

Loughborough University Institutional Repository

Towards accessible interactions with pervasive interfaces, based on human capabilities

This item was submitted to Loughborough University's Institutional Repository by the/an author.

Citation: ATKINSON, M.T. ... et al, 2010. Towards accessible interactions with pervasive interfaces, based on human capabilities. IN: Miesenberger, K. ... et al, (eds.). Computers Helping People with Special Needs, Part I: 12th International Conference, ICCHP 2010, Vienna, Austria, July 14-16, 2010, Proceedings. Lecture Notes in Computer Science; 6179. Berlin : Springer Verlag, pp.162-169.

Additional Information:

- This conference paper was published in “Computers Helping People with Special Needs, Part I” [© Springer Verlag Berlin / Heidelberg]. The original publication is available at <http://www.springer.com/computer/hci/book/978-3-642-14096-9>

Metadata Record: <https://dspace.lboro.ac.uk/2134/6519>

Version: Accepted for publication

Publisher: © Springer Verlag Berlin / Heidelberg

Please cite the published version.

This item was submitted to Loughborough's Institutional Repository (<https://dspace.lboro.ac.uk/>) by the author and is made available under the following Creative Commons Licence conditions.



For the full text of this licence, please go to:
<http://creativecommons.org/licenses/by-nc-nd/2.5/>

Towards accessible interactions with pervasive interfaces, based on human capabilities

Matthew Tylee Atkinson¹, Yunqiu Li¹, Colin H. C. Machin¹, and David Sloan²

¹ Department of Computer Science, Loughborough University, LE11 3TU, UK
{M.T.Atkinson,K.Y.Li,C.H.C.Machin}@lboro.ac.uk

² School of Computing, University of Dundee, DD1 4HN, UK
dsloan@computing.dundee.ac.uk

Abstract. We draw on the literature, established standards and practices, contemporary funded projects and ongoing proof-of-concept work to argue for accessibility advocates to capitalise on the current paradigm shift towards ubiquitous, distributed and collaborative information systems. We discuss the contemporary interaction trends and accessibility challenges that give rise to our proposal; describe and provide an argument for capability-based development over the current, more device-focused, approach and document ongoing proof-of-concept and funded development work in this area.

1 Introduction

One perennial barrier to improving the accessibility of ICTs is that of providing incentive for developers and content creators to make their work accessible. By social means, this may be partly addressed by recognising that accessibility barriers affect a range of people [1], including older users [2]. A further major challenge is that there is a growing requirement for accessibility to be appropriately managed in applications that involve collaborative and/or social interaction, including situations where an application may be used in company, for example older people's use of e-mail [3].

One cause of the classic barrier to accessibility provision is the perceived expense of its incorporation into development practices. Due to the present trends towards web-based and progressively more ubiquitous applications, however, these practices, are currently in flux, presenting an ideal opportunity for our field to develop techniques that can ease the development process for heterogeneous applications, whilst also mainstreaming accessibility techniques. This paper proposes our vision of such a technique, which is based on the notions of (a) adaptivity as a means to mainstream accessibility and (b) human capabilities as the primary, portable, measure of adaptation requirements.

Our contributions are: (1) a discussion of the disconnect between the trend towards ubiquity and the current methods of providing accessibility, leading to the need for portable adaptive accessibility being identified; (2) the overview of and justification for our proposed solution to the technical challenge and (3) a

description of how our techniques are being developed and tested as part of the Sus-IT³ project.

Adaptive systems raise serious matters of privacy, user interruption, security and ethics (e.g. do we have a responsibility to alert the user to suspected medical problems?) These are key matters and are being considered in parallel to the technical work described here [4].

2 Complementary and Closely-related Work

The need for improved accessibility is inherently a social one and social factors are key to the successful provision of accessibility—however, the scope of this paper is the technical aspects of the development process.⁴ It is important to keep in mind that adopting practices such as User-Centred Design or Design for All can motivate developers to improve accessibility and provide useful data on how this might be achieved. However these practices improve the design, not technical development processes—the latter being the focus here.

Some key complementary technical outputs include the growing array of highly-focussed accessibility adaptations being developed by researchers in the area. Examples include models of motor difficulties and compensation techniques (from key-debouncing to models of performance [5, 6]) and sensory impairments, such as Jefferson’s techniques for compensating for colour deficit [7]. The focussed nature of these projects—each addressing only one specific problem—makes them highly suitable candidates for becoming accessibility adaptations, as will be elaborated upon later.

Related work includes that of Gajos, who developed techniques to render a GUI corresponding to an abstract interface specification in which the concrete interfaces may be rendered on desktops, portables, PDAs and mobile telephones. These techniques were extended to support users with some kinds of vision or motor impairments—including combinations of the two [8].

IBM’s Web Adaptation Technology (WAT) [2] is able to make changes to system accessibility settings based on both detected problems (such as tremors whilst typing) and at the request of the user. The way accessibility settings are presented and applied is interesting, for two reasons: (1) the many accessibility-related settings buried deep within the operating system are brought into the WAT in-browser interface so that users can easily find and change them; (2) the settings are applied not just to web content but the application as well—i.e. a text size increase will affect the browser’s menus too. This is almost certainly what the user would wish for, going a long way to addressing the issues of user awareness of ATs (for one specific application), but the literature yields very few examples of this sort of approach.

³ <http://sus-it.lboro.ac.uk/>

⁴ Most workpackages in the Sus-IT project are focussed on social aspects, providing requirements, input data and validation for technical processes and artefacts developed as part of the project.

The ÆGIS project [9] is a current research effort that seeks to better deploy ATs on desktop and mobile systems. This project is focussed on providing an open framework and standards for deployment of ATs, which will provide a foundation for future work. The Sus-IT project seeks to develop reasoning processes that could run on top of such a framework, to direct deployment of the correct ATs for a given user at a given time, thus at least partially addressing users' lack of awareness of available ATs.

The tripartite ISO 24751, based on work from the IMS AccessForAll project, covers accessibility profiling of a learner and of an electronic learning object and describes how these might be used together in a learning management system to deliver appropriately-adapted resources for an individual. While it provides a useful reference in terms of accessibility profiling, its scope does not extend to methods of initial data population, nor of maintenance of accurate user profile information over time. Additionally, the adaptations it considers are couched in very machine-oriented terms (e.g. font enlargement, cursor size or double-click speed), which would not necessarily be portable to new device types.

3 Adaptivity as a Means to Mainstream Accessibility

Individuals' fluctuating capabilities can lead to a lack of awareness of accessibility needs and the appropriate solution for them [4]—so a system that can help close this gap by matching a person with appropriate accessibility accommodations at a given time can provide a valuable extension to existing accessibility support. While employing inclusive design practices is important, since perceptions of an interface and the situations in which applications are used are variable, no single interface can satisfy the needs of all users at all times.

Adopting standards can increase the likelihood that a given accessibility adaptation will be able to render any given content in a manner accessible to a given user, because standard methods for adaptation may be used (such as the DOM [2]). It is of key importance, however, to note that adopting such baseline standards is *just the beginning*—doing so provides a common platform for adaptations to build upon but is not the final step in ensuring accessibility, due to the dynamic diversity of users, devices and environments.

For example, it is far from always possible to anticipate the context in which information or applications may be accessed. Often, users would benefit from slight changes (adaptations) to the manner in which information is rendered, based on their capabilities and preferences, device and environment (which may change suddenly, e.g. in terms of light/noise levels, or familiarity), mandating a reactive adaptation to be executed (such as contrast changes, or redirection of output to other modalities).

In order to minimise unintended side-effects, adaptations would have to be focussed on solving one particular accessibility problem, be aware of specific users' characteristics *and* be applicable in a very localised fashion (e.g. be able to read aloud the content of long documents to a user with vision fatigue, but not the containing application or browser's menus or toolbar, if this is the user's

preference). This makes some traditional ATs unsuitable in their current form (often applicable to an entire output modality, rather than specific applications or parts thereof), and the adaptations developed through research, introduced above, seem more ideal. Infrastructure to allow adaptations to be targeted at specific parts of applications (akin to Gajos’ and others’ abstract UI systems described above) would be required to support this. In addition, the reasoning system would have to recognise that combinations of adaptations could produce “interference” and that the effects of this may vary for each user [10].

4 Capability-based Adaptive Accessibility

In order to reason about users, devices and adaptations, reliable methods and metrics are required to match users to devices—i.e. to relate device input and output functionality to the human capability requirements for providing or perceiving that input or output. Examples include optimising content displayed on the screen of a PDA for an outdoor environment when a substantial change in light level is detected (or switching to audio output, if necessary). Such metrics are also required for matching users to accessibility adaptations, including some degree of temporal flexibility, due to users’ changing needs over time. The diversification in terms of device form-factors and platforms mandates a portable solution (or already-learned user capabilities and preferences may not be transferable to new devices, applications or content). Currently the problems of reliable discovery and monitoring of user capabilities, as well as the interactions of potentially-conflicting accessibility adaptations, present formidable challenges.

However, there is a body of work that provides a basis for such reasoning. There exists an internationally-recognised standard for classifying human capabilities: WHO’s ICF [11]. Existing work by Billi uses the ICF to assist with modelling human-computer interaction scenarios [12]. The approach taken by Billi is to extend the existing classification in two ways: (1) Extending the human skills classification via an appendix to include more fine-grained details such as the user’s performance in clicking, dragging and moving the mouse, and (2) adding a device functionality appendix, in a similar vein to the extra human capabilities, which categorises device functions such as screen resolution, colour depth, whether joystick input is supported and whether voice input is supported.

It is argued that, although this approach could help improve modelling of interaction with—and therefore hopefully the accessibility of—*current* devices, it is not the correct approach for the medium- and long-term. This is because it is device-dependent. A more sustainable approach would be to express the additional classifications in terms of *human* functionality, as (a) human capability classifications change incredibly slowly over time⁵ and (b) expressing the capability requirements of device functions in terms of human capabilities allows the same classification to support new devices as they are developed. Consider the mouse scroll wheel as an example: this may eventually be phased-out in favour of

⁵ For the purposes of this work, the gamut of capabilities of the human species may be considered static.

multi-touch control surfaces. When this happens, the information “user can use the mouse scroll wheel” becomes irrelevant. It would be better to express that the user has a certain level of sustainable, repeatable bending in their finger(s), as this information can be used for reasoning independently of device class.

Expressing as much data as possible in terms of human capability requirements has another benefit: it allows any “interference” relationship between adaptations to be exposed within the reasoning process. For example, the use of larger text may have a positive effect on visual acuity, but—through enlarging the volume of information when rendered visually—it may impose a requirement for scrolling, which may burden the user’s motor capabilities. Each adaptation could be thought of as an overlay, or delta, to the user’s overall capabilities data structure (which would be similar to that prescribed by the ICF, though not necessarily fully-populated). Further, when two or more users are collaborating, the reasoning technique can consider the combination of each user’s requirements in the same way, effectively simply tightening the constraints on the process.

The promise of capability-based reasoning is to unify the process by which: (a) users can be matched to devices; (b) the rendering of information can be tailored for each user-device pair and (c) the application of accessibility adaptations can be orchestrated. However there are significant challenges to overcome. Using human capabilities as the basis for this reasoning requires a means to: (1) accurately assess and subsequently monitor user capabilities, as well as (2) the effects of the environment on a user and (3) the interference caused by the application of multiple adaptations.

The following sections provide a technical overview of our proposed capability-based reasoning process, particularly regarding how such a system can be made sufficiently passive (so as to avoid distracting users) and self-regulating, through the provision of metrics to enable performance and reliability to be assessed. A description of the user-visible adaptivity components is given in other work [4].

5 The User as the Expert

Our system seeks to map minor-to-moderate impairments brought about by capability change on the part of the user, environment or device to the appropriate accessibility adaptations. We begin with the basic assumption that the user’s actions can enable us to learn a great deal about their capabilities and preferred adaptations. Further, we acknowledge that: (a) in many cases it may not be possible to directly measure individual capabilities and (b) users may employ undetectable means to address accessibility problems (e.g. reading glasses).

Some basic features of our model follow. A user has a set of capabilities (as in classifications such as ICF). A computer system’s purpose is to transmit information to and receive information from the user in various modalities. Four are self-explanatory: audio, vision, motor and cognition. Our system also considers two further channels: time (as devices may impose constraints and users may have impairments that alter the time needed or available for certain interactions) and volume of information (because impairments and situations can alter

the amount of information a user may be capable of perceiving or processing). There is a 1:n relationship between a modality and the capabilities that operate over it (i.e. both visual acuity and colour perception are visual capabilities).

If, at any time, the requirements of the system, device and environment outweigh the level of capability afforded by the user, then an adaptation must be made. Figure 1 demonstrates such a situation and two alternative courses of action. According to the adaptations available, multiple courses of action may be available; a decision-tree algorithm similar to those used by Gajos could be used to select the most appropriate, based on a statistical model of the user’s preferred actions in similar situations (e.g. when web-browsing vs. when watching DVDs). As suggested by the figure, a typical system contains many points that can be monitored for such activity, thus providing us with a passive way to monitor for capability change.

The notion of change is of key importance here: the system could not always pose an accurate picture of the user’s capabilities in detail. Instead it should attempt to use monitoring techniques to gauge the *change* in capabilities over time. It is also important to recognise that some of the most profound adaptations (such as putting on reading glasses) will be undetectable to the system—and that the fact that an adaptation is made does not indicate that the user’s capability was outweighed by the requirements at the time. However, the information continuously gathered will still be of use in formulating a list of possible useful adaptations, from which the user can select after requesting help. Our system seeks to raise awareness of and augment the processes by which users may receive help, not to replace existing support methods.

Sometimes it may be necessary to interrupt the user to perform a more direct assessment of capability in order to calibrate the system. The approach we are currently trialling is to trade some rigour in the testing with making the tests into engaging “mini-games” for the user. We only require a moderate indication of capability (such as double-click speed or reading text size) but certainly do not wish to alienate users. It may also be necessary to present a choice of adaptations in a given situation when the computer cannot determine a “clear winner.” Finally, the user’s feedback after an adaptation has been applied—namely acceptance, rejection or a request for further help—would be taken into consideration in much the same way as it is in systems such as those discussed elsewhere [4].

6 Classifications

The reasoning process must be capable of mapping capability information into the appropriate adaptations (expressed as in ISO 24751). To assist with this mapping, a small number of classifications are being developed. Adaptations are classified on several dimensions; a brief example is given Table 1. Some potential side-effects (in terms of channels that may become restricted when an adaptation is applied) are given (which, in turn, may require supplementary adaptations to be made to counteract the side-effects—monitoring of users’ actions can indicate which course of action is preferred in different situations).




Fig. 1. Cloud representing the entire system, with particular adaptation hooks pointed out. The user has activated one of these and this has had consequences which could be dealt with in one of two ways.

Applications also need to be classified by the nature of their interaction and presentation styles, which can give us a general expression of expected user behaviour (e.g. in media players, little interaction is expected in comparison to web authoring software). If, after an adaptation is applied, the user's behaviour changes unexpectedly—providing much more input than expected, or switching to keyboard after historically having used the mouse—we may surmise that the adaptation has caused this extra burden (e.g. full-screen magnification would introduce a significant scrolling requirement). If the adaptation is rejected or further help is sought by the user, the system could take the side-effects into account in future and either avoid the problem adaptation or perhaps suggest a supplementary one (e.g. lower the volume of information presented) to counteract its side-effects. Work is ongoing in this area and will be described in due course.

Table 1. Adaptation types (brief excerpt from our classification).

Action	I/O Scope	Method	Modalities	Side-effects
Double-click speed	I Device	Parameter	M	T
Word document zoom	O App.	Parameter	V(Acuity)	I
System text size	O OS/App.	Parameter	V(Acuity)	I
Screen mag.	O Device	Filter	V(Acuity,Colour)	I

7 Conclusion and Further Work

We have presented part of our vision for mainstreaming accessibility through adaptations and the adoption of human capabilities as a benchmark for doing this. Work on the temporal aspects of the reasoning system is ongoing. Our reasoning process is being integrated into a pilot system, to be tested shortly. Longitudinal data gathering with users for approximately one year will be carried out towards the end of the project—the collected data on adaptation usage and prediction accuracy will be an additional output. This work was funded under the UK's New Dynamics of Ageing cross-council research programme.

References

1. Keates, S., Clarkson, P.J.: Countering design exclusion: bridging the gap between usability and accessibility. *Universal Access in the Information Society* **2** (October 2003) 215–225
2. Richards, J.T., Hanson, V.L.: Web accessibility: a broader view. In: WWW '04: Proceedings of the 13th international conference on World Wide Web, New York, NY, USA, ACM (2004) 72–79
3. Sayago, S., Blat, J.: Telling the story of older people e-mailing: An ethnographical study. *International Journal of Human-Computer Studies* **68**(1-2) (January 2010) 105–120
4. Sloan, D., Atkinson, M.T., Machin, C.H.C., Li, Y.: The potential of adaptive interfaces as an accessibility aid for older web users. In: W4A. (July 2010) *in press*.
5. Hurst, A., Hudson, S.E., Mankoff, J., Trewin, S.: Automatically detecting pointing performance. In: Proceedings of the ACM International Conference on Intelligent User Interfaces (IUI) 2008. (January 2008) 11–19
6. Biswas, P., Sezgin, T.M., Robinson, P.: Perception model for people with visual impairments. In: Proceedings of the 10th International Conference on Visual Information Systems (VISUAL 2008). LNCS 5188 (2008) 279–290
7. Jefferson, L., Harvey, R.: An interface to support color blind computer users. In: CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems, New York, NY, USA, ACM (2007) 1535–1538
8. Gajos, K.Z., Wobbrock, J.O., Weld, D.S.: Automatically generating user interfaces adapted to users' motor and vision capabilities. In: UIST '07: Proceedings of the 20th annual ACM symposium on User interface software and technology, New York, NY, USA, ACM (2007) 231–240
9. Korn, P., Bekiaris, E., Gemou, M.: Towards open access accessibility everywhere: The ægis concept. *Universal Access in Human-Computer Interaction. Addressing Diversity* **5614/2009** (July 2009) 535–543
10. Kurniawan, S.H., King, A., Evans, D.G., Blenkhorn, P.L.: Personalising web page presentation for older people. *Interact. Comput.* **18**(3) (2006) 457–477
11. World Health Organization: International classification of functioning, disability and health (ICF). <http://www.who.int/classifications/icf/> (2001)
12. Billi, M., Burzagli, L., Emiliani, P., Gabbanini, F., Graziani, P.: A classification, based on ICF, for modelling human computer interaction, for modelling human computer interaction. In: *Computers Helping People with Special Needs*, Springer Berlin / Heidelberg (2006) 407–414