Social Intellisense: A Task-Embedded Interface to Folksonomies

Scott Counts[1], Kristie Fisher[2], & Aaron Hoff[1]

[1]Microsoft Research, One Microsoft Way, Redmond, WA 98052
[2]University of Washington, Department of Psychology, Box 351525, Seattle, WA 98195

{counts, aaronho}@microsoft.com, kjfisher@u.washington.edu

Abstract

We present an application for accessing and creating socially constructed sets of information. Users store and retrieve information, such as bits of text, through the use of "keys" consisting of node names strung together using periods (e.g., "conferences.icwsm.dates") that point to user-added content ("May 23-26, 2010"). Furthermore, nodes can point to services on the web (e.g., "facebook.me.status" retrieves the user's current Facebook status). In short, Social Intellisense is a user interface mechanism for folksonomies that allows users to collaboratively "remap" web and user-added content via strings of their own creation. A study with our Social Intellisense-enabled email client showed that the paradigm is comparable or preferable to traditional web browser methods for retrieving certain information online. Issues and potential solutions to mismatches between a user's task at hand and collaboratively organized information are discussed.

Introduction

Social mechanisms for organizing information have proliferated in recent years. These often take the form of tagging systems in which users tag content with a word (e.g., tagging a photograph with the word 'sunset') and in the aggregate these tags form an organizational structure. Many notable websites, including CiteULike, Del.icio.us, and Flickr, utilize user-generated tags. Given that these systems organize information and are both social and somewhat informal, they are often referred to as "folksonomies", a neologism coined to capture the notion of user-generated, bottom-up creation of categorical structure.

Given the potential power of social tagging systems (see Shirky, 2005 for a thoughtful essay) and folksonomies generally, tools that can harness a collective group of people to organize digital information may prove critical as the sheer amount of information increases (Gantz et al., 2007). Especially as we move to units of information that are sub-document (e.g., "tweets"), novel user interface paradigms will be needed to support collaborative information organization and retrieval.

We created Social Intellisense (SI) as one attempt at such an interface. SI leverages the IntelliSense user interface concept[1] of hierarchically organized information accessed through strings of dot-connected words, with a drop down menu showing options at each consecutive node. This hierarchical structure distinguishes SI from tagging systems and yields a final product closer to an ontological organization. In this paper we present the SI system, along with results of a companion user study.

Background

As more of the information we access is on the web rather than our personal computers, the aforementioned tagging systems and other social information organization mechanisms are being utilized. A growing body of research literature has focused on analyses of existing socially-derived content structures such as del.icio.us (Golder & Huberman, 2006), Flickr (Marlow et al., 2006), and Dogear (Millen et al., 2006). These analyses are of sufficient scale to uncover real user behavior patterns highly important to the study of collaborative information organization more broadly. For example, Biddulph (2004) reports that some tags are used far more frequently than others, reflecting a common notion of appropriate words for tags among users. Golder and Huberman (2006) also report that more general tags tend to be used earlier than specific tags in the set of tags for a given URL. These findings imply that in social organizational systems users may adhere to a commonly understood structure for the information that starts general and moves to specific. The SI user interface leverages this notion in that users need only remember a general term (e.g., 'facebook' or 'wiki') to start the process of storing or retrieving a specific piece of information.

[1] http://en.wikipedia.org/wiki/Intellisense

The growth of information on the web and alternative organization mechanisms has not, however, replaced traditional hierarchical structures for organizing digital content. In fact, the file system on most personal computers provides a common example from daily life that has proven resilient to change over the past couple of decades. This resiliency is not without good reason, as hierarchical organization has its benefits. Folders provide "locations" that help users find content (Berlin et al., 1993), and in comparison to search, folders increase user control and help users better understand their information (Jones et al, 2005). Recent updates on the folder concept provide additional flexibility. For example, Facet Folders (Weiland & Dachselt, 2008) allow users to view files in hierarchical folder structures that can be rearranged based on desired metadata. Given these benefits of folders, SI also incorporates hierarchical structure to help users organize and navigate information.

Social Intellisense Prototype

The SI prototype was developed to bring together the benefits of hierarchical and social organization mechanisms into an interface embedded in the user's workflow. In the prototype, SI functionality is provided in our custom email client. SI is accessed simply by typing top-level hierarchy words followed by a dot (e.g., "flickr." or "conferences."), at which point a dropdown menu appears revealing items in the next level of the hierarchy (Figure 1a). The user simply continues typing or clicking words in consecutive drop-downs until an end node is reached at which point hitting enter inserts the content into the document. Words are auto-completed for faster typing. Thus, from an information access standpoint, the user is able to bring web and user-added content into her document without stopping typing.

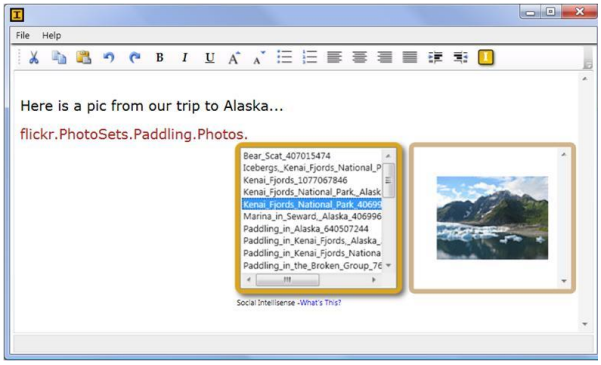

Figure 1a: Retrieving a photo from Flickr using SI. Hitting enter inserts this photo into the document. The drop-down menu is present at each level of the hierarchy, with a content preview at the lowest level.

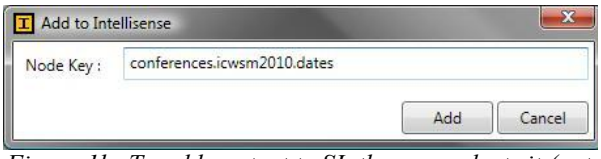

Figure 1b: To add content to SI, the user selects it (not shown), clicks the "I" icon at the right of the menu bar (see Figure 1a), then types the desired key.

Users can add content to a hierarchy by selecting the content they wish to add, such as a snippet of text or URL received in email, and then clicking a button at the right of the tool bar of the application. This opens a small dialogue box into which the user types the key representing the location desired for their selected content (Figure 1b). Currently we support adding content only into the non-service parts of the hierarchy. That is, a user can add to "username." or "conferences.", but not to "facebook.". Any part of a hierarchy that can be written to can also be overwritten and users can remove content and nodes by typing a minus at the end of a SI key and typing enter. Manipulations of hierarchies require appropriate permission: they must either be public or private and owned by the user. Private content is simply any content under a "private." node in a user's personal hierarchy (i.e., under "username.") which is tied to her corporate login. In terms of content types, to date we have focused on text and photos as they support a large number of scenarios.

We see SI as social in two ways. First, users can collaboratively create information hierarchies, and second, SI draws from existing social media information sources. Internet-based services currently supported include Facebook, Flickr, Twitter, Picasa, Wikipedia, del.ico.us, MSN Money, and any user-specified RSS feed. For services requiring authentication, SI logs users in as necessary.

User Study

Research Questions

RQ1 How does use of SI's web-based services compare to traditional web search and browse?
RQ2 What are the effects of a shared hierarchy whose organization is inconsistent with a user's current task?

Method

Participants Participants were 23 employees of our corporation. There were 19 men and four women, with an average age of 30.47 (SD = 5.66). They were recruited through a mass e-mail and were compensated with a $10 lunch coupon.

Procedure Participants were first given a 15 minute orientation on how to use SI, including inputting information, retrieving information, and using each of the web-based services built into the prototype. After the orientation, participants completed two tasks using the SI prototype.

Task 1: SI vs. Web The purpose of this task was to have users directly compare the experience of finding and inserting information using a web browser with that of using SI. Participants completed two form e-mails in the SI email client that were 250-275 words in length. Each email contained 10 "missing" items, which participants were asked to find and insert. These e-mails were meant to

simulate a correspondence with a friend, and the missing items were pictures and pieces of information that could be found online using Facebook (with each participant accessing his or her own account), Flickr (using an account that we created specifically for the study), Wikipedia, MSN Money, or a search engine. For one e-mail, participants were asked to use a browser to find and insert the items using copying and pasting functions, and for the other e-mail they were asked to use SI to find and insert the items. The method that participants were asked to use for each e-mail, and which method they used first, was counterbalanced across participants. Participants were given five minutes to complete each e-mail.

Task 2: Shared Hierarchy Problems may arise when a user must navigate a shared hierarchy of user-added content that was created and organized in part or entirely by other people. Even if the information is organized in a meaningful fashion, the organization may not be appropriate to any given user's task at hand. Task 2 was designed with this issue in mind. Participants were again given two consecutive form e-mails within the SI-enabled email client. Each e-mail was written to simulate a business correspondence from a corporate employee, and each was between 250-275 words in length. Again each e-mail contained 10 "missing" items that participants were asked to find and insert using SI. For one e-mail, they were asked to find the items using a hierarchy that was organized to correspond more directly to the task (henceforth called the "Good" hierarchy). For the other e-mail, participants were asked to find and insert the items using a hierarchy that, though organized logically, did not have an intuitive organization for the task at hand (henceforth called the "Bad" hierarchy). The two hierarchies were hand crafted by the experimenters to be sensible but more or less well matched to the task. As one example item from the task, participants were asked to retrieve company fourth quarter earnings where the appropriate leaf node "Earnings_Releases.Fourth_Quarter" could be found either under "Investor_Information." (good hierarchy) or under "Frequently_Requested_Info." (bad hierarchy). Condition orders were counterbalanced. Participants were given five minutes to complete each e-mail.

Results

RQ1: SI vs. Web (Task 1) SI was compared to browser-based web search in Task 1, in which participants were asked to find and insert information into email. Participants were able to complete more items when using SI, $F(1,21) = 13.18$, $MSE = 2.58$, $p < .01$, although they rated SI as less natural and intuitive to use than the web $F(1,21) = 5.21$, $MSE = 2.09$, $p < .05$. See Table 1 for complete results.

RQ2: Mismatch of shared hierarchy and user task (Task 2) As can be seen in Table 2, the Bad hierarchy led to poorer performance on most measures, lower ratings, and a higher user desire to change it. A repeated-measures ANOVA with hierarchy (Bad vs. Good), was performed on the data from Task 2. The organization of the hierarchy with respect to the task at hand made a significant difference in percentage of correct items retrieved $F(1, 21) = 10.25$, $MSE = .025$, $p < .01$, time needed to complete each item $F(1, 21) = 7.74$, $MSE = 161.40$, $p < .05$, and the average number of extra steps needed to find each item $F(1, 21) = 8.82$, $MSE = 2.12$, $p < .05$. There were also significant differences between the Bad and Good

Method	*Measure*	*Result*
SI	Total items completed*	7.13 (2.03)
	% items correct	68% (20%)
	How efficient? (Q)	5.30 (1.33)
	How natural and intuitive? (Q)*	4.96 (1.43)
Web	Total items completed*	5.39 (1.62)
	% items correct	63% (20%)
	How efficient? (Q)	4.74 (1.32)
	How natural and intuitive? (Q)*	5.87 (1.10)

*Table 1: Social Intellisense versus web search during Task 1. SDs in parentheses. (Q) denotes 7-point Likert scale questionnaire items. Non-specified numbers represent means. * = $p < .05$ significant differences between SI and web*

Hier.	*Measure*	*Avg. (SD)*
Bad	Total items inserted	6.57 (2.52)
	% correct*	59% (27%)
	Time to insert item (sec.)*	34.09 (18.01)
	Extra steps to complete each item*	1.80 (2.03)
	# items requiring extra steps	1.91 (1.76)
	How easy was it to find information? (Q)*	3.83 (1.30)
	How well-organized was the hierarchy? (Q)*	3.26 (1.36)
	How likely would you be to change the hierarchy? (Q)*	5.78 (1.44)
Good	Total items inserted	7.70 (2.18)
	% correct*	74% (23%)
	Time to insert item (sec.)*	23.73 (8.74)
	Extra steps to complete each item*	.47 (.45)
	# items requiring extra steps	.65 (.57)
	How easy was it to find information? (Q)*	5.74 (1.36)
	How well-organized was the hierarchy? (Q)*	6.22 (.95)
	How likely would you be to change the hierarchy? (Q)*	2.87 (1.87)

*Table 2: Performance and questionnaire responses from Task 2. (Q) denotes 7-point Likert scale questionnaire items * = $p < .05$ significant differences between good and bad hierarchies.*

hierarchies in how participants rated the ease of finding information $F(1,21) = 30.88$, $MSE = 1.42$, $p < .001$, how well-organized they thought the hierarchy was $F(1,21) = 70.60$, $MSE = 1.43$, $p < .001$, and how likely they would be to change it $F(1,21) = 54.37$, $MSE = 1.79$, $p < .001$.

Subjective Experience and Qualitative Feedback To assess subjective user experience, we examined survey questions and open-ended feedback collected at the end of our study. Most users had a preference for SI ($Pref_{SI}$=87%, $Pref_{Web}$=4.3%, $Pref_{None}$=8.7%) and thought that SI felt faster than cutting and pasting from the internet ($Longer_{SI}$=26.1%, $Longer_{Web}$=60.9%, $Longer_{none}$=13%).

The majority of the open-ended feedback boiled down to making SI more flexible, in that rather than linearly searching from the first order node to the leaf node, users suggested searching through the structure backwards or "sideways" through related nodes. This functionality is particularly relevant to the idea that whether or not a hierarchy is well organized is primarily a function of what it is being used for and is similar to faceted browsing. So called "sloppy" syntax approaches to command interfaces (e.g., Inky; Miller et al., 2008), also suggest an alternative for increasing flexibility.

Discussion

Results from the user study show that users were relatively facile with SI as first-time users. For example, in the first task users retrieved more information using SI than from the web, though we acknowledge that should SI contain as much information as accessible on the web generally, these results may be different (such a comparison was not possible to set up experimentally). However, despite users' web proficiency and the sophistication of search engines (reflected in the expected ratings of the web feeling more natural), it's important to point out that users actually preferred SI and felt it faster than the web.

Similar to known issues with group information management, such as everyone organizing and naming files and folders differently (the "vocabulary problem"; e.g., Furnas et al., 1987; Berlin et al., 1993; Rader, 2009), Task 2 highlighted the potential cost of information organized in a manner ill-suited to a user's task, something likely to happen at least occasionally when information is socially organized. Results from Task 2 showed that while users performed reasonably well, the negative effects of the "bad" hierarchy in terms of both usage (time and errors) and subjective impression (e.g., ease of finding information), were significant. In response to issues with incompatibilities between information organization and user task, we propose "pointer" functionality that allows users to map one key to another. For example, information under the key "workgroup.conferences.2010.icwsm.dates" added by person A could be pointed to by "userb.icwsm2010.dates" if person B preferred such a mapping. We hope this is a "best of both worlds" solution that lets users leverage the content and organization of others when appropriate while maintaining an organization that makes most sense to them and is best suited to their tasks.

Conclusion

As the amount of digital information continues to grow, social approaches to helping users manage and access that information within their computing contexts are likely to play an increasingly important role. Here we presented Social Intellisense as one possible interface for folksonomies that supports storage and retrieval of user-added and web service-based content in email. Users create, access, modify, and delete nodes in hierarchies that contain pieces of information such as text or photos, and this construction of the information space can be collaborative.

Our user study showed that SI was better than web-based methods in terms of speed for retrieving information in the context tested (Task 1). The results of Task 2 highlight difficulties with social organization systems and point to future efforts such as the proposed "pointer" functionality, that preserve the benefits of social contribution and organization while helping users maintain their own mental model of the information and keep it organized in a personalized and task-relevant way.

References

Biddulph, M. (2004). Introducing Del.icio.us. XML.com. http://www.xml.com/pub/a/2004/11/10/delicious.html

Furnas, G., Landauer, T., Gomez, L., & Dumais, S. The vocabulary problem in human-system communication. *Communications of the ACM, 30(11):* 964-971, 1987.

Gantz, J., Reinsel, D., Chute, C., Schlichting, W., McArthur, J., Minton, S., Xheneti, I., Toncheva, A., & Manfrediz, A. (2007). The Expanding Digital Universe. IDC whitepaper.

Golder, S., & Huberman, B. (2006). Usage patterns of collaborative tagging systems. *J. of Information Science, 32 (2)*,198-208.

Jones, W., Phuwanartnurak, A., Gill, R., Bruce, H. (2005). Don't take my folders away!: Organizing personal information to get things done. *In CHI '05 EA*.

Marlow, C., Naaman, M., boyd, d., & Davis, M. (2006). HT06, tagging paper, taxonomy, Flickr, academic article, to read. *In Proc. Hypertext and Hypermedia '06*.

Millen, D., Feinberg, J., & Kerr, B. (2006). Dogear: Social bookmarking in the enterprise. *In Proc. CHI '06*.

Miller, R., Chou, V., Berstein, M., Little, G., Van Kleek, M., Karger, D., & schraefel, m. (2008). *In UIST '08*.

Shirky, C. Ontology is Overrated: Categories, Links, Tags. www.shirky.com/writings/ontology_overrated.html

Weiland, M., & Dachselt, R. (2008). Facet folders: flexible filter hierarchies with faceted metadata. *In Proc. CHI '08*.