Using a Common Semantic Structure to Provide Comparable Contextual Models of Users and Technology

Matthew Bell, Colin Machin, and Roger Stone

Loughborough University, Loughborough, Leicestershire, LE11 3TU, UK
{m.j.bell,c.h.c.machin,r.g.stone}@lboro.ac.uk

Abstract. The accessibility solution that is appropriate for an individual in a given situation may be provided through variations in the choice of device, assistive technologies (AT) and adaptations used. Profiles can be created to represent users and technology, however, owing to trade-offs between profile specificity and transportability, there is currently no universally accepted method for creating profiles for holistic interaction.

This paper describes an approach which represents both user and technology in symmetrical (hierarchical) recursive profiles, using a vocabulary that moves from device-specific to device-agnostic capabilities. Through the use of semantic relationships, capabilities can be attributed —and accessibility comparisons made—at varying (appropriate) levels of granularity, using contextually comparable data.

Where accessibility problems are identified, they are described in terms of the gap between the capabilities of the user and technology, inherently providing a functional description of the support required. Speculative augmentation can then be used to evaluate different solutions in order to maximise accessibility for the individual.

Keywords: Evaluation of Accessibility, Usability, and User Experience, User and Context Modelling and Monitoring, Semantics.

1 Introduction

The accessibility solution that is appropriate for an individual in a given situation may be provided through variations in the choice of device, assistive technology (AT) and adaptations used (Sloan *et al.* 2010). Different devices are often used in different contexts or for different tasks owing to their individual characteristics. A desktop computer may be used at work or for creating documents, a Smart TV used at home or for consuming online media, and a tablet used 'on the move' or for checking emails. ATs can then be used to augment the capabilities of devices or users in order to improve their accessibility. Screen readers may be used to provide an audio version of a device's visual content and a hearing aid may be used to increase the range of sounds that a user can perceive. Finally, adaptations can be used as micro-ATs, providing personalisation that increases the range of users who are able to access content (Vanderheiden

C. Stephanidis and M. Antona (Eds.): UAHCI/HCII 2014, Part I, LNCS 8513, pp. 203–214, 2014.

2008). By moving closer to a screen a user can improve their ability to see its contents; the same result can also be achieved by increasing the size of the text on the screen. As they are intended for customisation, many adaptations are not labelled as accessibility options, however device-specific settings are increasingly being included in profiles describing accessibility needs.

In order to maximise accessibility, a model is therefore required that can compare between profiles representing the user and their device, as well as any ATs and adaptations they are using. The variety in interaction paradigms results in a need for specificity in order to capture the nuances of each particular device, interface or control. However, as users frequently interact with multiple technologies there is also a need for profiles that are generic enough to be transportable.

This paper will provide a framework for identifying accessibility problems by moving between device-specific and device-agnostic profiling through the logical structuring of profiles. The approach provides flexibility during both profile acquisition/maintenance and the matching process. Dynamic comparison of profiles at varying (appropriate) levels of granularity allows discovery of accessibility issues to be performed efficiently. In addition it is also possible for contextual information to be represented using the same structure, providing a measure of the applicability of data to a given situation.

2 Background

The presence (or absence) of accessibility issues between a user and a piece of technology can be predicted by simulating their interaction using representative models and profiles. Given the variety found within users, technologies and interaction paradigms, there is no universally accepted method of modelling any aspect of the simulation. This is due in part to the range of hierarchical levels and functional layers of abstraction for which simulations are created. Levels of abstraction describe the way that higher-level tasks can be built up from lower-level ones (e.g. hand-eye co-ordination). Layers of abstraction describe the way that data can be viewed in multiple frames of reference; with lower layers providing structural information and higher layers storing specific contextualised measurements.

Biswas & Robinson (2010) describes the issue of fidelity and the resulting trade-off between low-level (high-fidelity), and high-level (low-fidelity models). A similar problem is observable when contrasting between low-layer (low-fidelity) and high-layer (high-fidelity) models. The greater the fidelity of a simulation, the more accurately it is able to represent the situation under investigation. This accuracy is however gained at the expense of transportability, resulting in an inability to reuse the simulation to model other situations. Reducing the fidelity of the simulation increases its transportability, whilst decreasing its ability to accurately represent a situation.

As demonstrated by the generation of a glossary of terms in Peissner *et al.* (2012), this has resulted in a lack of consensus on the components that are required within a simulation. Some common components include models or profiles describing:

User: A representation of the user in terms of a number of variables describing his/her various characteristics.

Device: A representation of the device with which the user wishes to interact describing its various characteristics. The device may be a piece of hardware or software.

Task: A representation of the task to be simulated. It can be used to dictate the structure and contents of the other models by indicating the level of abstraction of the simulated interaction.

Environment: The context in which the interaction takes place. The environment can affect accessibility by interfering with interaction and its importance is dependent on the layer of abstraction at which the simulation is targeted.

At a high level, accessibility simulation is a comparison between user and technology (or task). In order for a comparison to take place, the respective profiles must be described in compatible vocabularies. In the field of Knowledge Representation a "domain model" contains the vocabulary and structural information necessary to describe a particular problem domain. A domain model for a generic accessibility problem will contain information about the user and technology in order to allow a comparison to be made. There are two main vocabularies that are used to facilitate this description: technology-focused preferences and human-focused capabilities. Competing standards have been developed that use each of the vocabularies respectively.

ISO 24751 deals with the provision of individualised adaptability and accessibility in the context of e-learning. It is a three-part standard, which describes the user and application profiles separately. Currently under revision, it is based on a fixed vocabulary of technology-focused needs and preferences. Preferences refer to device or software settings that can be specified to improve accessibility (e.g. font size). As they are device specific, they provide an accurate representation of the needs of the user in terms of the technology-focused preferences required in order for a device to be accessible. As above however, this specificity reduces the transportability of a user's profile owing to a need for mappings between all related preferences across devices.

ISO 24756 uses Common Accessibility Profiles (CAPs) to allow direct comparison between the needs and capabilities of users, systems and their environment (Fourney & Carter 2006). It views accessibility in terms of channels of communication that are facilitated by human-focused capabilities. CAPs are constructed from a series of Interacting Components which are able to either input or output via a fixed vocabulary of modalities (e.g. visual acuity). Unlike preferences, the approach provides a static vocabulary against which any device setting may be mapped (Atkinson *et al.* 2010). Human capabilities provide the transportability required to compare a single user profile against multiple device profiles. Their generalisability, however, comes at the expense of the extra effort required in terms of describing device settings and adaptations in terms of their corresponding human-focused capabilities.

It is possible to link the two vocabularies by modelling higher-level technology-focused preferences in terms of their lower-level human-focused capabilities. The resulting multi-level models lead to improved transportability while still allowing specific situations to be simulated. This type of simulation is called Hierarchical Task Analysis (HTA) and can be found in the GOMS family of models which break tasks down into a series of sub-tasks. Those sub-tasks can then be evaluated based on their constituent Motor, Cognitive and Perceptual elements.

Hierarchical structures can be found in models related to both accessibility and the personalisation of knowledge-based adaptive systems. Kaklanis *et al.* (2012) describes a framework for simulating accessibility evaluations. It uses hierarchical structures both within the organisation of its user-profile and during the simulations that are based on them. Data is firstly stored under a series of parameter types (Cognitive, Behavioural, Physical, Kinematic and Geometric), within those types a hierarchy is then developed based on the structure of the human body. During simulation, higher-level tasks are built from lower-level primitives, which are computed from the bio-mechanical data in the profile. Whilst the framework is hierarchical in nature, it is reliant on a fixed lowest level, from which higher levels are computed.

Sosnovsky & Dicheva (2010) describes the widespread use of hierarchical modelling within web-based adaptive systems. The article describes the Adaptive Web as focused on the personalisation of information and many of the systems are knowledge-based recommenders. The featured user models are all based on ontologies, which provide a hierarchical structure allowing the systematic storage and creation of new knowledge. Once again, higher-level knowledge can be inferred from its lower-level constituents. Through the use of stereotyping however, the models can also provide top-down inference.

Efforts have also been made to standardise the data structures used to expose profile data. Peissner *et al.* (2012) describes a number of formats based on XML and RDF. The Semantic Web is a collection of standards (including RDF) that allows data to be shared and reused across application, enterprise, and community boundaries[1]. By storing data in terms of its relationships to other data, networks can be formed to produce machine-readable information. Semantic technologies have already been used to produce an implementation of the CAP (Sala *et al.* 2011), and the CC/PP (Klyne *et al.* 2004) provides an RDF structure for defining the technology-focused accessibility needs of individuals and is therefore directly compatible with ISO 24751.

3 Approach

Two trends are visible in the previous section:

– The need for a standard vocabulary to describe both users and technology in order to expose profile information for comparison, and

[1] `http://www.w3.org/2001/sw/`

– The need to provide abstraction within profiles in order to cater for the diversity found in real life.

The approach taken builds on existing profiling techniques by providing a series of standard semantic relationships which allow both users and technology to be described using a symmetrical (hierarchical) recursive structure.

Taking the high-level view that accessibility simulation is a comparison between user and technology, there is little to differentiate between their resulting profiles. Rather than modelling them differently, they are both considered as 'actors' taking part in an interaction. Various elements within the interaction model can be standardised, as seen in figure 1. A vocabulary based on interaction capabilities allows accessibility to be described using a functional assessment of the abilities and requirements that are being compared. The abilities and requirements are stored in a common 'internal structure' and exposed via a standard 'external interface'. If a match is observed between the capabilities of the user and technology then interaction is possible using those matching capabilities. When a mis-match is observed, a description of the required assistance is provided in terms of the gap that must be bridged and the capabilities that are available to do so. Rather than relying on the presence of specific technology-based preferences or assistance targeted at a specific labelled impairment, a solution could involve re-routing interaction via alternate (matching) capabilities.

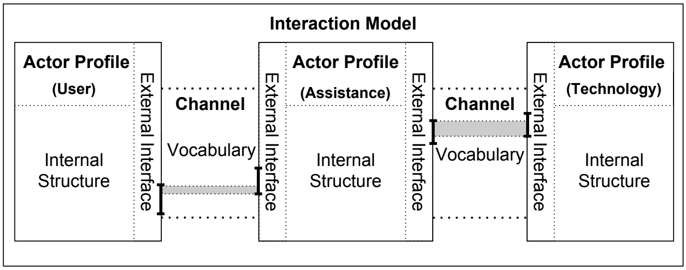

Fig. 1. Standardisation of Elements Within the Interaction Model

Where alternate matching capabilities are not available, assistance can be chosen based on its ability to bridge the gap between available capabilities. The use of a common internal structure allows capabilities to be linked internally and exposed in a predictable way via a standard external interface. The format is able to represent both users and technology (including hardware, software, assistive technology and adaptations) with chains of actors being used to represent the route that information takes. This provides the potential for automation as profiles can be stored in a machine-readable format allowing matching algorithms to be developed.

Profiles are also able to contain data at various levels and layers of abstraction[2]: from low-level human-focused capabilities to high-level technology-focused preferences, and from generic low-layer descriptions of abilities to specific high-layer contextualised ability recordings. This allows direct matching between profiles at the appropriate level of abstraction rather than requiring reference to a fixed (low) reference level as seen in existing models. In addition, a measure of the confidence of a match can be estimated based on the similarities between the highest layer at which information is provided. This approach reduces the rigidity suffered through the application of either of the themes mentioned above in isolation. Figure 2 shows the use of both levels and layers of abstraction within a profile as well as a channel of communication between two actors.

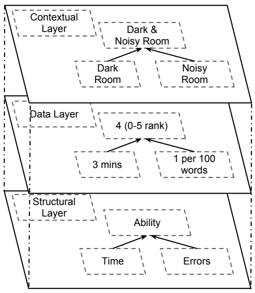

Fig. 2. Hierarchical Levels and Functional Layers of Abstraction

4 Design

The design of a framework built on the above approach will now be described. The framework consists of a series of components that are connected by a series of relationships.

4.1 Actors

Actors describe the entities (users and technologies) that are able to take part in interactions. Within the framework, profiles are built up using a nested series of actors, which may possess several sub-actors, describing their components. Sub-actors are actors in their own right (demonstrating the recursive nature of the framework) and can be viewed either as autonomous entities, separate to other actors at their level, or as constituents of higher-level actors. This mirrors the real-life construction of users, as demonstrated through the hierarchical VUMS profiles in Kaklanis *et al.* (2012).

[2] As describe in section 2.

The metaphor is even more useful when modelling technology, as individual actors can potentially be placed inside or separated from their higher-level actors, mirroring the modular nature of technology. A device like a desktop computer is actually a composite of several sub-devices (a screen, keyboard, mouse and speakers), each of which can be purchased separately and removed for use within another computer system. In the same way software is represented as an actor, which can be placed within numerous higher-level actors to represent its use within different devices.

An actor can also be used to represent the environment within which an interaction is taking place. Comparison between different environments (or the same environment at time times) can then be made based on their sub-actors. As an example, the environmental conditions of an open-plan office and a small home office may vary in terms of the number of people present, the type of light bulbs used, the weather outside the window and the ergonomy of furniture available.

This method of representation inherently identifies the abstract level at which comparisons should initially be made. For example, a simulation could be developed to describe a user interacting with a mobile phone. As the actors are defined at a high level, the initial comparison would focus on the high-level assessment of whether "the user can use the phone" (and whether the phone can be used by the user). The result of this comparison is a function of lower-level comparisons between the sub-actors of each actor (e.g. the user's eyes versus the phone's screen), which are in turn dependent on their sub-actors (e.g. the user's visual cones versus the phone's colour sub-pixels).

Through investigation of the sub-actors that are involved in an interaction, it is possible to identify which parts of the user and technology are responsible for its success or failure ('problem actors'). This form of assessment allows problem actors to be avoided, either by re-routing communication, or putting assistance in place to cope with the actor's deficiencies. For technology-based actors, replacement and upgrading are also possible.

4.2 Capabilities

Actors are able to interact with each other via capabilities, which are organised in the same hierarchical fashion as actors. Capabilities are described in terms of their direction, modality and a measure of their bandwidth, allowing channels of communication to be built and investigated between actors. An actor may have more than one capability, which allows it to take part in communication in either one or both directions. Processing functions provide links between capabilities allowing an actor to receive information via one capability and transmit it via another. An actor also inherits the capabilities of all of its sub-actors, (which are likely to be its own capabilities' sub-capabilities). A finger is an actor with capabilities to both input and output via the tactile medium. Output is provided through the capabilities of its muscles (sub-actors) to generate movement and input is provided through its nerves which are able to sense pressure, texture and temperature.

Capabilities share the same hierarchical structure as actors with higher-level capabilities being inferred by (reliant upon) the presence of lower-level capabilities. Where a mismatch is found between two capabilities, investigation of their sub-capabilities can indicate its source by identifying the constituent capabilities that are causing the problem. Where a gap is identified, its effects will propagate upwards, affecting the interaction of any higher-level capabilities that rely on it. As described previously, higher-level capabilities are likely to be task-focused with a vocabulary moving towards human-capabilities as focus is moved further down the capability hierarchy. Figure 3 depicts a low-layer (structural) view of a series of actors, sub-actors and their related capabilities. It provides a limited example of a number of pieces of technology that are all dependent on one of the capabilities provided by a finger (1-D Movement). At a higher-level, each capability would have data attached, describing a measurable ability and the context in which it was collected.

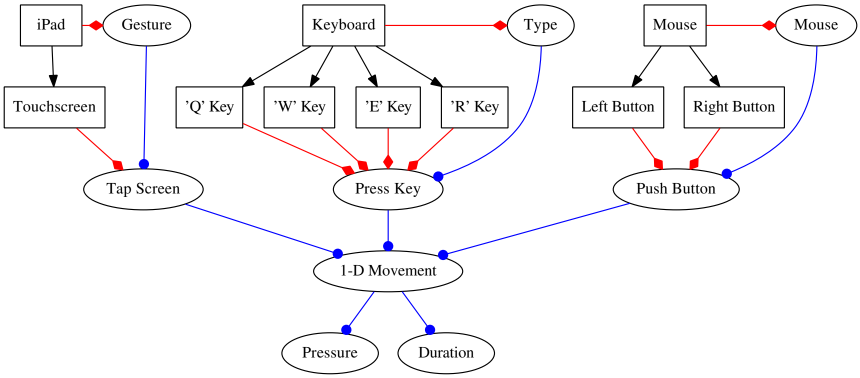

Fig. 3. A graphic representation of the hierarchical nature of physical actors and their capabilities (see figure 4 for key.)

4.3 Context

Capability data is constrained by the context within which it is collected (Peissner *et al.* 2012). User and device capabilities fluctuate throughout the day and their resulting abilities may be based on the presence of a number of factors. Multiple data items may therefore be stored against a single capability, providing information about an actor's capabilities in different contexts. Environments can be modelled using the same structure and vocabulary as any other actor. Interference (e.g. background light or noise) can be described as capabilities of the environment and exposed in a format compatible with a user's (or technology's) abilities to perceive them. As data is collected, information about its context can be stored by defining a relationship between the data and an actor representing its context. This again demonstrates the recursive nature of the framework.

5 Implementation

By defining flexible relationships rather than dictating rigid data structures, the (theoretical) design is not dependent on any specific technology or file format. The use of a graph-based representation (as seen in figure 4) does however bias implementation towards pattern-based representation and matching technologies. Development has been carried out using the semantic languages RDF and OWL for data storage and the graph-based logic language Prolog to provide reasoning capabilities. Figure 4 is a generic graph describing some of the relationships that have been defined and used in the framework.

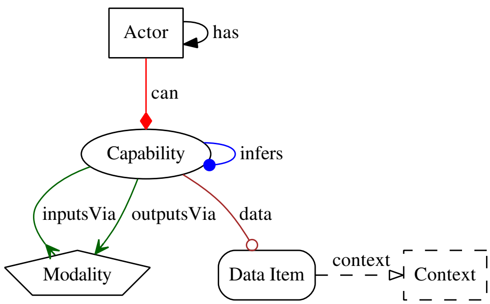

Fig. 4. A graph representing the elements of the framework and their relationships

Interaction is modelled in terms of a series of actors, with capabilities, that can be measured to provide data, which can be stored in context. Users, technology and context are all modelled as actors made up of one or more sub-actors (‘HAS’). An actor may also have capabilities allowing them to interact with each other (‘CAN’). Capabilities share the same hierarchical structure as actors with higher-level task-focused capabilities being inferred by lower-level human-focused ones (‘INFERS’). Capabilities are compared based on their modality (‘INPUTSVIA’ and ‘OUTPUTSVIA’) and can be measured to produce data that is collected in context.

Any two actors are able to sustain interaction when they have compatible capabilities with:

Opposite Directions: In order to communicate with each other one must transmit a message which is received by the other.
The Same Modality: Measured in the same units (or be convertible).
Overlapping Abilities: The two capabilities must share a common range within the modality.
Measured in Context: The ability measurements should be based on a common context (which is similar to that displayed in the current environment).

As the quality of a match is dependent on all of the above factors, adherence to the list can be used as a measure of the probability that two capabilities are compatible.

The framework has also been designed to support a modular implementation. The multi-layered approach taken to data storage allows the separation of vocabulary specification, the selection of data storage mechanisms and the development of matching algorithms. This provides the potential for different modules to be used based on the needs of a particular situation. Low-level capabilities related to transmitting or receiving sound could for example be described using a frequency response graph, with associated matching functions being used to predict their resulting accessibility.

This allows specific implementations of the framework to be developed to suit different interaction needs while providing the potential for integration through the definition of appropriate relationships. Gradual upgrades can be made available as vocabularies are extended, data storage needs change, algorithms are improved or improved context capturing is developed. Alternately "low-power" variants could be produced which restrict matching to lower layers, providing less information about the quality of the match or accuracy of the simulation.

6 Challenges and Further Work

This paper is focused on the description, storage and use of profile data, the success of the approach is therefore dependent on a number of other areas that are the subject of current research:

Semantic Web Technologies. are not yet fully realised, meaning there are still issues to be overcome before they can be widely adopted. A lack of accepted standards at the upper levels of the technology stack has resulted in the need for additional technologies to be used in order to achieve the functionality described at present. There are also outstanding questions regarding the scalability of graph-based data storage.

As they are researched further, many of these issues will disappear, assuming the further development and widespread use of semantic technologies. In addition, the identification of current deficiencies can be used to inform new research and development efforts. In terms of scalability, the modular and multi-layered design of the framework can be used to take advantage of pre-processing techniques that only present the data that is needed to evaluate a particular situation.

Data Acquisition. is a particular challenge with the problem being twofold: (1) initial population may require a bootstrapping procedure, and (2) information would need to be kept up-to-date. ISO 24756 for example, relies on the user (or other human agent) to create and update CAPs. However for mass adoption, the user cannot be relied upon either because (1) they may be unwilling, or (2) the information they provide may be unreliable (Godoy & Amandi 2005).

The framework is therefore reliant on the availability of agents that are able to populate user profiles. Given the nature of their task, current agents are highly specialised and tend to be developed as part of standalone systems (e.g. Hurst 2009). The intention of this framework is to allow agents to expose their data and as a result allow it to be used as part of a more holistic profile.

Context. can have a large effect on the usability of data, as described previously. At present the framework is able to model context and use it to determine the probability that a match is accurate. Where there is a lack of data to describe an actor's capabilities in a given context it is currently not possible to infer it (other than finding the closest possible match or using a stereotype). This concern can be addressed via inference back to human capabilities and the ability to model the environment as an actor. As interference (e.g. background light or noise) is stored as a capability in a format compatible with the users abilities to perceive them, it should be possible to determine the effect that a particular contextual capability has on an actor's capabilities. Known capabilities/contexts could then be used to extrapolated capabilities in other contexts.

Matchmaking Algorithms and Assistance Delivery Mechanisms. are, like data acquisition agents, are outside of scope given the focus of this research on data representation and storage. They are however required in order for the approach provided in this paper to be realised. Work aimed at developing a Global Public Inclusive Infrastructure is however providing research in this area (Vanderheiden & Treviranus 2011).

7 Discussion

The approach presented in this paper provides a flexible method of profiling users and technology, enabling direct matching between their capabilities at various hierarchical levels and functional layers of abstraction. The suggested relationships allow capabilities defined in various open and proprietary standards to be used together through attribution to an individual within a single profile. In addition contextual information can be included, providing a measure of applicability of data to a situation.

The accessibility of interaction between two actors can be assessed by matching their capabilities and the potential to include contextual information allows a measure of the accuracy of the resulting prediction to be provided. Where higher-level capability data is unavailable for a match to be made, inference can be used to recursively find a lower level at which matching constituent capabilities are available. As contextual data is stored in the same format as users and technology, contextual matching can also use similar algorithms.

Where a match is not possible owing to incompatible capabilities, an alternate route can be identified. Alternately, speculative augmentation of actors with representative sub-actors can used to identify the effect of different potential forms of assistance. As an example, both hearing-aids and text-to-speech functionality can be described in terms of the way they re-route and transform the existing channels of communication. Software can then be modelled on different devices by placing its highest representative actor within the actor describing each device.

References

Atkinson, M.T., Bell, M.J., Machin, C.H.C.: Towards ubiquitous accessibility: capability-based profiles and adaptations, delivered via the semantic web. In: Vigo, M., Abascal, J., Lopes, R., Salomoni, P. (eds.) Proceedings of the International Cross-Disciplinary Conference on Web Accessibility. W4A 2012, pp. 14:1–14:4. ACM, New York (2012)

Atkinson, M.T., Li, Y., Machin, C.H.C., Sloan, D.: Towards Accessible Interactions with Pervasive Interfaces, Based on Human Capabilities. In: Miesenberger, K., Klaus, J., Zagler, W., Karshmer, A. (eds.) ICCHP 2010, Part 1. LNCS, vol. 6179, pp. 162–169. Springer, Heidelberg (2010)

Biswas, P., Robinson, P.: A brief survey on user modelling in HCI. In: International Conference on Intelligent Human Computer Interaction, IHCI (2010)

Fourney, D., Carter, J.: A Standard Method of Profiling the Accessibility Needs of Computer Users with Vision and Hearing Impairments. In: Conference & Workshop on Assistive Technologies for People with Vision & Hearing Impairments (2006)

Godoy, D., Amandi, A.: User profiling in personal information agents: a survey. The Knowledge Engineering Review 20(11), 329–361 (2005)

Hurst, A.: Automatic assessment and adaptation to real world pointing performance. ACM, New York (2009)

Kaklanis, N., Moschonas, P., Moustakas, K., Tzovaras, D.: Virtual user models for the elderly and disabled for automatic simulated accessibility and ergonomy evaluation of designs. Universal Access in the Information Society, 1–23 (2012)

Klyne, G., Reynolds, F., Woodrow, C., Ohto, H., Hjelm, J., Butler, M.H., Tran, L.: Composite Capability/Preference Profiles (CC/PP): Structure and Vocabularies 1.0. W3C Recommendation, W3C (January 2004)

Peissner, M., Dangelmaier, M., Biswas, P., Jung, Y.M.C., Kaklanis, N.: White Paper: Virtual User Modelling Public Document. Tech. rept. 1.0. VUMS Cluster (June 2012)

Sala, P., Fernandez, C., Mocholí, J.B., Presencia, P., Naranjo, J.C.: Implementation of the ISO/IEC 24756 for the Interaction Modeling of an AAL Space. In: Stephanidis, C. (ed.) Universal Access in HCI, Part III, HCII 2011. LNCS, vol. 6767, pp. 210–219. Springer, Heidelberg (2011)

Sloan, D., Atkinson, M.T., Machin, C., Li, Y.: The potential of adaptive interfaces as an accessibility aid for older web users. In: Proceedings of the 2010 International Cross Disciplinary Conference on Web Accessibility (W4A). W4A 2010, pp. 35:1–35:10. ACM, New York (2010)

Sosnovsky, S., Dicheva, D.: Ogntological technologies for user modelling. Int. J. Metadata Semant. Ontologies 5(1), 32–71 (2010)

Vanderheiden, G.C.: Ubiquitous Accessibility, Common Technology Core, and Micro Assistive Technology: Commentary on "Computers and People with Disabilities". ACM Trans. Access. Comput. 1(2), 10:1–10:7 (2008)

Vanderheiden, G., Treviranus, J.: Creating a Global Public Inclusive Infrastructure. In: Stephanidis, C. (ed.) Universal Access in HCI, Part I, HCII 2011. LNCS, vol. 6765, pp. 517–526. Springer, Heidelberg (2011)