Individual Cognitive Abilities and Styles in HCI: Three Main Challenges and a Tiered Adaptation Model

Sven Bertel
Usability Research Group, Faculty of Media, Bauhaus-Universität Weimar
Bauhausstr. 11, 99423 Weimar, Germany sven.bertel@uni-weimar.de

ABSTRACT
For various tasks in human-computer interaction, measures of performance and emotion can be improved by adapting the user interface to a user's individual cognitive profile. Such tasks can be found, for example, with eLearning, information visualization, gaming, and human-computer collaboration in reasoning or problem solving (e.g. in design). Relevant factors within a cognitive user profile may include separate cognitive abilities, styles, and preferences, as well as personal characteristics of memory or attention. The three chief challenges for a successful adaptation to a user's cognitive profile lie (1) in establishing which of these factors are relevant for a given task-user pair, and to which extent, (2) in establishing how, based on (1), adaption may best be performed, and (3) in establishing a user's individual values for these factors. All three challenges possess aspects related to cognitive theory and user modeling, as well as quite practical aspects related to measuring a cognitive profile. This contribution will start out by, in turn, addressing research questions and methods related to the challenges. A tiered structure for cognitive user models will be subsequently sketched, through which adaptation can be based on parameters that are individualized to varying degrees, depending on how much is known about a user's individual cognitive profile.

Author Keywords
Interaction Technologies; Adaptation; Cognitive Factors.

ACM Classification Keywords
H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION
Some people are better than other people when it comes to processing information given to them as pictures. Some people outperform others when working with texts. The hypothesis that there exist inter-individual differences in how people process information on visual and verbal dimensions has been the focus of research for decades. For learning, for example, this has led to rather well established theories on different cognitive *abilities* (i.e. relating to differences in what people are capable of doing), different cognitive *styles* (i.e. relating to differences of modality in how they process and represent information), and different learning *preferences* (i.e. relating to differences in how they prefer information to be presented to them; cf. [21]). While eLearning studies into employing a learner's individual visual/verbal learning style for adapting learning material have so far shown mixed effects of adaptation on performance, they have shown clear benefits on intrinsic learner motivation [7, 2]. Would it not thus be useful if an eLearning system knew about a user's individual cognitive abilities and styles, and if it used such knowledge to automatically adapt how learning material gets presented?

A second example: A study by Keehner [17, 18] on effects of conflicting visual/haptic cues on surgeons' spatial scene understanding during laparoscopic surgery demonstrated that such effects depend on the surgeons' individual spatial abilities. One may conclude that, for surgeons with low general spatial abilities, specific types of visual/haptic cue mismatches should be avoided, as these lead to a decreased surgery-related scene understanding. Would it not be useful if the laparoscopic system knew about a surgeon's individual low spatial abilities and used such knowledge to prevent certain cue mismatches from being presented?

A third example: For human reasoning with spatio-temporal information (e.g., about distances/directions between objects or places) the construction process of corresponding mental scene models has been found to be frequently subject to preferences (e.g. [19, 15]) in that, among all possible consistent models, some are constructed more frequently (as well as faster and easier) than others. Parsimony seems to play a role, as models that reflect simple and typical configurations seem to be preferred [12], which can be explained by simple and commonly used mental processes [8]. Also, cultural left/right biases may have an effect on preferences [15]. Let us assume that we have a task that involves some spatial configuration and that is either shared by a human and a computational system, or in which a human user is assisted by a computational system. Would it not be useful if the system knew about such general human cognitive preferences, and adapted its actions accordingly (e.g., by occasionally drawing the user's attention also to some non-preferred model)?

It seems feasible to ground some of the adaptive behavior required for effectively addressing these examples onto general and general-purpose models of cognition, such as ACT-R [1], and to also, at least partially, evaluate adaptive behavior based on these. The third example is likely to be amenable to such an approach. For more specific cognitive faculties, use of more specific models of cognition may be required (e.g., for visuo-spatial reasoning, [8]), or of models which address specific classes of interactive tasks or domains (e.g., in-vehicle user interfaces [25], web-based navigation [11], or the design of built environments [6]). For purposes of either usability engineering or testing of adaptive user interfaces, a few approaches exist that incorporate user models along with interaction models, with the user models targeting general or specific user groups (for the latter, e.g. [24]). Programmable user models [28] are a good example of approaches which specifically try to bridge gaps of knowledge, methods, and philosophy of approach between the cognitive scientist and the interface designer, as they permit predictive interface evaluations based on implemented psychological theories. Even if an individual user's previous domain knowledge and the history or context of use can often be modeled, when it comes to inter-individual differences in cognitive processing, such as may be expressed by cognitive abilities or styles, such models address the average user of the targeted group. The models are specific with respect to task, context, expertise, or group, but are still general with respect to that group's cognitive profile.

It should seem obvious that adaptive user interfaces which will adequately address the first and the second example will need to draw on more individual parameters than can be derived from general models of cognition, even from task- or group-specific ones. For the individual user, these parameters may be assessed through various test instruments (for example, through tests of mental rotation or mental perspective taking performance for the respectively related, specific cognitive spatial abilities). While definitions of the terms *cognitive abilities*, *styles*, and *preferences* vary in the literature, I will, for the current purpose, simply rely on the definitions by Mayer and Massa [21] that were already provided in the first paragraph above. As far as other terminology is concerned, I will follow [23] and hold an *adaptive* system to be one that automatically modifies some of its characteristics to better fit a user's needs. To keep the discussion short enough for these few pages, the reader may simply envision such adaptive behavior to occur either on macro or micro levels (cf. the discussion by [20]) and either at design- or run-time.

Next, three chief challenges for user interface design connected to such an individualized approach will be outlined. The third section will sketch how a practical design compromise may look like that mediates between the desire (or even, need?) to know as much about a user's cognitive profile as possible and severe limitations in obtaining such knowledge. The contribution will conclude with a short discussion.

THREE MAIN CHALLENGES

As it turns out, the first two examples above are still both of a rather benevolent type for the interface designer interested in how cognitive factors (i.e., abilities, style or preferences) vary among his user base. This is because the cited studies all investigated relationships between certain tasks (learning, scene understanding) and specifically selected cognitive abilities or styles (visual/verbal, spatial). When starting with a given task and an individual user only, however, the first question is which of the currently known, separate cognitive factors may play a role in determining how that one user cognitively tackles the task. This question is already much harder to answer. It becomes harder yet when asking for specific effect sizes and directions, let alone when asking for possible interactions between the factors. Hunt [14] compared seven different cognitive skills; all of these may easily differ between users and may influence user performance. These were reading comprehension, vocabulary, grammar, mechanical reasoning, quantitative skills, mathematics achievement, and spatial reasoning. Some of these skills are more intercorrelated than others, with reading comprehension and spatial reasoning forming the maximally unrelated pair. While this list of skills serves to give a flavor of the variety one has to be able to deal with, it is by no means exhaustive.

The **first main, hard challenge** is thus to establish which factors are actually important for adaptive user interfaces for each given class of tasks, and which are not. Given that HCI specialists are usually no specialists in the cognitive sciences, and vice-versa, this challenge is best addressed interdisciplinarily. It must certainly be largely addressed incrementally. It seems that it will not be enough to simply identify relationships between cognitive factors on the one hand and measures of task performance or satisfaction on the other, but that one has to equally demonstrate that specific adaptation strategies formed on top of discovered inter-individual differences in users' cognitive profiles will be effective. Very likely, not all differences in profiles will be equally open to on operationalization for adaptive user interfaces.

A related, hard problem is created by the fact that many distributions of cognitive ability or style expressions are far from being uniform. Usually, some style expressions will be more frequently encountered than others within a targeted user group. For example, with cognitive learning styles, a majority of learners will be visual (74%), especially when sampling from populations in the natural sciences or engineering (e.g., [9, 2]). This may be of quite some importance for scenarios in eLearning or eTutoring. When a user interacts with an adaptive system, asking how often his cognitive profile is to be encountered within a user population will likely seem to be irrelevant to him. What likely will be relevant to him is that any adaptation which is to occur based on information about cognitive profiles will occur based on *his* specific profile. For systematically investigating interrelationships between cognitive factors and tasks, as well as for investigating how an adaptation should best occur based on a given user profile, frequencies of specific types of user profile are, however, far from being irrelevant. When a user type is encountered too infrequently, achieving reliable statistical comparisons between types may be difficult. As it seems to be no viable option to exclude those users with rare profile types from using an interface, or to at least decide to not try to provide them with adaptive interface behavior, other approaches need to found. A possible approach seems to lie in using mixed-method designs that combine conclusions based on quantitative observations for the frequent types with conclusion based on qualitative observations for the more infrequent types.

Should one be in the lucky position to have already determined which cognitive factors effectively influence a given user's performance on a specific task, and in which ways, the **second main challenge** lies in determining how adaptation of the interface should best be performed. In the case of the eLearning example above, this may be comparably straightforward, as one may choose the representational format of learning material to correspond to a user's individual cognitive profile. However, even with the example, it seems unclear if this would always be the best strategy. Depending on the specific learning goals and the context, it may be more important to train learners to instead better cope with material in formats that do not well match their cognitive profiles (see e.g. [9] for a discussion of this point for inter-individual differences of students' learning styles). The question thus becomes one of either adapting *to* a profile, or *against* it. This is particularly interesting in gaming applications, when system adaptation can take the form of a computational player either adapting to keep a human player in the game for as long as possible (and to, e.g., maximize fun or engagement) or to adapt to become a maximally strong adversary (see e.g. [27] for an example and a discussion).

Choices of either adapting to or against a user's individual cognitive profile may be limited for purely practical reasons for all sorts of asymmetric tasks which are not fully specified. Such tasks can be frequently encountered, for example, in computer-assisted design: here, usually only a subset of design requirements can be formally described (often, the more technical requirements, such as e.g. with regard to thermal insulation), while other requirements remain the exclusive province of a human designer (often, those regarding an aesthetical or a more holistic evaluation of the design). The result is a setting of an asymmetric human-computer collaboration, with shared initiative, in which both parties need to rather act as partners that observe, adapt to, and, ideally, anticipate the other's actions in order to be jointly effective [3].

Finally, the **third chief challenge** lies in making sure that a user interface has enough information about an individual user available to allow for an adequately precise establishing of that user's relevant cognitive profile. From a practical point of view, this likely poses the hardest of the three challenges. While test instruments exist for many of the various cognitive abilities and styles, administering them is often a rather lengthy process that, what is even worse, is often tiring and/or boring for the tested user. The same holds true for many of the established instruments that assess individual characteristics of a user's memory and attention systems. The best ways of measuring a user's working memory capacity are performance measures, that is, one tests how much one can cram into the memory system. However, running at full load will not only quickly tire a user out for further tests of cognitive ability or style, but also for any main tests of potentially adaptive systems that the HCI professional will be chiefly interested in. It thus seems highly impracticable to measure factors related a user's cognitive profile separately and over and over again for each task, interface, or tool. There will be neither enough time, cognitive user (or researcher) resources, or user motivation for such an endeavor.

In addition, individual expressions of abilities or style will be only largely indicative of a user's general traits in cognitive processing, though not necessarily for life and in all situations. Measurements and the factors that they reflect are each expected to vary to different degrees depending on the task or context (e.g. for learning styles, see [2]). Depending on the task at hand, it may be necessary to dynamically assess additional information about the individual user, for instance, to infer information about his current mental state (e.g., regarding current foci of attention). Different measurements of psychophysical parameters may be used to further inform, individualize, and situate more general cognitive user models, such as through parameters derived from gaze (e.g. [10], [4]), EEG, or skin conductance, etc. Where some of the criticism above was directed against cognitive user models that were too general (i.e., not individualized enough), the problem here is that information obtained about a user's individual cognitive profile via the established test instruments may be too general (i.e., not situation- or context-dependent enough) to serve as the sole base from which to derive parameters governing adaptive interface behavior. For example, for tasks or games that involve problem solving, one can assume reliable information about a user's current strategies and foci to likely be at least as important for a generation of effective adaptation strategies as reliable information about the user's individual cognitive profile, especially when such profile information has been obtained independent of task or context.

TIERED ADAPTATION

I have suggested above that a viable course of action to deal with different frequencies of cognitive user profiles in a user population may lie in combining largely qualitative measures applied to the infrequent types with largely quantitative measures used for the more frequent types. This may imply that different information obtained about a user's individual profile may be reliable to different degrees, depending on the methods through which it was acquired. It may also mean that adaptive interface behavior may have to be more or less assertive: less, if it is based on less reliable bits of information, more if it solidly grounds in reliable data. Such a graduated approach fits well with a situation in which obtaining any specific information about a given user-task pair (i.e., information about relevant cognitive factors which is either individualized or situated, or, better, both) is nearly always costly, with prices being chiefly paid in currencies of user fatigue or motivation. We thus need an approach that facilitates the striking of a practical compromise between a user interface designer's wish to know as much about a user's cognitive profile as is possible and strong practical limitations in obtaining such knowledge.

The model proposed here is based on a related sketch for situations of joint human-computer spatial reasoning and problem solving suggested by [5]. It conceptually extends [5] and is not limited to applications of spatial reasoning or problem solving. The model (see Fig. 1 for an illustration) consists of three tiers of user-related data, in which information available about the cognitive profile of an individual user is gradually refined from bottom to top level. The model is tiered, as whenever more specific information is lacking, it may be substituted by less specific information, albeit at a price, as we will see. General cognitive factors are such as may be obtained, for instance, through general models of human cog-

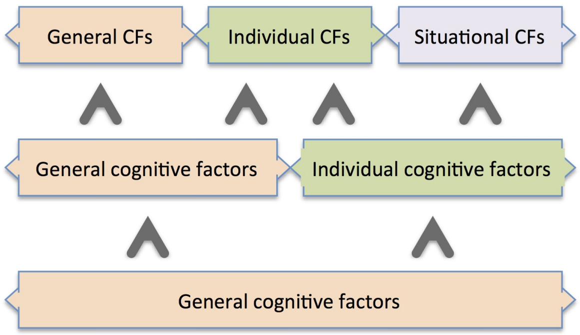

Figure 1: Sketch of the model with general, individual, and situational cognitive factors. Ideally, knowledge would exist up to the top tier for all factors relevant to adaption of a user interface behavior for a specific task. Where such is not available, less specific values may substituted from lower tiers.

nition, e.g. regarding general (i.e. averaged) characteristics of memory or attention, temporal characteristics of cognitive processing, or general preferences of mental model construction such as those described for the third example above. Individual cognitive factors of our modeled user are of the kind that can be obtained through test instruments, such as for various relevant cognitive abilities or styles. The first and the second of the examples discussed in the introduction would likely benefit from an adaptation rooted in general and individual cognitive factors. Situational cognitive factors, as shown on the top level, are such factors as can be derived based on live measurements of (e.g., various psychophysical) parameters about our user. With such measurements, one may attempt to address questions of the following kind: Which controls has the user gazed at over the course of the last five seconds? Is the user's current attention span likely to be lower or higher than his usual individual level? Are there any indications that his usual expressions of learning styles should be modified for the current task? etc. Situations related to the first two examples are easily conceivable during which the effectiveness of adaptive user interfaces may be improved by utilizing a selection of situational cognitive factors in addition to individual and general ones.

Through establishing the three levels, we get a number of interesting properties. First, information about factors on higher levels comes generally at higher costs than information about factors on lowers levels. Information obtained higher up is also more likely to be specific to a situation or task, and will be less well suited for drawing general conclusions about the cognitive profile of our user, or even about a group of users. General cognitive factors exist on all three levels, as only some of the more general assumptions and findings about a user's cognitive profile can and need be gradually refined along the way up. The same holds true for individual cognitive factors on levels two and three, and their relationship with situational cognitive factors. Last, and perhaps most importantly, we may use this model to sketch relationships between groups of factors which become useful when we do not have complete information about a user's individual cognitive profile (that is, virtually always). Ideally, our knowledge would extend up to the top tier for all of those factors that we judge to be relevant to effective adaption of a user interface for a specific task under consideration. Whenever such knowledge is incomplete, we may revert to using knowledge about factors on lower tiers, thus effectively using our model as one of graduated defaults. Such reversion to lower tiers will come at a price, of course, as we will lose some of our individualized or situated potential. In other words, whenever we move from higher to lower tiers, more specific values about cognitive factors and factor expressions will be substituted by less specific ones. It seems likely that one may often rather easily devise mechanisms of adaptive behavior that can reflect such changes in specificity, for example, by adapting a user interface less strongly (e.g., in less assertive and more subtle ways) whenever information about the individual user's cognitive properties is drawn from factors on less reliable tiers.

CONCLUSIONS

I have discussed three chief challenges that need be addressed to be able to productively harvest inter-individual differences in cognitive user profiles for generating effective, adaptive behavior of user interfaces. It seems that at least challenges (1) and (3) can only be adequately tackled inter-disciplinarily and through collaboration of user interface designers and cognitive scientists. In addition to incremental solutions for those two challenges, such collaboration should, at best, also result in a process of establishing a base of rules or best practices of which cognitive factors frequently relate to an individual user's task performance, and how, and of which test instruments are best to be used under which circumstances. The reason why I raise this point is that, naturally, we will very likely not see many user interface designers suddenly starting a training in the cognitive sciences. This would neither seem practicable, nor necessary. What is needed, however, is, first, an increased awareness among user interface designers for inter-individual differences rooted in their users' cognition and, second, approaches to designing adaptive user interfaces that permit scaling. This is to say that the approach needs to be able to scale from simple, recipe-like stages (à la "The ten most important rules for adapting your iOS app to your users' diversity in attentional resources", perhaps similarly simple and iconic as e.g. Shneiderman's *Golden Rules* [26] or Norman's *Design Principles* [22]) to much more detailed and focused stages in which specific cognitive factors will need to be addressed based on specific theories or (live) models of user's cognitive processes.

One should of course ask whether the three challenges that were raised here are the only challenges out there that currently keep us from creating effective adaptation of various user interfaces to the individual user's cognitive profile. The answer is no, of course. For the purposes of this contribution, I have tried to select and concentrate on the three challenges which I currently rate as the most urgent and difficult ones. Other, related challenges do, for instance, target questions of how changes of users' cognitive profiles over the course of a day, a task, or a lifetime, can be tracked, modeled, and reacted to, or of how users' cognitive abilities and styles interact with their emotional states. These are interesting questions, to be sure; however, I would strongly recommend embarking on a stepwise process, in which we tackle the most important challenges first, before moving on.

As a last point, I would like to argue that we are currently seeing a significant increase in the frequency of HCI settings that would benefit from a better adaptation of the involved interfaces and systems to the individual user's cognitive properties. Let me illustrate this point through two quick examples: First, eLearning. The use of MOOCs (Massive Open Online Courses) is currently on a rapid upswing, no matter if counted by the number of courses being offered or by attendance (surpassing 230,000 individuals per course for some courses, [16]). Also, participants are drawn from increasingly heterogeneous groups [13]. One possibly effective response to diminishing available instructor resources per participant may lie in constructing eLearing systems that more closely shadow the individual user's learning progress than is currently the common case, and that better adapt to it, similarly to how a good tutor would adapt material and methods to a student's progress. Such response would certainly benefit in quality if designers of those eLearning systems would know more about users' individual cognitive profiles as well as know better how to adapt interface behavior to these.

The second example is based on an extrapolation about the frequency of human-computer interactive systems that are asymmetric in the sense sketched above. The more computational tools we see that employ processes which remain partly opaque to the standard user (often because of reasons of data or process complexity, e.g. in applications of computer-supported design or big data analysis, or that employ techniques of data mining or machine learning), the more frequently HCI researchers and practitioners will need to address issues of human-computer collaboration and negotiation in which adequate cognitive user models will be key for effective interaction. My bet is that, at least for as long as we will continue to see an increase in the use of data-intensive applications, we will see an increase of asymmetric interaction settings that can greatly benefit from effectively adapting to users' individual cognitive profiles.

REFERENCES

1. Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., and Qin, Y. An integrated theory of the mind. *Psychological review 111*, 4 (2004), 1036.
2. Beckmann, J., Bertel, S., and Zander, S. Adaptive elearning based on individual learning styles - performance and emotional factors. In *Proc. Mensch & Computer 2014* (to appear).
3. Bertel, S. Show me how you act on a diagram and I'll tell you what you think (or: Spatial structures as organizing schemes in collaborative human-computer reasoning). In *2005 AAAI Spring Symposium: Reasoning with Mental and External Diagrams: Computational Modeling and Spatial Assistance*, AAAI (2005), 77–82.
4. Bertel, S. Towards attention-guided human-computer collaborative reasoning for spatial configuration and design. In *Foundations of Augmented Cognition (Proceedings of HCI International 2007, Beijing)*, Springer; Berlin (2007), 337–345.
5. Bertel, S. Supporting joint human-computer reasoning and problem solving: Blueprint of an architecture. In *Proc. of Spatial Cognition 2010 (Posters)*, D. N. Rapp, Ed., SFB/TR 8 Report No. 024-07/2010 (2010), 6.
6. Bertel, S., Barkowsky, T., and Freksa, C. Mental model-centered design for built environments. In *Proc. of EDRAMOVE Workshop on Movement and Orientation in Built Environments, EDRA 39th Annual Conference*, SFB/TR 8 Report No. 015-05/2008 (2008), 13–19.
7. Brown, E., Brailsford, T., Fisher, T., Moore, A., and Ashman, H. Reappraising cognitive styles in adaptive web applications. In *Proceedings of the 15th international conference on World Wide Web*, ACM (2006), 327–335.
8. Engel, D., Bertel, S., and Barkowsky, T. Spatial principles in control of focus in reasoning with mental representations, images, and diagrams. In *Spatial Cognition IV. Reasoning, Action, and Interaction*, C. F. et al., Ed., LNCS, Springer (Berlin, 2005), 181–203.
9. Felder, R. M., and Brent, R. Understanding student differences. *Journal of Engineering Education 94*, 1 (2005), 57–72.
10. Freksa, C., and Bertel, S. Eye movements and smart technology. *Computers in biology and medicine 37*, 7 (2007), 983–988.
11. Fu, W.-T., and Pirolli, P. Snif-act: A cognitive model of user navigation on the world wide web. *Human–Computer Interaction 22*, 4 (2007), 355–412.
12. Goodwin, G. P., and Johnson-Laird, P. N. Reasoning about relations. *Psychological Review 112*, 2 (2005), 468–93.
13. Ha, T.-H. Moocs by the numbers: Where are we now? `http://blog.ted.com/2014/01/29/moocs-by-the-numbers-where-are-we-now`, 2014.
14. Hunt, E. B. Verbal ability. In *Human Abilities: An information-processing approach*, R. Sternberg, Ed. New York: W.H. Freeman & Co., 1985, 31–58.
15. Jahn, G., Knauff, M., and Johnson-Laird, P. N. Preferred mental models in reasoning about spatial relations. *Memory & Cognition 35*, 8 (2007), 2075–2087.
16. Jordan, K. Mooc completion rates: The data. `http://www.katyjordan.com/MOOCproject.html`, 2014.
17. Keehner, M. Conflicting cues from vision and touch can impair spatial task performance: Speculations on the role of spatial ability in reconciling frames of reference. In *Spatial Cognition VI. Learning, Reasoning, and Talking about Space*. Springer, 2008, 188–201.
18. Keehner, M. Spatial cognition through the keyhole: How studying a real-world domain can inform basic science—and vice versa. *Topics in Cognitive Science 3*, 4 (2011), 632–647.
19. Knauff, M., Rauh, R., and Schlieder, C. Preferred mental models in qualitative spatial reasoning: A cognitive assessment of allen's calculus. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*, Erlbaum (1995), 200–205.
20. Leutner, D. Instructional-design principles for adaptivity in open learning environments. *Curriculum, plans and processes of instructional design: international perspectives* (2004), 289–307.
21. Mayer, R. E., and Massa, L. J. Three facets of visual and verbal learners: Cognitive ability, cognitive style, and learning preference. *Journal of educational psychology 95*, 4 (2003), 833.
22. Norman, D. A. *The design of everyday things*. Basic books, 2002.
23. Oppermann, R. *Adaptive user support: ergonomic design of manually and automatically adaptable software*. CRC Press, 1994.
24. Quade, M., Lehmann, G., Engelbrecht, K.-P., Roscher, D., and Albayrak, S. Automated usability evaluation of model-based adaptive user interfaces for users with special and specific needs by simulating user interaction. In *User Modeling and Adaptation for Daily Routines*, E. Martín, P. A. Haya, and R. M. Carro, Eds. Springer, 2013, 219–247.
25. Salvucci, D. D. Rapid prototyping and evaluation of in-vehicle interfaces. *ACM Transactions on Computer-Human Interaction (TOCHI) 16*, 2 (2009), 9.
26. Shneiderman, B. *Designing the user interface: strategies for effective human-computer interaction*, vol. 2. Addison-Wesley Reading, MA, 1992.
27. Wetzel, S., Spiel, K., and Bertel, S. Dynamically adapting an AI game engine based on players' eye movements and strategies. In *Proc. of Sixth ACM SIGCHI EICS 2014* (to appear).
28. Young, R. M., Green, T., and Simon, T. Programmable user models for predictive evaluation of interface designs. *ACM SIGCHI Bulletin 20* (1989), 15–19.