

The Design Space of Nonvisual Word Completion

Hugo Nicolau¹, André Rodrigues², André Santos², Tiago Guerreiro², Kyle Montague³, João Guerreiro¹

¹ INESC-ID, Instituto Superior Técnico, Universidade de Lisboa

² LASIGE, Faculdade de Ciências, Universidade de Lisboa

³ Open Lab, Newcastle University

hman@inesc-id.pt, afrodrigues@fc.ul.pt, arbsantos@fc.ul.pt, tjvg@di.fc.ul.pt,
kyle.montague@ncl.ac.uk, jpvguerreiro@gmail.com

ABSTRACT

Word completion interfaces are ubiquitously available in mobile virtual keyboards; however, there is no prior research on how to design these interfaces for screen reader users. In addressing this, we propose a design space for nonvisual representation of word completions. The design space covers seven categories aiming to identify challenges and opportunities for interaction design in an unexplored research topic. It is intended to guide the design of novel interaction techniques, serving as a framework for researchers and practitioners working on nonvisual word completion. To demonstrate its potential, we engaged blind users in an exploration of the design space, to create their own bespoke word completion solutions. Through this study we found that users create alternative interfaces that extended current screen readers' capabilities. Resulting interfaces are less conservative than mainstream solutions on notification frequency and cardinality. Customization decisions were based on perceived benefits/costs and varied depending on multiple factors such as users' perceived prediction accuracy, potential keystroke gains, and situational restrictions.

Author Keywords

Design space; Blind; Mobile; Text entry; Touchscreen; Word completion; Word prediction; Screen reader

CSS Concepts

• Human-centered computing → Accessibility → Empirical studies in accessibility

INTRODUCTION

Text entry is one of the most common tasks on smartphones, vital to browsing the web, sending emails, messaging or using social networks: it is unavoidable. Mobile keyboards often present word completion

Paste the appropriate copyright/license statement here. ACM now supports three different publication options:

- ACM copyright: ACM holds the copyright on the work. This is the historical approach.
- License: The author(s) retain copyright, but ACM receives an exclusive publication license.
- Open Access: The author(s) wish to pay for the work to be open access. The additional fee must be paid to ACM.

This text field is large enough to hold the appropriate release statement assuming it is single-spaced in Times New Roman 8-point font. Please do not change or modify the size of this text box.

Each submission will be assigned a DOI string to be included here.

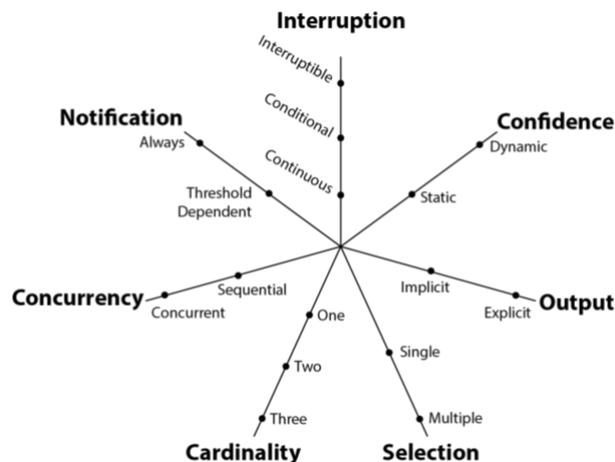


Figure 1. The design space of nonvisual word completion.

(suggestions) candidates that predict users' intended word as they type. These suggestions can potentially save keystrokes as they are always visible and displayed near the typing area. While sighted users can quickly scan the display for input feedback, content changes, and suggestion updates, blind people interact with touchscreen mobile devices in inherently distinct ways due to the one-dimensional and ephemeral nature of auditory feedback.

Although there is a large amount of work done on word completion interfaces [2,7,20,25,35,41], there is no prior research into how to design these interfaces for screen reader users. To fill this gap, we propose the first design space for nonvisual representation of word suggestions. The aim is to identify opportunities for future interaction designs, guide in the creation of novel interfaces, and spur research on the field.

Our design space for nonvisual representation of word completion covers a taxonomy of properties within seven categories: *notification*, *output*, *confidence*, *cardinality*, *concurrency*, *interruption*, and *selection*. In a first step, we analyze and deconstruct the typing process of blind users highlighting key challenges and opportunities that arise from the interaction between screen readers' capabilities and word completion systems. Such analysis served as a framework to build our design space. We detail the design space by describing each category and their possible

instantiations. Additionally, we discuss three different usages of the design space: (1) to analyze existing techniques and gaps in the literature, (2) to design innovative nonvisual representations of word completion suggestions and identify basic interaction possibilities, and (3) as a support tool that provides the building blocks for participatory design activities.

With the design space we offer an approach and perspective for designers, researchers, and practitioners to explore potential techniques arising from the combination of using screen readers and word completion systems. We strive to inspire readers to build upon the presented design space, aiming to help uncover opportunities to improve on existing techniques and generate novel solutions.

The contributions of this paper are two-fold: first, a design space for nonvisual representation of word completion suggestions, offering a new approach on how to think about this unexplored topic. Second, we show the potential of the design space to spur innovation by engaging screen reader users in the design of novel interfaces. Emerged solutions highlight not only the need for alternative word completion interfaces, but also current limitations of mobile screen readers. In terms of resulting interfaces, there were four main designs where participants consistently opted for more frequent notifications than mainstream interfaces. Customization choices were often dependent on personal and contextual factors such as perceived prediction accuracy, potential keystroke savings, cognitive demand, and situational restrictions.

RELATED WORK

We discuss related work in three fields of research: text input for blind people, word completion, design spaces.

Text Input for Blind People

Most smartphones already support nonvisual text input via built-in screen readers such as iOS VoiceOver or Android Talkback. Touchscreen screen readers enable an Explore by Touch approach by allowing users to drag their finger on the screen and having user interface elements (e.g. keys) read aloud as they touch them [22]. Although screen readers are effective in providing access to virtual keyboards, blind users still present significantly slower entry rates than their sighted counterparts [31]. While sighted users achieve mean entry rates of about 40 words per minute (WPM) [46], reported results for blind users are 4-5 WPM [4,31].

To address the mismatch in typing performance there have been many efforts to improve mobile text input for blind people [8,19,33,44]. Previous research has proposed alternative keyboards that leverage gestural interaction [19,44] and multitouch interaction [8], but with limited success in improving typing speed. Following the same approach of alternative keyboards, many Braille-inspired techniques have been presented over the last decade [4,27,34,39,40]. BrailleTouch [39] and Perkinput [4] were particularly successful in improving entry rates, with the

most proficient users reaching 32 and 22 WPM, respectively. Both techniques leverage multitouch capabilities of current touchscreen devices and allow users to type Braille characters by directly entering chords on the screen. These can be complemented with chord-based correction systems that reduce the number of errors [29].

Overall, much work has been done in the field of nonvisual text entry from understanding the fundamental challenges of interacting with touch-based screen readers [22,31,37,38] to novel keyboard designs [4,39] and speech input [3]. Despite the large amount of work done [4,8,19,33,39,40], nonvisual word completion interfaces remain unexplored.

Word Completion

Most virtual keyboards make use of word completion through a suggestion bar (Figure 2). For instance, the Android operating system presents three suggestions above the virtual keyboard while users are typing. When one of the word suggestions is above a confidence threshold, it turns bold: tapping on the space bar automatically accepts the suggestion and enters a blank space, i.e. auto-complete. Alternatively, users can tap one of the remaining suggestions at any given time or choose to ignore them. To undo the auto-completion action, they can backspace.

Screen reader users have a significantly different experience with word completion systems. Due to the inherently one-dimensional and ephemeral nature of auditory feedback, suggestion updates are given sequentially. In Android (v.8 'Oreo' + Gboard), the screen reader only reads aloud the auto-complete word (i.e. above a confidence threshold). Similarly, in iOS, VoiceOver updates are also restricted to the auto-complete word; however, notifications are given via earcon ('beep'), forcing users to explore the suggestion bar to get access to the suggested word.

Word completion systems aim to reduce the number of keystrokes needed from users to enter an intended word [20,41,42]. These have shown to help sighted people enter text more quickly and accurately. Word completion has shown to be useful for users with motor impairments; however, presenting suggestions can impose cognitive and motor costs that sometimes outweigh their benefits [20,25]. With the appropriate configurations, a system can offer both word completions and corrections for typing errors. Bi et al. [7] demonstrated that it is possible to simultaneously optimize a keyboard for both goals with correction accuracy rates of 8.3% and completion power of 17.7%.

Suggestions beyond word-level to sentence-level have also been investigated. Bridge and Healy [9] proposed GhostWriter-2.0, which supports users writing product reviews by suggesting short sentences mined from other reviews. Arnold et al. [2] investigated the use of phrase suggestions in composition tasks. The authors were particularly interested in having a system that provided valuable suggestions rather than just accurate predictions.

Sentence-level prediction was also used in other applications such as language translation [15], email responses [24], copy and paste tasks [45], while fixing typing errors [1] or as an AAC solution [23].

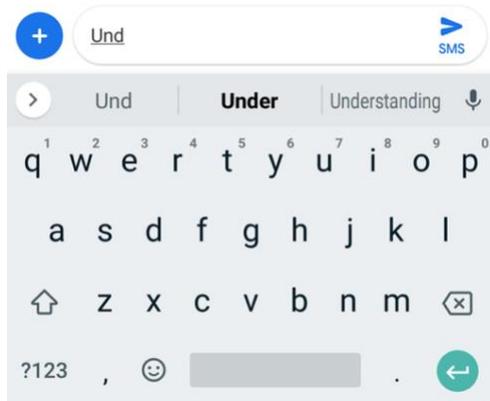
Quinn investigated the effect of visually presenting suggestions, demonstrating a trade-off between keystroke savings and typing speed, which was related to cognitive load [35]. Such cognitive load, introduced by word completion systems, has also been observed in a longitudinal study [12]. Despite its ubiquity, word completion has received little attention when used with screen readers. As nonvisual word completion is still largely unexplored, we outline a design space featuring the categories that can be explored by interface designers.

Design Spaces

Design spaces have been used in the field of human-computer interaction to understand and explore the potential of multiple technologies: from input devices [11,14] and smartphones [5] to shape-changing interfaces [26] and 3D printable interactivity [6].

Almost three decades ago, Foley et al. [14] showed that taxonomies are a useful way to organize knowledge about input devices and interaction techniques. Later, Card et al. [11] extended this work and proposed a design space to systematize the huge variety of input devices arising at that time. More recently, Kwak et al. [26] report on a design space and elicitation study for shape-changing interfaces. Ballagas et al. [6] survey the state of the art in 3D printing and proposed a design space in the form of a multidimensional box known as Zwicky box [47]. Another example is the work of Hirzle et al. [21], which provides a design space arising from the combination of head-mounted displays and 3D gaze.

An observant reader may note that previous research on design spaces emerged from a need to structure existing solutions. This approach is geared towards identifying gaps in the literature and families of successful solutions.



Although our design space can be used in such a way, the lack of previous literature presents a major challenge. Thus, our design space was mainly built to inspire others [28] and spur research in nonvisual word completion interfaces by offering a new approach to ideate interaction possibilities.

THE DESIGN SPACE

In the following, we analyze the typing process of screen reader users and identify the unique challenges that emerge from that experience. This analysis served as an underpinning framework to build a structured design space for nonvisual word completion interfaces. We then describe the design space's categories and values in more detail.

Deconstructing the Typing Process

Current word completion interfaces are not designed to support nonvisual interaction. These solutions rely on the users' ability to quickly glance at a word completion prompt while typing. Although suggestions are visually displayed and accessible on the screen, blind users may not be aware of the available completion options.

Figure 2 illustrates an example of a user typing the word "understanding" along with word completion updates. After typing the first three characters – "und" – there is a suggestion that meets the probabilistic confidence threshold of being the intended word. The user is *notified* visually as the suggestion turns bold. Apple's VoiceOver provides a simple earcon notification while Google's Talkback reads the word aloud via speech *output*. Unfortunately, the suggestion is not the intended word. Recent studies have shown that blind users spend a large amount of time correcting errors [30], thus accepting the wrong suggestion can be particularly damaging. This is especially relevant to VoiceOver users that over rely on the word completion system and can accept the suggestion (via space bar) without hearing it first.

On the other hand, screen reader users may ignore early suggestions by continuing typing, thus not benefiting from word completion. In Figure 2, the intended word is already

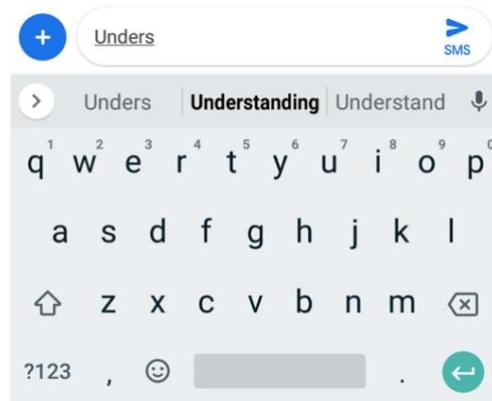


Figure 2. Example of word completion interface on Android 8.1 with GBoard. The user is typing the word "Understanding". Left – the user has type "Und", the auto-complete suggestion is "Under" with two less probable alternative suggestions of "Und" and "Understanding". Right - the user has type "Unders", the auto-complete suggestion is "Understanding" with two less probable alternative suggestions of "Unders" and "Understand".

available in the suggestion bar, just after three keystrokes. While sighted users have instant access to three suggestions, blind users are notified about, at most, a *single word*. Because auditory feedback is inherently *sequential*, having access to lower *confidence* suggestions means to intentionally stop the typing process and engage in a screen exploration task to *select* the suggestion. Moreover, users would perform this exploration without any guarantees of finding the intended word.

After six keystrokes (Figure 2 - right), there is a notification of the intended word – “understanding”. Ideally, the user hears the notification and enters a blank space to auto-complete the word, saving 7 keystrokes. However, if s/he continues to type, the suggestion output is *interrupted* by input feedback. This ephemeral nature of auditory feedback may result in users missing relevant notifications.

Categories and Values

Based on the previous analysis, we propose a taxonomy of properties relevant to non-visually interacting with word suggestions in text entry tasks. Although it is impossible to prove that taxonomies are complete – as technology evolves, so should the taxonomies of properties – the resulting design space aims to make researchers and designers aware of and help them to address challenges of future word completion interfaces for screen reader users. Our taxonomy includes seven categories: *notification*, *output*, *confidence*, *cardinality*, *concurrency*, *interruption*, and *selection*.

Notification. The notification category indicates when to notify users of word completion suggestions. Mainstream screen readers notify users when a suggestion has high probability of being the intended word, making them *threshold-dependent* notifications. Older versions of Talkback never notified users, forcing them to interact with the screen to access suggestions (*input-dependent*). On the other end of the spectrum, we may have a solution that *always* notifies users of word suggestions, emulating visual updates. *Threshold-dependent* designs can also resort to notification behaviors that are related with typing profile. For instance, notifications can occur when there is a keystroke gain of selecting the suggestion based on users’ typing speed or a fixed amount of characters.

Output. The output category indicates how users are notified of new suggestions. Output can be either *implicit* or *explicit*. VoiceOver uses *implicit* output since users are notified through an earcon, without knowing what the auto-complete suggestion is. Users can then explore the screen to access the suggestion or simply accept it and trust that the suggestion matches the intended word. On the other hand, Talkback reads aloud the auto-complete suggestion, making it *explicitly* visible. Both implicit and explicit can leverage multiple output modalities such as *spearcons* [43], *vibrotactile*, *haptics*, and *Braille displays*.

Confidence representation. The confidence category indicates whether confidence representation of suggestions is *static* or *dynamic*. Current representations are static, whereby no matter the level of confidence of the word prediction, the feedback is identical - contemporary screen readers behave similarly. In dynamic representations, the feedback is modified based on the level of confidence for the word completion. This approach can have multiple benefits, e.g., increasing volume can make users more aware of a suggestion that is a strong candidate; conversely, suggestions “whispered” can be less disruptive. Moreover, one can adjust other sound features, such as pitch or speed, to achieve similar results. Comparable behaviors can be mirrored in other modalities such as haptics.

Cardinality. The cardinality category indicates how many word completion suggestions are presented non-visually. Although there may be many suggestions visible on the screen, the cardinality category specifically indicates whether users are updated of *single* or *multiple* suggestions while typing. For instance, in the Google keyboard there are three available suggestions visible most of the time; still, Talkback only presents a single suggestion via speech feedback if above a confidence threshold.

Concurrency. The concurrency category indicates whether multiple suggestions are presented *sequentially* or *concurrently*. Screen readers’ output is sequential or one-dimensional. However, as visual representations of word completions, one can imagine using concurrent feedback to convey multiple suggestions [17,18]. Concurrency can take the form of binaural (left-right ear) or fully spatialized (3D) feedback. In both cases, it requires the use of headphones or even specialized hardware (e.g. head tracking technologies) to achieve the desired effect. Concurrency can also take place using multiple output modalities (e.g. audio and Braille displays).

Interruption. The interruption category indicates whether word completion feedback is *interruptible* (e.g. screen reader stops reading word suggestions as the user touches a new key) or *continuous*. Current screen readers are interruptible; if users touch a key when a suggestion is being read aloud the feedback is interrupted, as the reader assumes users want to continue typing without hearing the suggestion. Depending on typing speed, this can result in appropriate suggestions never being heard. However, one may wish that suggestions are continuously rendered along with input feedback. Continuous feedback should update when there are new notifications. It can be always active or based on a *condition*/threshold. For example, screen readers can keep reading suggestions that are above the auto-complete threshold, even when a new key is focused. It is important to highlight that interfaces should always be responsive to users’ input [32].

Selection shortcut. We assume that the suggestion list should be always available to select from using the de-facto method, e.g. explore by touch. Complementary to this, there

can be shortcuts to select suggestions, namely the most probable suggestion. This category represents the expressiveness of the shortcuts and how many suggestions from the list can be accessed from a direct action: *single* or *multiple*. While current screen readers only allow for selection of the most probable word suggestion via space key, it has been demonstrated to be useful to provide immediate access to additional less probable suggestions (e.g. motor impaired [13]). Selection shortcut mechanisms may include speech input, 2D/3D gestures, physical buttons, or novel keyboard layouts.

Usage of the Design Space

We propose three ways in which the design space can be leveraged by designers, practitioners, and researchers.

Views on the design space. The resulting design space can be filled with existing solutions, making it possible to visually express the main parameters of word completion representations as a 7-axis radar chart. Figure 3 shows an example of (the few) existing approaches, i.e. Talkback and VoiceOver. The ordinal values are mapped to these axes, such that the most informative values are placed further away from the center of the radar chart and the most restrictive are placed near the center of the chart. It is worth highlighting that as we are only using ordinal values, when comparing interfaces, the exact position on the axes is not as important as the relative position, thus these values are simply distributed evenly along the axes. Following these rules, when visualizing a word completion interface, any representation that entirely contains others, means it is more informative. However, it does not necessarily imply it is a better solution. Views of the design space can also be done using a table (e.g. Table 1), which help to identify promising families of solutions, as well as possible lack of techniques by the quantity of solutions in each category.

Ideation over the design space. Using the design space as an ideation tool is directed for those that aim to derive new interaction techniques and nonvisual representations of word completions. Each of the design categories and values can be used as the building blocks for novel technical solutions. Alternatively, and inspired by the work of Card et al.'s [11] operators within design spaces, one can choose to draw inspiration from existing solutions and manipulate them. Here, we start with a concrete set of values in each category and then replace one (or more) to originate a new word completion representation. An example would be to start with Android's solution and replace the *cardinality* parameter, enabling users to sequentially hear more than one suggestion. Furthermore, we could extend such technique by changing the *concurrency* parameter and allow simultaneous feedback. Moreover, each suggestion could be mapped to a position in space, originating tridimensional auditory *output*. We can imagine another technique where the output could still be tridimensional but sequential rather than simultaneous. These operations illustrate the potential of the design space to generate novel

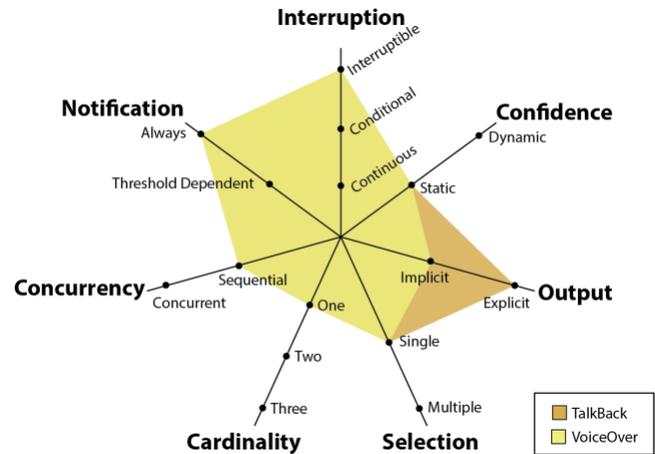


Figure 3. Illustration of both Talkback's and VoiceOver's nonvisual word completion interface.

techniques for nonvisual representations of word suggestions based on simple manipulations of its building blocks. Such manipulations can be extended to other operators using, for example, Raskar's idea hexagon [36] (e.g. generalization, fusion, opposite).

Elicitation of bespoke solutions. Our final usage of the design space, and less common in the literature, is to serve as a tool for participatory design activities. The set of categories and values can be used as "materials" when engaging users in design. These "materials" can then be easily mixed and manipulated to create and customize interfaces. The design space serves as an overall framework to expose possibilities, guide the design process, and collect feedback about bespoke solutions. In the next section, we present an example of such participatory activity with screen reader users.

EXPLORATION OF THE DESIGN SPACE

To demonstrate the potential of our design space, we took a participatory approach to design novel nonvisual word completion interfaces. We engaged screen reader users in exploring the design space enabling them to build personalized word completion representations. Such work yields a valuable contribution given the lack of knowledge regarding the expectations and needs of blind users towards mobile word completion interfaces. We aim to answer two main research questions: (1) Do screen reader users value alternative word completion interfaces? (2) What is the rationale for their personalization decisions?

Participants

We recruited 11 legally blind participants, 7 males, from a local training institution for visually impaired people. Participants' age ranged from 38 to 57 (M=45, SD=7). All participants have owned a smartphone for between 1 and 4 years and required a screen reader to interact with it. Although participants reported performing text entry tasks daily with their devices, only 4 reported using word completion. Those who did not use it, said it was either

because they did not know this feature existed (N=4) or did not see the benefit of using it (N=3).

Customizable Prototype

We modified Google's Android Open Source Project (AOSP) Keyboard¹, enabling us to augment the stock input method of most Android devices. The modifications made no layout or visual changes to the keyboard, i.e. target mapping remained unchanged. Instead, we augmented its settings capabilities to enable customization of word completion representations. Touch interaction also remained unchanged as participants could drag their finger on the keyboard to have keys read aloud and lift them to insert a character. They could also select suggestions from the top of the keyboard by focusing them and then double tapping. We also relied on the AOSP dictionary to retrieve the top three suggestions. The keyboard was developed to be customized and enable users to experience a set of attributes of the design space. To avoid fatiguing participants, we restricted the session length to 90 minutes. This timeframe allowed us to explore in considerable depth five out of the seven design categories (*notification*, *confidence*, *cardinality*, *concurrency*, and *interruption*).

Notification. Participants could choose when to be notified of word completion suggestions: *always* or based on the spellchecker *confidence*.

Confidence representation. The confidence levels of suggestions could be given through *static* or *dynamic* feedback; in the dynamic condition, volume was linearly mapped (0 to 100%) to spellchecker confidence.

Cardinality. Participants could receive *single* or *multiple* word completion suggestions. In the multiple condition, participants experienced up to three suggestions, emulating the visual interface.

Concurrency. Suggestions could be experienced either *sequentially* or *concurrently*. For concurrent feedback, we relied on Amazon Polly² to generate the audio with different voices for each suggestion [16,17]. Furthermore, the audio sources were set in different positions in space using the Spatial Audio API³ with a 250ms delay between words to help improve speech intelligibility [10]. Participants could decide the position in space for each suggestion based on its confidence value (i.e. right, center or left).

Interruption. Notification could be *interrupted* via touch input or *continuously* read. Additionally, participants could combine options by interrupting notifications only when they were below the auto-complete threshold; otherwise,

the suggestion would not be interrupted (herein referred to as conditional interrupt). Input feedback was rendered as mono audio using the screen reader's voice.

We did not explore variations in the following categories:

Output. Pilots studies showed that *explicit* feedback (speech) was preferred over implicit output (earcons). Thus, for this participatory design activity, users were notified via speech feedback (for auto-complete) as in Talkback. We solely used auditory feedback as this is still the most convenient and common output modality of screen readers. Alternative interface designs could leverage haptic devices such as refreshable Braille displays.

Selection shortcuts. The suggestion list was available on the top of the keyboard and could be selected through explore-by-touch. Tapping on the space bar selected a single (the most probable) auto-complete word. We did not explore shortcut techniques for multiple selection as we wanted to focus on novel word completion representations rather than new input techniques. However, we did invite participants to make suggestions on new selection interfaces.

Apparatus

We used the customizable prototype, previously described, running on a Xiaomi Redmi 3. The mobile device features a 5-inch capacitive touchscreen, running Android 7.1.2. All audio feedback was given either through Android Talkback (female voice) for default text input interactions or Amazon Polly (one male and two female voices) for concurrent word completion feedback. Participants were requested to use an Ozone Onda ST headset. All participants' touchscreen actions were logged through our application.

Procedure

At the beginning of the session, participants were informed they would be exploring a variety of word completion interfaces to find the one that best suited them. After completing a demographics and smartphone usage questionnaire, participants were asked to write sentences while exploring different values of the design space.

First, we described how the current default word completion interfaces work. Next, participants were asked to write a sentence with the default keyboard and Talkback behaviors to get familiarized with the device. We then prompted them to start exploring the design space. Namely, participants could personalize: *notification*, *confidence*, *cardinality*, *concurrency*, and *interruption* categories.

We suggested a set of sentences randomly selected from a corpus representative of the language (.98 correlation with the language character frequency). However, participants were also able to type freely. They were encouraged to use the same sentences to test each interface design.

For each personalized interface, participants were asked to share comments, (dis)likes, and improvements. We then invited the participants to experience interfaces that were theoretically more informative but also more cognitively

¹ <https://source.android.com>, accessed 4th April 2019.

² <https://aws.amazon.com/polly/>, accessed 4th April 2019.

³ <https://developers.google.com/vr/reference/ios-ndk/group/audio>, accessed 4th April 2019.

	Notification	Confidence	Cardinality	Concurrency	Interruption
P1	Always	Dynamic	3	Concurrent	Interruptible
P2	Always	Static	3	Sequential	Interruptible
P3	Always	Static	2	Concurrent	Interruptible
P4	Always	Static	1	-	Interruptible
P5	Always	Static	1	-	Interruptible
P6	Always	Static	2	Concurrent	Interruptible
P7	Always	Static	2	Sequential	Interruptible
P8	Always	Static	2	Sequential	Interruptible
P9	Always	Dynamic	3	Sequential	Conditional Interrupt
P10	Always	Dynamic	3	Sequential	Continuous
P11	Always	Static	3	Sequential	Interruptible

Table 1. Participants' preferences that resulted from participatory design session. The Concurrency dimension is not applicable to P4's and P5's interfaces as there was only a single word suggestion (no concurrent feedback).

demanding [16]. For more demanding styles, we increased the cardinality and/or concurrency.

All participants started by experiencing just an earcon when a suggestion was flagged to be auto-completed. Next, participants tried speech feedback on auto-complete as in Talkback. Afterwards, they experienced notifications after each key tap. Next, they tried concurrent speech with two suggestions. From this point onward, participants led the exploration of the design space based on their preferences and feedback. We guaranteed that participants engaged (at least once) with: one suggestion, multiple suggestions, dynamic confidence representation, sequential feedback, concurrent feedback, and the different interruption methods. The participatory process was similar to an optometry appointment where variations were presented back and forth until both participant and researcher were confident in the selection. The number of interfaces experienced by participants ranged from 6 to 21 ($M=10$, $SD=4.8$), mostly depending on whether they saw benefits in having access to multiple word suggestions, which would mean more interface possibilities.

Lastly, in the debriefing stage, participants were asked about their preferred combination of attributes and the rationale for their selection. We finished with an open-ended question where participants were encouraged to comment on the experience: *“what are your main thoughts on this experience with word completion systems?”* The total time for the session was between 60 and 90 minutes. Participants were compensated for their time with a 20€ gift card.

Data Analysis

We used thematic analysis to analyze participants' comments about their preferences following a grounded approach. One researcher analyzed the sessions' transcripts using an iterative and inductive coding process with open and axial coding to identify emergent themes in the data from the participatory design session.

FINDINGS

We examine participants' preferences (Table 1) and personalization process when engaged in exploring the design space.

Always provide feedback. Current nonvisual word completion interfaces notify users based on a confidence threshold – that is, they only receive word completion feedback when there is enough confidence that the suggested word matches the intended word. Screen readers take a conservative approach when notifying users at the cost of delaying valuable suggestions. Interestingly, all participants ($N=11$) preferred to receive suggestions after each character entered as long as they could interrupt the speech when starting to type a new character (P10 was the exception with continuous feedback).

“It [the interface] is good because it suggests the words I want to type. Therefore, if we are at the beginning of the word it is much easier [to complete]” (P7).

It is worth highlighting that participants started by trying the default interface behavior. The rationale for choosing more frequent updates was related with participants' perception of the benefits of having access to the intended word just after entering a few characters. Nevertheless, they were conscious about the potential costs and benefits of their choices, as illustrated by P1's and P10's comment:

“Not for short words, but for longer words we could save some characters” (P1). “In the first letter I don't pay much attention [to suggestions], in the second letter I stop a little bit [to listen], after the third I start paying much more attention” (P10).

This comment suggests that participants are aware of the limitations of word predictions and devise strategies to cope with its interaction costs depending on the number of entered characters.

Enable continuous typing. Participants reinforced the importance of allowing them to keep typing despite word completion feedback. Screen readers should be reactive to users' actions and give them the choice of hearing the suggestion and/or continuing to type. While nine showed a

preference for interruptible notifications, two participants preferred alternative interruption behaviors, which allowed them to explore the keyboard while listening to word suggestions. P9 chose continuous feedback just when suggestions were above the auto-complete threshold, while P10 did not want touch actions to interrupt notifications.

“I don’t mind listening to all suggestions, because I can keep typing. Sometimes I just ignore the last suggestions. I just continue to type” (P10).

Confidence representation is related to perceptual cost. After experiencing dynamic and static representations of confidence, most participants (N=8) preferred static feedback as illustrated by P6:

“They [the suggestions] are all at the same volume level, this way I can easily understand them” (P6).

Most participants valued all word completion suggestions. In the dynamic condition, suggestions with lower confidence levels were harder to hear. Participants commented that if they did not want to hear them, they could continue typing. Otherwise, they would like to clearly hear all feedback.

On the other hand, three participants chose the dynamic representation. Interestingly, these participants chose to hear three suggestions, indicating that more demanding interfaces may require dynamic feedback to alleviate cognitive load.

“I liked the dynamic feedback. Overall, it makes the last word less audible when is not related with what I wrote. That helps, it is less entropy” (P10). “Having three voices is interesting if I can associate different volume levels to the most correct” (P9).

Interface designers may consider other features to represent confidence (e.g. reading speed). It is worth noticing that participants had an indirect encoding of confidence as they positioned suggestions in space from left to right, even with sequential feedback, left being the top suggestion.

Number of suggestions is highly user-dependent. Preferred cardinality showed to be highly dependent on participants’ abilities and perceived benefit. Participants expressed tensions between potential benefits and ease of use. Two participants preferred just one suggestion. The main reason was that it was easier to discriminate between touch feedback and word completion suggestions:

“I prefer a single suggestion because then I have a voice giving suggestions and the other [voice] reading the keys I touch. With more suggestions I get confused and lose track of what I wrote.” (P5). “Having one suggestion makes it easier to understand it” (P4).

P5 felt that hearing more than one suggestion was cognitively demanding and would decrease her input performance. On the other hand, most participants (N=9) preferred multiple word completion suggestions:

“For those less experienced, maybe having a single suggestion is a good choice; however, there’s the risk of having to type the whole word to get the suggestion” (P9). “Three suggestions have benefits; I wouldn’t need to explore [the screen], I could decide [whether to select a suggestion]” (P1).

Although participants’ choices may be biased by a novelty effect, P1 predicted that having the same information as sighted users would allow him to make informed decisions on when to select word suggestions. It is worth highlighting that P1 already uses word completion with his smartphone. Thus, expertise may be playing a significant role. Other participants (N=4) felt that more than two suggestions read aloud was attentionally demanding:

“With three suggestions it is more complicated, we need to be super attentive. It is confusing, two suggestions are enough” (P6).

Word discrimination is the biggest drawback in concurrent suggestions. From the nine participants that chose multiple suggestions, six preferred sequential feedback while only three preferred concurrent feedback. Participants felt that having multiple suggestions read simultaneously decreased their ability to discriminate words, which could result in missing accurate word predictions:

“When reading aloud [suggestions] at the same time, I can’t hear properly, I can’t understand what they are saying” (P10).

Although we followed guidelines for concurrent feedback [16], we believe the similarity between word completions was an important factor. Even P1, who preferred concurrent feedback, highlighted the advantage of sequential feedback:

“Having the words read aloud in sequence is ‘safer’. I don’t think it happened - not understanding a word -, but I am not sure” (P1).

Context. Context of use was raised by some participants as an important factor. Three participants mentioned that the interface should be personalized to the situation, particularly if people are not using earphones:

“For this to work [concurrent feedback], we need to wear earphones. It would be good if it [the smartphone] gave me three suggestions in sequence when I’m not using earphones” (P1).

P4 goes even further and states that multiple suggestions are worthless in mobile contexts:

“We use our smartphones outside where there are cars and noise. Two suggestions don’t work. [Only] when we are seated and in a calm environment, but that’s not where we use our phones. One suggestion works well.” (P4)

It is clear that P4 is well aware of how environmental factors influence cognitive load and its effect on mobile

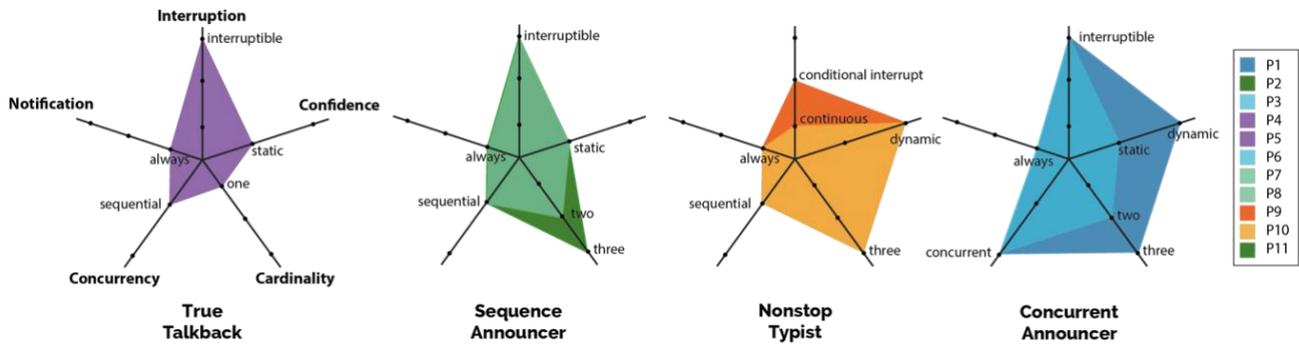


Figure 4. Interface designs that resulted from the participatory session overlaid on five categories of the design space.

usage. Moreover, he firmly believes that no one would be able to keep track of multiple suggestions unless they are in a controlled and calm environment. Such concerns highlight the need for user- and context-dependent interfaces. For example, P10 would like to personalize interface cardinality.

“If I was outside, I would prefer to hear just two suggestions. The third suggestion is usually useless. It is like a cost-benefit analysis, you see? The cost of hearing more [understandable sounds], and the benefit of having the third suggestion is not that much” (P10).

The need for faster selection. Although we did not engage participants in exploring the selection category, one participant felt the need for improved selection mechanisms. Interestingly, this participant preferred a single suggestion.

“Suggestions should be on bigger buttons. Also, when it gives the right suggestion, I should be able to do a double tap rather than search for an additional button” (P4).

DISCUSSION

In this section, we discuss major findings, explain the limitations of our work, and describe major challenges for future research.

Alternative Nonvisual Word Completion Interfaces

The design space (*notification, output, confidence, cardinality, concurrency, interruption, selection*) served as a structured guide in designing customized nonvisual word completion representations. The process showed to be appropriate to involve users in reflecting about possible interfaces that would fit their needs and expectations. Overall, the participatory design exercise showed that blind people have different preferences on how to engage with word completion systems, resulting in distinct nonvisual interfaces that go beyond converting text into sequential auditory feedback.

All personalized interfaces were less conservative than Talkback regarding notifications as participants preferred to be constantly updated on available suggestions. Moreover, analyzing participants’ choices, four main designs emerged (Figure 4).

First, the *True Talkback*: two participants created an interface similar to Talkback but less conservative than the current mainstream solution on when to provide notifications. This design can be seen as a familiar interface that aims to maximize the benefits of word completion by making the most probable suggestion visible.

Second, the *Sequence Announcer*: four participants preferred multiple suggestions (between 2 and 3) read sequentially. This design can be used with a single auditory channel. As users write, priority is given to input feedback, but when they stop, all suggestions are read aloud ordered by spellchecking confidence.

Third, the *Nonstop Typist*: two participants also preferred sequential feedback for suggestions but desired to hear them while exploring the screen to select characters. Users of the Nonstop Typist interface want to make sure they do not miss valuable suggestions, but they do not want to stop typing, thus the need for concurrent feedback.

Finally, the *Concurrent Announcer* – three participants personalized their interface to read the input and all suggestions concurrently. This interface provides quick access to word completion options in exchange for higher cognitive load. As users type, they can hear three concurrent and spatialized suggestions.

The Unattainable Ultimate Interface

Participants preferred constant feedback on word completion suggestions. Mostly, they wished to have quick access to information that is already available visually in the suggestion bar without the need to spend time exploring it. Prediction accuracy was largely low, and arguably because of it, users preferred to take control over whether they heard suggestions. Still, it was crucial to provide a mechanism to bypass notifications, empowering users to continue typing.

Participants were conscious of their design decisions, suggesting that they would adjust the typing behavior based on perceived benefit from word suggestions, e.g. they might type shorter words and use suggestions for longer words or ignore suggestions in the first written characters. For instance, P10 commented on waiting until entering the third letter to pay attention to suggestions, which is supported by

the threshold-dependent parameter under the *Notification* category. Nonvisual word completion interfaces could support threshold-dependent notifications other than prediction accuracy such as number of characters entered.

Most participants chose to be updated on multiple suggestions. The number of suggestions was dependent on multiple factors, including perceived prediction accuracy, potential keystroke savings, attentional demand, context, and available hardware.

Concurrency was mostly related with users' ability to discriminate word suggestions. Overall, participants' preferences were based on a judgement of benefit and cost, which could vary depending on individual abilities and contextual factors, i.e. an interface that works in a situation may not hold for all mobile contexts. For instance, in a calm environment in which users are devoting their full attention to the input task, the interface can maximize the cognitive load enabling higher entry speeds, while in mobile scenarios, attentional demand should be minimal. Similarly, when using two earbuds, the interface can leverage spatialized and/or concurrent feedback and switch to sequential feedback when only a single auditory channel is available.

Limitations

For the purposes of replicability, we used a publicly available dictionary (Google Android dictionary) to feed the word completion interfaces. Findings are constrained by the generally low theoretical letter savings of 21%. Prediction accuracy may have had an effect on perceived usefulness and trustworthiness when customizing interfaces. Still, it is expected that prediction systems will improve and change users' preference and perceived cost-benefit trade-offs.

Given lack of knowledge on the field, our goal was to shed light on how screen reader users engage with word completion interfaces using the design space as a tool to elicit feedback. There are no previous reports of how blind users leverage word completion. We report on the users' attitudes towards word completion, and desirable customization options that can guide future research. As with any participatory design activity, participants were restricted to use the provided materials. In this case, users had 2⁵ (32) design alternatives they could choose from. Although we never felt participants were restricted, as they were highly encouraged to comment and come up with new ideas, further iterations of this exercise with other design space categories and attributes can uncover novel insights.

Challenges for Future Research

In this section we describe two main challenges that can guide future research on nonvisual word completion.

Quantifying the design space. Finding the most efficient word completion interface was outside the scope of this paper. Nevertheless, we need a better understanding of the perceptual and cognitive costs of presenting nonvisual

suggestions. The proposed design space can serve as a tool to frame further research in understanding the effects of each individual attribute within the design space. Researchers should strive to minimize interaction costs and provide efficient interface designs; from improved notification mechanisms to faster selection techniques. Moreover, these efforts should leverage longitudinal studies to account for users' learning patterns.

Enable customization and adaptation. Customization and adaptation can be key in coping with users' needs and contextual challenges. This work informs on users' customization choices, but this capability needs to be incorporated into accessibility services in order to study if, and how, users leverage them in everyday use.

CONCLUSION

This paper is the first step towards improving the experience of screen reader users with word completion interfaces. We proposed a design space consisting of seven categories: *notification, output, confidence, cardinality, concurrency, interruption, selection.*

The design space is not limited to existing input devices or prediction algorithms. The goal was to present a design space directed towards researchers, designers, and practitioners that provides an approach on how to think about challenges and opportunities that emerge in nonvisual representation of word completions. The categories and values that we chose cover the current typing process of screen reader users, however we want to highlight that the framework can be extended.

We demonstrated the potential of our design space by engaging blind users in a participatory design session and empowering them to customize their interfaces. Participants created alternative solutions that go beyond current screen readers' capabilities, approximating their designs to the visual experience of their sighted counterparts with frequent notifications and multiple suggestion updates. Moreover, we show how individual abilities and contextual factors may influence their choices when designing word completion interfaces.

ACKNOWLEDGMENTS

We thank Fundação Raquel and Martin Sain in Lisbon (Portugal) and all participants. This work was supported by the Fundação para a Ciência e a Tecnologia (FCT) through scholarship SFRH/BD/103935/2014, INESC-ID research unit UID/CEC/50021/2019, LASIGE Research Unit UID/CEC/00408/2019, and project mIDR (AAC 02/SAICT/-2017, project 30347, cofunded by COMPETE/FEDER/FNR).

REFERENCES

1. Ahmed Sabbir Arif, Sunjun Kim, Wolfgang Stuerzlinger, Geehyuk Lee, and Ali Mazalek. 2016. Evaluation of a Smart-Restorable Backspace Technique to Facilitate Text Entry Error Correction. In *Proceedings of the 2016 CHI Conference on Human*

- Factors in Computing Systems* (CHI '16), 5151–5162. <https://doi.org/10.1145/2858036.2858407>
2. Kenneth C. Arnold, Krzysztof Z. Gajos, and Adam T. Kalai. 2016. On Suggesting Phrases vs. Predicting Words for Mobile Text Composition. *Proceedings of the 29th Annual Symposium on User Interface Software and Technology - UIST '16*: 603–608. <https://doi.org/10.1145/2984511.2984584>
 3. Shiri Azenkot and Nicole B. Lee. 2013. Exploring the use of speech input by blind people on mobile devices. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility - ASSETS '13*, 1–8. <https://doi.org/10.1145/2513383.2513440>
 4. Shiri Azenkot, Jacob O. Wobbrock, Sanjana Prasain, and Richard E. Ladner. 2012. Input finger detection for nonvisual touch screen text entry in Perkinput. In *Proceedings of Graphics Interface*, 121–129. <https://doi.org/2305276.2305297>
 5. Rafael Ballagas, Jan Borchers, Michael Rohs, Jennifer G Sheridan, James Foley, Victor Wallace, and Peggy Chan. 2006. The Smart Phone: A The input design space. *Shoot* 5, 1: 70–77. <https://doi.org/10.1109/MPRV.2006.18>
 6. Rafael Ballagas, Sarthak Ghosh, and James Landay. 2018. The Design Space of 3D Printable Interactivity. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2, 2: 1–21. <https://doi.org/10.1145/3214264>
 7. Xiaojun Bi, Tom Ouyang, and Shumin Zhai. 2014. Both Complete and Correct?: Multi-objective Optimization of Touchscreen Keyboard. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14), 2297–2306. <https://doi.org/10.1145/2556288.2557414>
 8. Matthew N. Bonner, Jeremy T. Brudvik, Gregory D. Abowd, and W. Keith Edwards. 2010. No-look notes: Accessible eyes-free multi-touch text entry. *Lecture Notes in Computer Science* 6030 LNCS: 409–426. https://doi.org/10.1007/978-3-642-12654-3_24
 9. Derek Bridge and Paul Healy. 2011. GhostWriter-2.0: Product Reviews with Case-Based Support. In *Research and Development in Intelligent Systems XXVII*, 467–480.
 10. Douglas S Brungart and Brian D Simpson. 2005. *Improving multitalker speech communication with advanced audio displays*.
 11. Stuart K. Card, Jock D. Mackinlay, and George G. Robertson. 2003. The design space of input devices. June: 117–124. <https://doi.org/10.1145/97243.97263>
 12. Girish Dalvi, Shashank Ahire, Nagraj Emmadi, Manjiri Joshi, Anirudha Joshi, Sanjay Ghosh, Prasad Ghone, and Narendra Parmar. 2016. Does Prediction Really Help in Marathi Text Input?: Empirical Analysis of a Longitudinal Study. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '16), 35–46. <https://doi.org/10.1145/2935334.2935366>
 13. Torsten Felzer, I Scott MacKenzie, and Stephan Rinderknecht. 2013. OnScreenDualScribe: A Computer Operation Tool for Users with a Neuromuscular Disease. In *Universal Access in Human-Computer Interaction. Design Methods, Tools, and Interaction Techniques for eInclusion*, 474–483.
 14. James D. Foley, Victor L. Wallace, and Peggy Chan. 1984. Human Factors of Computer Graphics Interaction Techniques. *IEEE Computer Graphics and Applications* 4, 11: 13–48. <https://doi.org/10.1109/MCG.1984.6429355>
 15. Spence Green, Jason Chuang, Jeffrey Heer, and Christopher D Manning. 2014. Predictive Translation Memory: A Mixed-initiative System for Human Language Translation. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (UIST '14), 177–187. <https://doi.org/10.1145/2642918.2647408>
 16. João Guerreiro and Daniel Gonçalves. 2014. Text-to-speeches: Evaluating the Perception of Concurrent Speech by Blind People. In *Proceedings of the 16th International ACM SIGACCESS Conference on Computers & Accessibility* (ASSETS '14), 169–176. <https://doi.org/10.1145/2661334.2661367>
 17. João Guerreiro and Daniel Gonçalves. 2015. Faster Text-to-Speeches: Enhancing Blind People's Information Scanning with Faster Concurrent Speech. In *Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility* (ASSETS '15), 3–11. <https://doi.org/10.1145/2700648.2809840>
 18. João Guerreiro, André Rodrigues, Kyle Montague, Tiago Guerreiro, Hugo Nicolau, and Daniel Gonçalves. 2015. TABLETS Get Physical: Non-Visual Text Entry on Tablet Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
 19. T Guerreiro, P Lagoá, H Nicolau, D Gonçalves, and J A Jorge. 2008. From tapping to touching: Making touch screens accessible to blind users. *IEEE MultiMedia*: 48–50.
 20. D Jeffery Higginbotham. 1992. Evaluation of keystroke savings across five assistive communication technologies. *Augmentative and Alternative Communication* 8, 4: 258–272. <https://doi.org/10.1080/07434619212331276303>

21. Teresa Hirzle, Jan Gugenheimer, Florian Geiselhart, Andreas Bulling, and Rukzio Enrico. 2019. A Design Space for Gaze Interaction on Head-Mounted Displays. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*: 1–12. <https://doi.org/10.1145/3290605.3300855>
22. S K Kane, J P Bigham, and J O Wobbrock. 2008. Slide rule: making mobile touch screens accessible to blind people using multi-touch interaction techniques. In *Proceedings of the 10th international ACM SIGACCESS conference on Computers and accessibility*, 73–80.
23. Shaun K Kane and Meredith Ringel Morris. 2017. Let's Talk About X: Combining Image Recognition and Eye Gaze to Support Conversation for People with ALS. In *Proceedings of the 2017 Conference on Designing Interactive Systems (DIS '17)*, 129–134. <https://doi.org/10.1145/3064663.3064762>
24. Anjuli Kannan, Karol Kurach, Sujith Ravi, Tobias Kaufmann, Andrew Tomkins, Balint Miklos, Greg Corrado, Laszlo Lukacs, Marina Ganea, Peter Young, and Vivek Ramavajjala. 2016. Smart Reply: Automated Response Suggestion for Email. In *Proceedings of the 22Nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*, 955–964. <https://doi.org/10.1145/2939672.2939801>
25. Heidi Horstmann Koester and Simon P Levine. 1994. Modeling the speed of text entry with a word prediction interface. *IEEE transactions on rehabilitation engineering* 2, 3: 177–187.
26. Matthijs Kwak, Kasper Hornbæk, Panos Markopoulos, and Miguel Bruns Alonso. 2014. The design space of shape-changing interfaces. June: 181–190. <https://doi.org/10.1145/2598510.2598573>
27. Sergio Mascetti, Cristian Bernareggi, and Matteo Belotti. 2011. TypeInBraille: A Braille-based Typing Application for Touchscreen Devices. In *Proceedings of the 13th International ACM SIGACCESS Conference on Computers and Accessibility*, 295. <https://doi.org/10.1145/2049536.2049614>
28. Meredith Ringel Morris, Jazette Johnson, Cynthia L Bennett, and Edward Cutrell. 2018. Rich Representations of Visual Content for Screen Reader Users. *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*: To Appear.
29. Hugo Nicolau, Kyle Montague, Tiago Guerreiro, João Guerreiro, and Vicki L Hanson. 2014. B#: Chord-based Correction for Multitouch Braille Input. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'14)*, 1705–1708. <https://doi.org/10.1145/2556288.2557269>
30. Hugo Nicolau, Kyle Montague, Tiago Guerreiro, André Rodrigues, and Vicki L Hanson. 2015. Typing Performance of Blind Users: An Analysis of Touch Behaviors, Learning Effect, and In-Situ Usage. In *ACM SIGACCESS Conference on Computers and Accessibility*. <https://doi.org/10.1145/2700648.2809861>
31. Hugo Nicolau, Kyle Montague, Tiago Guerreiro, André Rodrigues, and Vicki L Hanson. 2017. Investigating Laboratory and Everyday Typing Performance of Blind Users. *ACM Transactions on Accessible Computing (TACCESS)* 10, 1: 4:1–4:26. <https://doi.org/10.1145/3046785>
32. Jakob Nielsen. 1995. 10 usability heuristics for user interface design. *Nielsen Norman Group* 1, 1.
33. João Oliveira, Tiago Guerreiro, Hugo Nicolau, Joaquim Jorge, and Daniel Gonçalves. 2011. Blind people and mobile touch-based text-entry: Acknowledging the Need for Different Flavors. In *The proceedings of the 13th international ACM SIGACCESS conference on Computers and accessibility - ASSETS '11*, 179. <https://doi.org/10.1145/2049536.2049569>
34. João Oliveira, Tiago Guerreiro, Hugo Nicolau, Joaquim Jorge, and Daniel Gonçalves. 2011. BrailleType: unleashing braille over touch screen mobile phones. *Lecture Notes in Computer Science* 6946 LNCS, PART 1: 100–107. https://doi.org/10.1007/978-3-642-23774-4_10
35. Philip Quinn and Shumin Zhai. 2016. A Cost-Benefit Study of Text Entry Suggestion Interaction. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*: 83–88. <https://doi.org/10.1145/2858036.2858305>
36. Ramesh Raskar. 2012. Idea Hexagon. Retrieved from <https://de.slideshare.net/cameraculture/>
37. André Rodrigues, Kyle Montague, Hugo Nicolau, and Tiago Guerreiro. Getting Smartphones to Talkback: Understanding Smartphone Adoption Process of Blind Users. In *ACM SIGACCESS Conference on Computers and Accessibility (ASSETS'15)*.
38. André Rodrigues, Hugo Nicolau, Kyle Montague, Luis Carriço, and Tiago Guerreiro. 2016. Effect of Target Size on Non-visual Text-entry. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '16)*, 47–52. <https://doi.org/10.1145/2935334.2935376>
39. Caleb Southern, James Clawson, Brian Frey, Gregory Abowd, and Mario Romero. 2012. An Evaluation of BrailleTouch: Mobile Touchscreen Text Entry for the Visually Impaired. In *Proceedings of the 14th International Conference on Human-computer*

Interaction with Mobile Devices and Services (MobileHCI '12), 317–326. <https://doi.org/10.1145/2371574.2371623>

40. Daniel Trindade, André Rodrigues, Tiago Guerreiro, and Hugo Nicolau. 2018. Hybrid-Braille: Combining Physical and Gestural Interaction for Mobile Braille Input and Editing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI '18), 27:1--27:12. <https://doi.org/10.1145/3173574.3173601>
41. Keith Trnka, John McCaw, Debra Yarrington, Kathleen F McCoy, and Christopher Pennington. 2009. User Interaction with Word Prediction: The Effects of Prediction Quality. *ACM Trans. Access. Comput.* 1, 3: 17:1--17:34. <https://doi.org/10.1145/1497302.1497307>
42. Keith Trnka and Kathleen F McCoy. 2008. Evaluating Word Prediction: Framing Keystroke Savings. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies: Short Papers* (HLT-Short '08), 261–264. Retrieved from <http://dl.acm.org/citation.cfm?id=1557690.1557766>
43. Bruce N Walker, Amanda Nance, and Jeffrey Lindsay. 2006. Spearcons: Speech-based earcons improve navigation performance in auditory menus. In *Proceedings of the 12th International Conference on Auditory Display*.
44. Georgios Yfantidis and Grigori Evreinov. 2004. Adaptive blind interaction technique for touchscreens. *Universal Access in the Information Society* 4, 4: 344–353. <https://doi.org/10.1007/s10209-004-0109-7>
45. Shengdong Zhao, Fanny Chevalier, Wei Tsang Ooi, Chee Yuan Lee, and Arpit Agarwal. 2012. AutoComPaste: Auto-completing Text As an Alternative to Copy-paste. In *Proceedings of the International Working Conference on Advanced Visual Interfaces* (AVI '12), 365–372. <https://doi.org/10.1145/2254556.2254626>
46. Suwen Zhu, Tianyao Luo, Xiaojun Bi, and Shumin Zhai. 2018. Typing on an Invisible Keyboard. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI '18), 439:1--439:13. <https://doi.org/10.1145/3173574.3174013>
47. F Zwicky. 1967. The Morphological Approach to Discovery, Invention, Research and Construction. In *New Methods of Thought and Procedure*, 273–297.