this interpretation on the basis of the textual description of the frustrating incident. This check led to including 163 (88%) of the time-lost reportings – 130 were included unchanged and another 33 after resetting the hours to zero.

All three quality checks were performed independently by both authors, then disagreements were discussed until a consensus was reached. The agreement between the two authors' coding was assessed with Cohen's [16] Kappa. The Kappa values were .84, .67, and .84 for the first, second, and third quality check, respectively. That is, all three Kappa values were well above the threshold of .60 recommended by Lazar et al. [35] as indicating satisfactory reliability.

We acknowledge that resetting the hours to zero for some time-lost reportings is an interpretation and that introducing this interpretation possibly leads to underestimating the time lost. Therefore, we estimated the time lost in two different ways in the data analysis:

- *Time lost adjusted downward*: This estimate of the time lost consisted of resetting the hours to zero for 33 time-lost reportings in excess of an hour (as described above). It possibly underestimates the time lost and should, thus, be seen as a lower bound.
- *Time lost adjusted to maximum*: This estimate consisted of changing the time lost to one hour (i.e., the full reporting period) for the 33 time-lost reportings in excess of an hour. It possibly overestimates the time lost and should, thus, be seen as an upper bound.

In addition to ratings (e.g., of the level of frustration experienced), the participants also provided free-text responses. We used open card sorting [52] to organize these qualitative data into groups. That is, the responses were printed on paper cards and then the authors collaboratively read one card at a time, discussed it, and placed it next to the cards with similar content. Initially, the relations among neighboring cards were merely tentative but gradually groups started to emerge. When a group had materialized, we gave it a brief descriptive label. After all cards had been sorted, we walked through the unlabeled groups to arrive at a descriptive label. In this final labeling of the groups, a small number of cards were reassigned to other groups and some small groups were collapsed into more generic ones.

We used card sorting to analyze the responses to four questions. First, we sorted the 185 descriptions of frustrating incidents. This analysis resulted in 25 groups, each representing a category of frustrations experienced during the reporting period. Second, we sorted the 185 descriptions of how the participants resolved the incidents. This analysis resulted in six groups. Third, we sorted the 234 descriptions of what the participants in general considered a frustrating episode with computers.

During this analysis, many descriptions were split into multiple cards because the description covered several dimensions of frustration, for a total of 419 cards. These cards were sorted into 39 groups, which collectively described the dimensions in the participants' concept of frustration with computers. Fourth, we sorted the 234 descriptions of the participants' occupation. This analysis resulted in five occupation groups.

4 Results

In the following, we first describe the participant profile. Then, we analyze the quantitative data about time lost, frustration level, and frustration recurrence. Finally, we analyze the qualitative data about how the participants resolved the frustrations, what they were concretely frustrated about, and how they in general construed frustration.

4.1 Participant profile

Table 4 gives the profile of the 234 participants. They were an average of 28.6 years old ($SD = 9.0$). All participants were 18 years or older. Male participants were overrepresented (which is common in crowdsourced studies [42,43]). The participants' median daily use of computers was 6-8 hours. Our card sorting of the participants' occupation showed that the largest group was professionals, that only a modest number of participants were IT professionals, and that about one third of the participants were students.

The participants participated in the study on computers running Windows (83%) or Mac (13%); their web browsers were typically Chrome (75%) or Firefox (18%). The most frequent screen resolution was 1920×1080 (32%), 1366×768 (21%), or 1536×864 (15%). Thus, as intended, participants seemed to construe "the computer you normally use" to be a stationary or laptop computer.







Regarding the participants' rating of their computer experience and computer self-efficacy, Table 5 shows that they tended to be experienced and self-efficacious. For all but one question, the answers clustered in the upper half of the response scale. The only exception was the question about how the participants felt when they ran into a computer problem. For this question, the answers clustered around the middle of the response scale. That is, the ratings in Table 5 aligned well with the participants' substantial daily use of computers.

Table 4. Participant profile, $N = 234$ participants.

Classification	Categories	N	%
Age	10-19 years	18	8
	20-29 years	130	56
	30-39 years	56	24
	40-49 years	20	9
	50-59 years	9	4
	60-69 years	1	0
Gender	Female	77	33
	Male	155	66
	Other/Prefer not to say	2	1
Daily use of computers	0-2 hours	9	4
	3-5 hours	40	17

Occupation	6-8 hours	84	36
	9-11 hours	59	25
	12 hours or more	42	18
Occupation	Student	82	35
	Professional	100	43
	IT professional	25	11
	Unemployed	20	9
	Other	7	3

Table 5. Participants' computer experience and computer self-efficacy, $N = 234$ participants.

Question	Median	Distribution ^a
<i>Computer experience</i>		
Computers make me feel (1: very uncomfortable – 9: very comfortable)	8	
How experienced do you think you are when it comes to using a computer? (1: very inexperienced – 9: very experienced)	7	
<i>Computer self-efficacy</i>		
When you run into a problem on the computer you are using, do you feel? (1: anxious – 9: relaxed/indifferent)	5	
When you encounter a problem on the computer you are using, how do you feel about your ability to fix it? (1: helpless – 9: confident that I can fix it)	7	
When there is a problem with a computer that I can't immediately solve, I would stick with it until I have the answer (1: strongly disagree – 9: strongly agree)	7	
If a problem is left unresolved on a computer, I would continue to think about it afterward (1: strongly disagree – 9: strongly agree)	7	

^a Response distribution from 1 (left) to 9 (right), median in red.

4.2 Time lost

The participants reported an average of 0.79 ($SD = 0.92$) frustrating incidents during the one-hour reporting period. For 86 (37%) participants, the reporting period passed without frustrating incidents. The remaining participants reported one frustrating incident (131 participants, 56%), two frustrating incidents (9 participants, 4%), or three to seven frustrating incidents (8 participants, 3%). With increasing age, participants reported slightly more frustrating incidents, $r(234) = .13$, $p = .048$.

The time lost due to a single frustrating incident varied from 1 second to 1 hour. This time included the time spent trying to solve the problem and the time spent recovering from any lost work. Across the one-hour reporting period, the participants lost an average of between 6.63 and 12.15 minutes to frustration, corresponding to 11-20% of the reporting period, see Table 6. The lower bound resulted from resetting the hours to zero in 33 time-lost reportings that were in excess of an hour; the upper bound resulted from changing these 33 time-lost reportings to the full one-hour reporting period.

Gender affected downward-adjusted time lost, $t(230) = 2.32$, $p = .021$, but not adjusted-to-maximum time lost, $t(230) = 1.79$, $p = .075$. For downward-adjusted time lost, female participants ($M = 4.05$, $SD = 9.67$) lost less time than male participants ($M = 7.90$, $SD = 15.35$). Furthermore, the number of frustrating incidents reported by a participant correlated with both downward-adjusted time lost, $r(234) = .50$, $p < .001$, and adjusted-to-maximum time lost, $r(234) = .37$, $p < .001$. In contrast,

participant age did not correlate with either downward-adjusted time lost, $r(234) = .11, p = .101$, or adjusted-to-maximum time lost, $r(234) = .10, p = .131$. Similarly, the participants' daily use of computers and all the questions in Table 5 about their computer experience and computer self-efficacy did not correlate with either downward-adjusted time lost or adjusted-to-maximum time lost (r_s , all $ps > .228$).

Table 6. Time lost to frustrating incidents (in decimal minutes and as a percentage of the one-hour reporting period), $N = 234$ participants.

Category	Mean	Std. deviation	%
Time lost (adjusted downward for 33 incidents)	6.63	13.78	11
Time lost (adjusted to maximum for 33 incidents)	12.15	22.20	20

4.3 Task importance and level of frustration

The frustrating incidents happened during tasks with a median importance of seven and they caused the participants frustration at a median level of seven, see Table 7. For both task importance and frustration level, the participants' ratings clustered in the upper half of the response scale. With increasing task importance, the participants experienced a higher level of frustration, $r_s(185) = .59, p < .001$.

Task importance did not correlate with downward-adjusted time lost, $r_s(163) = .12, p = .121$, but correlated weakly with adjusted-to-maximum time lost, $r_s(163) = .17, p = .028$. That is, the data provided mixed support for a hypothesis about spending more time trying to solve and recover from important tasks. Frustration level correlated weakly with both downward-adjusted time lost, $r_s(163) = .26, p < .001$, and adjusted-to-maximum time lost, $r_s(163) = .25, p = .002$. That is, the data supported a hypothesis that longer-lasting incidents caused more frustration. In addition, frustration level correlated with only one of the questions in Table 5 about computer experience and computer self-efficacy: Participants who reported feeling more relaxed/indifferent when they ran into a problem on their computer experienced a slightly lower level of frustration during the incidents, $r_s(185) = -.15, p = .041$. That is, frustration level was unaffected by computer experience and largely unaffected by computer self-efficacy. There was no effect of gender on frustration level, $U(183) = 3426.00, p = .428$.

Table 7. Task importance and frustration level, $N = 185$ frustrating incidents.

Question	Median	Distribution ^a
How important was it for you to get it done? (1: not at all important – 9: extremely important)	7	
How frustrating was this episode for you? (1: not very frustrating – 9: very frustrating)	7	

^a Response distribution from 1 (left) to 9 (right), median in red.

Further support for the high level of frustration caused by the incidents was provided by the high percentage of incidents that had happened before or could happen again, see Table 8. Experiencing the same frustrating incident multiple times without knowing how to avoid it likely contributed to

drive the frustration level upward. It should, however, be noted that the three questions in Table 8 were optional and only answered for 107 (58%) of the 185 frustrating incidents.

Table 8. Consequence and recurrence of frustrating incidents, $N = 107$ frustrating incidents.







Question	Incidents for which participant answered ‘yes’	
	<i>N</i>	%
Will this episode, or how you felt during it, affect the way you subsequently use your computer? (yes/no)	23	21
Did the frustration happen before? (yes/no)	90	84
Could the episode you described happen again? (yes/no)	93	87

4.4 How did participants resolve the frustrations?

The participants were able to resolve a total of 74% of the 185 frustrating incidents, see Table 9. The largest category of resolved incidents consisted of those that were sorted out without the participants mentioning a specific strategy, except putting in extra effort. The other categories of resolved incidents involved waiting for the issue to resolve itself (e.g., “I patiently waited an extra 30 secs until the webpage loaded successfully”), repeating previous steps (e.g., “I restart the application and it helps. The error is gone”), taking another approach (e.g., “i had to use the mouse or trackpad instead to manually navigate between them”), and settling for a lower-quality outcome (e.g., “I’ve accepted that it wont look as good as I wanted it to”). Finally, the participants were unable to resolve 26% of the incidents. These incidents included deciding not to care (e.g., “It’s not important enough to me to care to resolve it - I just gave up.”), postponing further work on the problem to a later point in time (e.g., “I cannot resolve it at the moment. I will have to come back to it when I have time.”), and switching to a completely different activity to shake off the frustration (e.g., “I just shut down everything, going to sleep makes me feel better and more comfortable.”).

The median level of frustration tended to be lower for resolved than unresolved frustrations, thereby suggesting that being able to resolve an incident tempered the frustration somewhat. This suggestion was supported by the resolved incidents with a lower-quality outcome having a median frustration level in between the resolved and unresolved incidents (Table 9). Furthermore, the participants perceived that resolving an incident by taking another approach was as frustrating as being unable to resolve it. Taking another approach often involved reverting to manual procedures, which the participants possibly experienced as less satisfying than technological solutions.

Table 9. Categorization of how participants resolved the incidents, $N = 185$ frustrating incidents.

Category	<i>N</i>	%	Level of frustration ^a	
			Median	Distribution ^b
Resolved	53	29	6	
Resolved by waiting	28	15	6	
Resolved by repeating previous steps	26	14	6	
Resolved by taking another approach	20	11	7	
Resolved but with lower-quality outcome	10	5	6.5	
Unable to resolve	48	26	7	

^a The level of frustration was rated on a scale from 1 (not very frustrating) to 9 (very frustrating). ^b Response distribution from 1 (left) to 9 (right), median in red.

4.5 What were participants frustrated about?

Table 10 shows the distribution of the 185 frustrating incidents across the 25 categories that were the result of the card sorting. Collectively, the incidents spanned diverse issues ranging from slow systems, over corrupted contents and user slips, to privacy concerns and running out of charge. No single category accounted for more than 12% of the frustrating incidents. The level of frustration varied across categories from a median of three (noisy hardware and having forgotten how to do something) to nine (running out of charge), see Table 10. The median frustration level was above the midpoint of the scale for 20 (80%) of the categories.

Table 10. Categorization of frustrating incidents, $N = 185$ frustrating incidents.

Category	Class	N	%	Level of frustration ^a	
				Median	Distribution ^b
System is slow	Performance	23	12	7	---■---■---■---■---■
System froze temporarily	Performance	19	10	6	---■---■---■---■---■
System crashed	Performance	18	10	7	---■---■---■---■---■
Functionality not working	Performance	13	7	6	---■---■---■---■---■
Disruption in the flow of activities	Usability	10	5	5.5	---■---■---■---■---■
Input-device glitch	Usability	10	5	6.5	---■---■---■---■---■
Connectivity issue	Performance	9	5	7	---■---■---■---■---■
It is just hard	Usability	9	5	8	---■---■---■---■---■
Difficulty finding things	Usability	9	5	4	---■---■---■---■---■
Popups	Usability	9	5	6	---■---■---■---■---■
Corrupted contents	Utility	8	4	7	---■---■---■---■---■
Accidentally making a wrong selection	Usability	6	3	6	---■---■---■---■---■
Unsatisfactory contents	Utility	6	3	4	---■---■---■---■---■
Desired/absent functionality	Utility	5	3	4	---■---■---■---■---■
Poor information presentation	Usability	5	3	6	---■---■---■---■---■
Repeated login	Usability	5	3	7	---■---■---■---■---■
Computer just does not work anymore	Performance	4	2	7.5	---■---■---■---■---■
Uncertainty about payments	Utility	4	2	8	---■---■---■---■---■
Noisy hardware	Performance	3	2	3	---■---■---■---■---■
Cross-application tasks	Usability	2	1	5.5	---■---■---■---■---■
Gameplay is (too) hard	Utility	2	1	6	---■---■---■---■---■
Privacy concerns	Utility	2	1	7	---■---■---■---■---■
Not having the right version of documents	Utility	2	1	8.5	---■---■---■---■---■
Having forgotten how to do something	Usability	1	1	3	---■---■---■---■---■
Running out of charge	Performance	1	1	9	---■---■---■---■---■

^a The level of frustration was rated on a scale from 1 (not very frustrating) to 9 (very frustrating). ^b Response distribution from 1 (left) to 9 (right), median in red.

The three most frequent categories of frustration were about systems that were slow, froze, or crashed. For example, participants described that “The computer became slow and the game became less enjoyable”, “My pc is slow and when you try to open something like word or excel it froze for 1-2 min”, and “When I tried to open the folder with the dataset images, the computer just crashed on me and I lost some work that I forgot to save”. A total of 39 of the 60 incidents in these three categories had happened before and 38 were rated as likely to happen again. That is, recurrence was a frequent characteristic of these incidents, some of which received the maximum frustration-level rating. For example, a participant gave the following description of the frustration he experienced when his computer crashed while he was trying to deliver some work to a client:

The frustration made me to be emotionally, physically, psychologically, and socially down. I was dumbfounded for some time. I couldn't understand what was going on. I was angry and annoyed due to the occurrence to the extent I punched the wall. The anger couldn't just be appeased.

The next category concerned functionality that did not work. In some of these incidents, the participants used a system the way they normally did but, on this occasion, the functionality did not work. In other incidents, the non-working functionality was a recurrent issue for which the participants had developed a workaround. For example, one participant's mail client kept showing the most recent mail as read even before he had read it. This problem had occurred since the participant switched to Windows 10 and it also caused a synchronization issue with his phone, which consequently gave incorrect notifications about the number of unread mails. The participant had found a way to work around this problem but would prefer to be able to stop it from happening:

I resolve it by always clicking away to another email that has already been read. If I leave the new email highlighted then it stays in bold and doesn't clear as being read. I would like to fix it and stop it from happening but this is a workaround.

The next three categories concerned disruptions in the participants' flow of activities, input-device glitches, and connectivity issues. These incidents were system-triggered and contrary to the participants' intentions. For example, a participant was disrupted in her flow of activities because “the autocorrect changed a couple of words in my email content that I didn't want changed”. Several participants experienced input devices that did not register the participants' input correctly, including: “My mouse is having problems with its scroll wheel. Whenever I try to scroll up or down it doesn't always go the way I want it to go”. In addition, participants experienced frustrations connecting to the Internet or establishing connections among devices. As an example, a participant got frustrated when reconnecting his AirPods earphones to his computer:

My frustration stems from having had my AirPods connected to a different device the other day. Since my AirPods were then disconnected from my Mac it took some time to reconnect them again, as they had to pair via Bluetooth.

The two next categories involved tasks experienced as hard and difficulties finding things. In contrast to the previous categories, the system was not considered at fault. Still, the participants were frustrated about how difficult it was to accomplish what they wanted to do. A participant simply noted: “Im not very good at using this package and I cant do exactly what I want to”. Another participant turned the blame inward when finally finding a file: “Well I could not find a file, or so I thought, but is seems that I have renamed it so it's all okay I guess”.

The participants were also frustrated by popups, in particular advertisements. The popups delayed access to content, had to be closed manually, and were “absolutely annoying”. Skipping the advertisements as soon as it became possible to do so was not perceived as a satisfactory solution.

Similarly, participants' frustrations were not tempered by only being able to remove popups by selecting a remind-me-later option: "I clicked on the 'later' button but it only helps for a couple of hours and the message pops back up".

The remaining 56 incidents constituted 15 categories. While these categories were modest in size, they included three of the four categories with median frustration levels of 8 or more. These three categories were: running out of charge, not having the right version of documents, and uncertainty about payments. For example, a participant reported:

What I hate with a lot of these kind of services is that you do not get any kind of notification beforehand like: 'This is a Paid Service no matter what'. It says nothing until you are done with everything and all of what you did is just for nothing. There are more programmes and apps that does that which is very annoying and unnecessary.

A classification of the 25 categories into those concerning performance, utility, and usability showed an uneven distribution across these three quality attributes, see Tables 10 and 11. The frustration level did not differ across incidents concerning performance, utility, and usability, $H(2) = 3.27, p = .195$. Relatedly, the four categories with median frustration levels of 8 or more were spread across performance (1), utility (2), and usability (1).

Table 11. Frustrating incidents classified by quality attributes, $N = 185$ frustrating incidents.

Quality attribute	<i>N</i>	%	Level of frustration ^a	
			Median	Distribution ^b
Performance (i.e., whether system is responsive and reliable)	90	49	7	..-.-.##
Utility (i.e., whether system functionality matches user needs)	29	16	7-##
Usability (i.e., whether system is easy and satisfying to use)	66	36	6	..-.-.##

^a The level of frustration was rated on a scale from 1 (not very frustrating) to 9 (very frustrating). ^b Response distribution from 1 (left) to 9 (right), median in red.

4.6 How did participants construe computer frustration?

Before the one-hour reporting period, we asked the participants to describe what they in general considered a frustrating episode with computers. Fourteen participants provided what amounted to definitional statements, including that "a frustrating episode with computers is when they won't do what they are supposed to do". All these statements expressed that the participants were delayed or thwarted in attaining their goals. Some of the statements emphasized that the episode was unexpected, for example: "A frustrating episode with computers is when you think something in your mind and you want the computer to do that but for some unexpected reason it will not". In several cases, the unexpectedness concerned deviations from normal operations, for example: "A frustrating episode with computers in my opinion is a situation whereby my computer refuses to function in the usual way". Other statements emphasized that the user felt out of control, for example: "frustrating episodes are when there is an issue with the computer that is preventing me from doing what I need/want to be doing and I don't know to solve the problem". Several statements mentioned limited experience and competence in solving computer problems.

Our card sorting of all the descriptions of what the participants in general considered a frustrating episode with computers resulted in 39 groups that spanned 11 dimensions and a small ‘Other’ group, see Table 12. The user, task, and infrastructure comprised one dimension each; the eight remaining dimensions concerned the system. Overall, the dimensions tallied with the concrete incidents experienced by the participants during the one-hour reporting period (Table 10).

The three most frequently mentioned dimensions concerned the system and involved slow/unresponsive systems, hardware issues, and computer crashes. The only system dimension not represented by a concrete incident experienced during the reporting period was the one about dependencies/lack of interoperability. A participant described this dimension in the following way: “When you install a new piece of software and it then causes problems with the other applications or the devices on the computer, but you don’t know how or what it has done so it is very hard to fix”. The infrastructure dimension was the fourth most frequently mentioned dimension. It included descriptions that accorded with concrete incidents experienced during the reporting period, but it also added descriptions of frustration caused by unexpected automatic updates and viruses. The user dimension accorded with the concrete incidents, except for the addition of ergonomic issues. For example, the user dimension included the description: “The most frustrating things with computers are physical problems, such as back pain or problems with eyesight and tiredness”. Finally, the descriptions of the task dimension emphasized the loss of work due to unsaved data (e.g., “when i have completed some work and i think it is saved and it doesnt save and i lose my work”). This issue also appeared in several of the concrete incidents, but not as the feature that defined the incident categories.

Table 12. Dimensions in the participants’ conception of computer frustration, $N = 419$ descriptions (after splitting).

Dimensions	N	%
User	14	3
Cannot find stuff (4), forgotten step (3), ergonomics (2), not seeing cursor (2), programming: bugs that I cannot fix (2), long-filename search (1)		
Task	16	4
Data not saved/lost work (14), loosing game (2)		
Infrastructure	44	11
Unexpected automatic updates (21), virus (12), notifications/adds (11)		
System: slow, unresponsive	161	38
Freezing (60), lost connection (49), slow (48), online meeting failure (2), unstable (2)		
System: hardware issues	55	13
Hardware not working (51), dirty computer (2), low battery (2)		
System: crash	52	12
Crash (52)		
System: difficult to understand/use	18	4
Inexplicable features (7), misbehaving functionality (4), having to switch back and forth among apps (2), incomprehensible error codes (2), lack of control (2), new and unfamiliar (1)		
System: dependencies, lack of interoperability	17	4
Compatibility/dependencies (11), interoperability (4), synchronization failures (2)		
System: corrupted links and files	14	3
Downloads do not work (8), page not found (5), corrupted files (1)		
System: installation, defaults	12	3
Installing (11), persistent default apps (1)		

System: licenses, certificates, accounts	8	2
Accounts (2), cookies/invalid certificates (2), license (2), password (2)		
Other	8	2

5 Discussion

The conception of frustration expressed by the participants in our study corresponds with the formal definition of frustration outlined in the CUF studies (see Section 2.1). Like the formal definition, the conception expressed by our participants emphasizes that they are delayed or thwarted in attaining their goal, and often unexpectedly so. This correspondence confirms a key aspect of the empirical self-report approach and attests to the construct validity of the study.

5.1 Extent and level of frustration

Since the CUF studies [9,10,15,36,37] in 2003-2006, computers have become more mature and computer use more extensive. Comparing our results with those of the CUF studies, five findings stand out:

- *The time lost to frustration has dropped from 44-50% in the CUF studies to 11-20% in this study* (Tables 2 and 6). While the drop is gratifying, the time lost is still high. Imagine the consequences if our results generalize: For each five-day week of computer work, between half a day and a full day will be spent in frustrating episodes and recovering from lost work. This makes frustration a common user experience and a substantial cost for users and their employers.
- *The median level of frustration remains high.* Using the same nine-point scale, the median level of frustration is 7-8 in the CUF studies and 7 in this study (Tables 3 and 7). It is only the duration of the incidents that has lowered, their intensity is unchanged. Like Lazar et al. [36], we find that the frustration caused by an individual incident increased with increasing task importance and time lost.
- *Frustrating incidents still tend to recur.* In the CUF studies, the median frequency was that incidents recurred several times a month; in the present study, participants indicated that 84% of the incidents had happened before and that 87% could happen again (Tables 3 and 8). The high rate of recurrence shows that the cause of the incidents is often beyond the participants' control, thereby leaving them to experience and recover from the same problem repeatedly.
- *The number of unresolved incidents is as high or higher.* In the CUF studies, participants were unable to resolve 11% [37] and 16% [15] of the incidents, and they ignored the problem or opted for an alternative in a further 14% [37] and 11% [15] of the incidents. In the present study, participants were unable to resolve 26% of the frustrating incidents, took an alternative approach in 11% of them, and settled for lower-quality outcomes in 5% of the incidents (Table 9).
- *Computer experience and computer self-efficacy may not influence frustration.* The CUF studies yield mixed results on this issue. While Bessi re et al. [10] found that higher self-efficacy reduces computer frustration, Lazar et al. [36] found no effect of experience and self-efficacy on frustration level. We find that frustration level and time lost were unaffected by computer experience and largely unaffected by computer self-efficacy.

It appears that frustration will be high when time is lost on an important task, irrespective of the user's level of experience and self-efficacy. However, the incidents that cause frustration differ for novice and experienced users. For example, user errors are a more frequent cause of frustration among novice users [38], while freezing and crashing applications are a more frequent cause of frustration among

experienced users [41]. In addition, Novick et al. [41] found that more experienced users are more often able to resolve the problem that caused the frustration. This finding is however moderated by the high rate of recurrence in our study, which mostly involved experienced participants. The high rate of recurrence indicates that the participants may be able to resolve the individual incident but not to fix the underlying problem in a way that stops it from recurring. Several participants mention feeling out of control or having limited competence in handling the issues that cause their computer frustrations.

5.2 Causes of frustration

The participants are frustrated by incidents concerning poor performance, utility, and usability. First, *poor performance* is the most common cause of frustration. Half of the incidents concern this quality attribute. Slow, freezing, and crashing systems are prominent examples of frustratingly poor performance, thereby confirming previous studies [15,37,39,41,49]. In addition, slow and unresponsive systems are by far the most frequently mentioned dimension in the participants' conception of computer frustration. It contributes to slow performance that new applications tend to require still more network bandwidth and computer power, thereby quickly outdating users' current hardware. To counter this recurrent mismatch, some studies attempt to manage online delays by engineering the user's time perception. For example, Hong et al. [30] found that filling long delays with additional content that distracts the user from the passage of time is preferable to displaying information that alerts the user to the passage of time, such as progress indicators. These studies reduce frustration by manipulating perceived time rather than clock time. Poor performance also involves that specific functionality, internet connections, or the entire computer ceases to work, that the hardware becomes noisy, and that the computer runs out of charge. The incidents concerning performance are frustrating because the computer becomes less responsive and reliable.

Second, *poor utility* is the cause of one in six incidents. In these incidents, frustration results from a mismatch between system functionality and user needs. The modest number of incidents concerning this quality attribute may reflect that many participants were engaged in voluntary computer use and could, thus, choose an activity-application pair with a fair match between need and functionality. The reported mismatches concern corrupted, unsatisfactory, and outdated contents, desired but absent functionality, concerns about payments and privacy, and games that are too hard. Concerns about payment and privacy are new sources of frustration compared to the CUF studies [15,37], as are frustrations with computer games. These sources of frustration illustrate that our data include computer use during leisure activities. By spanning both students and professionals, our data also cover computer use related to education and work. Hertzum [28] found that frustration levels during leisure activities are at the lower end of the spectrum of frustration levels that users experience at work. He also found that higher frustration levels tend to co-occur with higher error rates and longer task completion times, thereby linking frustration to performance [29].

Third, *poor usability* is the cause of one in three incidents. These incidents concern disruptions in the flow of activities, input-device glitches, operations that are just hard, difficulty finding things, popups, accidentally making the wrong selection, poor information presentation, repeated logins, cross-application tasks, and having forgotten how to do something. These kinds of problems resemble the output from a usability test. They mainly pertain to two dimensions in the participants' conception of frustration: the user dimension and the dimension about systems that are difficult to understand and use. Several studies equate any frustrating incident with the presence of a usability problem [e.g., 14,39,41]. This way of defining usability problems links them directly to negative user experiences and broadens the scope of usability to all eleven dimensions in the participants' conception of computer frustration (Table 12). Contrary to the performance incidents, several participants described

the usability incidents in terms of the demands imposed on the user rather than the qualities absent in the system. This difference demonstrates how system and user characteristics interact in constituting frustrating incidents. In addition, the repeated finding of frustration caused by auto-corrections shows that facilities intended to alleviate frustration may, instead, add more of it [3,39].

Just as the incidents span the three quality attributes of performance, utility, and usability, the incidents are not restricted to one or a few applications. A wide range of software and hardware causes frustrating incidents [15,37]. This multiplicity complicates directed efforts to combat computer frustration because there is no single best place to start.

Note that some models of user satisfaction, such as the Kano model [32], suggests that as users' expectations raise, application features that used to be satisfying become neutral. It might also happen that features are expected and that their absence causes dissatisfaction. Might the same mechanism be at play with respect to frustrations, so that users' increased expectations of systems are simply not fulfilled? We do not think so. User satisfaction is the result of a complex inference before, during, and after interaction; frustrations result from having goals thwarted, see Figure 1. Thus, the processes appear to be different. Furthermore, while we might have come to expect to be able to do more with computers, many participants report frustrations that are mundane and feature simple tasks.

5.3 Consequences of frustrating incidents

Evidently, the incidents go beyond immediate frustration (see Figure 1). That is, they are not merely the adaptive redirection of attentional resources to features of the information environment that have temporarily become obstacles. They also involve the escalation of frustration to various emotional responses to the thwarted goal attainment and they have downstream effects on the quality and quantity of work products. Due to the frequency of computer frustration, these escalations and downstream effects have real consequences. We see three main consequences:

First, *frustration is integral to how people experience computers*. Computers are not solely experienced as empowering, fun, and supportive of creative expression; they are also the source of repeated frustration. This consequence is not the result of limited computer experience. On the contrary, the participants tended to be experienced with computers and used them for a median of 6-8 hours a day. While computer self-efficacy had virtually no influence on frustration level, we speculate that it may influence the downstream consequences of computer frustration. Specifically, frustrating incidents may reinforce low computer self-efficacy and cause people with low computer self-efficacy to approach computers with hesitation, shy away from them, or need support from colleagues and relatives to use them successfully. This way, frustrating incidents are a barrier to attaining universal usability [51].

Second, *a disturbing amount of time is lost to frustration*. The downstream consequences of this loss of time are enormous. People would, on a daily basis, have more time for productive work, enjoyable activities, and relaxation if they did not lose 11-20% of their computer time to frustration. This finding reminds us of Landauer's [34] critique of computers for not providing the productivity gains they are introduced to provide. While much has happened since Landauer's work, which was published in 1995, it appears that computer use still involves substantial amounts of unproductive time due to freezes, crashes, redoing lost work, and so forth. In addition, computers often introduce additional tasks that are experienced as frustratingly peripheral to the primary work. For example, the electronic health records introduced in hospitals have increased the time medical doctors spend on documentation tasks and reduced their face-to-face time with patients, thereby causing frustration at a level described as medical doctors hating their computers [22].

Third, *users are often delayed or obstructed in doing what they want their computer to help them do*. The delays and obstructions shift users' focus of attention from their activities to the computer. In 26% of the frustrating incidents, the participants were unable to resolve the incident and, thereby, prevented from resuming their activity. In the other incidents, the shift of attention broke the flow in the participants' activities and required them to regain their focus once they had resolved the incident. Several participants mentioned deadlines that became harder to meet and added further consequences to the incidents. This way, the frustrating incident foreshadowed subsequent frustrations in much the same way as the large number of frustrations that could happen again foreshadowed a continuing series of frustrating incidents.

5.4 Implications

The overarching implication of this study is that there is still an urgent need to make computer use a less frustrating experience. Meeting this need may seem a daunting task, given that frustrating computer experiences are diverse and in several respects as severe now as they were fifteen years ago. More specifically, we want to point out ten implications:

First, it is worth reiterating that the frustrating incidents are not caused by limited computer experience, low computer self-efficacy, or the like. Further support for this finding is reported by Kjeldskov et al. [33], who compared the usability problems experienced by the same users of the same system just after they started using it and after 15 months of daily use. While the overall number of experienced usability problems decreased, the decrease was modest and there was no decrease in the number of critical usability problems. Thus, additional experience and training will not do away with the frustrations. Only small gains can be expected from improving the user in one way or another.

Second, the categories and dimensions of frustration (Tables 10 and 12) provide a catalog that can be used in devising targeted efforts toward reducing computer frustration. Categories such as poor information presentation and absent/desired functionality appear solvable through increased use of user-centered design. Other categories, such as running out of charge, may appear mundane but strain current technological capabilities. Still other categories, especially popup ads, reflect conflicts of interest between users who want to get on with their activities and companies that want to promote their products. Diverse efforts are required to tackle the variety of reasons for the frustrations.

Third, slow performance is the single largest category of frustrating incidents. Faster hardware may reduce these problems but is probably a limited solution: More efficient hardware will likely be followed by functionality that demands more resources and by expectations for more immediate responses. This effect resembles Jevons' paradox [2] and suggests that other solutions are needed. The underlying issue appears to be insufficient attention to shielding slow performance from what the user experiences. Careful engineering of the user's time experience may reduce the perceived waiting time and, thereby, the experienced frustration [25,47]. The importance of such reductions is shown by Ramsay et al. [44], who document downstream effects of slow performance.

Fourth, frustrating incidents have downstream consequences. That is, they have consequences beyond the experienced frustration and time lost. For example, Ramsay et al. [44] report that frustratingly long download times carried over into less favorable assessments of the system content - once it became available. Relatedly, Zimmerman et al. [56] found that users who were subjected to a frustrating computer malfunction during the first task in an experiment performed worse on the second task compared to a control group that was not subjected to the computer malfunction during the first task. Our data also include examples of participants who needed to leave work and turn to household activities, such as cooking, to calm down after a frustrating incident. These findings illustrate the need for further studies of the downstream consequences of computer frustration.

Fifth, the participants indicated that most of the incidents were recurring problems. Longitudinal studies are required to validate this finding, which suggests that users are to a large extent living with recurring frustrations rather than resolving their computer frustrations. We speculate that recurring frustrations influence brand perceptions, foster workarounds, and contribute to burnout. These effects are aggravated by the general tendency for bad experiences, such as frustrations, to impress themselves more strongly on people than good experiences [7]. While workarounds are an adaptive response of taking action to mitigate problems, burnout is a distinctly maladaptive response where the recurrent presence of stressors results in a chronic feeling of energy depletion.

Sixth, the distinction between immediate frustration and emotional outcome (Figure 1) deserves further study because it accentuates that the absence or presence of an emotional outcome is an imperfect indicator of whether users were thwarted in attaining their goal. How often and under what conditions is immediate frustration accompanied by an emotional response? Hadlington and Scase [24] found more maladaptive emotional responses among users who scored high on fear of missing out, internet addiction, and the big-5 personality traits extraversion and neuroticism. Further work is needed to investigate how task characteristics, such as pace, and technology characteristics, such as mandatoriness, influence whether immediate frustration is accompanied by an emotional response.

Seventh, 37% of the participants reported no frustrations; this was not possible in Ceaparu et al. [15]. Apart from a slight age effect, our data do not explain differences in the number of frustrating incidents. Hadlington and Scase [24] suggest that the no-frustration participants may score higher on the personality traits of agreeableness, conscientiousness, and openness. Some participants may also be better than others at working around problems. Another explanation could be that the absence of frustration for some participants stemmed from the particular tasks they performed during the reporting period, rather than from their personality. Experience sampling [17] is one alternative methodology that might help explore these issues because it lends itself to longitudinal studies. Such studies would allow for investigating how computer frustration varies across a person's tasks and whether some people consistently experience no computer frustrations.

Eighth, the relation between frustration and usability should be analyzed further. Some studies equate experienced frustration with the presence of a usability problem [e.g., 14,39,41]. This practice defines usability very broadly, which makes sense in the context of usability testing. However, some of the reported frustrations are related to utility and, thus, clash with a traditional distinction between usability and utility. Are such frustrations indicators of usability problems or something else? Furthermore, the practice of equating usability problems with experienced frustration disregards the situations where the user's immediate frustration is not accompanied by an emotional outcome (see Figure 1). To detect these situations in usability tests, the evaluator must attend to whether the users' focus is shifted away from their goal to recover from breakdowns, irrespective of whether the shift is accompanied by the expression of frustration.

Ninth, the types of frustrations that participants report offer ideas for research in HCI and software development. It seems that operating systems could be developed to protect users much better from frustrations, in particular relating to crashes and unresponsiveness. It is also a fact that hardware decays. However, decaying hardware need not degrade the user experience. A gracious adaptation to decay should be possible, either manually as when one can select a display resolution that matches the strength of one's graphics card or automatically. Better feedback on the system state and better diagnostics might alleviate other types of frustration; however, it seems to us that such improvements would only matter for a small number of frustrations. Finally, participants describe trying again as a common strategy to work around frustrations that recur. Earlier work by Akers et al. [1] has shown

how usability problems may be predicted from logged user interface events. Similarly, research could explore the extent to which frustrations can be automatically diagnosed and alleviated in real-time.

Finally, future studies should continue to investigate the time lost to frustration. It appears especially important to validate the amount of time lost for other participant groups, with different methods, and over longer timespans. For example, computer users could be shadowed by expert observers, like in previous studies of communication loads and interruptions at hospitals [e.g., 53]. Or data spanning entire days may be collected using the day reconstruction method, like in previous studies of how working people spend their time [31]. Specifically, future studies should look more into the finding that female participants lost less (downward-adjusted) time to frustration than male participants but experienced similar levels of frustration. Is this finding due to situational factors or different gender attitudes to computers?

5.5 Limitations

Three limitations should be remembered in interpreting the results of this study. First, the participants are crowdworkers, who may not be representative of computer users in general. Specifically, male participants are overrepresented. In addition, the participants tend to be experienced with computers and have high computer self-efficacy; 79% of them use computers at least six hours a day. Second, we relied on participants to self-report their frustrations but acknowledge that self-reporting has limitations. Self-reporting enabled us to collect frustrating episodes during the participants' normal computer use but also meant that we did not have control over the use situation. In addition, the participants may have experienced frustrations they did not report. We aimed to make it straightforward to report frustrations but acknowledge that reporting meant extra work for the participants. That said, some participants reported even one-second frustrations. Furthermore, Ceaparu et al. [15] found near identical results for self-reported frustrations and frustrations collected by observing others. In contrast, Novick et al. [40] recommend observation over self-reporting to improve accuracy in frustration studies. Third, 79% of the participants are between 20 and 39 years of age. That is, older people are underrepresented relative to the age distribution of the general population. With increasing age, people likely experience more computer frustration because they undergo a gradual decline in mental and motor abilities. Increasing age is also associated with lower computer self-efficacy [45]. Collectively, these limitations suggest that our results more likely under- than overestimate the extent and impact of computer frustration.

6 Conclusion

Frustration is a common user experience. Across the 234 study participants, the average time lost to frustrating incidents was between 11% and 20% of the one-hour reporting period. Participants indicated that 84% of the incidents had happened before, that 87% could happen again, and that they were unable to resolve 26% of them. This high rate of recurrence and lack of control likely contributed to the frustration level, which was a median of 7 on a 9-point scale. The single largest source of frustration was slow performance, but it was followed by a variety of other issues, including systems that froze or crashed, functionality that did not work, disruptions in the flow of activities, input-device glitches, and many more. Collectively, these issues delay or obstruct users in doing what they want their computer to help them do and, thereby, make frustration integral to how they experience computers. Removing, or just reducing, the sources of frustration poses considerable challenges for computer science and human-computer interaction.

Acknowledgments

Research assistant Tor-Salve Dalsgaard made the tool for the participants to self-report the frustrations they experienced with their computer. We are grateful to the participants for sharing their computer frustrations.

References

- [1] David Akers, Matthew Simpson, Robin Jeffries, and Terry Winograd. 2009. Undo and erase events as indicators of usability problems. In *Proceedings of the CHI2009 Conference on Human Factors in Computing Systems*. ACM, New York, 659–668. DOI:<https://doi.org/10.1145/1518701.1518804>
- [2] Blake Alcott. 2005. Jevons' paradox. *Ecol. Econ.* 54, 1 (2005), 9–21. DOI:<https://doi.org/10.1016/j.ecolecon.2005.03.020>
- [3] Ohoud Alharbi and Wolfgang Stuerzlinger. 2022. Auto-cucumber: The impact of autocorrection failures on users' frustration. In *Proceedings of the Graphics Interface 2022 Conference*. Canadian Information Processing Society, Mississauga, ON.
- [4] Abram Amsel. 1992. Frustration theory - Many years later. *Psychol. Bull.* 112, 3 (1992), 396–399. DOI:<https://doi.org/10.1037/0033-2909.112.3.396>
- [5] Mara S. Aruguete, Ho Huynh, Blaine L. Browne, Bethany Jurs, Emilia Flint, and Lynn E. McCutcheon. 2019. How serious is the 'carelessness' problem on Mechanical Turk? *Int. J. Soc. Res. Methodol.* 22, 5 (2019), 441–449. DOI:<https://doi.org/10.1080/13645579.2018.1563966>
- [6] Ryan S. J. d. Baker, Sidney K. D'Mello, Ma Mercedes T. Rodrigo, and Arthur C. Graesser. 2010. Better to be frustrated than bored: The incidence, persistence, and impact of learners' cognitive-affective states during interactions with three different computer-based learning environments. *Int. J. Hum. Comput. Stud.* 68, 4 (2010), 223–241. DOI:<https://doi.org/10.1016/j.ijhcs.2009.12.003>
- [7] Roy F. Baumeister, Ellen Bratslavsky, Catrin Finkenauer, and Kathleen D. Vohs. 2001. Bad is stronger than good. *Rev. Gen. Psychol.* 5, 4 (2001), 323–370. DOI:<https://doi.org/10.1037/1089-2680.5.4.323>
- [8] Leonard Berkowitz. 1989. Frustration-aggression hypothesis: Examination and reformulation. *Psychol. Bull.* 106, 1 (1989), 59–73. DOI:<https://doi.org/10.1037/0033-2909.106.1.59>
- [9] Katherine Bessière, Irina Ceaparu, Jonathan Lazar, John Robinson, and Ben Shneiderman. 2003. Social and psychological influences on computer user frustration. In *Media Access: Social and Psychological Dimensions of New Technology Use*, Erik P. Bucy and John E. Newhagen (eds.). Routledge, New York, 91–103. DOI:<https://doi.org/10.4324/9781410609663>
- [10] Katie Bessière, John E. Newhagen, John P. Robinson, and Ben Shneiderman. 2006. A model for computer frustration: The role of instrumental and dispositional factors on incident, session, and post-session frustration and mood. *Comput. Human Behav.* 22, 6 (2006), 941–961. DOI:<https://doi.org/10.1016/j.chb.2004.03.015>
- [11] Mark Blythe and Andrew Monk. 2018. *Funology 2: From usability to enjoyment. Second edition*. Springer, Cham.
- [12] Kirsten Boehner, Rogério DePaula, Paul Dourish, and Phoebe Sengers. 2007. How emotion is made and measured. *Int. J. Hum. Comput. Stud.* 65, 4 (2007), 275–291. DOI:<https://doi.org/10.1016/j.ijhcs.2006.11.016>
- [13] Johannes Breuer and Malte Elson. 2017. Frustration-aggression theory. In *The Wiley Handbook of Violence and Aggression*, Peter Sturney (ed.). Wiley, New York, 40:01-40:12. DOI:<https://doi.org/10.1002/9781119057574.whbva040>

- [14] Anders Bruun, Effie Lai-Chong Law, Matthias Heintz, and Lana H A Alkly. 2016. Understanding the relationship between frustration and the severity of usability problems: What can psychophysiological data (not) tell us? In *Proceedings of the CHI2016 Conference on Human Factors in Computing Systems*. ACM, New York, 3975–3987. DOI:<https://doi.org/10.1145/2858036.2858511>
- [15] Irina Ceaparu, Jonathan Lazar, Katherine Bessiere, John Robinson, and Ben Shneiderman. 2004. Determining causes and severity of end-user frustration. *Int. J. Hum. Comput. Interact.* 17, 3 (2004), 333–356. DOI:https://doi.org/10.1207/s15327590ijhc1703_3
- [16] Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educ. Psychol. Meas.* 20, 1 (1960), 37–46. DOI:<https://doi.org/10.1177/001316446002000104>
- [17] Mihaly Csikszentmihalyi and Reed Larson. 2014. Validity and reliability of the experience-sampling method. In *Flow and the Foundations of Positive Psychology*, Mihaly Csikszentmihalyi (ed.). Springer, Netherlands, 35–54.
- [18] John Dollard, Neal E. Miller, Leonard W. Doob, Orval H. Mowrer, and Robert R. Sears. 1939. *Frustration and aggression*. Yale University Press, New Haven, CT.
- [19] Netia K. Dor-Shav and Mario Mikulincer. 1990. Learned helplessness, causal attribution, and response to frustration. *J. Gen. Psychol.* 117, 1 (1990), 47–58. DOI:<https://doi.org/10.1080/00221309.1990.9917772>
- [20] Joe S. Dumas. 1989. Stimulating change through usability testing. *ACM SIGCHI Bull.* 21, 1 (1989), 37–44. DOI:<https://doi.org/10.1145/67880.67884>
- [21] Nina Ferreri and Christopher B. Mayhorn. 2021. Individual differences in frustration and performance with online shopping activities. *Proc. Hum. Factors Ergon. Soc. Annu. Meet.* 65, 1 (2021), 148–153. DOI:<https://doi.org/10.1177/1071181321651054>
- [22] Atul Gawande. 2018. The Upgrade: Why doctors hate their computers. *New Yorker* November 12 (2018), 62–73. Retrieved from <https://www.newyorker.com/magazine/2018/11/12/why-doctors-hate-their-computers>
- [23] Donald Golden. 1980. A plea for friendly software. *ACM SIGSOFT Softw. Eng. Notes* 5, 4 (1980), 4–5. DOI:<https://doi.org/10.1145/1010884.1010885>
- [24] Lee Hadlington and Mark O. Scase. 2018. End-user frustrations and failures in digital technology: Exploring the role of fear of missing out, internet addiction and personality. *Heliyon* 4, 11 (2018). DOI:<https://doi.org/10.1016/j.heliyon.2018.e00872>
- [25] Chris Harrison, Zhiquan Yeo, and Scott E. Hudson. 2010. Faster progress bars: Manipulating perceived duration with visual augmentations. In *Proceedings of the CHI2010 Conference on Human Factors in Computing Systems*. ACM, New York, 1545–1548. DOI:<https://doi.org/10.1145/1753326.1753556>
- [26] Marc Hassenzahl. 2004. The interplay of beauty, goodness, and usability in interactive products. *Human-Computer Interact.* 19, 4 (2004), 319–349. DOI:https://doi.org/10.1207/s15327051hci1904_2
- [27] Morten Hertzum. 2010. Frustration: A common user experience. In *DHRS2010: Proceedings of the Tenth Danish Human-Computer Interaction Research Symposium. Computer Science Research Report #132*, Morten Hertzum and Magnus Hansen (eds.). Roskilde University, Roskilde, DK, 11–14.
- [28] Morten Hertzum. 2021. Reference values and subscale patterns for the task load index (TLX): a meta-analytic review. *Ergonomics* 64, 7 (2021), 869–878. DOI:<https://doi.org/10.1080/00140139.2021.1876927>
- [29] Morten Hertzum. 2022. Associations among workload dimensions, performance, and situational characteristics: A meta-analytic review of the task load index. *Behav. Inf. Technol.* 41, 16 (2022), 3506–

3518. DOI:<https://doi.org/10.1080/0144929X.2021.2000642>

- [30] Weiyin Hong, Traci J. Hess, and Andrew Hardin. 2013. When filling the wait makes it feel longer: A paradigm shift perspective for managing online delay. *MIS Q.* 37, 2 (2013), 383–406. Retrieved from <https://www.jstor.org/stable/43825915>
- [31] Daniel Kahneman, Alan B. Krueger, David A. Schkade, Norbert Schwarz, and Arthur A. Stone. 2004. A survey method for characterizing daily life experience: The day reconstruction method. *Science* 306, 5702 (2004), 1776–1780. DOI:<https://doi.org/10.1126/science.1103572>
- [32] Noriaki Kano. 1984. Attractive quality and must-be quality. *Hinshitsu - J. Japanese Soc. Qual. Control* 14, 2 (1984), 39–48.
- [33] Jesper Kjeldskov, Mikael B. Skov, and Jan Stage. 2010. A longitudinal study of usability in health care: Does time heal? *Int. J. Med. Inform.* 79, 6 (2010), e135–e143. DOI:<https://doi.org/10.1016/j.ijmedinf.2008.07.008>
- [34] Thomas K. Landauer. 1995. *The trouble with computers: Usefulness, usability, and productivity*. MIT Press, Cambridge, MA.
- [35] Jonathan Lazar, Jinjuan H. Feng, and Harry Hochheiser. 2017. *Research methods in human-computer interaction* (2nd ed.). Morgan Kaufmann, Cambridge, MA.
- [36] Jonathan Lazar, Adam Jones, Mary Hackley, and Ben Shneiderman. 2006. Severity and impact of computer user frustration: A comparison of student and workplace users. *Interact. Comput.* 18, 2 (2006), 187–207. DOI:<https://doi.org/10.1016/j.intcom.2005.06.001>
- [37] Jonathan Lazar, Adam Jones, and Ben Shneiderman. 2006. Workplace user frustration with computers: An exploratory investigation of the causes and severity. *Behav. Inf. Technol.* 25, 3 (2006), 239–251. DOI:<https://doi.org/10.1080/01449290500196963>
- [38] Valerie Mendoza and David G. Novick. 2005. Usability over time. In *Proceedings of the SIGDOC2005 Conference on Design of Communication*. ACM, New York, 151–158. DOI:<https://doi.org/10.1145/1085313.1085348>
- [39] Helena Mentis and Geri Gay. 2003. User recalled occurrences of usability errors: Implications on the user experience. In *Proceedings of the CHI2003 Conference on Human Factors in Computing Systems: Extended Abstracts*. ACM, New York, 736–737. DOI:<https://doi.org/10.1145/765891.765959>
- [40] David G. Novick, Edith Elizalde, and Nathaniel Bean. 2007. Toward a more accurate view of when and how people seek help with computer applications. In *Proceedings of the SIGDOC2007 Conference on Design of Communication*. ACM, New York, 95–102. DOI:<https://doi.org/10.1145/1297144.1297165>
- [41] David G. Novick, Baltazar Santaella, Aaron Cervantes, and Carlos Andrade. 2012. Short-term methodology for long-term usability. In *Proceedings of the SIGDOC2012 Conference on Design of Communication*. ACM, New York, 205–211. DOI:<https://doi.org/10.1145/2379057.2379097>
- [42] Agnieszka Piasna, Wouter Zwysen, and Jan Drahokoupil. 2022. *The platform economy in Europe: Results from the second ETUI internet and platform work survey. Working paper 2022.05*. European Trade Union Institute, Brussels.
- [43] Lisa Posch, Arnim Bleier, Fabian Flöck, Clemens M. Lechner, Katharina Kinder-Kurlanda, Denis Helic, and Markus Strohmaier. 2022. Characterizing the global crowd workforce: A cross-country comparison of crowdworker demographics. *Hum. Comput.* 9, 1 (2022), 22–57. DOI:<https://doi.org/10.15346/hc.v9i1.106>
- [44] Judith Ramsay, Alessandro Barbese, and Jenny Preece. 1998. A psychological investigation of long retrieval times on the world wide web. *Interact. Comput.* 10, 1 (1998), 77–86.

DOI:[https://doi.org/10.1016/S0953-5438\(97\)00019-2](https://doi.org/10.1016/S0953-5438(97)00019-2)

- [45] Kendra Reed, D. Harold Doty, and Douglas R. May. 2005. The impact of aging on self-efficacy and computer skill acquisition. *J. Manag. Issues* 17, 2 (2005), 212–228. Retrieved from <https://www.jstor.org/stable/40604496>
- [46] Pertti Saariluoma and Jussi P. P. Jokinen. 2014. Emotional dimensions of user experience: A user psychological analysis. *Int. J. Hum. Comput. Interact.* 30, 4 (2014), 303–320. DOI:<https://doi.org/10.1080/10447318.2013.858460>
- [47] Steven C. Seow. 2008. *Designing and engineering time: The psychology of time perception in software*. Addison-Wesley, Upper Saddle River, NJ. DOI:<https://doi.org/10.5860/choice.46-2138>
- [48] Brian Shackel. 1985. Human factors and usability - Whence and whither? In *Software-Ergonomie'85: Mensch-Computer-Interaktion*. 13–31.
- [49] Manoj K. Sharma, Shweta Sunil, B. N. Roopesh, Preeti M. Galagali, Nitin Anand, Pranjali C. Thakur, Priya Singh, S. J. Ajith, and Keshava D. Murthy. 2020. Digital failure: An emerging reason of anger expression among adolescents. *Ind. Psychiatry J.* 29, 2 (2020), 335–338. DOI:https://doi.org/10.4103/ipj.ipj_81_19
- [50] Saul Shiffman, Arthur A. Stone, and Michael R. Hufford. 2008. Ecological momentary assessment. *Annu. Rev. Clin. Psychol.* 4, (2008), 1–32. DOI:<https://doi.org/10.1146/annurev.clinpsy.3.022806.091415>
- [51] Ben Shneiderman. 2000. Universal usability. *Commun. ACM* 43, 5 (2000), 84–91. DOI:<https://doi.org/10.1145/332833.332843>
- [52] Donna Spencer. 2009. *Card sorting: Designing usable categories*. Rosenfeld, New York.
- [53] Rosemary Spencer, Enrico Coiera, and Pamela Logan. 2004. Variation in communication loads on clinical staff in the emergency department. *Ann. Emerg. Med.* 44, 3 (2004), 268–273. DOI:<https://doi.org/10.1016/j.annemergmed.2004.04.006>
- [54] Alexandre N. Tuch and Kasper Hornbæk. 2015. Does Herzberg's notion of hygienes and motivators apply to user experience? *ACM Trans. Comput. Interact.* 22, 4 (2015), article 16. DOI:<https://doi.org/10.1145/2724710>
- [55] Rongjun Yu. 2016. The neural basis of frustration state. In *Neuroimaging Personality, Social Cognition, and Character*, John R. Absher and Jasmin Cloutier (eds.). Academic Press, Amsterdam, 223–243. DOI:<https://doi.org/10.1016/B978-0-12-800935-2.00011-7>
- [56] Nichole K. Zimmerman, Everett Sambrook, and Jonathan S. Gore. 2014. The effects of a computer malfunction on subsequent task performance. *Behav. Inf. Technol.* 33, 9 (2014), 874–881. DOI:<https://doi.org/10.1080/0144929X.2012.733412>