Designing for Noticeability: Understanding the Impact of Visual Importance on Desktop Notifications

Philipp Müller
DFKI
Saarland Informatics Campus
Saarbrücken, Germany
philipp.mueller@dfki.de

Sander Staal
Institute for Visualisation and Interactive Systems
University of Stuttgart
Stuttgart, Germany
staals@student.ethz.ch

Mihai Bâce
Institute for Visualisation and Interactive Systems
University of Stuttgart
Stuttgart, Germany
mihai.bace@vis.uni-stuttgart.de

Andreas Bulling
Institute for Visualisation and Interactive Systems
University of Stuttgart
Stuttgart, Germany
andreas.bulling@vis.uni-stuttgart.de

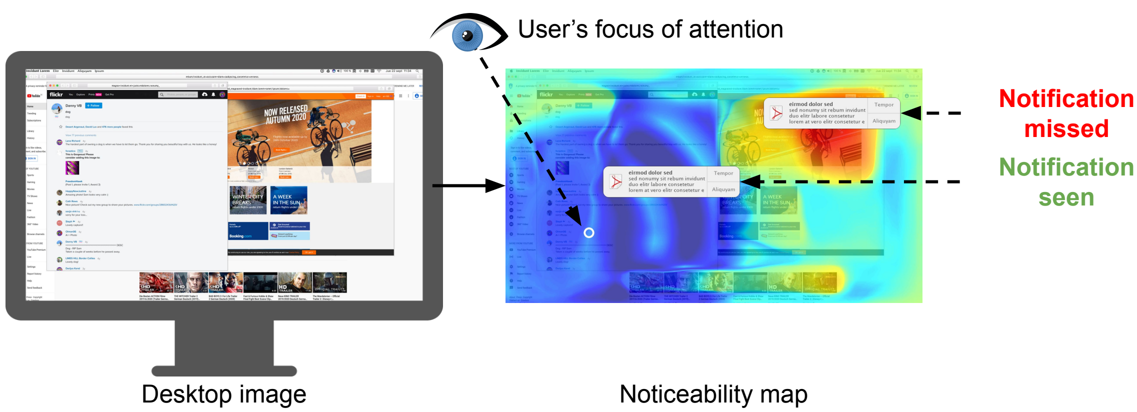

Figure 1: With data collected on realistically looking synthesised desktop images, we uncover the factors that impact noticeability of notifications. For a concrete desktop image and user attention focus, we build noticeability maps. These maps visualise the locations at which a notification is likely to be missed (red) or likely to be seen (blue).

ABSTRACT

Desktop notifications should be noticeable but are also subject to a number of design choices, e.g. concerning their size, placement, or opacity. It is currently unknown, however, how these choices interact with the desktop background and their influence on noticeability. To address this limitation, we introduce a software tool to automatically synthesize realistically looking desktop images for major operating systems and applications. Using these images, we present a user study (N=34) to investigate the noticeability of notifications during a primary task. We are first to show that visual importance of the background at the notification location significantly impacts whether users detect notifications. We analyse the utility of visual importance to compensate for suboptimal design choices with respect to noticeability, e.g. small notification size. Finally, we introduce *noticeability maps* - 2D maps encoding the predicted noticeability across the desktop and inform designers how to trade-off notification design and noticeability.

CHI '22, April 29-May 5, 2022, New Orleans, LA, USA

CCS CONCEPTS

• **Human-centered computing** → **Graphical user interfaces**; *User studies*; *HCI design and evaluation methods.*

KEYWORDS

notifications, attention, graphical user interfaces, dual task performance

ACM Reference Format:
Philipp Müller, Sander Staal, Mihai Bâce, and Andreas Bulling. 2022. Designing for Noticeability: Understanding the Impact of Visual Importance on Desktop Notifications. In *CHI Conference on Human Factors in Computing Systems (CHI '22), April 29-May 5, 2022, New Orleans, LA, USA.* ACM, New York, NY, USA, **??** pages. https://doi.org/10.1145/3491102.3501954

1 INTRODUCTION

Desktop notifications are widely used to notify users about incoming emails, upcoming calendar entries, or other relevant events.

To ensure that important information is actually noticed, notifications need to effectively attract and divert the users' attention from a primary task to a secondary task [44]. At the same time, notifications are part of the visual design of the user interface (UI) and are subject to aesthetic considerations. For example, designers may customise a notification's appearance in terms of size, placement, or opacity [46]. However, such design decisions can severely impair the user's ability to perceive notifications [24]. To create notifications that are not only in line with a designer's vision but also functional and noticeable, it is imperative for UI designers to understand the factors that impact noticeability.

Prior work studying the noticeability of desktop notifications has investigated factors including notification size, shape, color, movement, or opacity [16, 24, 28]. While this research has decomposed influences on noticeability in a highly controlled manner, it has two main limitations. First, these studies lack realism and have mainly used *simplified* or only few, carefully-selected desktop images as well as *highly abstract visual representations* of notifications [24, 28]. As a result, it remains unclear whether findings obtained in such artificial settings generalise to realistically-looking notifications placed on realistic desktop images. Second, prior work has not studied the *impact of the desktop's visual appearance (i.e. desktop background, icons and any applications) on noticeability.* It is well known that visual stimuli from the environment guide people's attention. To replicate these guidance effects, so-called saliency and visual importance models have been proposed [4, 11, 21]. It is therefore conceivable that the visual appearance of the desktop background also has a significant impact on noticeability and models of visual importance can be used to uncover these effects. One major obstacle that has so far prevented these studies is the lack of a dataset containing diverse and realistic desktop images.

Our work makes two original contributions to address these limitations. To study noticeability of desktop notifications in more diverse desktop environments, we introduce a software tool to efficiently synthesise a large number of realistic desktop images and notifications from three major operating systems. The synthesised desktop images contain realistic icons, task bars, diverse wallpapers, as well as variable arrangements of application windows. Our implementation of the software tool is publicly available[1]. Using these images, we conducted a 34-participant controlled user study in which participants were asked to detect notifications while performing the primary task of following a moving dot via the mouse pointer. While this primary task directed participants' attention to desktop locations across the entire screen, notifications appeared at random locations on the desktop interface and in different sizes, opacities, and aspect ratios. Using a state-of-the-art method to predict visual importance of desktop images [11], we show for the first time that the visual appearance of the background at the location of the notification has a significant impact on noticeability.

We analyse how visual importance interacts with major notification design factors investigated in previous work [24, 46], including opacity and size of notifications as well as the distance of the notification to the current attention focus of the user (i.e. the primary task location). This allows us to suggest how saliency-optimised display of notifications could be used to allow UI designers a larger degree of freedom for choices on these design factors that are sub-optimal with respect to noticeability. For example, a designer might want to display a notification with a low opacity value in order to integrate it aesthetically into the UI. While this is, in general, detrimental to noticeability, the notification can still be detected with high probability by taking visual importance into account when choosing the location at which the notification will be displayed.

We finally introduce *noticeability maps* - 2D maps that encode the expected noticeability at all locations on the desktop (see Figure 1). In contrast to current visual importance maps, our proposed noticeability maps are estimated from our study data and encode the interaction between visual importance and the users' current focus of attention (see Figure 2). These visualisations of noticeability values for different desktop regions provide an intuitive tool that can help designers to maximise noticeability of notifications, thereby increasing the degrees of freedom for aesthetic design decisions without sacrificing noticeability.

2 RELATED WORK

Our research is related to prior work on (1) understanding and optimising user interface notifications as well as (2) computational modelling of visual attention and visual importance in images.

2.1 Understanding and Optimising Notifications

Given the ever-increasing number of digital interfaces that generate an even larger number of notifications everyday [33], research on understanding and optimising notifications has surged in recent years. Early work has shown that while notifications are often perceived as a source of disruption [19] and distraction [17], a lack of notifications may lead to additional task switching [20]. It is well established, however, that excessive notifications and alerts have negative consequences and lead to inattention [25]. Consequently, significant research efforts have been spent on optimising when and how notifications are presented across a wide range of devices, from desktop computers [18], to mobile phones [29, 31, 33, 38], smartwatches [41], and even smart TVs [46] or virtual reality headsets [14]. Others explored multi-device settings – Weber et al. [47] studied how to best distribute notifications across multiple devices while Voit et al. [45] investigated how notifications were perceived in such scenarios.

Notifications on mobile devices have been studied extensively, especially concerning interruptibility of the user [1, 30, 32, 35], perceptibility of notifications [3, 9], as well as attentiveness [34] to notifications. Focusing on interruptibility, i.e. finding the opportune moment in time to interrupt users and deliver notifications, Poppinga et al. [35] developed a method using device-integrated sensors to predict when to display notifications. In addition to the device's context and sensors, Mehrotra et al. [30] proposed an approach that included the content of a notification for the same task. Further research investigated what factors influence perceptibility of notifications. For example, inserting visual elements or issuing notifications at specific times of the day leads to higher click rates [3]. Exler et al. [9] studied perceptibility of different notification types (e.g. ringtone, vibration, or LEDs) in different locations such as the user's pocket or on a table. Mehrotra et al. [31] studied

[1]https://github.com/sanderstaal/screenshot-synthesize

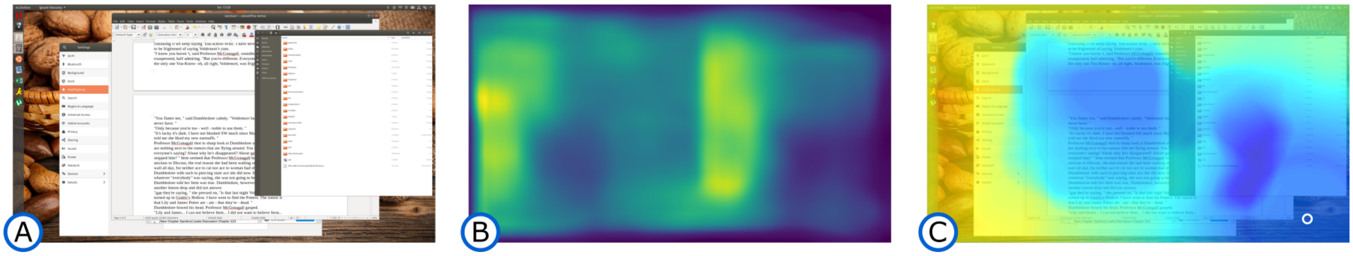

Figure 2: Realistic desktop image created with our tool (A), along with the corresponding visual importance map (B) and our proposed noticeability map (C). In contrast to visual importance maps that generally encode bottom up attention distributions, noticeability maps encode the likelihood of a notification to be detected while considering both bottom up (visual importance) and top down (the users' current focus of attention – white circle) features as well as the appearance and design of a notification. Visual importance map (B): Blue represents low visual importance, yellow and red represents high visual importance. Noticeability map (C): Blue is used for areas in which notifications will be highly noticeable and yellow to red is used for areas where a notification is expected to not be noticed (and thus should not be placed). Separate noticeability maps can be computed for different notification properties like opacity and scale.

different notification factors (e.g. sender-recipient relationship or alert modality) and their impact on response time and the users' ability to perceive notifications. A complementary task to predicting interruptibility is to predict attentiveness [34], which defines the level of attention paid towards a notification or message. Pielot et al. [34] used the smartphone sensors to build a random forest model to predict high or low attentiveness to notifications. While all the different facets of mobile notifications have been thoroughly explored over the years, desktop notifications and what makes them noticeable remains under explored.

In contrast to mobile devices – where optimising notifications is (mostly) about when and how to notify users, desktop environments are more complex and enable additional design considerations such as notification placement or different visual features. While the impact of notifications on users' interaction has been studied previously [18], the characteristics that make them (in)effective has only recently attracted research interest. Klauck et al. [24] provided first evidence of how different design properties, such as size, opacity, movement speed, or blink frequency, influence a notification's noticeability and distractivness. For example, their findings showed that a notification's size provides flexible control of noticeability relative to the gaze distance, while reducing the opacity can make notifications more subtle. Jones et al. [22] investigated shape-changing circuits as a way to provide notifications in the periphery. Mairena et al. [28] also investigated peripheral notifications, however, in their work, they studied the effect of different feature combinations (e.g. shape, color, or motion) and task interference. Another work analysed the effects of emphasis on simple scatter plots or visualisation [27].

While the above works provide a deeper understanding of some of the features that make visual desktop notifications more effective, the main limitation of prior works is the lack of realism and diversity in the appearance of notifications and the backgrounds on which notifications were presented. Visual stimuli in the environment and in complex UIs are known to guide user attention [4, 11, 21], however, prior works did not investigate this effect and its impact on noticeability. In our work, we rely on computational models of visual attention to study the impact of the desktop's appearance on notifications and their noticeability.

2.2 Computational Modelling of Visual Attention in HCI

Visual attention modelling (saliency modelling) is a core research area in computer vision [4] that aims to predict saliency maps that topographically encode the probability of visual attention over an image. Bottom-up models [13, 21] extract visual features only from the image while top-down models aim to incorporate task-related influences [36, 51].

Early work to use such models in HCI focused on web pages. Still et al. showed that bottom-up saliency maps correlated well with fixations during free-viewing of web pages [42]. Buscher et al. [6] proposed a method that leveraged both eye tracking data collected from 361 web pages and features from HTML to predict saliency of different page elements. A similar approach was taken by Shen and Zhao [40] who presented a computational saliency model that integrated both multi-scale low-level features and priors calculated from eye tracking data during web page viewing. They later improved their method to also include higher-level semantic feature representations, e.g., from object detections [39], Zheng et al. [52] presented a learning-based framework for predicting task-driven visual saliency on web pages whereas Li et al. [26] showed how an SVM trained on manually selected bottom-up and top-down factors could predict human visual attention while viewing web content. Bulling et al. used a visual saliency model in the context of gaze-based authentication to mask out salient image regions that would attract users' visual attention to improve the security of graphical passwords [5].

Methods that model attention on graphical user interfaces have only recently started to being investigated. Xu et al. proposed a spatio-temporal approach that used bottom-up user interface features as well as top-down information in the form of users' mouse and keyboard actions [50]. Gupta et al. developed a deep learning model based on an autoencoder to predict the saliency of mobile interfaces [15]. Another line of work introduced a learning-based method to predict "visual importance" for data visualisations [7], which encodes the relative importance of different visualisation or design elements. While ground truth saliency is typically collected from gaze data or approximated using mouse clicks and interfaces such as BubbleView [23], visual importance annotations require manual labelling of the visualisation elements that the annotators consider to be important (e.g. the title). As such, according to Bylinskii et al. [7], the importance scores are more uniformly distributed on the visualisation elements. This is beneficial for applications where the information from the visual stimuli is more structured such as in data visualisations or desktop images.

In our work, we used the recent Unified Model of Saliency and Importance (UMSI) [11] to understand and predict visual importance on desktop images. UMSI is especially suited four our application, as it is both able to predict visual importance on UI elements as well as slaiency on natural images, which are often part of the desktop background. As such, our work is the first to consider the context in which a notification is embedded (i.e. the desktop background including any applications) and its effect on noticeability.

3 SYNTHESISING REALISTIC DESKTOP INTERFACES

To systematically study noticeability of desktop notifications in realistic desktop environments, we needed a large and diverse collection of desktop images. However, to the best of our knowledge, there is no other work that provides such a dataset nor could we find publicly-available, high-resolution desktop screenshots, likely due to privacy concerns. For this reason, we decided to use an approach to automatically generate realistic desktop images. There exists only few prior work on this topic. For example, in SUPPLE [12], user interface rendering is modelled as a computational optimisation problem to generate UIs that meet the device's constraints while minimising user effort. Another tool is SpiderEyes [8], which is a system for designing attention- and proximity-aware collaborative interfaces for wall-sized displays. Based on the location and head position of up to four users, it automatically adapts the UI to support collaborative scenarios. Todi et al. [43] presented a tool that leveraged the user's browsing history to adapt a website's layout in a way such that the design looks familiar to each user. However, none of these works can be used to create realistically-looking desktop environments or notifications. To fill this gap, we developed a software tool to efficiently synthesise any number of realistic desktop images and notifications from three major operating systems: Microsoft Windows, macOS, and Ubuntu. Figure 3 shows a few examples generated using our tool. We describe the components of the synthesis tool below:

- *Desktop background.* We collected 67 different wallpapers for each operating system covering diverse motifs, styles, or colours. When generating a new image, the tool randomly samples a wallpaper and then places a random number of shortcut icons on it. Icons can be placed either in a grid pattern or randomly, similar to what most operating systems provide. If icons are placed in a grid alignment, they will be placed in blocks, close to one another to mimic realistic desktop setups. Our tool contains 100 common application icons collected from the Internet.
- *Menu and task bar.* The graphical interfaces of each operating system typically contain menu or task bars such as the top bar and the dock in macOS. Our tool synthesises these bars and populates them with items. Each text in these bars is substituted with random words and the number of items per section is varied, meaning that the tool can create both densely and sparsely-filled task bars. For cases where the real-world counterparts contain icons, the tool samples random icons from the same collection used to generate desktop shortcuts. Each of the three operating systems supported by our tool have a menu bar where frequently used applications and tasks are displayed (like the dock in macOS). As an additional customisation, our tool randomly selects some of these items and adds a 'highlighted' effect to them, which is used by the system to indicate currently open applications.
- *Applications.* Once the tool generated the desktop background and all required menu and task bars, the desktop image is randomly populated with a number of open application windows. It can also happen that no open window is added to the image. We collected 150 screenshots from 100 commonly used websites, some of which inspired by prior work [40]. We captured each screenshot twice, once in a maximised, widescreen browser window (3831 x 1933 px) and once in a restricted-sized browser window (1500x1500 px). This allowed us to capture different responsive designs of a website and hence obtain more diverse visual appearances. When generating new application windows, the tool randomly selects either a maximised, full screen window or restricted-sized windows that only cover parts of the background. Each website screenshot is enclosed in a proper browser application window: MS Edge for Windows, Safari for macOS, and Firefox for Ubuntu. A random string is generated to display the URL within the browser window. In addition to the website collection, we added screenshots of other common applications, such as Minesweeper or File Explorer, for each operating system (46 for Windows, 6 for macOS, and 24 for Ubuntu). Similarly to website screenshots, the tool generates a window which encloses the application screenshot in the same design as used by the OS. Due to the lack of diverse and natural high-resolution screenshots of websites or applications available on the Internet, all screenshots were manually captured by an experimental assistant.
- *Notifications.* Besides desktop images, the tool can be used to generate random but realistically-looking notifications in the same style of the corresponding operating system. For increased diversity, we also covered notifications that differed from official style guides[2]. Our tool first randomly samples one out of six (4:15, 3:4, 25:16, 35:14, 38:9, and 42:7) possible aspect ratios for the notification, where 35:14 is closest to the actual ratio used by these operating systems. All notifications contain a randomly generated text of variable length. In addition, the notification body may also contain an icon (from the desktop shortcuts collection) or an action button (only for Window and macOS).

[2] E.g. macOS Notifications https://developer.apple.com/design/human-interface-guidelines/macos/system-capabilities/notifications/

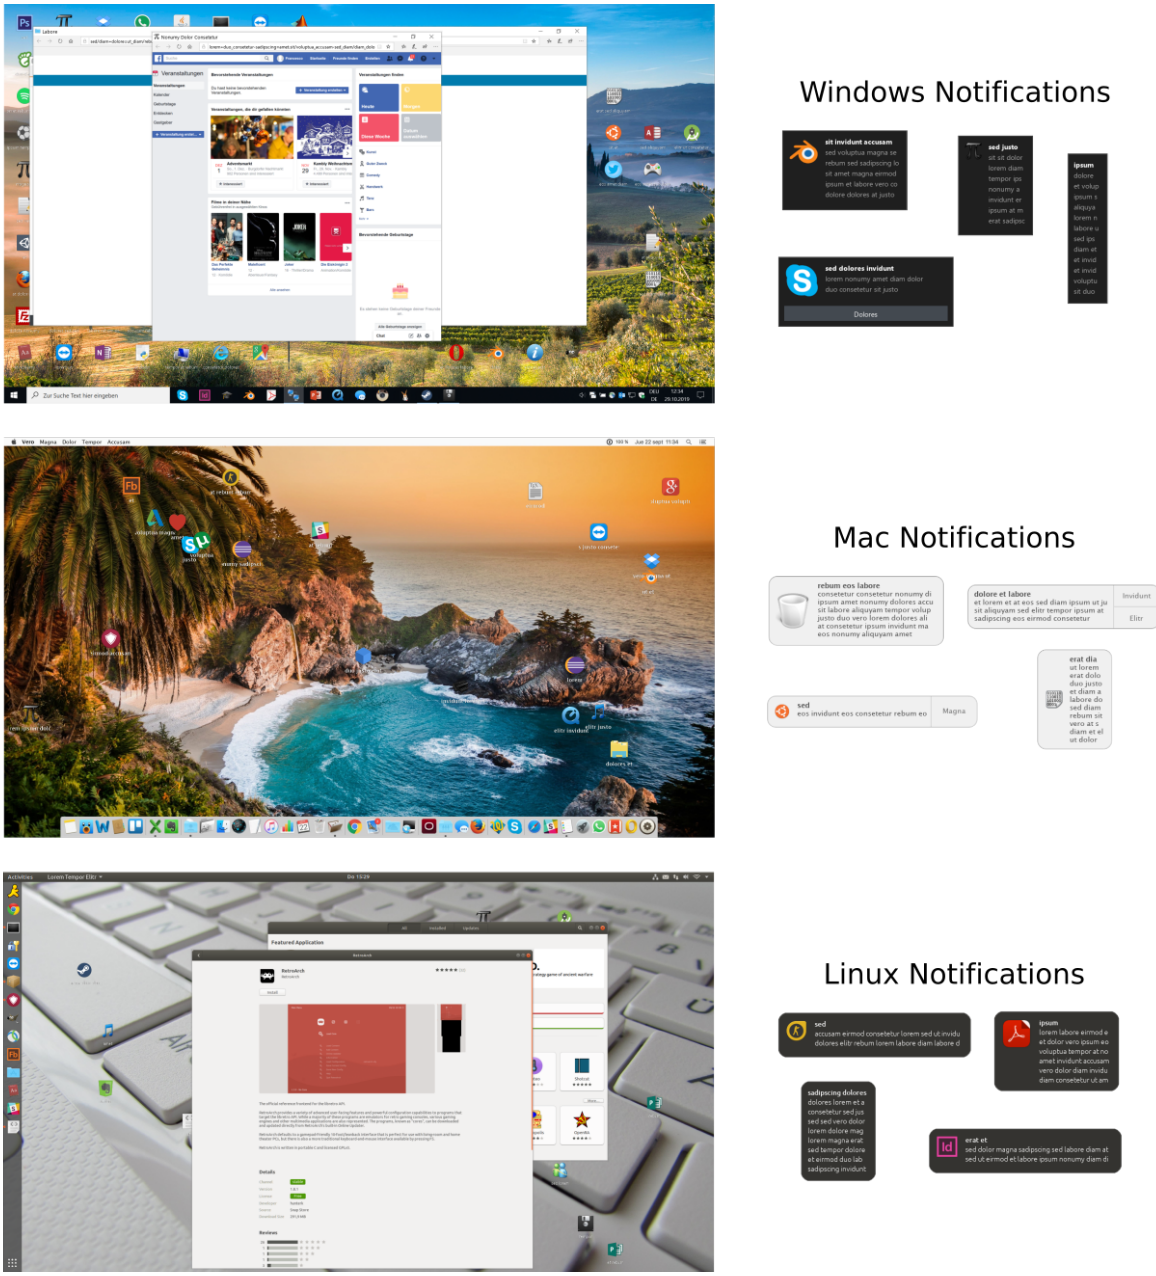

Figure 3: Using our proposed software tool, we are able to synthesise realistically-looking desktop images (left) and notifications (right). The tool can generate desktop images for three different operating systems that contain realistic icons, task bars, diverse wallpapers, and different application windows.

The final desktop image is realised by merging all of the individual components, i.e., the desktop background, menu, task bar, and applications. The outcome is a realistically-looking desktop image of 1920x1080 px. For further realism, we used the default font from each operating system and randomly placed a mouse icon on the image. This image formed the basis on which notifications were shown during the user study.

4 USER STUDY

We used our tool to synthesise 300 different desktop images and notifications, 100 for each supported operating system. Using these desktop images, we conducted a controlled user study in which participants had two tasks. Similarly to previous work [24], we employed a primary task in which participants had to follow a

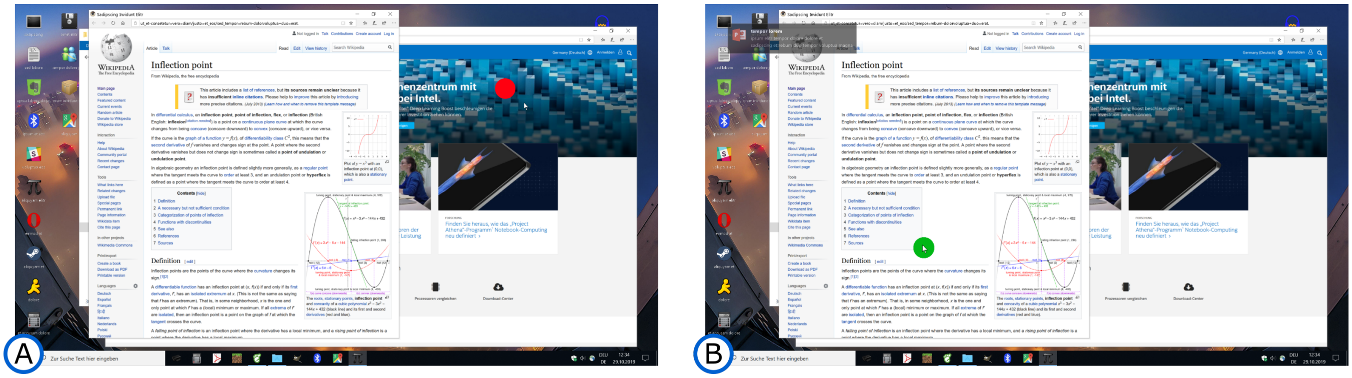

Figure 4: Example image participants saw during the user study. As a primary task, participants were instructed to follow a moving green dot with their mouse pointer. The dot moved along a random zig-zag path. If the mouse pointer was outside of the moving dot, its colour was shown in red (A). As a secondary task, participants were asked to detect notifications appearing on the screen. Figure (B) shows a notification with medium opacity in the upper left corner. Participants were instructed to press the space bar whenever they detected a notification.

moving dot with the mouse pointer in order to manipulate participants' focus of attention. While concentrating on the primary task, the secondary task involved detecting notifications appearing at random locations. Combining these two tasks results in a diverse set of configurations of user attention and notification locations.

4.1 Apparatus

The user study was conducted remotely and online in a web browser, and all participants could use their own personal computer for it. We restricted the study to participants using either a monitor or a laptop screen (no mobile or tablet devices allowed). The study was performed in a full screen browser window and participants were asked to not close this window during the experiment. In case the generated desktop image (1920 x 1080 px) did not fit the resolution or aspect ratio of a participant's screen, the image was automatically resized and padded with black borders where necessary. The largest resolution used by our participants was 2560 x 1440 px (used by five participants) and the smallest resolution was 1228 x 691 px. Most of our participants (18) used our default resolution of 1920 x 1080 px.

4.2 Experimental Procedure

In line with previous work [24] on noticeability of notifications, we employed a dual-task design. As in [24], participants followed a moving dot with the mouse pointer (primary task) and confirmed the appearance of notifications by pressing the space bar (secondary task). Each participant performed several sessions of this task. In detail, for each session a random desktop image from a random operating system (Linux, Mac or Windows) was sampled and served as the background on which a small coloured dot moved along a random zig-zag path. The path was generated by successively connecting random on-screen locations. For the primary task, participants were instructed to follow the moving dot with the mouse pointer as well as they could. If the mouse pointer was outside the moving target, the colour of the dot changed from green to red (see Figure 4A). We chose this task as users are required to focus their attention on the moving dot and the location can be used as a proxy to the users' focus of attention. The path followed a zig-zag pattern to (a) make sure participants could not predict the path and were required to continuously concentrate on the task and (b) to sample maximally diverse attention points from the screen.

As a secondary task, participants were asked to detect notifications appearing on the screen by pressing the space bar (see Figure 4B). Notifications were randomly displayed in time and location on the screen with varying opacity (between 20% and 100%), sizes, and aspect ratios. All these parameters were chosen randomly to collect a diverse set of configurations and prevent bias effects in participants' response behaviour. We especially did not include a separate condition with notifications placed in the upper right or lower right corner of the screen (as is default in major operating systems). Any difference between such a condition and randomly placed notifications might be due to bias effects resulting from the participants being able to clearly distinguish this as a special condition. Each notification was displayed for 2.5 seconds, where during the first and last 0.5 seconds a fade-in (or fade-out) effect was used. If a participant confirms a notification, the notification immediately vanished from the screen. No matter whether the participant confirmed a notification or not, the application showed the next notification after a random time of 2 to 5 seconds.

Every session lasted two minutes. Upon completion of a session, participants were awarded a score reflecting how well they performed in the two tasks. Their score was displayed on a global (anonymised) leader board, where they could compare their performance with other participants from the study. We used this leader board as a motivation, to encourage participants to repeat and perform multiple such sessions in our study. At the beginning of the user study, we asked each participant to play a least six such sessions. We measured that the time needed for a participant to complete six sessions is approximately 20 minutes. Participants were free to take breaks between these sessions, where they could also close and re-open the full screen window. Any data collected

Figure 5: Example of visual importance extraction. Given an input image A we extract a visual importance map (B) using the method proposed by Fosco et al. [11]. This visual importance map is subsequently normalised by the average visual importance screenshots from the particular operating system (C) to produce the final importance map (D).

during the study was completely anonymised and could not be linked to any participant.

4.3 Participants

We recruited 42 participants via local university mailing lists that recorded 248 sessions in total. Out of those, 14 sessions were invalid because participants quit the full screen mode during a data collection session. Out of the 42 participants, 34 (20 male, 13 female, one unspecified) finished at least six sessions. Based on a demographics survey, the ages of the participants ranged from 18 to 52 (M=26.09, SD=6.15). 24 participants used a desktop monitor, while the remaining 13 participants were using a laptop screen. Four participants reported suffering from visual impairments that were fully corrected by glasses or contact lenses. None of the participants reported deficiencies in color perception. 25 participants considered themselves as having a good technical expertise. 25 participants used a Windows operating system on their private computer, 5 participants used macOS, and 4 participants used a Linux distribution. The operating system installed on the computer did not influence which synthesised desktop image was shown to the participants.

5 RESULTS

Of the 34 participants in our study, one participant had to be removed from further analyses because of detecting less than half of the presented notifications (indicating that the participant did not focus on the task adequately), resulting in 33 remaining participants. In the following, we first describe how we extracted visual importance at notification locations from desktop images. Subsequently, we present analyses on the connections between noticeability and notification design factors. Finally, we introduce *noticeability maps* that encode the expected likelihood of detection for notifications presented at different locations. For statistical analysis, we use non-parametric tests due to violated normality assumptions in some cases and report median $\mu_{1/2}$, mean μ, and standard deviation $\hat{\sigma}$.

	VI	Distance	Size	AR	Opacity
Detected					
Median	0.94	0.32	0.0125	2.79	160.3
Mean	0.93	0.32	0.0125	2.84	160.3
SD	0.09	0.02	0.0006	0.28	8.576
Missed					
Median	1.24	0.41	0.0108	2.93	125.0
Mean	1.31	0.40	0.0107	2.99	121.9
SD	0.45	0.05	0.0010	0.84	16.43
Wilcoxon					
T	16	17	12	242	0
p	< 0.001	< 0.001	< 0.001	0.49	< 0.001

Table 1: Median, mean, and standard deviation of notification design factors for detected as well as missed notifications. The design factors are the visual importance of the background at the notification location (VI), the distance of the notification from the current locus of attention, the notifications' size, its aspect ratio (AR), and its opacity. Additionally, we report results of two-sided Wilcoxon signed-rank tests comparing detected and missed notifications for each factor (n=33). With Bonferroni correction, p-values smaller than 0.01 can be considered statistically significant.

5.1 Extracting Visual Importance on Desktop Images

To extract visual importance on the generated desktop images we use the recent state-of-the-art method for visual importance prediction across graphic design types by Fosco et al. [11]. The advantage of this method over other approaches is its ability to predict visual importance on natural images as well as on graphical designs and mobile user interfaces. This fits our purpose, as the desktop interactions we study contain images along with graphical designs and user interface elements. In order to better capture which regions in the desktop image stand out relative to the expected visual importance distribution on desktop images, we subtract the average visual importance computed over all desktop images generated for a given operating system from each single visual importance map for the corresponding operating system (see Figure 5 for an example). To compute the visual importance of a desktop image at the location where a notification is placed, we take the average visual importance score of the area covered by the notification.

5.2 Effects of Visual Importance on Noticeability

To check whether visual importance at the location of a notification is connected to the probability of the notification being detected, we conducted a Wilcoxon signed-rank test on the visual importance at the notification location with the two conditions "notification detected" ($\mu_{1/2}$ = 0.94; see Table 1 for more information) and "notification not detected" ($\mu_{1/2}$ = 1.24). This test reached significance (T=16; p<0.001; n=33). For a more detailed picture on how visual importance is connected to the noticeability of notifications, we

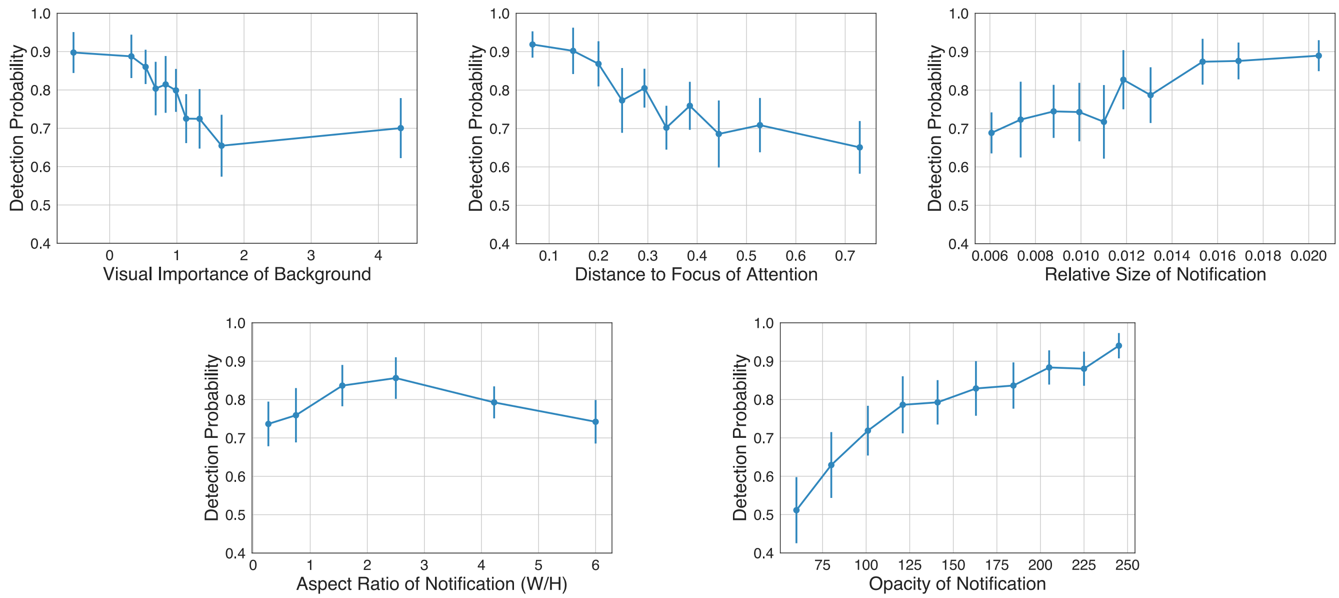

Figure 6: Influence of different factors on the probability for a user to detect (i.e. notice) a notification. Error bars indicate 95% confidence intervals. From top left, in clockwise order, the individual factors are: The visual importance of the desktop image at the location where the notification is placed, the distance of the notification to the current focus of user attention, the relative size of the notification on the screen, the opacity of the notification ranging from 0 (fully transparent) to 255 (fully opaque), and the aspect ratio of the notification expressed as width divided by height.

binned visual importance values into ten equally sized percentiles (10%,20%,...,100%) and plot the probability of notifications being detected for each visual importance bin (see top left of Figure 6). The figure shows a clear decrease in the detection probability with higher visual importance scores with a high plateau for low visual importance scores at around 0.9 detection probability and a low plateau between 0.65 and 0.7 for high visual importance scores.

We did not include a condition with standard notification locations (upper right for Mac/Linux, lower right for Windows) in our study design because such a "special" condition could easily bias participants' responses. In contrast to randomly placing notifications, sessions with fixed notification locations could easily be identified by participants as being different. Hence, participants could anticipate where a notification will show up and thus always notice it. Nevertheless, it is still possible that notifications placed in standard locations have different noticeability. To test this hypothesis, we analysed the detection probabilities of notifications that were placed (by chance) in standard locations. We defined a notification to be in a standard location for macOS and Linux if it was placed in the top right corner of the screen. More specifically, we checked whether the top right corner of the notification was both within the uppermost and the rightmost 20% of the screen. For Windows, due to a different standard location for notifications, we checked for the lower right corner of the screen. In total, the study had 117 notifications presented in standard locations. For these notifications, the median detection probability across all 33 participants was 0.83 ($\mu = 0.75$; $\hat{\sigma} = 0.29$). For notifications presented in any other location, the median detection probability was 0.78 ($\mu = 0.79$; $\hat{\sigma} = 0.09$). A two-sided Wilcoxon signed-rank test for detection probability as dependent variable was not significant (T=235; p=0.42; n=33). Therefore, our analysis does not indicate that whether a notification is presented in a standard location influences detection performance. Another hypothesis was that prior experience in using a specific OS can influence what participants perceived as the usual, standard location of a notification. To investigate this possibility, we focused on Windows users, which constituted the largest group in our participants. For these users, a total of 85 notifications were placed in the lower right corner of the screen, which is the default for Windows. The median detection probability for these notifications was 0.80 ($\mu = 0.74$, $\hat{\sigma} = 0.28$), whereas outside this area, we observed a median detection probability of 0.81 ($\mu = 0.80$, $\hat{\sigma} = 0.10$). A two-sided Wilcoxon signed-rank test comparing these conditions was not significant (T=112; p=0.45; n=23). Given the few notifications placed by chance in such default locations, it is difficult to draw any general conclusions. Results tend to indicate that placing notifications in standard locations did not have a strong influence on detection performance.

5.3 Influence of other Factors

We conducted Wilcoxon signed-rank tests to check whether the detection of a notification is connected to the additional design factors investigated in our study (see Table 1). These tests revealed that detected notifications are significantly closer to the current focus of attention, that they are significantly larger, and that their opacity is significantly higher. For a more detailed picture on how

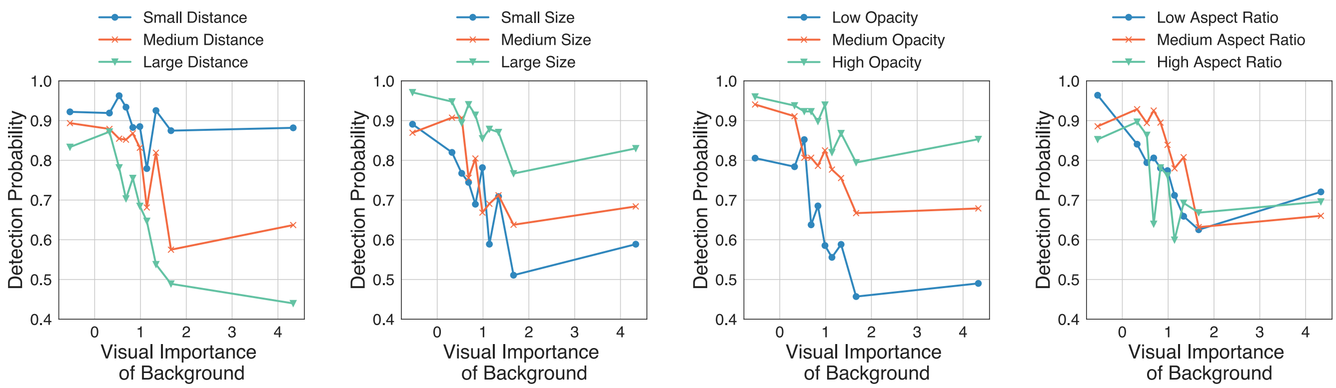

Figure 7: Interplay of the effects of visual importance of the background on the detection probability of notifications with other factors. From left to right: distance of the notification to the current focus of user attention; size of the notification; opacity of the notification; aspect ratio of the notification.

noticeability is related to these design factors, we binned the factors in the same way as we did for visual importance (see Figure 6) and computed detection probability for each bin. Only aspect ratio is an exception as due to the study design six distinct aspect ratios exist. Several factors have a clear, and generally monotonous, impact on noticeability. The distance of the notification to the current focus of attention (i.e. the distance to the moving dot) is inversely related with the detection probability. When notifications are placed very close to the current attention focus, the detection probability is above 0.9, but it decreases to below 0.7 for notifications appearing far away on the screen. A further important factor influencing users' ability to detect notifications is opacity, where low-opacity notifications only reach a detection probability of 0.5 while high opacity notifications are detected with more than 0.9 probability. The proportion of the screen covered by a notification is also positively connected to its noticeability, ranging from 0.7 detection probability for small notifications to almost 0.9 for large notifications. The relation between aspect ratio (notification width divided by notification height) and noticeability is less obvious. The medium aspect ratios appear to be most noticeable, with a peak at an aspect ratio of 35:14. This could indicate a preference for aspect ratios that are close to what users typically are confronted with in their daily interactions.

While our study was not designed to investigate the impact of the type of operating system on the noticeability of notifications, we can still make use of our data in an exploratory numerical analysis. The operating system was sampled randomly at the beginning of each session. As a result, the number of users being exposed to at least one session of a given operating system differs (n_{Linux} = 30, $n_{Windows}$ = 26, n_{Mac} = 25). The detection probability for notifications in the Mac OS ($\mu_{1/2} = 0.71$; $\mu = 0.71$; $\hat{\sigma} = 0.13$) was lower than in Linux ($\mu_{1/2} = 0.88$; $\mu = 0.84$; $\hat{\sigma} = 0.12$) or Windows ($\mu_{1/2} = 0.84$; $\mu = 0.81$; $\hat{\sigma} = 0.11$). The lower noticeability of Mac notifications might be a result of their bright colour which creates less contrast on many backgrounds.

5.4 Interplay of Visual Importance with other Factors

We analyse how visual importance of the desktop image at the notification location interacts with the other factors (Figure 7). Each plot shows the dependence between visual importance and detection probability for high, medium, or low values on the respective other factor. We partitioned the data into high, medium, and low for each of these factors by using the highest third, middle third, and lowest third of the data. In contrast to the aspect ratio of the notification, we can observe a clear interaction effect with visual importance for opacity of the notification, distance of the notification to the attention focus, and notification size. This interaction is similar for all three factors. In general, if the other factor challenges noticeability (e.g. low opacity, or large distance), the effect of visual importance on noticeability is especially strong. If the other factor makes detection easy, the effect of noticeability is less pronounced. For example, the range of probability scores resulting from different visual importance values is below 0.15 for notifications displayed at a small distance to the current focus of attention, but close larger than 0.4 for notifications displayed at a large distance from the current focus of attention.

These results indicate that visual importance could effectively be used to offset the detrimental effects on noticeability resulting from a small notification size, a notification with low opacity, or a notification that is far away from the current focus of attention.

5.5 Noticeability Maps

In order to present the effects of different factors on noticeability in an intuitive way, we propose the concept of noticeability maps (see Figure 8). These noticeability maps encode the expected likelihood of detection for notifications when presented at different locations on the desktop. To construct these maps, we first compute the detection probability for different combinations of visual importance scores and distances of the notification to the attention task based on the data observed in our study. To strike a balance between accuracy and robustness, we bin visual importance scores and distances in six bins each and compute the detection probability for each of the resulting 36 combinations. Using linear interpolation between bins, we obtain a function f mapping from visual importance and distance to the attention focus to a detection probability. To create a noticeability map for a given desktop image and attention focus, we compute for each pixel its visual importance and its distance to the attention focus. Using f we obtain an estimate of the detection probability at this pixel.

Figure 8 B and C show noticeability maps for the same desktop image but for different locations of the attention focus (indicated by the white circle). Regions of the desktop that are assigned a high estimated noticeability are at the blue end of the color spectrum. Regions with low estimated noticeability are at the red end of the color spectrum, effectively warning designers that notifications placed in these regions are unlikely to be noticed. Apart from distance and visual importance, the noticeability maps also encode the interaction between these two factors. This is visible when comparing Figure 8 B and C. Regions in the upper left show large differences in noticeability when they are far away from the attention focus (see Figure 8 B). In contrast, when these regions are close to the focus of attention, in addition to a generally higher noticeability, smaller differences are present (see Figure 8 C).

Finally, Figure 8 E and F shows the noticeability for high and low opacity notifications, respectively. Here, we restrict the data used in the computation of the noticeability map to the 50% of notifications highest or lowest in opacity. While it is much more challenging to place noticeable notification with low opacity, the noticeability map reveals several locations at which even a notification with low opacity is likely to be detected. A similar pattern can be observed when contrasting notifications of large size with those having a small size (not shown in the Figure).

6 DISCUSSION

6.1 On the Synthesis Tool

In this paper, we introduced a tool for the synthesis of realistic desktop images that we will make available as an open-source implementation upon acceptance. Our tool allowed us to conduct the first study on the noticeability of notifications displayed on *realistic* desktop images. This is a significant step over prior work that has focused on *simplified* or carefully-selected desktop images as well as *highly abstract visual representations* of notifications [24, 28]. With our tool, we could not only solve the problem of a missing dataset containing realistic desktop images, but it also provides additional benefits. In contrast to a dataset of images, our tool allows full control over the synthesis process, allowing researchers to experimentally vary the created images. This enables psychologists to e.g. study how findings obtained with simple shape- and colour stimuli [10, 48] translate to realistically looking desktop environments. Furthermore, due to the generative approach, a semantic segmentation of the images is directly available. Work on natural images suggests that this feature can be highly useful when conducting research on human attention prediction [49]. Our tool can be easily adapted to the needs of researchers by adding their own applications, icons, or desktop images. In the future, we plan to extend this tool to incorporate additional characteristics of user interactions. While desktop images are largely static, some dynamic elements exists, e.g. website banners or embedded videos. To study such interaction scenarios, we plan to extend our tool with dynamic elements. Furthermore, we also plan to adapt our tool to synthesise realistic mobile user interfaces.

6.2 On the Impact of Background on Noticeability

Further, we are first to show that visual importance of the location at which a desktop notification will be presented is a strong prior for the probability that the notification will be detected by the user. Our results show that this connection between visual importance and noticeability is inverse – higher visual importance of the background at the notification location results in lower noticeability. Uncovering the precise mechanism behind this connection is subject to future research. At this point, we can only speculate on possible explanations. A low-level explanation of this effect could be that the visual importance measure we used is correlated with visual clutter [37]. This is plausible as high visual importance is usually assigned to regions in the images containing a large degree of clutter. A large degree of visual clutter, in turn, could make it more challenging for notifications to "stand out" from their surrounding. A possible higher-level explanation could be that users do not expect notifications to appear at locations that already contain UI elements like application windows or opened websites. Such UI elements often get assigned a high visual importance, leading to the inverse connection between visual importance and noticeability.

While we analysed the interactions of visual importance with several other factors of notification design (e.g. size, opacity, or aspect ratio), our study was not designed to investigate the influence of personal characteristics. Relevant characteristics to address in future work include users commonly used operating systems, their age, and their profession. Another interesting direction for future work will be to investigate the impact of different primary tasks. Especially if tasks are associated with specific UI elements, this might impact the noticeability of notifications beyond a task-agnostic notion of visual importance.

6.3 On the Relevance for UI Designers

Our work suggests that by taking visual importance into consideration, UI designers can optimize placement of notifications for maximum noticeability. Visual importance can even serve to offset the effects of choices for other factors that are sub-optimal from a noticeability perspective. For example, Figure 7 shows the effects of visual importance on noticeability for different opacity values. Noticeability is considerably worse for low opacity notifications compared to high opacity notifications. However, when placing the notification at a location with low visual importance, a detection probability of 0.8 can be achieved. This is comparable to the detection probability of high opacity notifications placed at locations with high visual importance. Visual importance can be used in a similar fashion to offset effects of distance between the current focus of attention and notification location, as well as notification size. Thus, taking visual importance into account, designers have more possibilities to create notifications that integrate into the user interface aesthetically without sacrificing noticeability.

Using visual importance to guide notification placement requires the appearance of the desktop to be known. This is the case in two

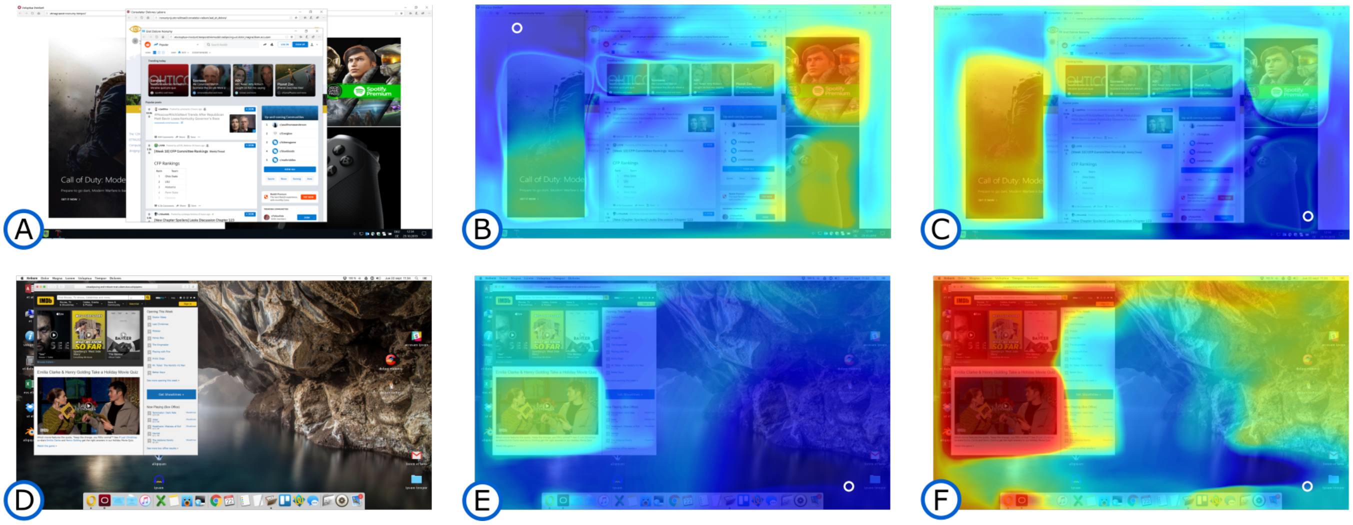

Figure 8: Examples of noticeability maps. Red indicates low-noticeability regions, blue indicates high-noticeability regions. The left column shows the effect of different locations of user attention on noticeability. A: Input image. B: Noticeability map for attention focus on the upper left of the desktop (white circle). C: Noticeability map for attention focus on the lower right of the desktop (white circle). The right column shows the effect of different opacity values on noticeability. D: Input image. E: Noticeability map for high opacity notifications and attention focus on the lower right of the desktop. F: Noticeability map for low opacity notifications and attention focus on the lower right of the desktop.

main scenarios. First, if the appearance of an application is mainly static, visual importance can guide designers in choosing where notifications should be commonly displayed in this application. Second, when the actual screen content is known at runtime, visual importance could be used to optimise notification placement dynamically. Future operating systems could either automatically place notifications of different applications, or offer access through an API to noticeability maps, giving application developers and designers fine-grained control over noticeability. Our analysis on the interaction between visual importance and the distance to the current focus of user attention indicates that by taking the users' attention into account, the quality of the proposed placement options can be increased. Users' focus of attention can be estimated either through dedicated eye tracking equipment or computational methods that e.g. analyse interactive behaviour [2, 50].

As an intuitive visualisation of the impact of different factors on noticeability for a concrete desktop image, we proposed noticeability maps (see Figure 8). These noticeability maps can be used by designers to better understand how different notification parameters play out in a concrete user interface. Designers can additionally simulate users' focus of attention in order to understand which notification placements lead to sufficient noticeability for the likely locations of user attention (e.g. a text entry field). While already useful for manual optimisation of notifications, the noticeability maps also point the way towards automatic means of optimising notifications that we are planning to explore in future work.

7 CONCLUSION

In this work we presented a novel tool to synthesize realistically looking and diverse desktop screenshots and notifications. We used these images to conduct the first study on the noticeability of desktop notifications displayed on realistically looking desktop images. We found that visual importance of the desktop at the location where a notification is placed is inversely related to its noticeability. We analysed how this effect interacts with other influences on noticeability including notification size, opacity, and distance to users' attention focus. We discuss how optimising notification placement with respect to visual importance can allow for a larger degree of freedom for design choices on other factors while still maintaining high noticeability. Finally, we introduce the concept of noticeability maps that can be used to visualise the effects of different factors on noticeability of notifications in a concrete desktop interface and thereby act as a guideline for designers. Taken together, our synthesis tool and findings open up an exciting new avenue in notification research and represent an important step towards future smart notifications that are automatically optimized for noticeability.

ACKNOWLEDGMENTS

P. Müller was funded by the German Ministry for Education and Research (BMBF), grant number 01IS20075. M. Bâce was funded by a Swiss National Science Foundation (SNSF) Early Postdoc.Mobility Fellowship. A. Bulling was funded by the European Research Council (ERC; grant agreement 801708).

REFERENCES

[1] Christoph Anderson, Judith S. Heinisch, Sandra Ohly, Klaus David, and Veljko Pejovic. 2019. The Impact of Private and Work-Related Smartphone Usage on Interruptibility. In *Adj. Proc. of the ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proc. of the ACM International Symposium on Wearable Computers*. ACM, New York, NY, USA, 1058–1063. https://doi.org/10.1145/3341162.3344845

[2] Ioannis Arapakis and Luis A Leiva. 2020. Learning efficient representations of mouse movements to predict user attention. In *Proc. of the International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, New York, NY, USA, 1309–1318. https://doi.org/10.1145/3397271.3401031

[3] Ravit Avraham Bahir, Yisrael Parmet, and Noam Tractinsky. 2019. Effects of Visual Enhancements and Delivery Time on Receptivity of Mobile Push Notifications. In *Extended Abstracts of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–6. https://doi.org/10.1145/3290607.3312993

[4] Ali Borji and Laurent Itti. 2013. State-of-the-Art in Visual Attention Modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 1 (Jan. 2013), 185–207. https://doi.org/10.1109/TPAMI.2012.89

[5] Andreas Bulling, Florian Alt, and Albrecht Schmidt. 2012. Increasing the Security of Gaze-Based Cued-Recall Graphical Passwords Using Saliency Masks. In *Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 3011–3020. https://doi.org/10.1145/2207676.2208712

[6] Georg Buscher, Edward Cutrell, and Meredith Morris. 2009. What do you see when you're surfing? Using eye tracking to predict salient regions of web pages. In *Proc. of the of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 21–30. https://doi.org/10.1145/1518701.1518705

[7] Zoya Bylinskii, Nam Wook Kim, Peter O'Donovan, Sami Alsheikh, Spandan Madan, Hanspeter Pfister, Fredo Durand, Bryan Russell, and Aaron Hertzmann. 2017. Learning Visual Importance for Graphic Designs and Data Visualizations. In *Proc. of the Annual ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 57–69. https://doi.org/10.1145/3126594.3126653

[8] Jakub Dostal, Uta Hinrichs, Per Ola Kristensson, and Aaron Quigley. 2014. SpiderEyes: Designing Attention- and Proximity-aware Collaborative Interfaces for Wall-sized Displays. In *Proc. of the ACM International Conference on Intelligent User Interfaces*. ACM, New York, NY, USA, 143–152. https://doi.org/10.1145/2557500.2557541

[9] Anja Exler, Christian Dinse, Zeynep Günes, Nadim Hammoud, Steffen Mattes, and Michael Beigl. 2017. Investigating the Perceptibility Different Notification Types on Smartphones Depending on the Smartphone Position. In *Proc. of the ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proc. of the ACM International Symposium on Wearable Computers*. ACM, New York, NY, USA, 970–976. https://doi.org/10.1145/3123024.3124560

[10] Charles L. Folk and Roger Remington. 1998. Selectivity in distraction by irrelevant featural singletons: evidence for two forms of attentional capture. *Journal of Experimental Psychology: Human perception and performance* 24, 3 (1998), 847. https://doi.org/10.1037/0096-1523.24.3.847

[11] Camilo Fosco, Vincent Casser, Amish Kumar Bedi, Peter O'Donovan, Aaron Hertzmann, and Zoya Bylinskii. 2020. Predicting Visual Importance Across Graphic Design Types. In *Proc. of the Annual ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, USA, 249–260. https://doi.org/10.1145/3379337.3415825

[12] Krzysztof Gajos and Daniel S. Weld. 2004. SUPPLE: Automatically Generating User Interfaces. In *Proc. of the ACM International Conference on Intelligent User Interfaces*. ACM, New York, NY, USA, 93–100. https://doi.org/10.1145/964442.964461

[13] D. Gao and N. Vasconcelos. 2007. Bottom-up saliency is a discriminant process. In *Proc. of the IEEE International Conference on Computer Vision*. IEEE, New York, NY, USA, 1–6. https://doi.org/10.1109/ICCV.2007.4408851

[14] S. Ghosh, L. Winston, Nishant Panchal, Philippe Kimura-Thollander, Jeff Hotnog, Douglas Cheong, G. Reyes, and Gregory D. Abowd. 2018. NotifiVR: Exploring Interruptions and Notifications in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics* 24 (2018), 1447–1456. https://doi.org/10.1109/TVCG.2018.2793698

[15] Prakhar Gupta, Shubh Gupta, Ajaykrishnan Jayagopal, Sourav Pal, and Ritwik Sinha. 2018. Saliency Prediction for Mobile User Interfaces. In *Proc. of the IEEE Winter Conference on Applications of Computer Vision*. IEEE, New York, NY, USA. https://doi.org/10.1109/WACV.2018.00171

[16] Carl Gutwin, Andy Cockburn, and Ashley Coveney. 2017. Peripheral popout: The influence of visual angle and stimulus intensity on popout effects. In *Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 208–219. https://doi.org/10.1145/3025453.3025984

[17] Eric Horvitz, Andy Jacobs, and David Hovel. 1999. Attention-Sensitive Alerting. In *Proc. of the Fifteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 305–313.

[18] Shamsi T. Iqbal and Brian P. Bailey. 2008. Effects of Intelligent Notification Management on Users and Their Tasks. In *Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 93–102. https://doi.org/10.1145/1357054.1357070

[19] Shamsi T. Iqbal and Eric Horvitz. 2007. Disruption and Recovery of Computing Tasks: Field Study, Analysis, and Directions. In *Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 677–686. https://doi.org/10.1145/1240624.1240730

[20] Shamsi T. Iqbal and Eric Horvitz. 2010. Notifications and Awareness: A Field Study of Alert Usage and Preferences. In *Proc. of the 2010 ACM Conference on Computer Supported Cooperative Work*. ACM, New York, NY, USA, 27–30. https://doi.org/10.1145/1718918.1718926

[21] Laurent Itti, Christof Koch, and Ernst Niebur. 1998. A Model of Saliency-Based Visual Attention for Rapid Scene Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 11 (Nov. 1998), 1254–1259. https://doi.org/10.1109/34.730558

[22] Lee Jones, John McClelland, Phonesavanh Thongsouksanoumane, and Audrey Girouard. 2017. Ambient Notifications with Shape Changing Circuits in Peripheral Locations. In *Proc. of the ACM International Conference on Interactive Surfaces and Spaces*. ACM, New York, NY, USA, 405–408. https://doi.org/10.1145/3132272.3132291

[23] Namwook Kim, Michelle Borkin, Krzysztof Z Gajos, Aude Oliva, and Frédo Durand. 2017. BubbleView: An Interface for Crowdsourcing Image Importance Maps and Tracking Visual Attention. *ACM Transactions on Computer-Human Interaction* 24 (2017), 36:1–36:40. https://doi.org/10.1145/3131275

[24] Michaela Klauck, Yusuke Sugano, and Andreas Bulling. 2017. Noticeable or Distractive?: A Design Space for Gaze-Contingent User Interface Notifications. In *Proc. of the ACM SIGCHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1779–1786. https://doi.org/10.1145/3027063.3053085

[25] Kostadin Kushlev, Jason Proulx, and Elizabeth W. Dunn. 2016. "Silence Your Phones": Smartphone Notifications Increase Inattention and Hyperactivity Symptoms. In *Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1011–1020. https://doi.org/10.1145/2858036.2858359

[26] J. Li, L. Su, B. Wu, J. Pang, C. Wang, Z. Wu, and Q. Huang. 2016. Webpage saliency prediction with multi-features fusion. In *Proc. of the IEEE International Conference on Image Processing*. IEEE, New York, NY, USA, 674–678. https://doi.org/10.1109/ICIP.2016.7532442

[27] Aristides Mairena, Martin Dechant, Carl Gutwin, and Andy Cockburn. 2020. A Baseline Study of Emphasis Effects in Information Visualization. In *Proc. of the Graphics Interface Conference*. ACM, New York, NY, USA, 13 pages. https://doi.org/10.20380/GI2020.33

[28] Aristides Mairena, Carl Gutwin, and Andy Cockburn. 2019. Peripheral notifications in large displays: Effects of feature combination and task interference. In *Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300870

[29] Abhinav Mehrotra, Robert Hendley, and Mirco Musolesi. 2016. PrefMiner: Mining User's Preferences for Intelligent Mobile Notification Management. In *Proc. of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, New York, NY, USA, 1223–1234. https://doi.org/10.1145/2971648.2971747

[30] Abhinav Mehrotra, Mirco Musolesi, Robert Hendley, and Veljko Pejovic. 2015. Designing Content-Driven Intelligent Notification Mechanisms for Mobile Applications. In *Proc. of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, New York, NY, USA, 813–824. https://doi.org/10.1145/2750858.2807544

[31] Abhinav Mehrotra, Veljko Pejovic, Jo Vermeulen, Robert Hendley, and Mirco Musolesi. 2016. My Phone and Me: Understanding People's Receptivity to Mobile Notifications. In *Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1021–1032. https://doi.org/10.1145/2858036.2858566

[32] Martin Pielot, Bruno Cardoso, Kleomenis Katevas, Joan Serrà, Aleksandar Matic, and Nuria Oliver. 2017. Beyond Interruptibility: Predicting Opportune Moments to Engage Mobile Phone Users. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 91 (Sept. 2017), 25 pages. https://doi.org/10.1145/3130956

[33] Martin Pielot, Karen Church, and Rodrigo de Oliveira. 2014. An In-Situ Study of Mobile Phone Notifications. In *Proc. of the International Conference on Human-Computer Interaction with Mobile Devices & Services*. ACM, New York, NY, USA, 233–242. https://doi.org/10.1145/2628363.2628364

[34] Martin Pielot, Rodrigo de Oliveira, Haewoon Kwak, and Nuria Oliver. 2014. Didn't You See My Message? Predicting Attentiveness to Mobile Instant Messages. In *Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 3319–3328. https://doi.org/10.1145/2556288.2556973

[35] Benjamin Poppinga, Wilko Heuten, and Susanne CJ Boll. 2014. Sensor-Based Identification of Opportune Moments for Triggering Notifications. *IEEE Pervasive Computing* 13 (2014), 22–29. https://doi.org/10.1109/MPRV.2014.15

[36] Vasili Ramanishka, Abir Das, Jianming Zhang, and Kate Saenko. 2017. Top-Down Visual Saliency Guided by Captions. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, New York, NY, USA, 3135–3144. https://doi.org/10.1109/CVPR.2017.334

[37] Ruth Rosenholtz, Yuanzhen Li, and Lisa Nakano. 2007. Measuring visual clutter. *Journal of vision* 7, 2 (2007), 17–17. https://doi.org/10.1167/7.2.17

[38] Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, Martin Pielot, Dominik Weber, and Albrecht Schmidt. 2014. Large-Scale Assessment of Mobile Notifications. In *Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 3055–3064. https://doi.org/10.1145/2556288.2557189

[39] Chengyao Shen, Xun Huang, and Qi Zhao. 2015. Predicting Eye Fixations on Webpage Withan Ensemble of Early Features and High-Level Representations fromDeep Network. *IEEE Transactions on Multimedia* 17 (11 2015), 1–1. https:

//doi.org/10.1109/TMM.2015.2483370

[40] Chengyao Shen and Qi Zhao. 2014. Webpage Saliency. In *Proc. of the European Conference on Computer Vision*. Springer International Publishing, Switzerland, 33–46. https://doi.org/10.1007/978-3-319-10584-0_3

[41] Alireza Sahami Shirazi and Niels Henze. 2015. Assessment of Notifications on Smartwatches. In *Proc. of the ACM International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*. ACM, New York, NY, USA, 1111–1116. https://doi.org/10.1145/2786567.2794338

[42] Jeremiah Still and C.M. Masciocchi. 2010. A saliency model predicts fixations in web interfaces. *CEUR Workshop Proceedings* 617 (01 2010), 25–28.

[43] Kashyap Todi, Jussi Jokinen, Kris Luyten, and Antti Oulasvirta. 2018. Familiarisation: Restructuring Layouts with Visual Learning Models. In *Proc. of the ACM International Conference on Intelligent User Interfaces*. ACM, New York, NY, USA, 547–558. https://doi.org/10.1145/3172944.3172949

[44] Niels van Berkel, Anja Exler, Martin Gjoreski, Tine Kolenik, Tadashi Okoshi, Veljko Pejovic, Aku Visuri, and Alexandra Voit. 2020. UbiTtention 2020: 5th International Workshop on Smart & Ambient Notification and Attention Management. In *Adj. Proc. of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, New York, NY, USA, 431—-435. https://doi.org/10.1145/3410530.3414616

[45] Alexandra Voit, Dominik Weber, and Niels Henze. 2018. Qualitative Investigation of Multi-Device Notifications. In *Proc. of the ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*. ACM, New York, NY, USA, 1263–1270. https://doi.org/10.1145/3267305.3274117

[46] Dominik Weber, Sven Mayer, Alexandra Voit, Rodrigo Ventura Fierro, and Niels Henze. 2016. Design Guidelines for Notifications on Smart TVs. In *Proc. of the ACM International Conference on Interactive Experiences for TV and Online Video*. ACM, New York, NY, USA, 13–24. https://doi.org/10.1145/2932206.2932212

[47] Dominik Weber, Alireza Sahami Shirazi, and Niels Henze. 2015. Towards Smart Notifications Using Research in the Large. In *Proc. of the ACM International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*. ACM, New York, NY, USA, 1117–1122. https://doi.org/10.1145/2786567.2794334

[48] Dirk Wentura, Philipp Müller, and Klaus Rothermund. 2014. Attentional capture by evaluative stimuli: Gain-and loss-connoting colors boost the additional-singleton effect. *Psychonomic Bulletin & Review* 21, 3 (2014), 701–707. https://doi.org/10.3758/s13423-013-0531-z

[49] Juan Xu, Ming Jiang, Shuo Wang, Mohan S Kankanhalli, and Qi Zhao. 2014. Predicting human gaze beyond pixels. *Journal of vision* 14, 1 (2014), 28–28. https://doi.org/10.1167/14.1.28

[50] Pingmei Xu, Yusuke Sugano, and Andreas Bulling. 2016. Spatio-Temporal Modeling and Prediction of Visual Attention in Graphical User Interfaces. In *Proc. of the ACM SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 3299–3310. https://doi.org/10.1145/2858036.2858479

[51] J. Yang and M. Yang. 2017. Top-Down Visual Saliency via Joint CRF and Dictionary Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 3 (2017), 576–588. https://doi.org/10.1109/TPAMI.2016.2547384

[52] Quanlong Zheng, Jianbo Jiao, Ying Cao, and Rynson Lau. 2018. *Task-Driven Webpage Saliency*. Springer International Publishing, Switzerland, 300–316. https://doi.org/10.1007/978-3-030-01264-9_18