What Else and Where Else: Two Worthwhile Questions for an Information Interface

Aran Lunzer and Yoshiki Ohshima

This material is based upon work supported in part by the National Science Foundation under Grant No. 0639876. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Presented at NordiCHI '12, October 14-17, 2012, Copenhagen, Denmark. Proceedings of the ACM 978-1-4503-1482-4/12/10. 2pp.

VPRI Memo M-2012-008

Viewpoints Research Institute, 1209 Grand Central Avenue, Glendale, CA 91201 t: (818) 332-3001 f: (818) 244-9761

What Else and Where Else: Two Worthwhile Questions for an Information Interface

Aran Lunzer and Yoshiki Ohshima
Viewpoints Research Institute
1209 Grand Central Avenue, Glendale, CA 91201, USA
aran@acm.org, yoshiki@vpri.org

ABSTRACT
We make a case for widespread support, in interfaces for retrieving and laying out information items, for two simple operations on the items being displayed. The operations are predicated on the fact that the calculations determining how items should be laid out, for example into a ranking, are often highly sensitive to criteria that, for a given application, are somewhat arbitrary. Subtly different criteria can lead to dramatically different results. Our suggestion is to support users in understanding such differences among multiple result displays, with a 'What else?' operation that shows which items appear in a given location in any of those displays, and conversely a 'Where else?' that shows other locations where a given item can appear.

Author Keywords
Information displays; result comparison; exploration

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation: User interfaces—*Graphical user interfaces*

INTRODUCTION

'Comparisons are odious', as people have been saying to each other since at least the 15th century. Recent trends in information presentation appear to have forgotten this truth. To be clear, not *all* comparisons are problematic: the kind that are under suspicion are those that, by focussing on limited features of the available candidates, cast some as winners and some as losers without taking into account their other characteristics, good or bad. So how problematic is it when a search engine delivers a page offering the 'most relevant' of two million wildly heterogeneous Web pages, or a magazine publishes a survey suggesting the 'ten best science-fiction films', or 'top fifty liveable cities'?

The fact is that we are constantly faced with choices, and even if it were theoretically possible to gather and rank candidates by marshalling and comparing all their details, in trying to do so we would simply run out of time in the day. Resorting to simplifications, such as rankings derived for us by agencies we trust, is the only practical approach. But how good are the comparisons that are being made on our behalf? A Web search engine might take into account dozens of item features, along with other factors such as a user's current location, interaction history and even shopping records, but given the huge number of items being filtered even a tiny change in how these factors are weighted can lead to completely different results. In spite of this, users commonly place great trust in search engines. For example, Hearst [1, p. 136] writes: 'Studies and query logs show that searchers rarely look beyond the first page of search results. If the searcher does not find what they want in the first page, they usually either give up or reformulate their query.' Furthermore, at least one study has revealed an expectation that even within a search-results page the higher results are better; subjects were observed choosing the items in higher positions even when they were less relevant to the task [4].

When it comes to film surveys, people know that any given magazine's suggestions depend heavily on the writers' particular criteria and tastes. Someone who cares about finding good films would do well to compare the results of surveys from additional sources, whose criteria will be different but potentially just as valid. Similarly, Web searches based on subtly different criteria or weightings can yield results that are different yet equally interesting. The goal of the current work is to encourage such exploration beyond the limits of a single result display, by making it easy for users to probe the similarities and differences between alternatives.

PROBING ALTERNATIVE DISPLAYS

The operational details for requesting multiple alternative yet related result displays depend on the particular application, and are beyond the scope of this discussion. We assume the existence of an environment that makes such requests straightforward, for example based on subjunctive-interface techniques [2, 3] in which alternatives are represented as distinct scenarios that can be viewed and manipulated in parallel. We then propose enabling the user to probe these alternatives with the following two contextual operations on the displays:

What else could I be seeing here?

This question, posed at a particular location in the display (such as the top entry in a list, or an instrument used for displaying some measurement), should bring up a view of the various alternative values that appear at that location under the set of scenarios currently being considered.

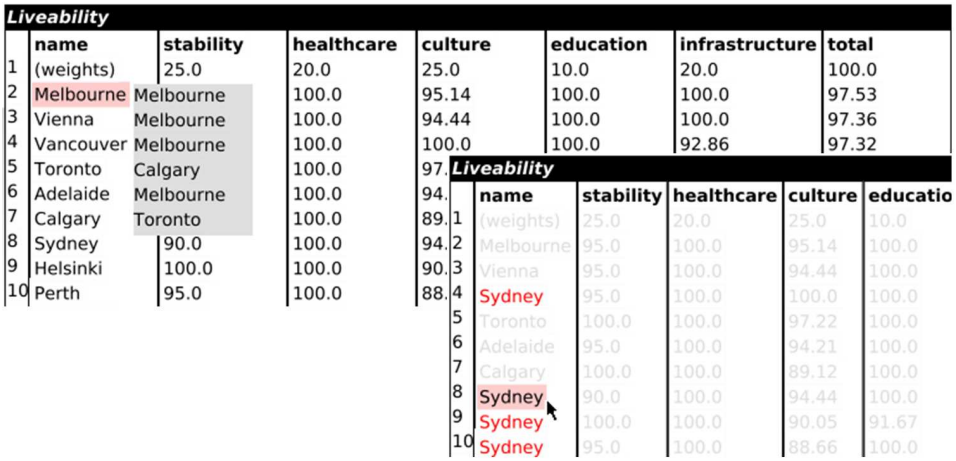

Figure 1. The 'What else?' and 'Where else?' operations at work

Where else could I be seeing this?

This is the complementary question: it should support the user in seeing where else in the display the value being pointed to (which could be a graphical object or a piece of text) would appear in the alternative scenarios.

APPLICATION TO A WEIGHTED RANKING

The figure above shows an example of asking the 'What else?' and 'Where else?' questions on cells within a sorted table of cities from the Economist Intelligence Unit's Global Liveability Report (**`http://www.eiu.com/site_info.asp?info_name=The_Global_Liveability_Report`**). In the standard report, each city is scored on five criteria that are combined, with the weights shown in the top row of the table, to give an aggregate score by which the cities are then ranked. A natural question to ask is whether the ranking is sensitive to the choice of weights. In our interface this can be explored by setting up multiple scenarios, each with a different weighting scheme: in the figure are six scenarios representing the original scheme and, for each of the five criteria, a scheme in which that criterion is weighted as zero.

The result of the 'What else?' question on Melbourne, sitting in top spot in the list, is a menu showing that Melbourne occupies this location in all but two of the scenarios. At this point, clicking on the entry for Calgary, for example, would switch the interface to a presentation of the ranking in which, thanks to a different set of weightings, Calgary came out top.

Below and right, the user has invoked the 'Where else?' query on Sydney, in 7th place (row 8). Three other positions in the list are highlighted. A user who wants to understand the circumstances under which Sydney would rank 3rd can click on that entry and be taken to the corresponding results. Note that a highlight might represent more than one scenario: in this example it happens that there are two scenarios in which Sydney occupies 8th place. In such cases, the interface must help the user to see the differences between those scenarios.

Here the small data set and the relatively bold differences between scenarios make for a rewarding exploration of the alternatives. Situations involving subtler differences are likely to be frustrating for users unless we add markup to highlight display locations that are worth exploring. Conversely, in situations in which the scenarios differ grossly it would be hard for a user to make sense of the differences without being able to probe larger regions, such as multiple cells at a time.

CONTINUING DEVELOPMENT

Our current experiments are based on a personal-computing platform, itself still under development, in which support for creating, maintaining and viewing multiple scenarios is built in as a fundamental feature of the platform. The key enabling technologies are an implementation of Worlds [5]—a mechanism for running isolated (e.g., experimental) calculations in parallel—coupled with an interface architecture based on a side-effect-free reactive language. Moving on from the operations demonstrated here, our goal is to make it straightforward for developers of information interfaces to support users in exploring variations—not just in ranked-list retrievals, but in a wide variety of application and display types.

ACKNOWLEDGEMENTS

We thank our colleagues and collaborators for fruitful discussions, and the reviewers for their suggestions. This material is based upon work supported in part by the National Science Foundation under Grant No. 0639876.

REFERENCES

1. Hearst, M.A. *Search User Interfaces*. Cambridge University Press, 2009.
2. Lunzer, A. and Hornbæk, K. Subjunctive interfaces: Extending applications to support parallel setup, viewing and control of alternative scenarios. *ACM Trans. Comput.-Human Interact.* 14, 4, Article 17 (January 2008), 1–44.
3. Lunzer, A. and Hornbæk, K. Subjunctive interfaces for the Web. Chapter in Cypher, A., Dontcheva, M., Lau, T. and Nichols, J.W. (eds.) *No Code Required: Giving Users Tools to Transform the Web*, Morgan Kaufman (2010), 267–285.
4. Pan, B., Hembrooke, H., Joachims, T., Lorigo, L., Gay, G., and Granka, L. In Google we trust: Users' decisions on rank, position, and relevance. Journal of Computer-Mediated Communication, 12(3), article 3 (2007). **`http://jcmc.indiana.edu/vol12/issue3/pan.html`**
5. Warth, A., Ohshima, Y., Kaehler, T. and Kay, A. Worlds: Controlling the Scope of Side Effects. In *Proc. ECOOP 2011*, LNCS 6813, Springer (2011), 179–203.