Can adaptive systems participate in their design? Meta-adaptivity and the evolution of adaptive behavior

Alexandros Paramythis

Johannes Kepler University,
Institute for Information Processing and Microprocessor Technology (FIM)
Altenbergerstraße 69, A-4040 Linz, Austria
alpar@fim.uni-linz.ac.at

Abstract. This paper discusses the opportunities arising from the employment of meta-adaptivity as a tool that can facilitate the design of adaptation. The discussion is structured around a specific example problem in the domain of adaptive course delivery systems. The paper builds upon this example to argue that meta-adaptivity is a viable solution to the "ground up" design of adaptive systems, and may be especially suited for cases where there is little empirically validated evidence to support design decisions.

1 Introduction

As any researcher or practitioner in the field of adaptive hypermedia systems will readily admit, designing adaptive behavior is hard. Some of the difficulties stem from the lack of established and proven practices in this respect; others are due to the lack of sufficient empirical evidence to support the design of adaptation. Furthermore, although adaptive systems are inherently interactive systems, the design approaches, methods and tools available from the field of human-computer interaction do not suffice in themselves in addressing the aforementioned problems, as they have not been devised to design systems that dynamically change their behavior to suit the (changing) requirements of individuals. At the same time, increased attention is being paid to adaptation meta-data, and, specifically, their generation and incorporation into the adaptation design / authoring process. This is both spurred by the goal of achieving meta-adaptivity *per se*, and a natural consequence of the introduction of semantic web technologies for the attainment of adaptivity in modern adaptive systems (which, to a large extent, provide an "enabling" layer for meta-adaptivity). In the context of this paper, the term "meta-adaptivity" will be used to refer to the capacity on the part of an adaptive system to observe, assess and modify its own adaptive behavior, towards a (set of) adaptation design goal(s). Furthermore, we will be exclusively focusing on a specific type of meta-adaptivity, namely self-regulation (in short, self-regulation assumes the capability for an adaptive system to perform self-evaluation and learn from that) [2]. This paper argues and provides an example of how meta-adaptivity offers a design tool, or vehicle, that potentially facilitates the design of adaptive systems.

2 Example adaptive system

This section presents an example design case study of an adaptive system. The basis of our exemplary system design is a simple, yet quite popular, adaptive function in the area of adaptive hypermedia systems: the annotation of links within learning content. For our needs we will assume a system that exhibits characteristics common to a large range of adaptive systems in the field (e.g., NetCoach [5]). The system's most important features are as follows:

– The system's domain model is a small, course-specific ontology comprising: learning concepts, corresponding modules / pages, and semantic relations between these.
– The user model is an overlay model (over the domain), with a number of discrete, user-specific "states" with respect to each of the concepts in the domain model.
– Individual user models are updated on the basis of observable user activities.
– Adaptation logic is expressed through adaptation rules, such as in the case of [4].
– Further to the above, the user model may contains other user attributes, some of them explicitly provided by the user, and others inferred from user activity.
– Based on the user model, the system can decide on recommendations regarding the future visits of different modules / pages. It is exactly on the basis of these recommendations that link annotation is being considered in the context of our example.

The design question at hand is whether and in what way to annotate links to concepts or modules to convey the semantics of the system's recommendations. Although there is a considerable body of research on this question, for the purposes of this example, we will assume that the system's designer has no empirical evidence to support the considered design alternatives. The alternatives themselves are encapsulated in five different strategies as far as link annotation is concerned: *Strategy A*: *No annotation*. This can be considered the base-line strategy, and would simply involve not exposing the user to the system's recommendations. *Strategy B: Annotation using different link colors*. In this strategy different colors are used directly on the links to signify system recommendation. *Strategy C: Annotation using bullets of different colors*. This is very similar to strategy B above, with the exception that the colors are applied externally to the links. Annotations (i.e., the bullets) are dynamically added to the document. *Strategy D: Annotation using custom icons*. A variation of strategy C above, with the bullets replaced by icons intended to carry more semantic information. *Strategy E: Link hiding*. This strategy involves hiding (but not disabling) links [1], for which the system's recommendation is that the user is not yet ready to visit them.

Given the strategies above, the design question at hand is which one(s) to use, and for a given user and context of use. Note that the strategies are not necessarily mutually exclusive. Also note that it would be desirable to identify situations that might justify a transition from one strategy to another.

3 Evolution of adaptive behavior

The designer starts out with no evidence about when and under what conditions to use each strategy, or whether, indeed any one strategy is "better" than all the rest.

Each strategy has obvious trade-offs as far as flexibility and user control over the navigation process is concerned. The designer's goal however, is clear: students should encounter concepts that they are not "ready" for as little as possible, and this should be achieved with the least possible restrictions on interaction / navigation. Continuing with the example introduced above, we will look at three potential iterations that the design process could have gone through.

The first step of the design would involve the encoding of the strategies as sets of adaptation actions (e.g., as in [3]). Since the designer has no evidence regarding the applicability of the different strategies, these cannot be directly assigned to adaptation logic. The system would need to be able to recognize these strategies and apply them (separately or combined) in more or less a trial-and-error fashion. The design information that already exists, and can be conveyed to the system, is which strategies are mutually exclusive, and which ones can be applied in combination. Given these constraints, and a suitably encoded representation of the design goal stated earlier as self-regulation metric, the system is then ready to undergo the first round of user testing.

We will assume that the results of the first round of testing do not yet suffice for building a comprehensive body of adaptation logic to guide the system's adaptive behavior. They do, however, provide enough evidence for the following:

- Eliminating strategies (and combinations thereof) that do not seem to meet the desired design goal under any circumstances. In the context of the ongoing example we can assume this includes strategy B and all its combinations (e.g., because changing links' colors seems to be confusing for users).
- Categorizing and providing a tentative "ranking" of the remaining combinations, based, respectively, on their design / interaction semantics, and on the rate of success they have exhibited during the first round of testing.

The aforementioned categorization and ranking process, might result in something like the following in the case of our example: *Category I*: Includes only strategy A and corresponds to absolute freedom in navigation, with no system assistance / guidance whatsoever. *Category II*: Includes the uncombined strategies C and D and corresponds to absolute freedom in navigation, but this time with explicit system assistance / guidance. *Category III*: Includes all combinations that include strategy E and corresponds to the application of restrictions on the navigation, to enforce a path through the learning material. The preceding categorization and ranking is, obviously, only one of several possibilities. It does, nevertheless, serve to demonstrate the following points:

- Although it is a ranking, it is not obvious in which "direction" it should be applied. For example, should the system start with the most "liberal" (in terms of navigation freedom) category and move to the more "restrictive" one?
- Applying such a ranking incorporates two concepts that may need to be extricated and made explicit: the concept of the "default" category of strategies that might be applicable for a new user; and the concept of a "fallback" category that gets applied when none of the available categories / strategies has the desired effect.

For our example, we will assume that the designer has opted to use the ranking in the order presented above (i.e., "liberal" to "restrictive"), and to let the default and fallback categories be the first and last ones respectively. With these additional constraints, the system would be ready for a second round of user testing.

The introduction of additional structure in the adaptation design space effected in the previous phase, along with more results from user testing, based on that structure, can be expected to finally provide detailed enough results to start building more concrete and comprehensive adaptation logic around the alternative strategies.

According to results from related research in the literature, this phase might result in user model-based adaptation logic along the following lines:

– For users that are novice or unfamiliar with the knowledge domain of the learning material, the more restrictive category (III) of strategies would be applicable.
– Within Category III, a ranking between strategies would be possible, such as: (i) strategy E – link hiding, no explicit recommendations by the system, (ii) combination D+E – link hiding and icons to “explain” the rationale behind provided guidance, and (iii) combination C+E – same as previous, but with less visual clutter.
– Category II (uncombined strategies C or D) would be reserved for users that are sufficiently familiar with the system and the recommendation mechanism, and Category I for users that already have familiarity with the knowledge domain, or exhibit behavior aimed at circumventing constraints applied on their navigation freedom.

Please note that the above adaptation logic is only exemplary in nature and might differ significantly from the actual results one might get with a specific system and learning material. Also note that, although the example case study is ending here, there is no reason why in real-world settings this would be the last design iteration.

4 Adaptive system design revisited

Based on the design evolution outlined in the previous section, we will now move on to an overview of the meta-adaptive facilities being utilized behind the scenes, and their effects on the design process.

To start with, the basis of the design iterations has been the derivation of new knowledge regarding their suitability for different users (or contexts of use) given the overall design goal / self-regulation metrics. This knowledge, in its simplest form, is derived by applying alternative (combinations of) adaptation strategies, and assessing the extent to which the self-regulation metrics are satisfied, always in connection to the current user’s model. Knowledge derivation, then, is achieved by analyzing all recorded cases where a particular strategy has had similar results, and identifying common user model attributes among the respective users. This is the core of the “learning” facilities in the context of self-regulated adaptivity, and their output could be expressed in various forms, including for instance as preliminary adaptation logic, intended to be reviewed, verified and incorporated into the systems by the designers. Although rather straightforward, the above step may already suffice to provide valuable input to the design process. For example, it should be capable to identify strategies that are not suitable for any (category of) users, in any context of use. This was assumed to be the case in the elimination of strategy B and all its combinations in the previous section.

A second set of capabilities alluded to in the previous section is the categorization, or “clustering” of adaptation strategies, as well as their “ranking”. Categorization can take place mainly along two dimensions: (a) The system can try to identify strategies that have similar effects with respect to the self-regulation metrics, given sufficiently similar user models; the output of this process would be a provisional clustering of strategies, based on their "cause and effect" patterns. (b) The system can try to identify the differentiating subsets of user models that render some strategies more effective than others. These dimensions give, respectively, two semantically rich measures of similarity and differentiation of adaptation strategies. When sufficient meta-data about the user model itself is available, the system can combine that with the measures to provide provisional rankings of alternatives within categories.

Before closing it is important to note that the example in the previous section, as well as the analysis in this section, only assume three types of analytical assessment capabilities on the part of a self-regulating adaptive system. Although these are by far not the only ones possible, they are already adequate for the type of design support put forward in this paper. Space constraints do not permit going into a discussion of additional possibilities, or of the requirements placed on the adaptation infrastructure by this type of self-regulation; interested readers are referred to [2] for details on the later.

In conclusion, this paper has presented a case for the use of meta-adaptivity as a facilitator in the design of adaptive systems. The applicability of the proposed approach is of course not universal: it requires that a self-regulating adaptive system (or infrastructure) is already operational, and that the cost of authoring the alternative adaptive behaviors is not prohibitive. It is also mainly intended for cases where there exist several alternative adaptive behaviors, with little or no empirical evidence as to their suitability for different categories of users, or different "states" of a single user. Within these confines, however, it is argued that meta-adaptivity does not represent only the next logical step in the evolution of adaptive systems, but also a potentially irreplaceable tool in how we design such systems in the future.

References

[1] Brusilovsky, P. (1996). Methods and techniques of adaptive hypermedia. *User Modeling and User-Adapted Interaction*, 6(2-3), 87-129.

[2] Paramythis, A. (2004). Towards Self-Regulating Adaptive Systems. In Weibelzahl, S., & Henze, N. (Eds.), *Proceedings of the Annual Workshop of the SIG Adaptivity and User Modeling in Interactive Systems of the German Informatics Society (ABIS04)*, Berlin, October 4-5 (pp. 57-63).

[3] Paramythis, A., & Stephanidis, C. (2005). A Generic Adaptation Framework for Hypermedia Systems. In Chen, S. Y., & Magoulas, G. D. (Eds.) *Adaptable and Adaptive Hypermedia Systems* (pp. 80-103). Idea Group, Inc.

[4] Stephanidis, C., Paramythis, A., Zarikas, V., & Savidis, A. (2004). The PALIO Framework for Adaptive Information Services. In A. Seffah & H. Javahery (Eds.), Multiple User Interfaces: Cross-Platform Applications and Context-Aware Interfaces (pp. 69-92). Chichester, UK: John Wiley & Sons, Ltd.

[5] G. Weber, H.-C. Kuhl, and S. Weibelzahl. Developing adaptive internet based courses with the authoring system NetCoach. In S. Reich, M. Tzagarakis, and P. de Bra (Eds.), *Hypermedia: Openness, Structural Awareness, and Adaptivity*, (pp.226-238), Berlin, 2001.