Spring 6-2021

Distributed Cognition as a framework for accessible design

Darley Sackitey
Darley.Sackitey.21@Dartmouth.edu

Distributed Cognition as a Framework for Accessible Design

by

Darley Sackitey

A Senior Honors Thesis

Presented to the Faculty of the Department of Cognitive Science

Dartmouth College

In partial fulfillment of the requirements for

the degree of Bachelor of Arts in Cognitive Science

Principal Reader: Professor Lorie Loeb

Second Reader: Professor Jonathan Phillips

Hanover, New Hampshire, USA

May 2021

Abstract

In order to navigate the world, humans have historically made tools that allowed them to exploit their environment in ways beyond their natural physical and mental capacities. This continues to be true in the information age. Being able to use digital tools in this age affords the individual agency to influence and participate in the world and so it is critical that this capacity is equally available to all people. Current standards of accessibility attempt to create accommodations for overlooked populations in the form of guidelines, but these rely on a flawed approach to accessibility that is surface-level and lacks understanding of the populations to which it caters. However, distributed cognition offers an approach through which we can understand users through their interaction strategies with technology and through which we can understand technology as an extension of the mind. This paper explores the short-comings of current approaches to accessible design, how distributed cognition has been used to describe human-computer interactions and how this can be extended to understand accessible design. A case study of how a distributed cognition approach may transform accessible design for individuals with Autism Spectrum Condition is developed and discussed.

Keywords: distributed cognition; accessible design; human-computer interaction; autism

Introduction

These days, it is difficult to imagine everyday life without technology. We rely on screens on a daily basis to communicate with each other, learn, access health services and even for entertainment. As the world becomes increasingly digitised, it becomes imperative to ensure that the use of technology is accessible to all people and that the design of these technologies accommodates the unique needs of different user groups in society. To this end, accessible design has been developed as an attempt to ensure that the needs of individuals that differ from typical users are not overlooked in the design of technology. Accessibility is described by how much a product can be used by the widest range of user groups to satisfy the needs of as many users as possible, in different contexts of use (International Organisation for Standardisation, 2008). Current approaches to satisfying this condition focus on identifying needs that may be overlooked for different populations and then making design changes to accommodate these needs, often through proposed guidelines. Identifying the needs of populations often occurs through usability testing where users are observed using a product in controlled environments, in order to identify what tasks they struggle with (e.g. by tracking the number of errors made) and what changes can be made to minimise these struggles. While this approach leads to the generation of tangible practical 'rules' through which the user experiences of people with disabilities or other communities with atypical needs may be improved, it falls short in providing an explanation of the underlying mechanisms that create friction between users and a piece of technology. They also tend to over-emphasise users' perceived deficits and attempt to compensate for them through guidelines and overlook the unique strengths individuals may have that could transform the way they interact with technology. 26 years after the first compilation of web accessibility guidelines, it may be time to revisit what it means for technology to be accessible.

Distributed cognition provides a potential lens through which accessible design can be reimagined. By emphasising the collaborative cognitive work that occurs between people and their environment, distributed cognition allows us to shift the priorities of accessible design from usability to the extension of human capabilities through integration with technology. It allows us to identify how internal and external representations may give rise to specific interaction strategies, and how these strategies may vary across users with different strengths and weaknesses. In this way, accessibility becomes about understanding how users navigate the world and how to design technology which is flexible enough to accommodate different interaction strategies. This paper explores how the theory of distributed cognition can be used to describe human-computer interactions and provide a reframing of accessibility that allows us to understand how users' cognitive and physical capacities may influence the approaches they take in constructing mental models and developing interaction strategies, as well as how information structures and design may make room for or exclude some interaction strategies.

This paper has five sections. Section I considers accessibility/accessible design as it is currently defined and practiced and identifies the shortcomings of current systems of thought on the topic. Section II focuses on distributed cognition and the unique approach it takes to framing human cognition. Section III is concerned with how distributed cognition has been used as a framework under which human-computer interactions can be explored and modelled. Section IV considers how work linking distributed cognition and human-computer interaction can be extended to understand accessibility in interface design and how this may transform our approaches to and conceptualization of accessibility. Section V applies a distributed cognition approach to accessibility issues faced by autistic individuals.

Section I: Accessibility

In human-centered design, accessible design is "design focused on principles of extending standard design to persons with some type of performance limitation to maximize the number of potential customers who can readily use a product, building or service" (International Organisation for Standardisation, 2008). The general idea behind this definition is that the successful use of technology should not be exclusive to individuals with normative needs, and that users with physical or cognitive disabilities should be adequately accommodated either by designing universally usable products, adapting interfaces to users or developing standardised interfaces that are compatible with wearable technology (or other specialised products for people with disabilities) (Persson, Åhman, Yngling, & Gulliksen, 2015). Usability is described as "effective, efficient and satisfying" design and is generally specific to users and contexts of use (Dattolo & Luccio, 2017). Consequently, design may be usable but not accessible. This also means that we can redescribe accessible design as design that is usable/can be used by anyone in any context.

To achieve the objectives of accessible design, designers often refer to sets of guidelines and recommendations that describe common standards that make web or mobile content accessible. This approach generally reduces the pursuit of web accessibility to satisfying a list of recommendations prescribed by a standard—designs that are able to check off more boxes are considered more accessible, and those are lacking are considered less accessible. Often, usability tests in which users are instructed to carry out specific tasks in a controlled environment are conducted to understand how users interact with technology and to observe pain points. For example, Al-Wakeel et al. evaluate mobile applications for people with autism by asking children to use an application and following up with questionnaires that evaluate their experience as well as collecting eye-tracking data and use this information to generate recommendations to improve usability (Al-Wakeel, Al-Ghanim, Al-Zeer, & Al-Nafjan, 2015).

In general, the emphasis in accessibility has been on creating tangible concrete guidelines that are easily measurable. The most popular standard for accessible web design is the WCAG, which spells out principles, guidelines, success criteria, and advisory techniques for achieving accessible design. The document describes the main principles of accessibility as: perceivable, operable, understandable, and robust. Perceivable means that information must be presented in a way that users can register through appropriate senses, and guidelines to achieve this center around text formatting, creating alternative ways for users to access information (e.g., subtitles), and increasing clarity of content. Operable design allows users to easily navigate information and is achieved by adherence to guidelines around allowing time for users to consume content, ensuring all navigation can be done on a keyboard, avoiding seizure-inducing content, making navigation clear, etc. To be understandable, content must be readable, predictable, provide error states, etc., and guidelines for robustness focus on making sure that different user groups can interpret content, and increasing compatibility to users (e.g., considering people who use wearable technology). These guidelines provide a neat framework under which designers can create and measure design, which can potentially transform the user experiences for often overlooked populations and allow more people to successfully use technology.

Guideline 1.1

Text Alternatives: Provide text alternatives for any non-text content

Success Criteria	WebAIM's Recommendations
1.1.1 Non-text Content (Level A)	☐ Images, form image buttons, and image map hot spots have appropriate, equivalent alternative text. ☐ Images that do not convey content, are decorative, or contain content that is already conveyed in text are given empty alternative text (alt="") or implemented as CSS backgrounds. All linked images have descriptive alternative text. ☐ Equivalent alternatives to complex images are provided in context or on a separate linked page. ☐ Form buttons have a descriptive value. ☐ Form inputs have associated text labels. ☐ Embedded multimedia is identified via accessible text. ☐ Frames and iframes are appropriately titled.

Figure 1: Guideline for perceivability based on WCAG 2.0 (WebAIM, 2021)

These guidelines rely on the assumption that accessibility is simply usability universalised; however, being able to use a website does not necessarily make it accessible to an individual. In fact, the relationship between usability and accessibility is not well established – according to an investigation by the Disability Rights Commission, there is no clear relationship between websites that conform with the WCAG and the levels of performance and satisfaction for people with disabilities (Disability Rights Commission , 2004). In addition, the WCAG is only so useful for identifying accessibility issues, and at best may address only about half of usability issues faced by disabled users (Rømen & Svanæs, 2012; Power, Freire, Petrie, & Swallow, 2012). This suggests that these guidelines provide an incomplete account of accessibility, and do not even address all usability issues.

There is clearly a gap between the impact accessibility guidelines are expected to have and what the reality of the experience of disabled people with technology is. I propose that this is because these standards are often an afterthought to design and do not fundamentally transform designs to accommodate disabled persons. Accessible design standards fall short in describing and addressing true accessibility in several ways:

1. They give an incomplete description of accessibility by over-emphasising usability
2. They do not provide insight on the mechanisms that cause a failure in users' adoption of technology.

3. Guidelines are rigid and do not adapt quickly enough to technology changes, or to all circumstances of use.

4. Usability tests as a measure of accessibility focus' focus on singular isolated tasks does not account for interaction between tasks, how usage changes with time, how users integrate with a product, and may overvalue efficiency

5. Accessibility standards take a deficit approach towards user and focus and may over-emphasise their weaknesses without considering how their strengths may also transform the interaction they have with technology.

Because accessibility standards tend to focus almost exclusively on usability, they also overlook nuances about how users' mental models may differ and how this may influence how they interact with software. These guidelines over-emphasise the presentation of content on screens at the expense of understanding how this content interacts with users' natural strategies for problem-solving and interacting with the world. Consider blind people for example— for users to navigate a page successfully, they must form mental models, and to do that they must understand the information groupings. Blind people will rely on different sensory input to form these mental models (Leuthold, Bargas-Avila, & Opwis, 2008), and so will have very different strategies for navigating an information space from sighted people (Savidis & Stephanidis, 1998). Consequently, a design is not accessible to blind people simply because it can be used by blind users (e.g., websites that are compatible with screen readers are not necessarily accessible to blind users).

Accessibility must consider more than users' ability to access information on screens, but also how the underlying information structure fits with users' mental models and enhances their natural interaction strategies. According to Power et. al., content not found where expected was one of the prominent problem areas faced by disabled users when navigating sites that conformed with WCAG 2.0 guidelines. From their study, users were often unable to find information where they expected it to be, for example, on a museum website, users followed a link to an object in a collection, expecting to find information about where the object was displayed but were unable to find this information. Even when users were able to find information they were looking for, it was not through following their natural logic because the pages did not match their mental models of the site architecture (Power, Freire, Petrie, & Swallow, 2012). Since mental models are built from a user's internal representation of the external world, they shape the expectations users have for interactions and consequently the ways in which they interact with the world. A more complete approach to accessibility would consider accessible design as design that gives rise to functionally equivalent mental models to all users, such that people with different mental models may still have similar ease in carrying out different tasks across different contexts of use. This perspective accommodates flexibility in the different forms of mental models differently abled individuals may rely on. In this way, we can shift the focus of accessibility from designing interfaces that are technically usable to designing interfaces that leverage the ways in which users make sense of the world to communicate information.

By virtue of being standards, accessibility guidelines oversimplify accessibility and lack explanatory value. Whilst several accessibility guidelines exist, there is little fundamental understanding of why and/or when these guidelines work or do not work. Accessibility standards are often applied to designs with little recognition of what gave rise to the accessibility issues that they attempt to remedy. Accessibility recommendations are built from usability studies and based on surface observations of user friction. We say that the user was unable to carry out a task because e.g., they were not able to understand a button label, and so we recommend that buttons are labelled with clear language, but we do not consider that perhaps buttons are simply not the suitable method of navigation for the user. This kind of information cannot be revealed by user testing which occurs in a hyper-controlled environment and relies largely on observing user behaviour. In this scenario, we understand the failed interaction between the user and technology to be the software's inability to meet an accessibility guideline (labelling buttons) instead of a more fundamental mismatch between information structure and users' mental models.

Inaccessible or unusable design is often attributed to bad design, but what exactly was misunderstood by the design or badly executed is often hard to identify beyond what accessibility guidelines were met and which were not. This reduces the pursuit of accessible design to identifying surface features of web and mobile technologies and deciding whether or not users are able to perceive and/or interact with them. In this respect, accessible design as it stands acts as a Band-Aid solution to problems that may be much more intrinsic than the colour of text on screens or the presence or absence of subtitles. For example, the WCAG suggests providing captions for live audio content in synchronised media as a way of achieving perceivability; however, it does not detail how this step relates to differently abled people or why captions specifically are suited towards addressing their unique needs. While this does not minimize the practical usefulness of these guidelines— captions may still have positive impact on accessibility for some populations— it is important to acknowledge their limitations and avoid overestimating how much they tell us about what it means for design to be accessible.

Further, accessibility guidelines are generally too rigid to accommodate different varieties of use-cases or the rapidly changing landscape of digital content and interfaces. Considering that the objective of accessibility is to develop interfaces that different users can use in a variety of contexts, it seems counterintuitive that we rely on a set of guidelines to achieve this goal, particularly because anticipating the different contexts of use to include is almost impossible. Even within the subset of people with disabilities, it is difficult to make generalisations about how individuals may interact with an interface, and what modifications may be beneficial for users. The needs of a blind user may differ from that of an autistic user, and even within the subset of blind users, different users may have different needs. It is difficult to account for all of these needs in a short recommendation, especially given that they may change in different contexts for use. In different contexts, some guidelines will be more useful than others, and guidelines that may be helpful for some users in some contexts may be subtractive and maybe even detrimental for other users in the same context. Because a large number of use cases and user types exist, to achieve accessibility goals with current approaches to accessibility theoretically requires developing infinite guidelines.

The inflexibility of accessibility guidelines makes them unsuited for accommodating the changes between the interactions between users and technology over time. Because usability emphasises the experience on first use, it does not acknowledge how technology use may change over time and how these changes may require different accommodations. For example, when people first start using a piece of technology, their immediate accessibility needs may revolve around understanding the information structure of the software, but as they continue to use it, they may adapt their behaviour to further include the technology in their lives, and consequently their accessibility needs may begin to become more about how they may achieve this adaptation more seamlessly. It is also the case that different interface systems require different accessibility accommodations—the needs of blind people on 2D interface will be different from 3D interfaces, and as these trends change with the advancement of technology, it is important to identify overarching frameworks that can describe and direct accessible design for these different platforms.

Usability tests are also a limited measure of accessibility, since they do not recreate the contexts of use in which users may rely on or encounter a product. Usability studies often occur in a controlled environment, where participants are issued clear instructions to carry out individual tasks without a real-world context. Because people may use the same product in different ways in different situations, these studies do not allow us to identify how accessibility needs may vary across contexts of use. To avoid confounds and develop targeted insights about usability, it makes sense that these studies investigate tasks individually. However, in our use of technology, prior interactions influence subsequent ones, because they set our expectations for subsequent interactions—a user would be quite surprised if the navigation for a product changed with every interaction, even if the navigation was completely new to them on first use. The results of one interaction may also give rise to new interactions, and usability tests are unable to investigate these interrelations due to their controlled nature.

Also, because usability tests are designed to be very focused, they are effective at identifying major usability issues for normative users, or issues that the majority of a user group may encounter but this tends to side-line less popular issues other users might face, which is counter-productive given the aims of accessibility. Indeed, accessible design is in the details that are present or absent, the small choices that accumulate to the exclusion of user groups and relying solely on usability testing does not afford a complete measure of this. Regardless, usability is useful because a product cannot be accessible if it is not first usable; however, it is not a complete measure of accessibility. For accessibility guidelines to capture a more complete picture of accessibility, they would have to rely on measures beyond these tests and consider methods that incorporate time and every-day contexts in understanding users' interactions such as diary studies or participatory design methodologies.

Current standards of accessibility assume a deficit approach to solving usability issues for people with disabilities. This means that disabled users are considered almost solely in relation to the ways in which they fall short of a normative standard. For example, accessibility guidelines may only consider blind people in terms of their blindness or people with attention disorders in terms of their difficulties with attention. The accessibility principles of the WCAG are often tied to what users cannot do, such that an action only becomes an accessibility issue when users fail to carry it out. For example, to be understandable, content has to be readable, and this implies understandability only becomes an issue when readability is not satisfied. The subtle implication of this is that the reason a design may not be understandable is that the user cannot read content, which undermines alternative ways of understanding content and overlooks the chance that users may be able to perceive meaning in other ways. Also, guidelines are built from usability studies which are designed to look for problems that users face and ignore the points of success. They do not aim to understand users, they aim to identify failures and 'fix' them, while overlooking the unique strengths of these individuals. The direction of focus thus becomes reducing the number of usability issues, and not necessarily maximizing the overall experience of users, even though that may still be achieved sometimes. Accessibility principles should focus more broadly on understanding the ways that differently abled people use technology, and not only the problems they encounter when doing so (Power, Freire, Petrie, & Swallow, 2012).

Since they take a deficit approach, accessibility guidelines also inherit the biases of holding a normative user interaction as the standard to achieve. Deficits only exist in relation to an accepted standard that often tends to be normative, and by relying on this, accessibility guidelines end up producing recommendations that center around nudging the interactions of differently abled people towards 'normal' interactions. This is evident in the earlier example of understandability as a principle of accessibility. The first recommendation for understandability in the WCAG is making text readable, and all subsequent guidelines assume a normative standard for reading and readability, with references to factors like pronunciations, abbreviations, definitions, etc. In fact, different ways of making meaning from content (text or otherwise) are not discussed, as if these guidelines can be easily adapted to people with disabilities. These recommendations assume that differently abled people would use a piece of technology similarly to typically abled people, but a blind person does not collect information about the world in the same way as a sighted person with their eyes closed might because they may have different mental models of the world. This is because the lived experience of a blind person is not only about their inability to see. Consequently, the difference between the accessibility needs of a blind and sighted person cannot be reduced to accommodations that only address the blind person's ability to see—for example, audio transcriptions of screens. True accessibility must take a holistic approach to understanding how differently abled individuals interact with technology—their successes, failures, and unique patterns of interaction—in order to generate principles are able to leverage both the strengths and weaknesses of differently abled individuals.

It is clear that in our pursuit of accessible design, standardizing user experiences has emerged to be a primary focus. Whilst this emphasis on clean concrete principles like the WCAG2.0 has been useful in corporate design, this approach does not describe the underlying mechanisms that make a piece of technology accessible or not. Without a theoretical framework to understand why these guidelines work, it becomes difficult to identify how they might be inappropriate for different contexts of use, and when they might fail the users they aim to accommodate. By focusing on observable points of friction and cataloguing them into principles of accessibility, we employ a bottom-up approach to solving accessibility problems. However, to build a complete picture of what it means to achieve accessibility, it is important to combine these with top-down approaches that contextualize observed human-computer interactions. It is important to step beyond the concrete and into the abstract systems behind the surface principles to improve our understanding of accessibility.

II. Distributed cognition

> "Humans create their cognitive powers in part by creating the environments in which they exercise those powers" (Hutchins, 1995)

Distributed cognition is a perspective on cognition that proposes that cognitive processes occur in conjunction with the environment and are not bound to the human brain. Classical computational theory of mind describes the mind as a computational machine that manipulates a finite set of symbols to build more complex ideas; connectionist models of the mind suggest that the mind acts as a pattern matcher instead. Unlike both of these models, distributed cognition considers how individuals interact with each other and the environment to carry out cognitive tasks. It considers cognitive capabilities as shared between the individual and the environment within which they exist (Hutchins, Distributed cognition, 2000). What sets distributed cognition apart from other theories of mind is that the unit of analysis of cognitive events expands beyond the individual and is not bound to an individual brain. It considers that cognitive activity is situated in our environments such that elements of our environments are computational mediums that allow us to complete cognitive tasks.

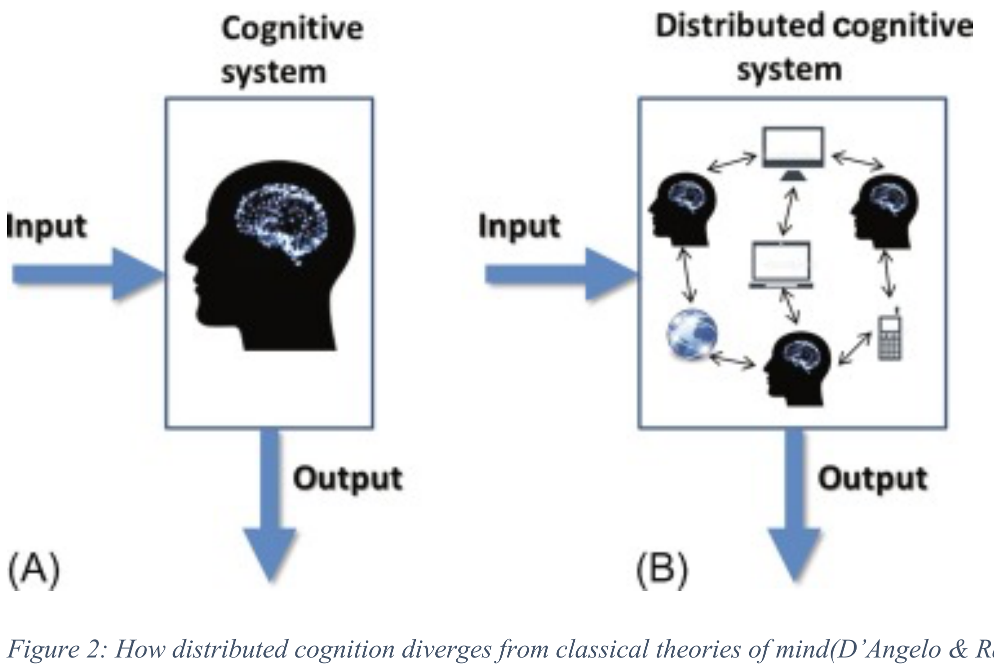

Figure 2: How distributed cognition diverges from classical theories of mind(D'Angelo & Rampone, 2018)

Distributed cognition describes two kinds of representations that facilitate the interactions between individuals and their environment – internal representations which exist within the mind of the individual, and external representations which exist in the environment (Zhang & Norman, 1994). When asked to compute a complex arithmetic calculation, most people will resort to using a paper and pen—in this process, they set up the problem in physical form through external representations (written symbols) and manipulate these symbols in a specific order to complete the task. Whilst they rely on memory in carrying this out, the process does not occur entirely in their head (Giere & Moffatt, 2003); in fact, for many people, it would be close to impossible to complete this task mentally. In this example, the pen and paper can be considered cognitive artefacts that facilitate completion of the arithmetic task. Cognitive tasks can be distributed between the members of a group; for example, in collaborative work environments (Rogers & Ellis, 1994). They may also be distributed across internal and external representations as in solving arithmetic on paper or distributed across time such that the results of earlier events transform the properties of subsequent events (Zhang & Patel, 2006). Under this approach to cognition, a central challenge is understanding the distribution, transformation, and propagation of information across different components of the cognitive system and how they affect the performance of the system as a whole.

A common misconception about distributed cognition is that it simply amplifies already existing human capabilities such that cognitive tasks are off-loaded onto cognitive artefacts, without necessarily transforming the cognitive processes of humans. This perspective that the environment is an amplifier of cognitive abilities, fails to recognize the ways in which cognitive processes are structurally transformed by the different cognitive artefacts available to the individual. Cole and Griffin dispute the idea that the consequent improvement of human capacities from using the environment as a cognitive tool results from a process of amplification. According to them, 'amplification' describes an intensification of something without changing its basic structure, which is not a complete picture of how individuals interact with the environment to complete cognitive tasks (Griffin & Cole, 1980). For example, the tasks of writing something down and reading it later is not in fact an amplification of memory because the set of functional skills that are used to carry out the task of remembering in this scenario are completely different. Similarly, the processes underlying arithmetic calculation on paper would vary vastly from those relying on the manipulation of a physical or mental abacus.

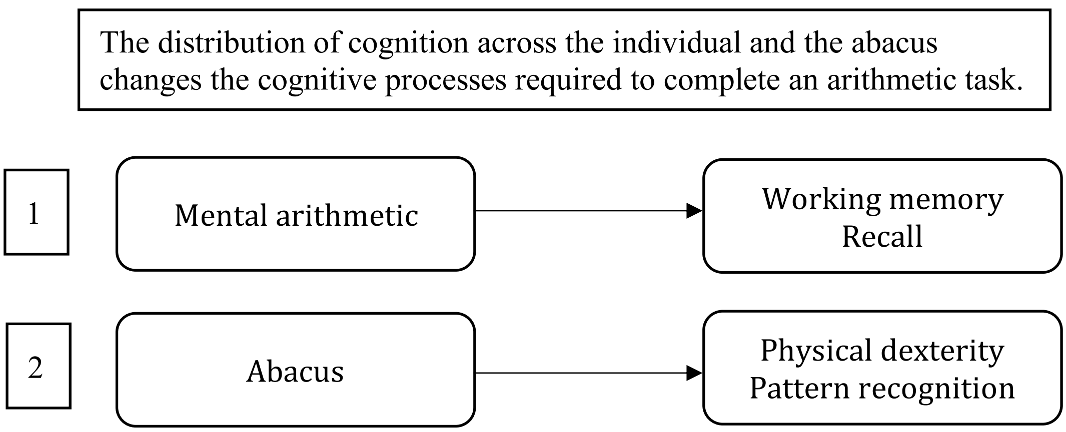

Figure 3: Sharing cognitive work with the environment transforms the way we do tasks

Eve … esses may be faster and easier to use because of the artefacts in relies on. It follows that cognitive artefacts that take advantage of skills that humans are good at—pattern-matching, object manipulation, mental simulation of simple dynamics (Hutchins, Distributed cognition, 2000)—will allow for more efficient completion of cognitive tasks. The distributed cognition perspective is powerful because it allows us to think about the ways in which cognitive artefacts may be designed to transform cognitive tasks into forms that leverage human strengths and weaknesses.

III. Distributed cognition as a framework for human-computer interactions

Human-computer interaction (HCI) explores the relationships that we form with technology in our everyday use of devices. According to Sinha, Shahi and Shankar, human-computer interaction is "concerned with the design, evaluation and implementation of interactive computing systems for human use and with the study of major phenomena surrounding them". It focuses specifically on the interaction between one or more computational machines and is concerned with:

1. The design of computer interfaces
2. Methods of implementing interfaces
3. Evaluating interfaces
4. Developing new interfaces and interactions
5. Developing theories and models of interaction (Sinha, Shahi, & Shankar, 2010).

Distributed cognition provides a perspective that can be useful for addressing issues 1, 3 and 5 especially. Within a distributed cognition framework, we can think of computers as cognitive artefacts with which individuals are able collaborate to accomplish cognitive tasks. For example, when we use a calculator to compute arithmetic, we are offloading some of the memory requirements of doing the same task on paper onto the calculator since we only need to remember what the numeric symbols mean. The focus on the interaction between individuals and cognitive artefacts makes it a useful lens through which we can evaluate human-computer interaction.

Under a distributed cognition lens, the individual as a 'user' can be reimagined. Traditional models of human-computer interactions consider technology as a tool to be used by the individual, which centers the cognitive work in the individual. However, the relationship between the individual and the technology they use is bi-directional: people are transformed by technology as they use it, and in turn they may transform the ways in which it is used. For example, computer games have been found to transform classroom learning and improve student's performance (Miller & Robertson, 2010), and while social media was initially created with the purposes of connecting friends, it has quickly become a marketplace for different goods in some cases (e.g. art markets on Twitter) because of the ways that people use it. Over time, people become more and more integrated with the technology they use as they adapt to it and adapt it to their needs. This is reflected in social trends around the use of technology—e.g. checking emails several times per day has become second nature for many people in recent time. Distributed cognition highlights this integrative process where technology is more than an aid towards carrying out a task, but also transforms our cognitive capacities. This allows us to think about HCI beyond the design and implementation of interfaces, but also in terms of the ways in which human cognitive capacities may be elevated with technology.

Distributed cognition provides a framework for approaching empirical and conceptual research problems in HCI. Distributed cognition allows us to identify and understand when and how humans exploit their physical environments to do cognitive work, which can tell us how to design digital environments that allow individuals to employ their human strengths in their use of technology. As discussed earlier, the properties of physical objects influence the cognitive strategies that are employed by individuals when they collaborate with their environments to perform cognitive tasks, and this principle applies to digital environments. Cognitive ethnographies, as used in distributed cognition research, reveal how people switch their attention between the properties of a representation and the properties of the thing being represented. This reveals opportunities for complex interactions that support different kinds of cognitive work (Hollan, Hutchins, & Kirsh, 2000).

So far, designs and implementations of interfaces within HCI has focused significantly on replicating real-world cognitive artefacts. For example, skeuomorphic design implements user interface elements like buttons that replicate their visual presentation in real-world scenarios through using shadows and 3D designs. This strategy is effective, because it builds on users' previous experiences and the mental models they have already developed in previous non-digital contexts. These strategies focus on the value that digital representations have by virtue of the fact that they signal to some other familiar object but overlook how the properties of the representation themselves may be exploited by users to do work (Hollan, Hutchins, & Kirsh, 2000). Hazlehurst's study on fishermen reveals how distributed cognition methodologies such as cognitive ethnographies allow us to understand the ways individuals balance their attention between the properties of external representations and the properties of what they represent. Swedish fishermen rely on a false-colour sonar display which shows fish populations as flecks on a screen to coordinate boats (Hazlehurst, 1994). In their speech, these fishermen refer to flecks, sprinkles and fish, in their interpretations of the display, and often treat the flecks themselves as fish as in "that fleck is dense enough to set the net upon" (Hollan, Hutchins, & Kirsh, 2000). The flecks in this context have the unique property of colour that allows fishermen to make judgments about when to set their net that the fish being represented cannot afford to these fishermen. The ways that human agents integrate with their environment is important for the design of human-computer interfaces, and distributed cognition provides the language and theoretical framework under which to analyse these interactions.

Distributed cognition allows us to develop models of human-computer interactions. Wright, Harrison and Fields use distributed cognition to develop the distributed information resources model. This model attempts to characterise information structures relevant to the control of action and describe how these structures may be used as resources for action. The model considers abstract information structures and their representations as distinct properties of an interaction system.

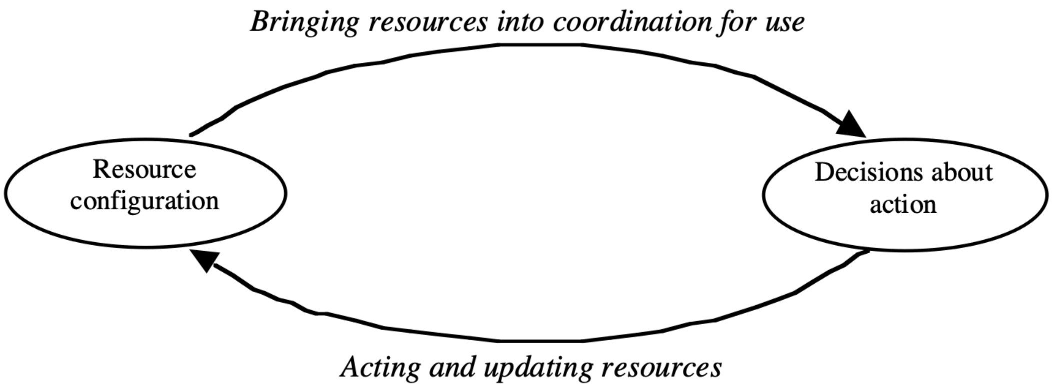

Figure 4: The cycle of interaction in the resources model (Wright, Fields & Harrison, 2000)

The abstract level of analysis allows for investigations of how the structural characteristics of information are tied to its ability to act as a kind of resource, and consequently allows for comparisons of structure between different information representations (Wright, Fields, & Harrison, 2000). For example, information structures that rely on pictorial information may enhance a piece of information's capacity to act as a resource for data visualisation, whereas text-based structures may hamper this capacity. At the representational level, the details of how information structures are distributed across people and the environment, as well as the form of these representations is considered. In the distributed information resources model, information processing occurs through a cyclic model of interaction in which the line between users and their environment is blurred. According to Wright, Harrison and Fields, actions are constrained by the configuration of internal or external representations, and when an action occurs, this configuration changes which then informs the next action, creating a cycle of interaction as illustrated in Figure 1 (Wright, Fields, & Harrison, 2000).

The ways in which different representation configurations are able to influence action are described as interaction strategies. These strategies are inextricable from the resource configurations because they presuppose specific configurations so that they are useful to the individual. In turn, the interaction strategies available to individuals are constrained by the configuration of information in their environment. The resources model thus applies the ideas of cognition as collaboration with the environment to model human-computer interactions. This allows us to identify realms of investigation through which these interactions can be better understood such as the strategies people use and how this relates to the configurations of both their internal and external representations.

IV: Distributed Cognition as a Framework for Accessible Design

While there exists plenty literature on distributed cognition and accessible design, little research has been done to explore how the two may be related. In the same way that distributed cognition may allow us to deconstruct human-computer interaction, it can also be instrumental in providing a framework through which accessible design can be considered. Building off the work of Wright, Harrison and Fields, the resources model approach to human-computer interaction can be applied to framing accessibility in a novel way. Under this model, the primary considerations of accessibility become the interaction strategies that individuals have access to, and how these relate to their internal and external representations. Successful accessible design would focus on creating interaction experiences that provide the optimal configuration of resources to allow users the opportunity to carry out the interaction strategies they are most comfortable with. The focus then switches to how information in the environment can be arranged to the maximum benefit of users, instead of identifying accommodations that can be added on to already standing information architectures as a Band-Aid to issues that users may struggle with.

Further, a distributed cognition approach to accessibility would transform the ways in which accessible design is evaluated. The pivot in focus from isolated task analysis to systems of interaction transforms the ways in which effective accessible design is evaluated. Under this approach, it becomes important to first identify which resources are available to users by way of internal and external representations and what their preferred interaction strategies are. With this information we can determine what the idea configuration of internal and external representations might be for users. The more accessible a piece of design is, the closer it's information structure would be to this configuration. To collect this kind of information, ethnographic methods of data collection may be more suitable than hyper-controlled experiments in which performance on a task-level is recorded. Usability testing in this framework consequently becomes about understanding how people interact with a piece of technology 'in the wild', and the ways their actions are limited by, or advanced by, the information structure within which they operate. A distributed cognition approach to accessibility transforms the goal of accessible design from accommodating users in a system that may not be optimal for them, to identifying the ways in which technology acts as a collaborative tool through which users carry out cognitive tasks.

Fundamentally, distributed cognition tells us about how humans are integrated with their environment and how our cognitive capacities extend beyond the brain, but is embodied within the environments in which we find ourselves. From this standpoint, what makes a piece of design truly accessible is how seamlessly it integrates with the user as a sort of 'mind extension'. The idea is not necessarily to reproduce cognitive capacities that humans can carry out within their brain on a larger scale, but to transform the space of interaction strategies available, in a way that considers and takes advantage of both strengths and weaknesses of the individual. Because accessibility standards are still based on normative interaction strategies, they focus on creating information structures that allow the individual to interact with a piece of technology in the way that most closely resembles what a 'normal' interaction is conceived to be. Using distributed cognition as a framework overhauls this approach and opens up opportunities to experimenting with and creating new interaction strategies that are more specific to the individual and consequently allow for smoother integration between the individual and a piece of design.

Thinking about accessibility in terms of integration and distributed cognition transforms approaches to understanding users by shifting towards ethnographic studies over usability studies. To understand how individuals integrate technology into their daily lives and the contexts in which they use these technologies, it is important to observe individuals ‘in the wild’, so to speak, which the controlled environments of usability studies do not allow for. On the other hand, ethnographic studies as used in developing distributed cognition theories allow for observing users in chaotic environments, over longer periods of time, which allow us to identify the different contexts of use in which users depend on a digital product to carry out specific tasks, as well as the different interaction strategies that users may implement different interaction strategies. Whilst ethnographic studies do not negate the usefulness of usability tests, they provide a context within which to interpret the results of usability tests and give deeper insights into the collaborative dynamic between individuals and technology in accomplishing cognitive tasks.

Applying a distributed cognition approach to accessibility allows us to transform our understanding and practice of accessible design in the following ways:

1. Distributed cognition gives us a framework to understand the mechanisms that make interactions work and how integration between person and technology can be disrupted or enhanced.
2. Distributed cognition provides a non-deficit approach to designing for people with disabilities
3. The language of distributed cognition allows us to categorise types of interactions and identify their relationships with the environment.

Distributed cognition allows us to understand how human-computer interactions may be disrupted and how accessibility issues arise. As raised earlier, current accessibility standards do not provide explanatory value to why some standards may work or not, and when they may fail or succeed. With distributed cognition, we are able to identify what resources are relevant in an interaction, and how they relate with the user in different ways to understand why some specific decisions about information structure may be more beneficial than others. It also provides us a framework through which we can analyse existing principles and guidelines, to form a more complete understanding of what it means for design to be accessible. For example, the principle of perceivability can be understood in terms of how internal representations may interact with external ones. External representations that are not perceivable do not form part of the resources available to the user and so cannot influence the action of users which may influence the interaction strategies available to the user. What guidelines overlook is that lack of perceivability may also create a set of interactions that may actually be optimal for some users, and so the guideline might not always be beneficial. Framing accessibility in terms of distributed cognition— for example, perceivability becomes ensuring that resources for optimal interaction strategies are made available to users, and this might not necessarily be through perception— allows for more flexibility in tailoring experiences to users.

Further, the focus on the integration between the individual and their environment allows for analysis of user behaviour in terms of their strengths and weaknesses, which provides a more complete view than the deficit approach that current accessibility guidelines employ. Deficit approaches to accessibility can be harmful because they focus solely on the identification of user points of weakness, and on creating solutions that attempt make up for these weaknesses. Whilst this may lead to the generation of helpful recommendations for differently abled users, it does not address true accessibility for these users since it fails to consider users in their entirety. Deficit approaches do not result in the best outcomes for users and may affect their ability to integrate with the design over time and use it smoothly in daily life. This partially explains why as raised earlier, accessibility guidelines only address 50% of disabled users' dissatisfaction with technologies.

Under a distributed cognition approach to accessibility, the focus is on how different users interact with technology in different contexts, and not so much what their natural abilities or disabilities are, even though these may influence the interactions under scrutiny. In this way, there are no 'disabled' users per se, because the categorisation criteria of consequence become the users' interaction strategies, which describes the information resources they may require to successfully use some digital product to achieve some task. Considering the environment as an embodied collaborator in carrying out cognitive tasks (instead of as an accessory to human thinking) shifts the focus of accessibility on how that relationship between human and environment may be enhanced. This allows us to consider both the ways in which the environment may create new ways for humans to overcome their weaknesses in carrying out tasks, as well as how their strengths may transform the sorts of strategies that would be helpful in facilitating that interaction.

Distributed cognition provides a language for describing the relationships between users which provides new opportunities for categorising and investigating human-computer interactions. Being able to call interactions by name allows us to categorise them and consequently investigate when and how different strategies are employed by individuals. Further, it allows us to develop new categories of users beyond their abilities or disabilities but by the interactive strategies on which they depend, which may transform the ways in which user groups are created in design practice. For accessible design, this is particularly important because while people's abilities and disabilities may influence their interaction strategies, it may be the case that these are not the only relevant factors that influence this, and user groups based on interaction strategies may even be more diverse than we might expect. Of course, this may also not be the case, but distributed cognition provides us with the language that allows us to frame the research question that investigates this. To investigate an issue, we must first name it, and distributed cognition allows us to conceptualise new ways of organising design practice by providing the vocabulary which facilitates this.

Accessibility is only going to become a more critical facet of design as the world becomes increasingly digitised. Consequently, it is important that we start taking the steps towards understanding the mechanisms that describe the way we interact with technology, and how designs may favour some interaction strategies over others. When we use distributed cognition as a framework for accessibility, we are able to transform the focus of the field onto the integration of users with different technologies, and how the configuration of environmental resources enable or undermine this integration process. The language of distributed cognition allows us to identify the ways in which users' internal representations and external representations interact to influence action in ways that may not be immediately obvious. In this way, not only are we are able to develop a more complete understanding of what it means for technology to be accessible, we are also able to identify more precisely what has failed to function properly when a piece of technology has poor accessibility.

Section V: A distributed cognition account of accessibility for autistic people.

"If you've met one person with autism, you've met one person with autism"

—Dr. Stephen Shore

Autism Spectrum Condition (ASC) is a neurodevelopment condition commonly associated with unique sensory processing, attention, social cognition and executive function capacities. According to the DSM V, autism is described by deficits in social communication and interaction (American Psychiatric Association, 2013), and categorises people with autism from level 1 to 3, depending on the degree of social support they are perceived to need. Even though 2.2% (Center for Disease Control and Prevention, 2006) of adults and 2.5% (Organization for Autism Research, 2019) of children in the US are diagnosed with autism, there is neither strong consensus on the causes of autism or complete understanding of the characteristics of autism. This can be partially attributed to the diverse presentation of the condition in individuals. Categories like high-functioning or low-functioning are used to create high-level generalisations of individuals with ASC, but even within these categorizations, individuals have very diverse cognitive capacities. This diversity within the community of people with ASC makes accessible design for people with ASC a complex affair.

While the presentation of autism varies wildly, some themes can be identified in the literature on the unique capacities of autistic individuals. For the purposes of this paper, we will consider central coherence, attention and working memory. Research shows that autistic individuals exhibit biases towards local processing (Happe & Frith, 2006). People with ASC's attentiveness to local and featural information may impact their ability to 'experience wholes without full attention to the constituent parts' (Kanner, 1943) and this has even been used as diagnostic criteria in the DSM V (American Psychiatric Association, 2013). The weak central coherence theory of autism draws on this to describe autism as a cognitive style. Further, attentional peculiarities have been consistently reported by people with ASC. Autistic people may exhibit over-selective attention where they are hyper-focused on specific stimuli, but may also exhibit abnormally broad focus of attention which may lead to hyperstimulation by stimuli (Allen & Courchesne, 2001). Studies also show that autistic individuals are able to maintain sustained attention in some contexts (Buchsbaum, et al., 1992) and may have difficulties with disengaging attention

(Casey, Gordon, Mannheim, & Rumsey, 1993) as well as shifting attention (Courchesne, et al., 1994).

Autism has also been associated with poorer performance on working memory tasks, particularly those requiring cognitive flexibility and planning (Kercood, Grskovic, Banda, & Begeske, 2014). Whilst studies show that autistic individuals perform similarly to typically developing individuals on verbal working memory tasks, although they exhibit difficulties in spatial working memory tasks (Williams, Goldstein, Carpenter, & Minshew, 2005). In addition, the empathising-systemising account of autism suggests that individuals with ASC have a high ‘systemising quotient’ which means that they are driven to analyse and construct rule-based systems more than typically developing individuals (Baron-Cohen, 2009). Whilst the theory itself may not completely explain autism, the behaviour it describes has been observed in autistic people. People with ASC perform worse than typically developing individuals on tasks with weaker ‘rule constraints’ (Ciesielski & Harris, 1997) and they are more likely to perform worse on open-ended tasks where no explicit strategy is implied by the instructions given because they tend to explore fewer spontaneous strategies (White, Burgess, & Hill, 2009). The cognitive profile of an individual with ASC influences the ways in which they interact with the world, and consequently has implications for the ways they interact with technology. Research examining the details of people with ASC's interaction with technology is fairly limited. However, our understanding of the cognitive profile of people with autism allows us to hypothesise about possible barriers they may face in adopting technology. The previous discussion suggests that people with ASC may experience difficulties with technology in at least two domains:

1. Making sense of digital content
2. Navigating 'open-ended' information structures

The ways in which information is presented on screens may transform the ways that people are able to interact with it and the kind of work they can do with the information. For example, users may struggle more with text-based content than photo-based content within the context of data visualisation. For autistic people, their sensitivity to different stimuli and the ways that they direct their attention may mediate how accessible information on screens is to them. It follows that if they experience issues with disengaging attention, it would be harder for them to engage with some kinds of content over others. For example, a bias towards photographic content may imply that their attentional resources may be allocated to photographic information, which means that they might miss other kinds of important information that may be presented in other forms such as text. Individuals with ASC have been found to spend larger proportions of time on images, leaving less time for texts in reading tasks due to their atypical attention patterns (Yaneva, Temnikova, & Mitkov, 2015). This is regardless of whether or not the image is relevant to the text, and so having decorative imagery which may make screens more engaging for typical users may actually impair reading comprehension for people with ASC. On the other hand, this also means that opportunities for presenting data in photographic ways may allow for better integration between autistic individuals and technology since it would be taking advantage of their natural methods of collecting information about the world.

Further, people with ASC's bias for local processing as described by the weak central coherence theory may influence the ways in which they pick up meaning from interfaces. In design practice, gestalt principles are a popular way of embedding semantic information into the organisation of a page. The gestalt principles of similarity, continuation, closure, closeness, figure, and symmetry/order describe how the mind makes sense of objects in relation to each other when they are perceived. For example, the gestalt principle of closure is what allows us to perceive Figure 5 as two triangles and three circles, instead of three lines and three pie shaped objects. These principles rely on being able to see the larger picture but may not apply well for people with ASC who have a local processing bias (and so may perceive this image as 6 disparate parts). According to Brosnan et al., autistic people use gestalt grouping principles significantly less than their typically developing counterparts (Brosnan, Scott, Fox, & Pye, 2004). This means that designs that rely on these principles may be conveying information in a way that is inaccessible to autistic individuals, which may hamper their ability to successfully integrate with these products.

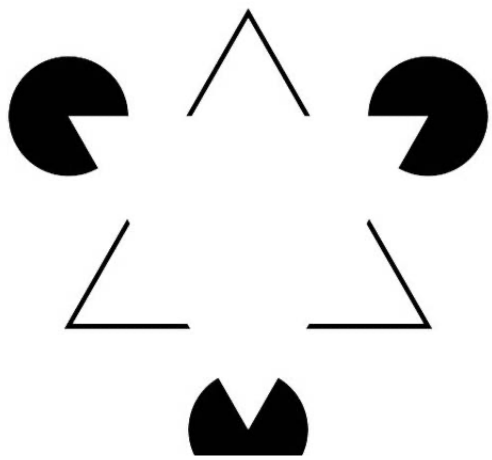

Figure 5: Gestalt principle of closure

Based on their bias towards systemizing according to the E-S theory, it is reasonable to expect that people with autism may face issues in navigating more ‘open-ended’ information structures. In scenarios where people with ASC have to use a product to carry out tasks that are not clearly defined, they may struggle—for example, trying to find an email with software that relies on implicit gesture-based interactions such as swiping for filtering and sorting functionalities may be frustrating for individuals with ASC. This also means that autistic individuals’ propensity for rule-based and systematic interactions can be leveraged to produce interfaces that integrate better with them. For example, the inclusion of gamification elements has been found to have a general positive effect for autism interventions (Camargo, Barros, Brancher, Barros, & Santana, 2019). This may be because games generally come with explicit sets of rules, and definitions of possible/impossible behaviours, which makes it easier for people who have ASC to navigate and consequently integrate with.

These potential points of user friction are by no means novel, and a number of approaches have been proposed towards addressing these. Like most approaches to accessibility, these have centered around the creation of guidelines and standards to specify accessibility. The guide for easy-to-read information proposes guidelines for making text content accessible to people with learning disabilities, which include suggestions like supporting text with images, placing images on the left and text on the right, using large font, etc. (Freyhoff, Hess, Kerr, Tronbacke, & Van Der Veken, 1998). Further, Britto and Pizzolato developed web guidelines specifically targeted for people with autism based on a bibliographic study of interface design research.

They identify ten categories within which recommendations for interfaces fall—customization, visual and textual vocabulary, engagement, redundant representation, multimedia, feedback, navigability, affordance, system status and interaction with touch screen and suggest that the most critical of these are those that address autistic people's comprehension of visual and text information. They also go further to provide specific guidelines that address the different needs specified by these categories (Britto & Pizzolato, 2016). These approaches to accessibility for autistic individuals are useful because they allow us to tangibly elevate the user experience of these users. These guidelines are practical and concrete but fall short in considering the interaction strategies of autistic individuals since they are based on usability studies, and they do not factor how the mental models of autistic people may be different. The rigidity of guidelines is particularly limiting in considering autistic populations because there is no one-size-fits-all approach that can cater to all autistic individuals because of the varied ways it presents in individuals.

On the other hand, distributed cognition allows us to make sense of the potential points of user friction by reframing pain points for people with ASC in terms of information structures and interaction strategies. Through a distributed cognition lens, we can understand how the challenges autistic people may face with technology through their interaction strategies. People with ASC can be described as having cognitive styles that transform their interaction strategies when using technology, which may give rise to the issue with technology adoption raised earlier.

Autistic people's reliance on explicit structure and instructions to carry out tasks suggests that they may have difficulties with flexibly trying different interaction strategies when navigating interfaces, which would create frustrations when interacting with information structures that rely on some level of user inference. This explains why guidelines such as labelling icons can be helpful for people with ASC—where typical individuals may be reasonably be able to understand icons, people with ASC may struggle with the extra level of inference required to make sense of unlabelled icons. Since they are biased towards systemising, this also means that their interactions with technology will be influenced by this: when faced with uncertainty, they may choose systemising strategies of problem-solving (Craig, Grossman, & Krichmar, 2017). In the case that their mental models are in line with the information structure of the interface, this will result in successful interactions; however, a mismatch may lead to errors and frustration. So, if buttons were not labelled but were located in parts of the screen that corresponded with the autistic individual's mental model where location and function are related, then user friction may be avoided.

We can also consider their bias for local processing in terms of interaction strategies and representations. Since autistic people tend to process information at a local level, their internal representations may be structured with an emphasis of these local details at the expense of more overarching gestalt information. Because interaction strategies may be influenced by these representations, people with autism may face difficulties using interfaces that do not allow flexibility in interaction strategies to accommodate their internal mental models. For example, screens that use physical closeness of kinds of information as the only indicator of category may be inaccessible to autistic individuals who may not be able to infer that items in the same location belong to the same category. This idea also extends further into understanding how attentional peculiarities of autistic people may influence their interaction with content on interfaces. Hyper-focusing on some elements e.g., photos over others, means that while the same information may be presented to both typically developing and autistic individuals in similar ways, their attentional patterns make it such that the information actually available to these different user groups may be different, which may give rise to different interaction strategies for navigating interfaces. Thus, using distributed cognition as a framework allows us to identify how representations and interaction strategies facilitate successful user experiences.

Being able to understand accessibility issues of people with ASC in terms of interaction strategies is particularly important because of the variety of ways that autism presents in people. What may work for some individuals may not work for others, and so guidelines may result in positive outcomes for some subsets of autistic people whilst being detrimental for others. This also allows us to generate segmentations of user groups within the group of people with ASC based on their interaction strategies—for example, rather than grouping users by their attention deficits, we can group them by the kinds of information they depend on in navigating interfaces, e.g., photo vs text etc. This way, we can identify the unique resources that different sub-groups of autistic people may need to satisfactorily collaborate with a piece of technology. For example, while photographic content may be helpful for autistic individuals with attention disengagement issues, they may also potentially create overwhelming sensory input for people with sensory processing issues. What makes this perspective so useful is that it can apply beyond accessibility for disabled people and nudge us more in the direction of accessibility for all people. Even within typically developing user groups, there are likely different interaction strategies that different people rely on. Designing products with these strategies in mind allows us to get closer to the goal of designing products that are universally accessible.

Beyond contextualising current approaches to accessibility, distributed cognition's value as a theoretical framework is in transforming the kinds of questions and opportunities we search for when thinking about accessibility. While current approaches ask, 'what is wrong and how can we fix it?', distributed cognition asks 'how do people understand the world and how do they collaborate with it?' In the case of individuals with autism, one opportunity this creates is the examination of the temporal component of user experiences. Whilst we can reasonably infer from usability testing that people with autism might struggle with vague iconography, we do not learn about how this struggle may change over time and how temporal changes to the information structure may affect how people with autism integrate with a piece of technology. Technology is far from static and so it is important to account for the role that time may play in understanding accessibility. Applying current understandings of the experience of autism suggests that beyond having an explicit information structure, individuals with autism may benefit from systematic changes in a piece of technology as they use it and may struggle with integrating with interfaces that change dramatically in unpredictable ways over time. For example, an application that introduces feature updates incrementally instead of changing several features at once may be more accessible to people with ASC. This suggests that the ways that interface changes are managed over time is important to achieving accessible design.

Further, with a distributed cognition approach to accessibility for people with autism no longer becomes about only the individual, but also the support networks on which they depend. Since many people with ASC still depend on their parents, it becomes necessary to consider how cognition is distributed not only across the autistic individual and a piece of technology, but also how it is distributed across the individual, technology, and their parents. For a piece of technology to be accessible in this context, it has to also consider the interactions between the individual with ASC and their caretakers and the interaction strategies that the caretaker may employ in understanding an interface. For example, non-verbal autistic children may collaboratively use speech production technology, and if parents are unable to navigate the interface, it may have an indirect negative impact on accessibility for the autistic child. With traditional approaches to accessibility, the user is considered as an isolated entity, which means that this interrelationship between parents, children and technology is overlooked. Because distributed cognition research methods do not isolate the participant but observe their behaviours in different contexts, this allows us to identify the different people or other technologies that fit within the distributed network of the individual, which allows us to develop designs that people with ASC may integrate better with.

Limitations

While the distributed cognition perspective lends us increased flexibility in defining and practicing accessible design, the majority of the discussion in this paper has been theoretical, and so is limited by lack of empirical backing. To test this theory and its implications, it would be especially important to conduct research (e.g. cognitive ethnographic studies) on the different interaction strategies that people use when collaborating with technology to carry out tasks. This will allow us to identify if there are overarching trends among different user groups and what environmental resources different groups may rely on to carry out tasks. Another limitation of this perspective is that it is difficult to develop a single metric of integration for different products and users. Because integration with digital products may be influenced by a variety of factors, including users' interest in the product, it is difficult to isolate unique factors that may affect this for different users. A practical downside of this approach is also that ethnographic research occurs over longer periods of time and is more resource intensive than usability studies, which makes this approach harder to implement in real-world contexts.

Conclusion

The central claim of distributed cognition is that we interact with our environment to do cognitive work. Whilst this claim may not completely encompass human cognition, it provides a useful framework through which we can understand the role of technology in our lives. According to this perspective, technology acts as a cognitive artefact, an extension of our minds that affords us capacities beyond our natural human abilities. Whilst this idea has been applied to understanding HCI, its implications for accessibility have not been considered in the current literature. Applying a distributed cognition framework to accessible design allows us to transform our understanding of accessibility by emphasising integration and focusing on the ways that information structures transform the ways people can interact with an interface. Whilst this does not render accessibility guidelines useless by any means, it allows us to apply them within appropriate contexts and broadens the scope of accessible design so that we can approach the larger goal of universal design.

Acknowledgments

I would like to thank Professor Lorie Loeb for her guidance throughout my college career and her in development of this paper. I would also like to thank Professor Jonathan Phillips for his advice and feedback in assembling my different ideas on design and distributed cognition into this paper. Finally, I would like to thank my close friends and family for supporting me and being a sounding board for my thoughts.

References

International Organisation for Standardisation. (2008, July). *Ergonomics of human-system interaction — Part 171: Guidance on software accessibility* . Retrieved from International Organisation for Standardisation: https://www.iso.org/standard/39080.html

Persson, H., Åhman, H., Yngling, A., & Gulliksen, J. (2015). Universal design, inclusive design, accessible design for all: different concepts—one goal? On the concept of accessibility—historical, methodological and philosophical aspects . *Universal Access in the Information Society*, 505–526.

Al-Wakeel, L., Al-Ghanim, A., Al-Zeer, S., & Al-Nafjan, K. (2015). A Usability Evaluation of Arabic Mobile Applications Designed for Children with Special Needs — Autism. *Lecture Notes on Software Engineering*.

Dattolo, A., & Luccio, F. L. (2017). Accessible and Usable Websites and Mobile Applications for People with Autism Spectrum Disorders: a Comparative Study. *EAI Endorsed Transactions on Ambient Systems*.

Disability Rights Commission . (2004). *The Web Access and Inclusion for Disabled People.* Retrieved from https://www.city.ac.uk/__data/assets/pdf_file/0004/72670/DRC_Report.pdf

Rømen, D., & Svanæs, D. (2012). Validating WCAG versions 1.0 and 2.0 through usability testing with disabled users. *Universal Access in the Information Society*, 375–385.

Leuthold, S., Bargas-Avila, J. A., & Opwis, K. (2008). Beyond web content accessibility guidelines: Design of enhanced text user interfaces for blind internet users. *International Journal of Human-Computer Studies*, 257-270.

Savidis, A., & Stephanidis, C. (1998). The HOMER UIMS for dual user interface development: Fusing visual and non-visual interactions. *Interacting with Computers*.

Power, C., Freire, A., Petrie, H., & Swallow, D. (2012). Guidelines are only half of the story: accessibility problems encountered by blind users on the web. *CHI '12: Proceedings of the SIGCHI Conference on Human Factors in Computing System*, (pp. 433–442).

Hutchins, E. (2000). Distributed cognition. In *International Encyclopedia of the Social and Behavioral Sciences.* Elsevier Science.

Giere, R., & Moffatt, B. (2003). Distributed Cognition:: Where the Cognitive and the Social Merge. *Social Studies of Science*, 301-310.

Rogers, Y., & Ellis, J. (1994). Distributed Cognition: An Alternative Framework for Analysing and Explaining Collaborative Working. *Journal of Information Technology*.

Zhang, J., & Norman, D. A. (1994). Representations in Distributed Cognitive Tasks . *Cognitive Science* .

Zhang, J., & Patel, V. L. (2006). Distributed cognition, representation, and affordance. *Pragmatics and Cognition*.

Griffin, P., & Cole, M. (1980). Cultural amplifiers reconsidered. *The social foundations of language and thought*.

Sinha, G., Shahi, R., & Shankar, M. (2010). Human Computer Interaction. *International Conference on Emerging Trends in Engineering and Technology, ICETET.* Goa: IEEE.

Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*.

Wright, P., Fields, R., & Harrison, M. (2000). Analyzing human-computer interaction as distributed cognition: the resources model. *Human-Computer Interaction*.

Miller, D. J., & Robertson, D. P. (2010). Using a games console in the primary classroom: Effects of 'Brain Training' programme on computation and self-esteem. *British Journal of Educational Technology*.

American Psychiatric Association. (2013). *Diagnostic and Statistical Manual of Mental Disorders, 5th Edition: DSM-5.*

Center for Disease Control and Prevention. (2006). Retrieved from Center for Disease Control and Prevention: https://www.cdc.gov/ncbddd/autism/data.html

Organization for Autism Research. (2019, January 03). *NEW STUDIES ESTIMATE AUTISM PREVALENCE AT 1 IN 40* . Retrieved from Organization for Autism Research: https://researchautism.org/new-studies-estimate-autism-prevalence-at-1-in-40/

Hutchins, E. (1995). *Cognition in the Wild.* MIT press.

Happe, F., & Frith, U. (2006). The Weak Coherence Account: Detail-focused Cognitive Style in Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders*.

Kanner, L. (1943). Autistic disturbances of affective contact. *Nervous Child*, 217-250 .

Kercood, S., Grskovic, J. A., Banda, D., & Begeske, J. (2014). Working memory and autism: A review of literature. *Research in Autism Spectrum Disorders*, 1316-1332.

Williams, D. L., Goldstein, G., Carpenter, P. A., & Minshew, N. J. (2005). Verbal and Spatial Working Memory in Autism . *Journal of Autism and Developmental Disorders*.

Allen, G., & Courchesne, E. (2001). ATTENTION FUNCTION AND DYSFUNCTION IN AUTISM. *Frontiers in Bioscience*.

Buchsbaum, M., Siegel, B. V., Wu, J. C., Hazlett, E., Sicotte, N., Haier, R., . . . Sabalesky, D. (1992). Brief report: Attention performance in autism and regional brain metabolic rate assessed by positron emission tomography. *Journal of Autism and Developmental Disorders*, 115-125.

Casey, B. J., Gordon, C. T., Mannheim, G. B., & Rumsey, J. M. (1993). Dysfunctional attention in autistic savants. *Journal of Clinical and Experimental Neuropsychology*, 933-946.

Courchesne, E., Townsend, J., Akshoomoff, N. A., Saitoh, O., Yeung-Courchesne, R., Lincoln, A. J., . . . Lau, L. (1994). Impairment in shifting attention in autistic and cerebellar patients. *Behavioral Neuroscience*, 848-865.

Baron-Cohen, S. (2009). Autism: The Empathizing–Systemizing (E-S) Theory . *The Year in Cognitive Neuroscience 2009*, 68-80.

Ciesielski, K. T., & Harris, R. J. (1997). Factors related to performance failure on executive tasks in autism . *Child Neuropsychology*.

White, S. J., Burgess, P. W., & Hill, E. L. (2009). Impairments on "open-ended" executive function tests in autism . *The International Society for Autism Research*.

Yaneva, V., Temnikova, I., & Mitkov, R. (2015). Accessible Texts for Autism: An Eye-Tracking Study. *ASSETS '15: Proceedings of the 17th International ACM SIGACCESS Conference on Computers & Accessibility*, (pp. 49-57).

Brosnan, M. J., Scott, F. J., Fox, S., & Pye, J. (2004). Gestalt processing in autism: failure to process perceptual relationships and the implications for contextual understanding. *J Child Psychol Psychiatry*, 459-69.

Camargo, R. M., Barros, R. M., Brancher, J. D., Barros, V. T., & Santana, M. (2019). Designing Gamified Interventions for Autism Spectrum Disorder: A Systematic Review. *Joint International Conference on Entertainment Computing and Serious Games*, (pp. 341-352).

Freyhoff, G., Hess, G., Kerr, L., Tronbacke, B., & Van Der Veken, K. (1998). *Make it Simple. European Guidelines for the Production of Easy-to-Read Information for People with Learning Disability.* Retrieved from Cornell University Library: https://ecommons.cornell.edu/handle/1813/76527

Britto, T. C., & Pizzolato, E. B. (2016). Guidelines of Interaction and Interface Design for People with Autism Spectrum Disorder. *Conference: Ninth International Conference on Advances in Computer-Human Interactions - ACHI 2016.*

Craig, A. B., Grossman, E., & Krichmar, J. L. (2017). Investigation of autistic traits through strategic decision-making in games with adaptive agents. *Scientific Reports*.

Phillips, B., & Zhao, H. (1993). Predictors of assistive technology abandonment. *Assist Technol*.

Ringland, K. E., Wolf, C. T., Boyd, L. E., Baldwin, M. S., & Hayes, G. R. (2016). Would You Be Mine: Appropriating Minecraft as an Assistive Technology for Youth with Autism. *ASSETS '16: Proceedings of the 18th International ACM SIGACCESS Conference on Computers and Accessibility.*

WebAIM. (2021, Februrary 26). *WebAIM*. Retrieved from WebAIM's WCAG 2 Checklist: https://webaim.org/standards/wcag/checklist

D'Angelo, G., & Rampone, S. (2018). *Security and Resilience in Intelligent Data-Centric Systems and Communication Networks.* Academic Press.