

Commentary

Discussion of the Article “Website Morphing”

Andrew Gelman

Departments of Statistics and Political Science, Columbia University, New York, New York 10027,
gelman@stat.columbia.edu

The article under discussion illustrates the trade-off between optimization and exploration that is fundamental to statistical experimental design. In this discussion, I suggest that the research under discussion could be made even more effective by checking the fit of the model by comparing observed data to replicated data sets simulated from the fitted model.

Key words: discussion; Web; morphing

History: Received: November 17, 2008; accepted: November 17, 2008. Published online in *Articles in Advance* February 24, 2009.

“Website Morphing” (Hauser et al. 2009) is a fun paper to discuss. In particular I like the summary of technical challenges at the beginning: The website must (1) morph based on relatively few clicks, (2) learn which characteristics are best for which customers, (3) use prior information, and (4) be implemented in real time. So often in applied statistics we find ourselves analyzing a data set in isolation or attacking a large problem with limited data; it is refreshing to see such clear goals, and I think I will try to be focused in this way in my future research projects.

At a more technical level, the trade-off between exploitation and exploration is a standard issue in statistical design of experiments: for exploitation (which we call optimization) we want the best solution based on our current state of knowledge; for exploration (in statistical terms, inference) we want the “leverage” that comes from using design points that can be far from practical. The particular tools developed in this article should be useful for other problems such as computerized adaptive testing in educational applications.

My main suggestion for improving this work is to make it more open-ended by incorporating simulation-based model checking, also called posterior predictive checks (as described, for example, in Chapter 6 of our Bayesian statistics text, Gelman et al. 2003). The idea is to use the fitted model to simulate several replicated data sets and then to compare these to the actual data. The comparisons can be graphical (for example, using scatterplots or time-series plots

of individual users’ patterns of clicks) and numerical (for example, mean squared errors or log likelihoods of observed responses given predicted probabilities from the model). This open-ended exploration gives a framework for the continuous improvement of details of the proposed method as it is being used while keeping the existing inference and optimization structure.

One advantage of a fully Bayesian model (in computer science terminology, a generative model) is that it can be used immediately to simulate replicated data. I recommend that the authors of this article—and the many readers of this journal who will surely want to apply the methods described to their own problems—take advantage of the generative property to explore the ways in which their model fits, and does not fit, reality. Such graphical and numerical diagnostics could be performed using the data gathered from the application of the fitting and optimization procedure to real users.

Acknowledgments

The author thanks the Columbia University Applied Statistics Center, the National Science Foundation, the National Institutes of Health, and Yahoo! Research for partial support of this work.

References

- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Rubin. 2003. *Bayesian Data Analysis*, 2nd ed. Chapman and Hall, London.
- Hauser, J. R., G. L. Urban, G. Liberali, M. Braun. 2009. Website morphing. *Marketing Sci.* 28(2) 202–223.

Commentary

Discussion on “Website Morphing” by
Hauser, Urban, Liberali, and Braun

John Gittins

Department of Statistics, Oxford University, Oxford OX1 3TG, United Kingdom, gittins@stats.ox.ac.uk

These comments are a tribute to an impressive paper and suggestions for clarification of some fairly minor issues.

Key words: website morphing; Gittins index

History: Received: November 13, 2008; accepted: November 17, 2008. Published online in *Articles in Advance* February 24, 2009.

I would like to congratulate Hauser et al. (2009) on their paper. It brings together expertise in website and software design, psychological profiling, marketing research, statistical modeling, hidden Markov models, and adaptive Bayesian procedures in an impressive, practical, and effective way.

My book (Gittins 1989) on what became known as Gittins indices is now 19 years old, and it is very pleasing to have the opportunity to acknowledge one of the best papers since then that applies that theory. I particularly like the simple notion of an expected Gittins index, which Hauser et al. (2009) put to very effective use.

In recent years I have left these developments to others, so I have just a few detailed comments.

1. I would like to avoid confusion as to how the index G_{rmm} is, in principle, to be derived from Equation (1): The Gittins index for morph m with known cognitive segment r is the value of G_{rmm} for which the two terms inside the curly brackets on the right-hand side of Equation (1) are equal.

2. As the authors say, index values depend importantly on the discount parameter a . Note that the authors use $a = 0.999999$ in Table 1, Figure 3, and Appendix 3.

3. I am puzzled by the authors' statement on page 7, “It is a valid fear that such exploration might lead to costly false morph assignments more so than a null strategy of one website for everyone.” In expectation, and with suitable priors and discount parameter, the index strategy with morphing is optimal and beats any strategy of one website for everyone. The simulation is still interesting, of course, but the authors might wish to elaborate.

4. The authors have told me that their calculations of index values for values of a that were very close to 1.0 were very time consuming. With luck, any future calculations could be shortened by making use of the approximations for such a values that are set out in §7.4 of my book. It would be interesting to know, for example, whether the limiting behaviour described in the display (7.15, p. 169) was observed.

References

- Gittins, J. C. 1989. *Allocation Indices for Multi-Armed Bandits*. John Wiley & Sons, London.
- Hauser, J. R., G. L. Urban, G. Liberali, M. Braun. 2009. Website morphing. *Marketing Sci.* 28(2) 202–223.

Commentary

Discussion of “Website Morphing”

Hal Varian

School of Information, University of California, Berkeley, Berkeley, California 94720,
hal@ischool.berkeley.edu

Website morphing seems to be a useful technique, with applications beyond matching cognitive style.

Key words: website; morphing; discussion

History: Received: November 5, 2008; accepted: November 6, 2008. Published online in *Articles in Advance* February 24, 2009.

I think that the concept of website morphing might be a bit broader than that presented by Hauser et al. (2009). Their definition says that morphing matches the look and feel of a website to cognitive styles, but the way they learn about the cognitive style is by observing the click stream, so a more general definition would be that morphing matches the look and feel of a website to the clickstream.

Cognitive style could be one *interpretation* of what is driving click behavior, although it is not the only interpretation. Suppose, for example, that gender is an important characteristic to an advertiser, and that men click on different things than women. After a few clicks, one might be able to predict the gender of the visitor with reasonable accuracy and morph the website to appeal to the appropriate gender.

Experimentation is, of course, widely used by websites to compare different user interface designs. Google Analytics, for example, offers a “website optimizer,” which allows for sophisticated experimentation. What is novel about the approach of Hauser et al. (2009) is that it explicitly recognizes that one size does not necessarily fit all, so the interface optimization should be conditioned on other aspects of observed behavior, such as the prior clickstream.

Their paper is primarily concerned with the case of a single session. However, in practice, state can be stored in cookies and preserved across sessions, meaning that each visitor can have a personalized website. Cookies managed by an ad network also persist across URLs, so that the set of ads that one sees when viewing a major website will generally be different than the set of ads someone else sees, simply because each has a different click and impression history.

Of course, one runs the risk of misclassification, so it is a good idea to allow the user to reset his or her

history to a blank slate. Otherwise, viewers might well be frustrated if the classification is wrong. According to a *Wall Street Journal* story (Zaslow 2002):

Mr. Iwanyk, 32 years old, first suspected that his TiVo thought he was gay, since it inexplicably kept recording programs with gay themes. A film studio executive in Los Angeles and the self-described “straightest guy on earth,” he tried to tame TiVo’s gay fixation by recording war movies and other “guy stuff.”

“The problem was, I overcompensated,” he says. “It started giving me documentaries on Joseph Goebbels and Adolf Eichmann. It stopped thinking I was gay and decided I was a crazy guy reminiscing about the Third Reich.”

Morphing a user interface based on personalized interaction history is not uncommon. In fact, Microsoft Office reorders menu options depending on how frequently a particular operation is performed. This is, admittedly, a trivial example compared to changing the entire the look and feel of a website.

What is attractive about the model presented in the paper of Hauser et al. (2009) is that it delivers a rigorous, well-thought-out strategy that encompasses experimentation, optimization, and personalization. I suspect that there will be other exciting applications of these techniques.

References

- Hauser, J. R., G. L. Urban, G. Liberali, M. Braun. 2009. Website morphing. *Marketing Sci.* 28(2) 202–223.
- Zaslow, J. 2002. If TiVo thinks you are gay, here’s how to set it straight. *Wall Street Journal* (November 26), http://online.wsj.com/article_email/SB1038261936872356908.html.