University of Maryland's
Human Computer Interaction Lab
27TH ANNUAL SYMPOSIUM

MAY 27-28, 2010

design for the world

Presentation Summaries, Technical Report Abstracts, & Other Information

HCIL is a partnership of the University of Maryland Institute for Advanced Computer Studies (UMIACS) in the College of Computer, Mathematical and Physical Sciences (CMPS) and the College of Information Studies - Maryland's iSchool

www.cs.umd.edu/hcil

Special thanks to our Champion Symposium Sponsors:

Special thanks to our Symposium Sponsors:

Special thanks to our University of Maryland Sponsors:

College of Computer, Mathematical and Physical Sciences
College of Information Studies – Maryland's iSchool
Department of Computer Science
The Institute for Systems Research
Dingman Center for Entrepreneurship
Maryland Institute for Technology in the Humanities
Department of Electrical and Computer Engineering

Human-Computer Interaction Lab

University of Maryland

May 27th, 2010

Welcome to the 27th Annual HCIL Symposium!

Thank you for joining us for the Human-Computer Interaction Lab's 27th annual Symposium. This year you will hear diverse talks from the HCIL's faculty, staff, and students on cutting-edge research which focus on three important areas of HCI work: communities, text/translation, and search. Following these talks, we will continue our long-standing tradition of welcoming you to our lab for demos and posters (and of course birthday cake)! We hope you will be staying the next day as well, May 28, 2010, to enjoy a wide variety of tutorial and workshop topics that range from iPhone development to consumer health informatics.

The Human-Computer Interaction lab has a long, rich history of transforming the experience people have with new technologies. From understanding user needs, to developing and evaluating those innovative technologies, the lab's faculty, staff, and students have been leading the way in HCI research and teaching. We believe it is critical to understand how the needs and dreams of people can be reflected in our future technologies. To this end, the HCIL develops advanced user interfaces and design methodology. Our primary activities include collaborative research, publications, and the sponsorship of open houses, workshops and symposiums. This year we have added a new lab activity, honoring our "HCI Heroes" with our first annual HCIL Awards. During the Symposium we will award these special people for the work they do.

In closing, it's an exciting time of growth for the lab and we are happy you can be here to share it with us. We look forward to hearing your thoughts and feedback on the symposium, our research, and our on-going lab programs. It is with your help, we can continue to change the world of HCI for many years to come.

Sincerely,

Dr. Allison Druin

Director, Human-Computer Interaction Lab and Associate Professor
College of Information Studies (Maryland's iSchool)
Institute for Advanced Computer Studies (UMIACS)

Allison Druin ♦ allisond@umiacs.umd.edu ♦ http://www.umiacs.umd.edu/~allisond/
College Park, Maryland 20742 ♦ (301) 405-7406

Human-Computer Interaction Lab

U n i v e r s i t y o f M a r y l a n d

The Human-Computer Interaction Lab (HCIL) was established in 1983 as an interdisciplinary effort within the Institute for Advanced Computer Studies. Today HCIL participants include faculty, staff, and students from the Department of Computer Science, the College of Information Studies – Maryland's iSchool, the Department of Psychology, the College of Education, the Department of English, The Maryland Institute for Technology in the Humanities, The College of Engineering, The Institute for Systems Resaerch, and the Dingman Center for Entrepreneurship in the Robert H. Smith School of Business, at the University of Maryland, College Park.

This booklet contains all Symposium presentation summaries, posters and a sampling of technical report abstracts that have come out since our 2009 Symposium.

Please see our web site (www.cs.umd.edu/hcil) for a complete list and on-line versions of these and other papers.

If you would like more information, please contact:

Human-Computer Interaction Lab
2117 Hornbake Bldg, South Wing
College Park, Maryland 20742
(301) 405-2769
hcil-info@cs.umd.edu

Special thanks to our Research Sponsors:

Amazon Web Services
Carnegie Hall
Center for Integrated Transportation Systems Managemement (USDOT-RITA)
Department of Defense
Defense Advanced Research Projects Agency (DARPA)
Google
IBM
Intel Corp.
Institute of Museums and Library Sciences (IMLS)
Lockheed Martin
Microsoft Research
Mongolia Ministry of Education, Culture and Science
National Center for Health Statistics
National Institutes of Health
National Science Foundation (NSF)
National Endowment for the Humanities (NEH)
Office of the National Coordinator for Health Information Technology
Pacific Northwest National Laboratory
Samsung
Sesame Workshop
Tibco/Spotfire
U. S. National Park Service
Washington Hospital Center

Table of Contents

27th Annual Symposium
May 27, 2010

8:15am Sign-In - and Coffee, Tea, Breakfast

9:00am Welcome

Jenny Preece
Dean, College of Information Studies - Maryland's iSchool

Allison Druin
Director of HCIL

Jennifer Golbeck
Co-Director, HCIL

9:20am Keynote Talk

Keynote: The promise of Zoomable User Interfaces
Ben Bederson

SESSION I: Communities – 9:50am
Session Chair: *Jenny Preece, Dean, College of Information Studies - Maryland's iSchool*

Self-Promotion in 140 Characters: The Use of Twitter by Congress
Jennifer Golbeck, Justin Grimes, Anthony Rogers

Analyzing Social Media Networks with NodeXL
Derek Hansen, Cody Dunne, Ben Shneiderman

ManyNets: An Interface for Multiple Network Analysis and Visualization
Manuel Freire, Awalin Nabila, Miguel Rios, Jennifer Golbeck, Catherine Plaisant, Ben Shneiderman

SHORT BREAK – 10:40am

SESSION I continued: 11:00am

New Design Methods with Children: Layered Elaboration
Greg Walsh, Mona Leigh Guha, Beth Bonsignore, Beth Foss, Evan Golub, Leshell Hatley

Designing Social Musical Technologies at Carnegie Hall
Quincy Brown, Greg Walsh, Allison Druin, Ben Bederson, Jennifer Golbeck

SESSION II: Text and Translation – 11:35am
Session chair: *Catherine Plaisant, Associate Director of Research, HCIL*

Human-Computer Collaborative Translation
Chang Hu, Ben Bederson, Philip Resnick

Finding Entries in an On-line Arabic Dictionary
Sarah Wayland, C. Anton Rytting, David Zajic, Timothy Buckwalter, Jason White, Corey Miller, Jeffrey Carnes, Nathanael Lynn, Paul Rodrigues, Michael Maxwell, Evelyn Browne

LUNCH 12:10 – 1:10pm

SESSION II continued – 1:10pm

iOpener Workbench: Tools for Rapid Understanding of Scientific Literature
Cody Dunne, Ben Shneiderman, Bonnie Dorr, Judith Klavans

CrowdFlow: A Human-Computer Hybrid Cloud Computing Model
Alex Quinn, Jimmy Lin, Ben Bederson, Tom Yeh

SESSION III: Search – 1:45pm
Session chair: *Michael Pack, Director, CATT Laboratory*

How Children Search Online at Home
Allison Druin, Beth Foss, Evan Golub, Leshell Hatley

Finding Temporal Patterns in Electronic Health Records
Krist Wongsuphasawat, David Wang, Catherine Plaisant, Ben Shneiderman

Analyzing Trends in Science & Technology Innovation
Ben Shneiderman, Ping Wang, Yan Qu, Cody Dunne

Demonstrations & Posters
2:45pm to 4:00pm – HCIL Lab, 2117 Hornbake Bldg, South Wing

StoryKit: Designing a Mobile Application for Story Creation By Children And Older Adults
Alex Quinn, Beth Bonsignore, Ben Bederson, Allison Druin

International Children's Digital Library (ICDL)
Anne Rose

Sikuli: Using GUI Screenshot for Search and Automation
Tom Yeh

Educational Simulation for Computing and Information Ethics
Ken Fleischmann, Russ Robbins, and Al Wallace

Perpetual Connectivity, Selective Participation - on the Mobile-Social Practices of Students and Faculty
Dana Rotman

Evaluating Visual Analytics Systems with Contests
Swetha Reddy, Catherine Plaisant

Guidelines for a Costume Designers Workbench
Rachael Bradley

The Regional Integrated Transportation Information System
Michael VanDaniker

Exploring and Visualizing Crashes
Darya Filippova

Dependency Graphing of Incident Datasets
Jon Gilmour

Virtual Incident Management Training through MMOGs
Walter Lucman, Mic Couture

4D, Wide-Area Temporal and Spatial Data Visualization
Drew Lund

Toque: Designing a Cooking-Based Programming Language For and With Children
Sureyya Tarkan, Allison Druin, Evan Golub, and Vibha Sazawal

New Interaction Styles and Paradigms in Educational Environments
Ana I. Molina, Francisco Jurado and Manuel Ortega

Lifelines2: Search, Group, and Comparison of Personal Records
David Wang, Catherine Plaisant, Ben Shneiderman

Immersive People, Immersive Games, and Instruments to Measure Them
Kent Norman

Item analysis of usability tasks
Susan Campbell

Comparing Small Graphs
Benjamin Smith

Investigating the Impact of Design Processes on Children
Mona Leigh Guha, Allison Druin, Jerry Alan Fails

A Dual-Display Reading Device in the Literature Classroom
Nick Chen

Usability Testing of Web-based Multimedia Health Tutorials by Older Adults
Bo Xie

Collaborative Identification and Annotation of Government Deep Web Resources
Pengyi Zhang, Yan Qu, Chen Huang, Paul T. Jaeger, John Wells, W. Scott Hayes, James E. Hayes, and Xin Jin

Presentation Summaries

The Promise of Zoomable User Interfaces

Benjamin B. Bederson
Human-Computer Interaction Lab
Department of Computer Science, UMIACS & iSchool
Contact: bederson@cs.umd.edu

Expanded version at HCIL Tech Report HCIL-2002-21
http://www.cs.umd.edu/localphp/hcil/tech-reports-search.php?number=2009-21

Zoomable User Interfaces (ZUIs) have received a significant amount of attention in the 17 years since they were introduced. They have enjoyed some success, and elements of ZUIs are widely used in computers today, although the grand vision of a zoomable desktop has not materialized. This paper describes the premise and promise of ZUIs along with their challenges. It describes design guidelines, and offers a cautionary tale about research and innovation.

The essential problem that Zoomable User Interfaces (ZUIs) aim to solve is a fundamental one – that there is more information than fits on the screen. The common solutions to this problem are, roughly, scrolling, linking & searching, along with denser representations (i.e., information visualization). Zooming, like fisheye displays, is an instance of the latter – a kind of information visualization that aims to take advantage of human spatial perception and memory. ZUIs place documents in two-dimensional space at any size, enabling (and requiring) animated spatial navigation to move among documents.

I have identified three key characteristics that have attracted people's attention over the years. The promise of ZUIs comes largely from the following general expectations.

Engaging: The animation is visually attention grabbing. It takes advantage of human visual perception abilities.
Visually rich: There are more degrees of freedom to visually structure objects, and thus they offer the potential of great creative expression.
Lure of simplicity: The fact that you find information by looking for it in a place implies a promise of simplicity that will solve our organizational and information retrieval problems.
But the potential benefits of ZUIs are sometimes mirages. ZUIs can be engaging, but they also make some people feel physically sick. They can be visually rich, but only if the author knows how to take good advantage of the platform's capabilities. Finally, the promise of simplicity often falls short. While human visual perception does make it easy to see where one is navigating, the reality is that it places a heavy load on short term memory to remember where in space you just were and where things are. And the requirement of human memory to know how space is organized means that ZUIs don't scale up very well. ZUIs are often motivated by the physical world and how people like laying papers out on their desk. But no one wants all of their papers on their desk. It is much more common to have only a relatively small number of papers that are actually being worked with.

So this talk discusses the potential of ZUIs along with their challengs along with a number of design guidelines that should be considered when designing ZUIs.

PAPERS FOR FURTHER READING

1. Bederson, B. B. and Hollan, J. D. 1994. Pad++: a zooming graphical interface for exploring alternate interface physics. *In Proc. of the 7th Annual ACM Symposium on User interface Software and Technology UIST 1994*. ACM, New York, NY, 17-26.
2. Bederson, B. B., & Meyer, J. (1998). Implementing a Zooming User Interface: Experience Building Pad++. Software: Prac. and Experience, 28(10), pp. 1101-1135.

SESSION I: Communities – 9:50am
Session Chair: *Jenny Preece, Dean, College of Information Studies - Maryland's iSchool*

Self-Promotion in 140 Characters: The Use of Twitter by Congress
Jennifer Golbeck, Justin Grimes, Anthony Rogers

Analyzing Social Media Networks with NodeXL
Derek Hansen, Cody Dunne, Ben Shneiderman

ManyNets: An Interface for Multiple Network Analysis and Visualization
Manuel Freire, Awalin Nabila, Miguel Rios, Jennifer Golbeck, Catherine Plaisant, Ben Shneiderman

Break 10:40am – 20 mins

SESSION I continued: 11:00am

New Design Methods with Children: Layered Elaboration
Greg Walsh, Mona Leigh Guha, Beth Bonsignore, Beth Foss, Evan Golub, Leshell Hatley

Designing Social Musical Technologies at Carnegie Hall
Quincy Brown, Greg Walsh, Allison Druin, Ben Bederson, Jennifer Golbeck

Use of Twitter by the US Congress

Jennifer Golbeck, Justin Grimes, Anthony Rogers
College of Information Studies
Contact: {jgolbeck,jgrimes1,azgul}@umd.edu

Twitter is a microblogging and social networking service with millions of members and growing at a tremendous rate. With the buzz surrounding the service have come claims of its ability to transform the way people interact and share information, and calls for public figures to start using the service. In this study, we were interested in the type of content that legislators are posting to the service, particularly by members of the United States Congress. We read and analyzed the content of over 6,000 posts from all members of Congress using the site. Our analysis shows that Congresspeople are primarily using Twitter to disperse information, particularly links to news articles about themselves and to their blog posts, and to report on their daily activities. These tend not to provide new insights into government or the legislative process or to improve transparency; rather, they are vehicles for self-promotion. However, Twitter is also facilitating direct communication between Congresspeople and citizens, though this is a less popular activity.

Data was collected in two parts. First, we collected the 200 most recent posts (also called "tweets") for each Congressperson listed on Tweet Congress on February 6, 2009 (69 people total). Due to technical restrictions at the time of the initial data collection, 200 was the maximum number of retrievable tweets per user account. When users had fewer than 200 tweets, we collected all the tweets in their histories. It is important to note that the overwhelming majority of Congressional Twitter users did not have more than 200 tweets at this time. There were, however, a few Congressional Twitter users that did have more than 200 tweets. These users were very early adopters and power users and often had thousands of tweets. We choose not to attempt to include these tweets as this we would have led to the overall data being skewed by a few users. The tweets were analyzed to provide insights into the content and types of posts. All together, we found 4,959 tweets from Congressional users as of that date. This cross-sectional, exhaustive sample of tweets was used for the bulk of our analysis. We also conducted a follow-up analysis, repeating the same procedure for all tweets posted in June 2009 and August 2009 (1,516 tweets) to see if the patterns of usage had changed or remained the same.

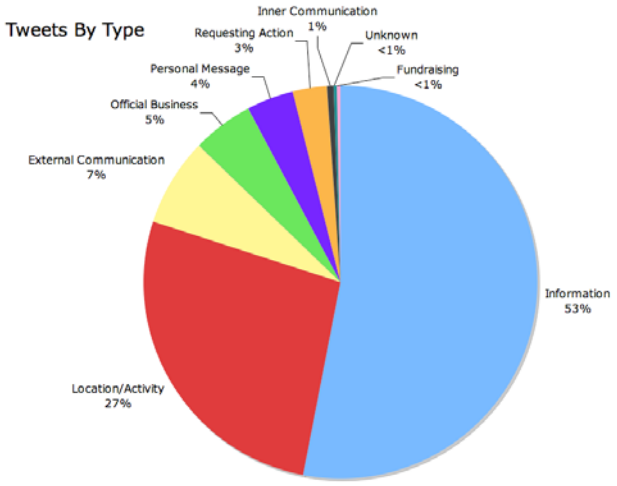

Figure 1: Types of Tweets by Congresspeople

Their tweets were coded into the following classes: Informational, Locations and Activities, Official Business (activities on the floor of Congres), Internal Communication (messages to staffers or others in Congress), External Communication (messages to people outside Congress), Personal Messages, Requests for Action, and Fundraising. Two researchers coded each tweet and we achieved an inter-rater reliability of 93.3%.

We found that informational posts are most common, accounting for just over half of all posts. The majority of these contained links that the Congressperson was sharing. Locations and Activities posts followed next, making up 27% of posts, and these detailed the daily activities of the Congressperson off the floor. While the other types of messages were less common, interesting Internal and External Communication debates and discussions can be seen across the data. A small secondary study of posts by members of Parliament in the U.K. match these results very closely, indicating both their accuracy and broader applicability.

Advocates of Congressional Twitter use argue that it can be used to increase transparency and improve communication. From this data, we draw several conclusions that directly address these points. Twitter does indeed provide a forum in which direct communication between Congresspeople and their constituents is supported. However, this is unlikely to scale, since there is a limit to how much meaningful personal communication one person can undertake. Until that threshold is crossed, Twitter is a venue that can facilitate this type of interaction.

On the other hand, Congresspeople are using Twitter primarily for outreach, not for improving transparency. While there are certainly limits to what can be communicated in 140 characters, we found the content of the tweets does little to improve insight into the activities of Congress, improve governmental transparency, or educate the readers about legislation or issues. From this perspective, the Twitter forum is not being used for new types of Congressional communication.

There is much future work to be done both on Twitter generally and in the context of government in particular. A similar analysis to what we presented here performed with a broader target group is important for establishing baselines of use habits. Within the government context, we see several spaces of future research. As mentioned above, metrics for evaluating impact will be necessary as research progresses. There are some existing services that measure how often posts are retweeted, which measures the reach of the posts. New metrics to measure impact will also be necessary. Finally, examining the connections and communication patterns between citizens and Congresspeople, particularly with respect to location may be useful. This can help Congresspeople focus their attention to their constituents as Twitter grows and to understand the perspectives being shared with them through the service.

PAPERS

1. Jennifer Golbeck, Justin Grimes, Anthony Rogers. 2010. Twitter Use by the US Congress, *to appear in JASIST.*

Notes

Analyzing Social Media Networks with NodeXL

Derek Hansen[1], Cody Dunne[2], Ben Shneiderman[2]
[1]iSchool, [2]Department of Computer Science
Contact: dlhansen@umd;{cdunne, ben}@cs.umd.edu
http://nodexl.codeplex.com/

Social media services, such as Facebook, Twitter, and Wikis have enabled new forms of collaboration and interaction in nearly every imaginable human endeavor. And we have only begun to realize the potential of technology-mediated social interaction. Despite numerous success stories, we must remember the countless failures due to social and technical factors. How can we support practitioners in their efforts to cultivate meaningful and sustainable online interaction?

One promising strategy is to provide tools and concepts that help practitioners make sense of social media data. There is precedence to this approach in the development of sophisticated, yet fairly intuitive website analytics tools such as Google Analytics. These tools help non-programmers understand website traffic patterns so they can make more informed design decisions. We envision an equivalent set of social analytics tools to help social media analysts and community administrators make better decisions based on their in-depth understanding of social participation and relationships.

Social network analysis (SNA) provides a set of concepts and techniques for making sense of social data through quantifiable metrics and network visualizations. These complement basic metrics of social participations used in current tools (e.g., number of posts; membership duration) and reveal the patterns in the network that result from social interactions. SNA concepts provide an effective vocabulary to characterize important relational properties of network members, as well as entire network structures. However, SNA also adds significant complexity and imposes obstacles for analysts that lack technical skills. Tools such as Pajek and UCINET have made SNA possible for those with sufficient drive and technical know-how, such as intelligence analysts, computer scientists, and social science doctoral students.

With the prevalence of social media network data, there is a great opportunity to make the powerful concepts and techniques of SNA accessible to a much wider audience of community analysts, participants, and designers. Doing so is hardly trivial, leading to the need to address the pressing research question: How can the complex, sophisticated set of SNA techniques be supported in an intuitive manner? To address this question, we have been working as part of a team of researchers funded by Microsoft Research to develop NodeXL, an open source add-in for the widely used spreadsheet application Excel 2007 (see http://nodexl.codeplex.com). It provides a range of basic network analysis and visualization features [1], that we have refined over time based on user studies [2]. NodeXL uses a highly structured workbook template that includes multiple worksheets to store all the information needed to represent a network graph. Network relationships (i.e., graph edges) are represented as an "edge list", which contains all pairs of entities that are connected in the network. Complementary worksheets contain information about each vertex and cluster. Data importers allow users to grab networks and user data from popular social media networks such as Twitter, YouTube, Flickr, and email. Visualization features allow users to display a range of network graph representations and map data attributes to visual properties including shape, color, size, transparency, and location (Figure 1).

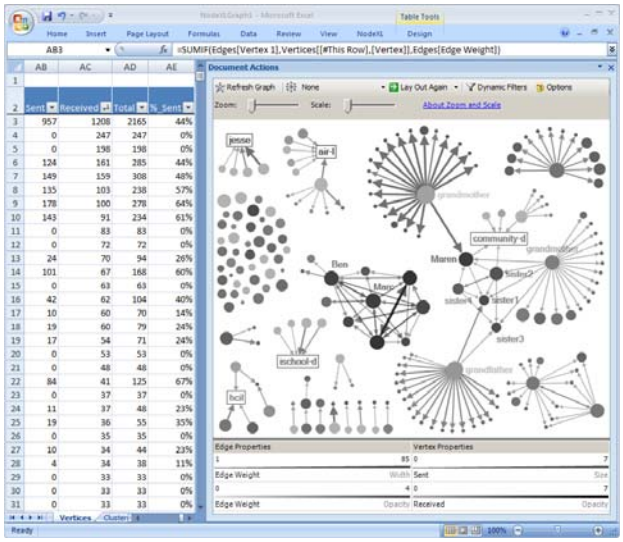

Figure 1 – Derek's personal email network shown in NodeXL, with data represented in the spreadsheet (left) and a network graph (right). Network and social metrics are mapped onto visual attributes including size, opacity, and edge thickness to highlight important people and relationships. Nodes are positioned to identify clusters such as the NodeXL team shown in the middle.

While there is a need to improve the functionality, efficiency, scalability, and usability of NodeXL, there is also a need to understand how non-technical experts can use NodeXL (and related tools) to understand community interaction. Over the past 2 years we have been studying the process through which non-technical practitioners and students can use NodeXL to make sense of online community data [3] (see Figure 2). Our experience has shown that with minimal training, non-technical students can use NodeXL to generate meaningful graphs and insights from online communities. The tight integration of visualization and data proved key to teaching, understanding, and applying network concepts for beginners [2,3]. NodeXL is now used to teach SNA in dozens of classes around the globe. The ability to automatically capture data was essential for non-programmers, as were the layout algorithms. Our studies have led to important improvements in the tool such as the inclusion of a legend and "binning" of isolate edges, as well as identify additional priorities for network analysis tools such as improved layout algorithms, support for grouping, and working with multi-modal data.

We have also been developing usage scenarios that highlight how SNA can be applied to a variety of social media networks in order to derive actionable insights. This work stands in contrast to the majority of SNA work by computer scientists and computational social scientists who characterize the mathematical properties of social media networks, but fail to speak to community administrators trying to gain practical insights. Many of these insights are captured in a forthcoming Morgan-Kaufmann book titled "Analyzing Social Media Networks with NodeXL: Insights from a Connected World" by Derek Hansen, Ben Shneiderman, and Marc Smith. The book introduces social media networks, social network analysis, and NodeXL to those unfamiliar with these concepts and then applies them to a number of social media including email, discussion forums, Twitter, YouTube, Flickr, wikis, and websites.

This work has been funded by Microsoft Research, with significant contributions from Natasa Milic-Frayling (Microsoft Research Cambridge), Marc Smith (Connected Action Consulting Group), Tony Capone (Microsoft Research), Eduarda Mendes Rodrigues (University of Porto), and Jure Leskovec (Stanford University). University of Maryland HCIL contributors include Ben Shneiderman, Derek Hansen, Cody Dunne, Dana Rotman, Elizabeth Bonsignore, Udayan Khourana, and Puneet Sharma.

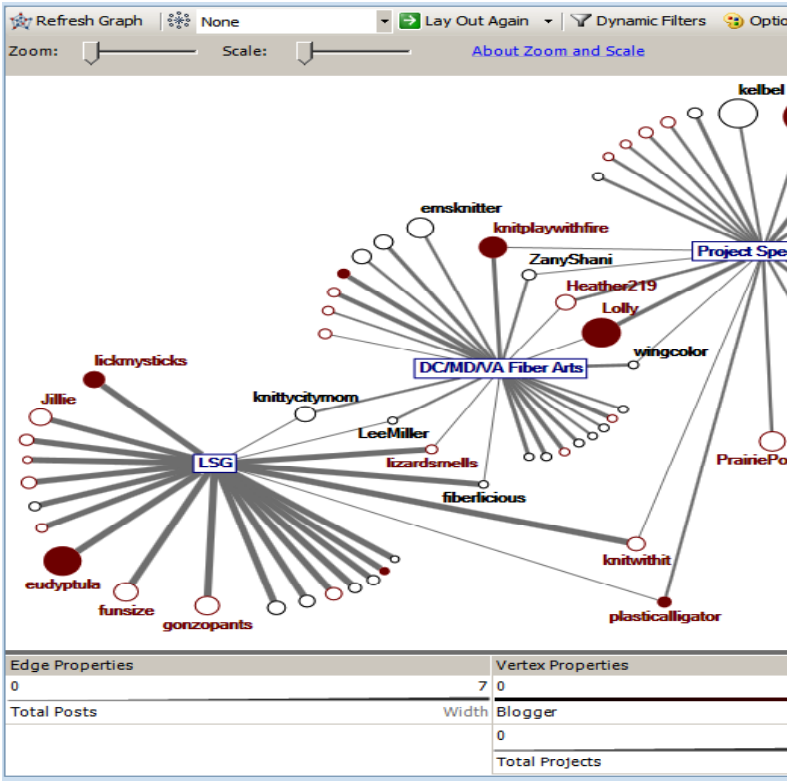

Figure 2 – A network created by Rachel Collins, a student new to SNA in Derek Hansen's Communities of Practice class. Bimodal network connecting three Ravelry groups (i.e., forums) represented as blue text boxes to contributors represented as circles. Edge width is based on number of posts (with logarithmic mapping). Vertex size is based on number of completed Ravelry projects. Maroon vertices have a blog and solid circles are either community moderators or volunteer editors. The network helps identify important boundary spanners (e.g., those connected to multiple groups) as well as compare groups.

PAPERS

1. Smith, M., Shneiderman, B., Milic-Frayling, N., Rodrigues, E. M., Barash, V., Dunne, C., et al. Analyzing Social (Media) Network Data with NodeXL. Forthcoming In *Proc. C&T 2009*.

2. Bonsignore, E.M., Dunne, C., Rotman, D., Smith, M., Capone, T., Hansen, D.L. & Shneiderman, B. (2009), "First steps to NetViz Nirvana: evaluating social network analysis with NodeXL", In SIN '09: Proc. international symposium on Social Intelligence and Networking. IEEE Computer Society Press.

3. Hansen, D., Rotman, D., Bonsignore, E., Milic-Frayling, N., Rodrigues, E., Smith, M., Shneiderman, B. (2009), Do You Know the Way to SNA?: A Process Model for Analyzing and Visualizing Social Media Data. In University of Maryland Tech Report: HCIL-2009-17.

Notes

ManyNets: An Interface for Multiple Network Analysis and Visualization

Manuel Freire, Awalin Sopan, Miguel Rios,
Catherine Plaisant, Jennifer Golbeck, Ben Shneiderman
Department of Computer Science
Contact: manuel.freire@uam.es,{awalin, mrios, plaisant, golbeck, ben}@cs.umd.edu
www.cs.umd.edu/hcil/manynets

ManyNets offers a powerful new approach that enables network analysts to work on multiple networks simultaneously. Several thousand networks can be presented as rows in a tabular visualization, and then inspected, sorted and filtereds in different ways. In a network table, each row is a network while the columns represent the featurs of those networks; it can be a network metic (e.g, node count) or a distribution of values throughout the whole network (e.g, distribution of degree of nodes). Cell visualizations and interactive column overviews allow analysts to assess the distribution of attributes within particular sets of networks. Details, such as traditional node-link diagrams, are available on demand. The node-link diagram can also show several networks merged into a single one.

ManyNets allows analysis of separate networks by loading them on a single table and then manipulating them as necessary. For example, we can load networks of Facebook students groups from different universities and compare these networks. Drawing these networks would take a long time, as some of them have more than a hundred thousand edges. The resulting pictures would also be difficult to interpret, and even more difficult to compare. But from ManyNets table we can easily compare the attributes of these netwroks.

We can also analyze parts of a single network by dividing it into smaller components according to neighborhood, motif, cluster, attribute values etc. For example, we can find 'user types' by looking at the local neighborhoods within social network by splitting it into all its ego-networks. Again, we can compare larger neighborhoods to one another by splitting the large network into its connected compoenets. MnyNets interface also provides flexible ways to split network according to a specific network feature, e.g, creation of time-sliced subnetworks to observe periodic trends.

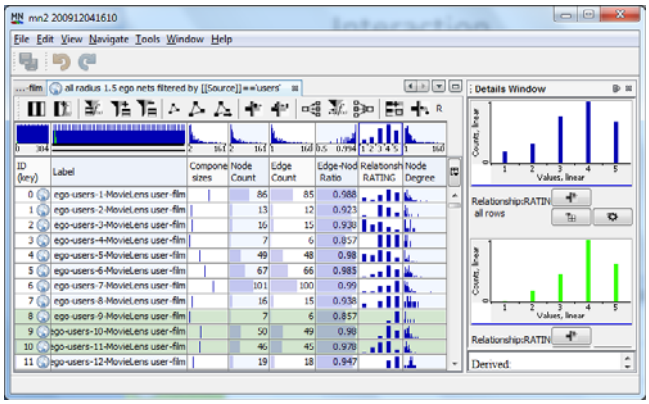

Figure 1: The MovieLens network split into user-film ego-networks, i.e. each row is a network with a user and the movies rated by him. Each user rates many films, with scores 1 to 5. The networks' size, node count, edge count etc. are shown in columns. The distribution of ratings and node degree appear as histograms. The histogram overview of the whole Rating column for all networks appears in the right side pane and the same for only the selected networks 8,9,10 is drawn just below it in green bars. The graph table of MovieLens, entity tables 'users' and 'films', relationship table 'user-film' are also shown in separate tabs.

Current network visualization tools such as NodeXL (node-link representation), and Zame (matrix representation) provide ways to visualize only a single network and show only the network attributes that contain a single value with limited user interactions. Especially the network generation task is not easy within the interface provided by them. Through ManyNets, entity-relationship schema can be used to load the tables and generate the networks. Its multi-level table manipulation system has entity and relationship tables in the lowest level. Entities are stand-alone, can be used as nodes where relationships relate two entities and are mapped to edges. The networks defined by the schema are presented as rows in the network table. Inside each network we have node and edge tables; nodes come from entities

and edges come from relationships. In ManyNets, we can mix multiple entities, relationships in a network resulting multi-modal network.

Useful visual analysis is only possible with proper user interactions. In ManyNets, columns can be removed, sorted and rearranged in the table. Moreover, we can add new columns from menu and via user-specified expression. Filtering rows from the table is also possible via selecting rows in table or even with user defined expression. The option to provide user defined expression is also available for filtering nodes and edges from networks. Details are provided on demand as tooltips and in side-pane as various types of overviews.

Tables often have many rows to fit on the screen and in such cases scrolling up and down is needed. Again, if the data within a column exhibits a certain trend, or if two columns are correlated, users might benefit from examining an overview of one or more columns. Our column overviews are interactive configurable visual overview of a whole column with provision to view in separate window. A network feature may consist of distribution of values, in this case each cell of the column contains a distribution, and we can represent the distribution in each cell by color coded heat map or by histogram. The overviews of such distribution columns can be represented by an aggregated histogram or stacked heat maps of all the cells. We can sort the heat map overview according to various distribution metrics that helps to find out similar patterns and outliers in the distributions. We can also perform cluster-based sorting of the distributions using complete-linkage agglomerative clustering, and then rearrange the resulting dendrogram using the optimal leaf ordering algorithm.

Another advanced feature of ManyNets is its ability to derive new relationships and generate new networks from them. Users can access and visualize the complete schema and define relationship resulting in new relationship table available to create new networks. Connections between two entities can be created if they share destination or they are reachable via a path within the network. As the provisions to create network from relationship and dividing existing network are integrated with this tool, analysts do not need to process the data from outside for these purposes and perform analysis using only ManyNets.

In conclusion ManyNets provides a way to look at many networks at once, even starting from a non-network dataset. It also facilitates to reveal patterns and outliers within network attributes with interactive exploratory search. It is designed in specifically for network analysts and helps them build and explore networks where they can split, filter, rank, visualize and synthesize many networks using this interface.

PAPERS

1. Sopan, A., Freire, M., Plaisant, C., Golbeck, J., Shneiderman, B. (April 2010) Distribution Column Overviews in Tabular Visualizations. *HCIL Tech Report* HCIL-2010-01. Under submission. (2010).
2. Freire, M., Plaisant, C., Shneiderman, B., Golbeck, J. ManyNets: An Interface for Multiple Network Analysis and Visualization. *Proc. of ACM CHI 2010*

Development page

Application download, datasets, manual

tangow.ii.uam.es/mn/

Academic page

Publications, demo videos

www.cs.umd.edu/projects/hcil/manynets/

Acknowledgements

- Partial support from Lockheed Martin
- Manuel Freire is supported by Fulbright Scholarship

Notes

Layered Elaboration: A New Technique for Co-Design with Children

Greg Walsh, Allison Druin, Mona Leigh Guha, Elizabeth Foss, Evan Golub, Leshell Hatley, Elizabeth Bonsignore, Sonia Franckel

Human-Computer Interaction Lab, University of Maryland
College Park, MD 20742 USA
gwalsh@umd.edu, allisond@umiacs.umd.edu

As technology for children becomes more mobile, social, and distributed, our design methods and techniques must evolve to better explore these new directions. This paper reports on "Layered Elaboration," a co-design technique created to support these evolving needs. Layered Elaboration allows design teams to generate ideas through an iterative process in which each version leaves prior ideas intact while extending concepts. Layered Elaboration is a useful technique as it enables co-design to take place asynchronously and does not require much space or many resources.

THE FOUNDATIONS OF LAYERED ELABORATION

The concept for Layered Elaboration has its roots in storyboarding for interactive media [5], paper prototyping [8], and annotation tools [6]. Storyboards and drawings have been used as a method in participatory design research for decades [2, 3, 7]. Paper prototyping has been used for game design [4] and interface testing [1]. The technique described in this paper builds upon these and is named *Layered Elaboration* because it enables design groups to add ideas presented by others with layers of transparencies while encouraging design team members to expand on those earlier ideas.

Layered Elaboration is different from other techniques because few lend themselves to revisiting design ideas and adding upon them in a straightforward manner. Many design techniques require the permanent change of original items as additional designers modify them. Layered Elaboration enables co-designers to add and modify ideas without permanently destroying the original through the use of transparent materials.

LAYERED ELABORATION IN PRACTICE

Layered Elaboration was originally conceived during the design of a motion-controlled, history-based, instructional video game [9]. An adult team member created storyboards representative of one module in the game. We wanted a way to add to and extend the ideas without having to recreate or destroy the artifacts from this initial phase of design. Adding a layer of overhead transparencies to the storyboard fulfilled this goal.

Figure 1. Using Layered Elaboration to design a game.

Our team made modifications to the technique based on those initial experiences. The modified technique emphasized collaboration and elaboration between groups. Our next use of Layered Elaboration occurred when our intergenerational design team was tasked with developing a technology that would help other children be more environmental-conscious in three different locations. Our team split into three smaller groups, each consisting of two children and at least one adult. Each team was assigned one of the three locations.

Increasing Participant Ownership

To begin the session, each team was given drawing paper and markers to create a design. The groups were assigned one of those topics and given fifteen minutes to create their designs.

After the time was up, the groups got together in the middle of the room for a "stand-up meeting" to rapidly move along the process and as an interim debriefing. At the meeting, each group explained their design and answered any questions that the other groups might have had. Once a group presented, a transparent overlay was added.

Inter-Group Collaboration

In order to solve some of the challenges of the initial implementation of Layered Elaboration, the groups exchanged designs to foster inter-group collaboration and communication. With the new designs in hand, the groups began elaborating on the ideas presented. By using the transparent overlay and markers, the design partners were able to draw on the previous group's work without permanently destroying it.

After another ten-minute period, the groups got back together in the middle of the room for another stand-up meeting. Each group presented what they added or changed to the previous design. After each group presented, another transparent overlay was added to the design and prepped with registration points. The groups then had one more opportunity to add to the design.

After the final design period, all of the groups sat down and discussed the final designs.

LESSONS LEARNED AND FUTURE WORK

As with any co-design technique, Layered Elaboration has strengths and challenges. Its strengths include:

- The ability to add to and modify the initial storyboard without permanently damaging or altering it.
- The ability of the design team to stack the transparency overlays over the original storyboard to see common trends or “hot spots” in the different groups' feedback.
- The portability as a co-design tool. Instead of needing a large physical space for low-tech prototyping, the stackable storyboards are no larger than a clipboard.
- The relatively rapid, iterative nature of the technique allows a number of design partners to provide input and ideas in a short amount of time.
- The cost of the materials is low.

We have also identified a few challenges with the modified technique, which we intend to address in the future. They include:

- Not all team members paid attention to the other groups as they presented in the stand-up meetings, which led to confusion and less elaboration.
- The washable markers used by the design teams were inadequate for writing on transparent overlays because they smudged; permanent markers were required.

This technique is useful when non-destructive design annotation, limited space, and evolutionary artifacts are design requirements. We would like to use this technique with a geographically distributed audience, so, we developed a web-based, co-design tool called DisCo.

Using the lessons learned from previous experiences, and trying to address the needs of a geographical distributed co-design audience, we designed and implemented DisCo to facilitate Layered Elaboration and support creative expression. DisCo is built with Flash, PHP, and JavaScript and enables users to collaborate asynchronously on designs.

Our initial work with DisCo has provided us with a foundation to develop new, distributed co-design tools.

ACKNOWLEDGEMENTS

We would like to thank the U.S. National Park Service, the University of Maryland’s College of Information Studies, and, as always, our child design partners.

REFERENCES

[1] Bailey, B. P., Biehl, J. T., et al. (2008). Adapting paper prototyping for designing user interfaces for multiple display environments. Personal Ubiquitous Comput., 12(3), 269–277.

[2] Brederode, B., Markopoulos, P., et al. (2005).pOwerball: the design of a novel mixed-reality game for children with mixed abilities. Proc. IDC '05: Proceedings of the 2005 conference on Interaction design and children, 2005. 32–39.

[3] Moraveji, N., Li, J., Ding, J., O'Kelley, P., and Woolf, S. (2007). Comicboarding: using comics as proxies for participatory design with children. Proceedings of the SIGCHI conference on Human factors in computing systems; 2007.

[4] Ollila, E. M. I., Suomela, R., et al. (2008). Using prototypes in early pervasive game development. Comput. Entertain., 6(2), 1–17.

[5] Orr, K. L., Golas, K. C., et al. (1994). Storyboard development for interactive multimedia training. Journal of Interactive Instruction Development, 6(3), 18–31.

[6] Reside, D. (2007). The AXE Tool Suite: Tagging Across Time and Space. Proceedings of Digital Humanities, 178–9.

[7] Scaife, M., Rogers, Y., Aldrich, F., and Davies, M. (1997). Designing for or designing with? Informant design for interactive learning environments. Proceedings of the SIGCHI conference on Human factors in computing systems; 1997.

[8] Snyder, C. (2003). Paper prototyping. Burlington, Massachusetts: Morgan Kaufmann, .

[9] Walsh, G. (2009). Wii can do it: using co-design for creating an instructional game. Proceedings of the 27th international conference extended abstracts on Human factors in computing systems; 2009.

Notes

Facilitating Cross Cultural Awareness in High School Students Using Co-Design Methodologies

Quincy Brown[1], Greg Walsh[1], Allison Druin[1], Christopher Amos[2], Sarah Johnson[2], Misty Tolle[2]
College of Information Studies[1], Weill Music Institute at Carnegie Hall[2]
Contact: {qbrown, gwalsh}@umd.edu, allisond@umiacs.umd.edu

The belief that children can be effective and valued design partners versus mere testers or end-users is the foundation of the Cooperative Inquiry design process [2]. Cooperative Inquiry adapts methods from cooperative design, participatory design, and contextual inquiry in a manner that enables children and adults to work as partners, in multi-generational teams, in the design of new technologies [4]. The methods and techniques employed during the design process enable designers, regardless of age, to brainstorm and articulate ideas that might not be revealed or expressed with more traditional usability testing approaches.

The Cooperative Inquiry process and methods have been used to effectively design software including mobile device applications and drawing tools, as well as hardware including robots and interactive environments [1-3, 5, 6]. In this research we extend the use of participatory methods and techniques to the design of curricula and concert performances. In collaboration with the Weill Music Institute at Carnegie Hall we explore the ways in which cooperative inquiry methods, traditionally used for technology design, can impact curricular design and concert performances.

The Weill Music Institute at Carnegie Hall offers a variety of education and community programs throughout the New York City metropolitan area, including school programs for grades K-12. One program, Carnegie Hall Cultural Exchange, offers students and teachers the opportunity to connect their classrooms to the world. The 2009-2010 season of the Cultural Exchange program includes global awareness forays into the Music of India, New York Jazz, and Music of Mexico. The program involves approximately 2000 students from 25 high schools in New York City, New Delhi, and Mexico City.

In multigenerational design teams, we partner with high school students in New York City to elicit their ideas for music visualizations, concert performance enhancements, and social networking site design. We employ lo-tech prototyping techniques such as bags of stuff and big paper to enable the students to brainstorm and articulate their design ideas in visible, tangible and personally expressive ways. We also facilitate the critique and evaluation of existing social networks as well as concert performances. The outcomes of this research include a music visualization lighting segment during a live concert performance, shown in figure 1, and social network site design modifications to focus the social interactions on students' common music interests. This work will be extended in the 2010-2011 Cultural Exchange program to include additional methods to support asynchronous design with students from different countries.

Figure 1 – Music visualization segment during live performance at Carnegie Hall

ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant #0937060 to the Computing Research Association for the CIFellows Project. The Sony Corporation is the proud leading sponsor of Carnegie Halls' Online Exchange Communities. Carnegie Hall Cultural Exchange is sponsored by MetLife Foundation. Carnegie Hall Cultural Exchange is supported, in part, by the Bureau of Educational and Cultural Affairs of the US Department of State. Delivery of the Weill Music Institute's programs to a nationwide audience is supported, in part, by the US Department of Education and by an endowment grant from the Citi Foundation. Cultural Exchange concert experiences are supported, in part, by The Tishman Fund for Education through Technology and by an endowment grant from the Citi Foundation.

REFERENCES

1. Bederson, B. B., Quinn, A. and Druin, A. Designing the reading experience for scanned multi-lingual picture books on mobile phones. In *Proceedings of the Proceedings of the 9th ACM/IEEE-CS joint conference on Digital libraries* (Austin, TX, USA, 2009). ACM.
2. Druin, A. Developing New Technologies for Children with Children. In *Proceedings of the CHI 1999* (Pittsburgh, Pa, USA, May 15-20, 1999). ACM Press.
3. Druin, A., Montemayor, J., Hendler, J., McAlister, B., Boltman, A., Fiterman, E., Plaisant, A., Kruskal, A., Olsen, H., Revett, I., Schwenn, T. P., Sumida, L. and Wagner, R. Designing PETS: a personal electronic teller of stories. In *Proceedings of the Proceedings of the SIGCHI conference on Human factors in computing systems: the CHI is the limit* (Pittsburgh, Pennsylvania, United States, 1999). ACM.
4. Guha, M. L., Druin, A., Chipman, G., Fails, J. A., Simms, S. and Farber, A. Mixing Ideas: A New Technique for Working with Young Children as Design Partners. In *Proceedings of the Interaction Design and Children* (College Park, Maryland USA, 2004). ACM.
5. Hourcade, J. P., Bederson, B. B., Druin, A. and Taxen, G. KidPad: collaborative storytelling for children. In *Proceedings of the CHI '02 extended abstracts on Human factors in computing systems* (Minneapolis, Minnesota, USA, 2002). ACM.
6. Sherman, L., Druin, A., Montemayor, J., Farber, A., Platner, M., Simms, S., Porteous, J., Alborzi, H., Best, J., Hammer, J., Kruskal, A., Matthews, J., Rhodes, E., Cosans, C. and Lal, A. StoryKit: tools for children to build room-sized interactive experiences. In *Proceedings of the CHI '01 extended abstracts on Human factors in computing systems* (Seattle, Washington, 2001). ACM.

Notes

SESSION II: Text and Translation – 11:35am
Session chair: *Catherine Plaisant, Associate Director of Research, HCIL*

Human-Computer Collaborative Translation
Chang Hu, Ben Bederson, Philip Resnick

Finding Entries in an On-line Arabic Dictionary
Sarah Wayland, C. Anton Rytting, David Zajic, Timothy Buckwalter, Jason White, Corey Miller, Jeffrey Carnes, Nathanael Lynn, Paul Rodrigues, Michael Maxwell, Evelyn Browne

12:10pm – Lunch

SESSION II continued: 1:10pm

iOpener Workbench: Tools for Rapid Understanding of Scientific Literature
Cody Dunne, Ben Shneiderman, Bonnie Dorr, Judith Klavans

CrowdFlow: A Human-Computer Hybrid Cloud Computing Model
Alex Quinn, Jimmy Lin, Ben Bederson, Tom Yeh

Monotrans: Human-Computer Collaborative Translation

Chang Hu[1,2], Benjamin B. Bederson[1,2] and Philip Resnik[1,3,4]
[1]Department of Computer Science, [2]Human-Computer Interaction Lab,
[3]Institute for Advanced Computer Studies, [4]Department of Linguistics,
University of Maryland
Contact: {changhu, bederson}@cs.umd.edu, resnik@umiacs.umd.edu
www.cs.umd.edu/hcil/monotrans

An enormous potential exists for solving certain classes of computational problems through rich collaboration between humans and computers. Take translation for example. Humans alone are expensive and can be surprisingly slow; despite significant recent advances, machine translation (MT) remains a crucial problem and fully automated high quality translation remains a distant dream for the vast majority of the world's language pairs. Usable translation quality can sometimes be obtained by statistical MT systems, but only for a minority of language pairs, and only in use cases where sufficient training text is available and the material being translated is reasonably similar to the material on which the system was trained.

Using the Web to reach non-professional human translators holds promise, and there has been some initial success with distributing translation over a crowd of bilingual users. However, compared to the total user population, the potential translator population is still small. For example, while Wikipedia currently has about 75,000 active contributors, there are fewer than 800 translators. With a much larger number of potential human helpers who speak only the source or target language, but not both, it seems natural to ask whether some combination of machine translation with volunteer *monolingual* speakers could result in high quality translation.

We propose a rethinking of the translation problem to bring together translation technology and human-computer interaction, producing a framework for translation that will exploit imperfect technology and limited human abilities in tandem to achieve capabilities neither can achieve alone.

The core of this framework is Monotrans, an iterative protocol in which the human participants work together to make sense of machine translations, and introduce redundant information to make their intended meanings clearer (Figure 1). This protocol makes it possible to detect and correct some translation errors, and to at least identify some passages that have errors even if they are not correctable given the available information. For example, "has cheeseburger" is a detectable error, even if it is not clear whether the intended meaning was "has cheeseburgers" or "have a cheeseburger". Back-translating a refinement and carrying along redundant information, e.g. a *picture* of multiple cheeseburgers, might help convey which of those alternatives the English speaker guessed, presenting the opportunity for confirmation or further correction.

Research Prototype

We built a research prototype as a multi-user web application (Figure 2). When a user logs in, the UI displays a book page in the user's language. Every sentence in the

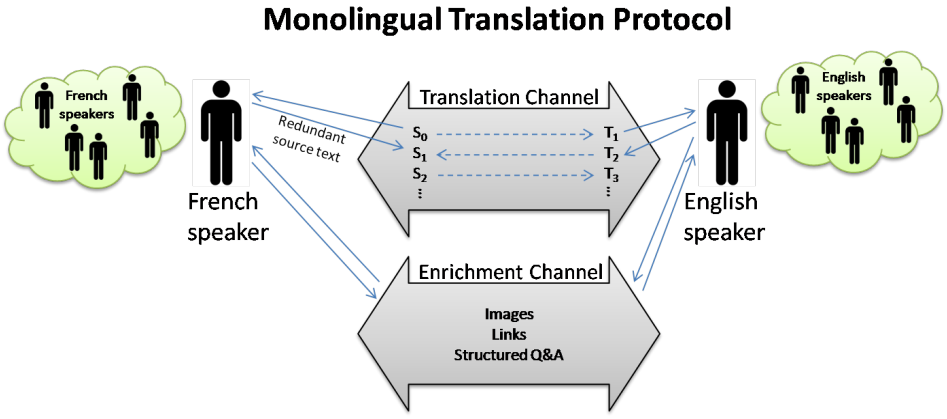

Figure 1 – Round-trip protocol. Dashed arrows show machine translation and solid arrows show human editing.

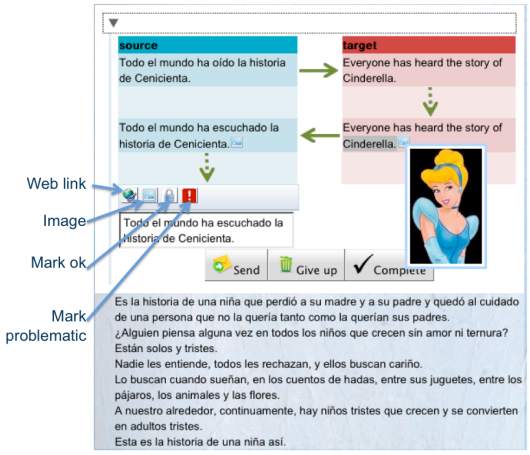

Figure 2 – In this prototype UI showing a page, a sentence is expanded, with an image attached to a word.

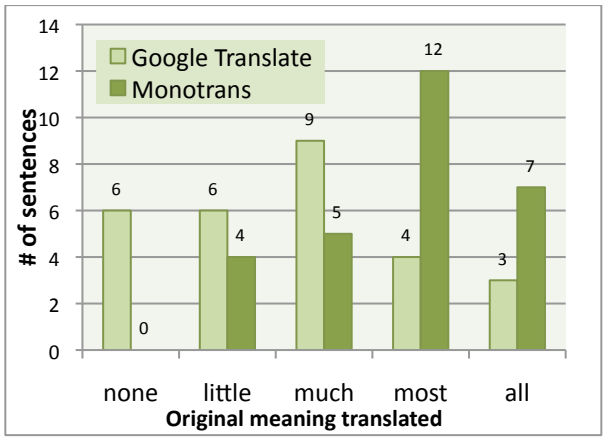

Figure 3 - Number of sentences in each adequacy category, rated by a professional translator.

page is displayed with the most up-to-date translation hypothesis (or corresponding back-translation). The user can navigate through all available pages with navigation controls, or expand a sentence translation to edit it.

When a sentence is expanded, the UI shows the most up-to-date translation hypothesis with all previous translation hypotheses of this sentence and a rich editor where the sentence can be edited and annotated.

The rich editor currently includes the following elements for the enrichment channel, aimed at enhancing redundancy and communicating shared context.

- Image annotations
- Web link annotations
- Annotation of correct and incorrect parts of a sentence

Word level alignments necessary to perform annotation projection can be obtained from our own machine translation engines. Also, some machine translation services make word alignment information publically available to researchers along with the translation hypotheses. The Google Translate Research API (open to university research projects) is one such example.

Preliminary Results

We used MonoTrans to translate part of a children's book from Russian to Chinese. Chinese and Russian are commonly spoken languages in the world. However, they make good experimental candidates because they are very different from the perspective of linguistic typology.

In the experiments, two Russian speakers and four Chinese speakers formed four pairs to use the prototype. (One Russian speaker participated three times, with different content.) The participants were all native speakers of one language and had no knowledge of the other. They were all computer-literate and fluent speakers of English. While most of the participants were computer science students and researchers, none of them work in the area of machine translation directly, and none of them were familiar with the details of this project. They were not linguists or linguistic students.

Participants worked on 6 pages (a total of 44 sentences) and finished translating 28 of them. This works out to approximately seven sentences per hour between any given pair of participants. It is about five times faster than the earlier "Wizard of Oz" experiment. With a standard rating procedure, sixteen of the 28 sentences translated with the prototype were rated as fully fluent and nineteen sentences of the 28 were rated as mostly or fully translated, by a professional translator not connected with the project.

The shift in adequacy is especially notable among these results. Completely inadequate MT outputs (none of the meaning preserved) dropped from 6 to 0. This means that the protocol helped the target language participants understand at least some of the meaning even when the original MT output quality was especially low and they had little to go on. In a coarse-grained way of thinking, if the adequacy rating could be categorized so that {none, little}=*bad* and {most, all}=*good*, then there would be a drop in *bad* (meaning) from 12 to 4 out of the 28, and there would be an increase in *good* from 7 to 19 of 28. That represents a factor of roughly 3 in each of the desired directions.

PAPERS

1. Chang Hu: Collaborative Translation by Monolingual Users. *CHI Extended Abstracts 2009*: 3105-3108
2. Chang Hu, Benjamin B. Bederson and Philip Resnik: Translation by Iterative Collaboration between Monolingual Users. *Graphics Interface 2010, Ottawa, Canada*

Notes

Finding Entries in an On-line Arabic Dictionary

Sarah C. Wayland, C. Anton Rytting, David Zajic, Timothy Buckwalter, Jason White, Corey Miller, Jeffrey Carnes, Nathanael Lynn, Paul Rodrigues, Michael Maxwell, Evelyn Browne
Center for Advanced Study of Language
Contact: {swayland, crytting, dzajic, tbuckwalter, jwhite, cmiller, jcarnes, nlynn, prodrigues, mmaxwell, ebrowne}@casl.umd.edu
http://www.casl.umd.edu/node/525

In this paper, we describe a new Arabic spelling correction tool called Arabic Did You Mean….? designed specifically to help non-native learners look up words in Arabic electronic dictionaries.

Mistakes that hinder successful dictionary lookup can stem from a few underlying causes:

1. Difficulty discerning phonemic contrasts (such as /q/ and /k/),
2. Difficulty discriminating visually similar letters (such as غـ and فـ),
3. Incorrectly reconstructing a citation form of an inflected word (such as trying to find the singular of مباريات by looking up the non-word مبارية).
4. Words read on Arabic-language opinion forums and blogs, or heard in audio or video webcasts are usually in dialect, and as such have different orthography, morphology, and lexical content from the Modern Standard Arabic commonly taught in foreign language classrooms.
5. If the language learner doesn't have access to Arabic keyboards, they sometimes make errors based on an unfamiliar or unintuitive Romanization scheme.
6. Users sometimes make simple typographical errors, such as hitting a key adjacent to the one intended.

Because Arabic language learners can generate a wide range of alternate inputs, a tool that suggests spelling corrections has the potential to make it easier for the learner to find unfamiliar words in existing lexicons.

We have created a spelling corrector for Arabic dictionary lookup that accepts input in a Romanization system known as the Standard Arabic Technical Transliteration System, or SATTS, verifies whether or not the query matches a citation form[1] in a bilingual Iraqi Arabic-to-English dictionary (Woodhead & Beene, 2003), and suggests similar citation forms the user may have intended.

Following Beesley (Beesley, 1998; Beesley & Karttunen, 2003) the underlying spelling correction algorithm relies on a weighted finite state transducer (FST) that assigns weights based on confusion matrices representing the various error types, including sound confusions, visual (Arabic letter) confusions, and keyboard proximity errors. These confusion matrices can be modified in order to adapt the spell-checker to the error types associated with different Arabic dialects, keyboard layouts, or even the listening errors of students who speak a native language other than English.

As we know of no existing spelling error corpus for Arabic, we simulated error data by generating a corpus using a Noisy-Channel model for error production. For the noise model, we learned the kinds of errors a nonnative speaker could make from a corpus of transcribed Arabic speech elicited from learners of Arabic during an imitation task (Sethy, et al., 2005). The elicitations were common MSA greetings and simple conversations uttered by native speakers of Levantine Arabic and Iraqi Arabic. We evaluated our system by comparing it with a baseline based on a standard spell-checking method, known as Levenshtein (1966) distance, which assumes no language-specific knowledge. Our system got a significantly higher Mean Reciprocal Rank (MRR) score on a test corpus of Arabic query strings and intended words than the baseline Levenshtein version ($\Delta t = 10.95$, $p=0.0001$). Both sound-based confusions and variant spellings contributed to the improved MRR score (Rytting et al., in press.)

Once we determined that our approach was reasonable, we needed to create a user interface that would display the results in a way that was meaningful to English speakers looking up words in Arabic (see Figure 1, below.) After interviewing a number of Arabic language learners at a variety of skill levels, we determined that users should be able to enter their queries using either native Modern Standard Arabic (MSA) or SATTS.

[1] A *citation form* is the form of a word that heads a lexical entry and is alphabetized in a dictionary.

These potential users told us that the output of our spelling correction algorithm should list not only the **citation form** and the type of **inflection** of the word (e.g., plural, citation, etc.), but also the **root form**[2] of the word, as well as its **part-of-speech** (POS). In addition, they requested that the system list the **definitions** associated with the root form.

The citation forms returned by the spelling correction algorithm have a rating associated with them that reflects the weights returned by the FST. We translate these weights into an integer rating system that ranges from one to five, with one being the least likely match, and five being the most likely match. These **ratings are reflected as the number of stars** in the user interface. Perfect matches are highlighted in yellow. Because rankings are based on the citation forms and not the root forms, a single root form may have more than one citation form associated with it; each of these citation forms may have a different ranking. Thus, entries are grouped by root, with the order of the roots determined by the rating of its highest-ranking citation form.

Other columns show the inflected form of the root as returned by the FST in a variety of output formats, including Modern Standard Arabic (**MSA**), **SATTS**, **Buckwalter** (another Romanized form of Arabic), **GU Phonetics** (the sound-based form listed in the Iraqi Arabic dictionary published by Georgetown University), and **SAMPA** (a computer-readable phonetic script based on 7-bit printable ASCII characters known as the Speech Assessment Methods Phonetic Alphabet). A dialog box allows the user to **show or hide any of the columns**.

To limit the output, users can ask the system to show only **verbs** or **non-verbs** (a distinction that is more meaningful to Arabic speakers than to non-Arabic speakers). In addition, although the default output is ordered by rank, we allow users to **sort** by whatever column makes sense to them.

Lastly, our users wanted to be able to download the results into a text file for further processing in other program. To facilitate that, we allow them to **download the results** into a comma separated value (CSV) text file.

We are in the process of finalizing our design for this tool; the next step is to collect feedback and usage information comparing Arabic dictionary lookup when users have access to the "Did You Mean...?" tool, and when they do not.

We have plans to extend our system to allow lookup of non-citation word forms, such as inflected and morphologically complex forms, as well as named entities.

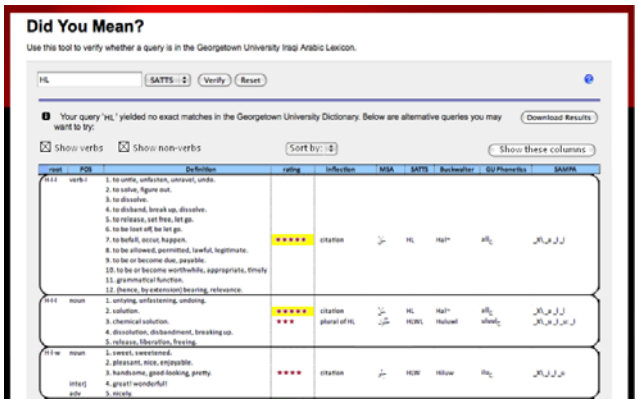

REFERENCES

1. Beesley, K. R. (1998). Arabic morphology using only finite-state operations. In *Proceedings of the Workshop on Computational Approaches to Semitic Languages*, 50-57, Montreal, Quebec.
2. Beesley, K., & Karttunen, L. (2003). *Finite state morphology*. Stanford, CA: CSLI Publications.
3. Center for Advanced Study of Language, University of Maryland. (2008). *Iraqi Arabic corpus*. [Data file].
4. Erwin, W. M. (2004). *A short reference grammar of Iraqi Arabic*. Washington, DC: Georgetown University Press.
5. Levenshtein, V. I. (1966). "Binary codes capable of correcting deletions, insertions, and reversals." *Soviet Physics Doklady*, 10, 707–710.
6. Rytting, C.A., Buckwalter, T., Wayland, S.C., Hettick, C., Rodrigues, P., Zajic, D., (in press). Spelling Correction for Dialectal Arabic Dictionary Lookup. ACM Transactions on Asian Language Information Processing.
7. Sethy, A. Mote, N., Narayanan, S. and Johnson, L. (2005). "Modeling and automating detection of errors in Arabic language learner speech," in *Proc. of Eurospeech, Interspeech*, Lisbon.
8. Standard Arabic Technical Transliteration System. (2010). Retrieved February 03, 2010, from *Wikipedia.org*: http://en.wikipedia.org/wiki/SATTS
9. Woodhead, D. R., & Beene, W. (2003). *A Dictionary of Iraqi Arabic*. Washington, DC: Georgetown University Press.

[2] The *root form* of an Arabic word is the base template form (three to five consonants) that, by the application of strict morphonemic rules, decline into noun and verb forms reflecting gender, plurality, voice, and other aspects of the word's meaning.

Notes

iOpener Workbench: Tools for Rapid Understanding of Scientific Literature

Cody Dunne, Ben Shneiderman, Bonnie Dorr & Judith Klavans
Department of Computer Science
Contact: {cdunne, ben, bonnie}@cs.umd.edu, jklavans@umd.edu
tangra.si.umich.edu/clair/iopener

Scholars and scientists today have the challenging task before them of keeping up with the advances within their rapidly expanding fields. The growing number of publications each year combined with ever-increasing cross-disciplinary research make following all the branches of research in a field and identifying key papers challenging. Rapid generation of accurate survey articles can help ease these difficulties, providing readers with concise overviews tailored to their needs. Students, educators, scientists, and government decision makers all have vastly different levels of expertise and requirements for learning about scientific fields.

The goal of iOpener (Information Organization for PENning Expositions on Research) is to generate readily-consumable surveys of different scientific domains and topics, targeted to different audiences and levels. We've created an infrastructure for automatic summarization of research domains that links bibliometric lexical link mining, summarization techniques, and visualization tools. Part of this is the *iOpener Workbench*, a new tool which presents the academic literature for a field using many different modalities: lists of articles, their full texts, automatic textual summaries, and visualizations of the structure of the citation network – all shown in Figure 1.

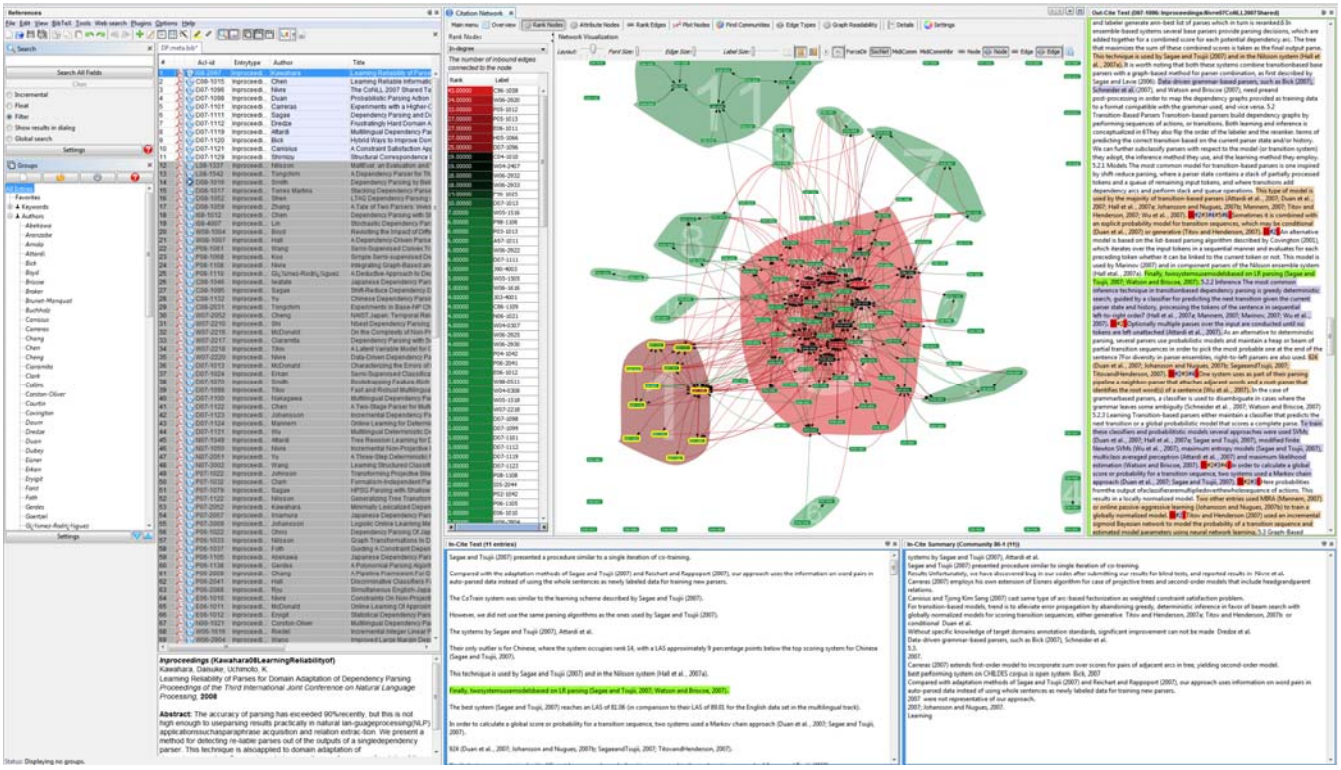

Figure 1 The *iOpener Workbench* combines (clockwise from top) a citation network analysis tool (SocialAction), full-text with linked citations, multi-document summaries of clusters, the text of incoming citations, and reference manager (JabRef),

The iOpener Workbench is partially an integration of two powerful existing tools – the SocialAction network analysis tool[1] and the JabRef reference manager[2]. SocialAction provides us with powerful network analysis capabilities including force-directed citation network visualization, ranking and filtering papers by statistical measures, scatterplots of paper attributes and statistics, categorical and numerical range coloring, and automatic cluster detection (shown using convex colored hulls in Figure 1). Using visualizations of the citation network we can easily find unexpected trends, clusters, gaps and outliers. Additionally, visualizations can immediately identify invalid data that is easily missed in tabular views.

JabRef supplies all the features one would expect from a reference manager, including searching using simple regular expressions, automatic and manual grouping of papers, DOI and URL links, PDF full text with annotations, abstracts, user generated reviews and text annotations, and many ways of exporting. It integrates with Microsoft Word, OpenOffice.org, and BibTeX, which allows quick adding of citations to discovered articles when writing survey papers.

These tools are linked together to form multiple coordinated views of the data. Clicking on a node in the citation network selects it and its corresponding paper in the reference manager, displaying its abstract, review, and other data associated with it. Moreover, when clusters of nodes are selected their papers are floated to the top of the reference manager for easy perusal. The inverse is true as well, with any paper, group, or search term selected in the reference manager highlighting the corresponding nodes in the network.

There are other coordinated views that provide the user with other aspects of the field. When any node or cluster is selected, the In-Cite Text window displays the text of all incoming citations to the paper(s), i.e. the whole sentences from the citing papers that include the citation to the selected paper(s). These are displayed in a hyperlinked list that allows the user to select any one of them to show their surrounding context in the Out-Cite Text window. This window shows the full text of the paper citing one of the selected papers, with highlighting showing the selected citation sentence as well as any other sentences that include hyperlinked citations to other papers. The last view is the summary window, which can contain various multi-document summaries of a selected cluster. Using automatic summarization techniques, we can summarize all of the incoming citations to papers within that cluster, hopefully providing key insights into that research community.

The iOpener Workbench integrates these many components in order to provide a tool that supports rapid understanding of scientific literature. Users can analyze the network of citations between papers, identify key papers and research clusters, automatically summarize them, dig into the full text of articles to extract context, make annotations, write reviews, and finally export their findings in many of document authoring formats. We hope this infrastructure will enable users to generate readily-consumable surveys of scientific fields.

Acknowledgements

This work has been partially supported by the National Science Foundation grant "iOPENER: A Flexible Framework to Support Rapid Learning in Unfamiliar Research Domains", jointly awarded to UMd and UMich as IIS 0705832.

PAPERS[3]

1.Aris, Aleks, Ben Shneiderman, Vahed Qazvinian, Dragomir Radev. Visual Overviews for Discovering Key Papers and Influences Across Research Fronts. JASIST 60(11): 2219-2228 (2009).

2.Mohammad, Saif, Bonnie Dorr, Dragomir Radev, Melissa Egan, Ahmed Hassan, Pradeep Muthukrishnan, Vahed Qazvinian, David Zajic. Using Citations to Generate Surveys of Scientific Paradigms. In NAACL-HLT 2009.

3.Qazvinian, Vahed and Dragomir R. Radev. Scientific paper summarization using citation summary networks. In COLING 2008, Manchester, UK, 2008.

4.Radev, Dragomir R. and Joseph, Mark and Gibson, Bryan and Muthukrishnan, Pradeep. A Bibliometric and Network Analysis of the Field of Computational Linguistics. Journal of the American Society for Information Science and Technology (JASIST).

5.Shneiderman, B., Research Agenda: Visual Overviews for Exploratory Search, Position paper for National Science Foundation Workshop on Information Seeking Support Systems, Chapel Hill, NC, June 26-27, 2008.

[1] www.cs.umd.edu/hcil/socialaction

[2] jabref.sourceforge.net

[3] This work is not yet published. Related papers are listed below.

Notes

CrowdFlow: Integrating Machine Learning with Mechanical Turk for Speed-Cost-Quality Flexibility

Alexander J. Quinn[1], Benjamin B. Bederson[1,2,3], Tom Yeh[3], Jimmy Lin[1,2]
University of Maryland
Human-Computer Interaction Lab
Department of Computer Science[1] || iSchool[2] || Institute for Advanced Computer Studies[3]
College Park, MD 20742 USA
Contact: aq@cs.umd.edu

Many problems involving visual perception, language understanding, or other human abilities can now be solved by computers. This allows the tasks to be done faster and more affordably than humans could do. Furthermore, it means they can be done on-demand, enabling computers to monitor national intelligence information streams, search video footage for missing persons, and perform a wide variety of other services of importance to society. The primary disadvantage is that machines are still not as accurate as humans on many tasks.

The field of human computation - alternately referred to as crowdsourcing [1, 2], distributed human computation [7], or collective intelligence [5]- is concerned with methods and systems for distributing small, independent tasks to anonymous workers connected by large online networks. This enables faster turnaround time and reduces the overhead of hiring humans to do the work in an office. Furthermore, it is typically much less expensive than hiring employees. Human computation employs a variety of paradigms, including games with a purpose (GWAPs) [9, 10] and online task markets, such as Amazon Mechanical Turk (AMT) [4].

Although such systems provide more flexibility than in-person workers, they are still more expensive than running the same jobs with a computer. Furthermore, they human effort in cases where the computer's accuracy is modestly good. Perhaps most importantly, they lack the ability to tune the speed, cost, and quality to the situation's needs.

CrowdFlow is our general toolkit (currently implemented as a Python library) built to solve this problem. The framework blends the flexibility of AMT with the speed and affordability of machine learning systems.

An especially powerful element of this framework is that the humans and machines benefit from one another in a kind of *man-computer symbiosis* [3]. Machines benefit from the human-created training data, which helps boost their accuracy. Human workers benefit from having the machine results as a starting point. When the machine results are correct, the worker need only verify that fact.

The user of CrowdFlow specifies a desired speed-cost-quality tradeoff. The system then allocates tasks to humans and machines in a way that attempts to fulfill the user's specification. By estimating system performance, we can describe a tradeoff space gamut within which it is possible for the human manager to manipulate the system.

Speed may be expressed as a time limit for completing the job. Similarly, cost is the maximum the user is willing to pay to Turkers and/or for any cloud-based computational resources. (The latter is planned but not currently implemented in our toolkit.) Quality is measured relative to some satisfaction criteria the user provides. It could be a heuristic that combines multiple criteria.

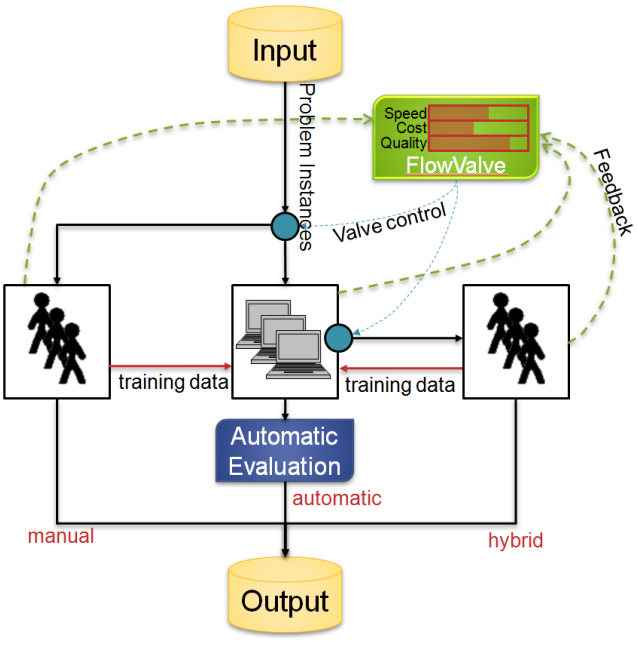

Figure 2. CrowdFlow conceptual model.

Making CrowdFlow work will entail several challenges. First, we must be able to estimate, within some confidence interval, the accuracy of the human workers, even if there is no ground truth. This is doubly important when you consider that without the humans' judgment, it is impossible to estimate the machine's accuracy. Also, the system must keep updating the accuracy estimates and automatically adjust the allocation of tasks as needed. Second, we need to specify a generic architecture that is both easily usable by developers of specific problems while still providing significant value. Third, we need to support a wide range of human capabilities in this context.

We aspire to keep CrowdFlow as general as possible. It should be applicable to problems with these properties:

- Solvable by humans but at unacceptable speed or cost.
- Solvable by computers but with unacceptable quality.
- Divisible into small, independent subtasks.

In particular, we expect it to be especially useful for tasks with this additional optional property:

- Cannot be solved by machine learning algorithms previously trained for other problems.

CrowdFlow currently uses Turkers in two roles (although we envision using a much richer model in the future):

- Worker answers the question from scratch.
- Fixer corrects a machine-created answer.

Which role is used depends on the specifics of the domain and the machine learning system. If the cognitive cost of fixing an incorrect result is low, and/or if the machine's accuracy is high, then the fixer role is preferred. This gives Turkers the benefit of the machine results, reducing their effort, and potentially reducing the required price and/or time required to get the work done. Otherwise the worker role is used. Currently, the user decides, but ultimately, we envision that the valve will decide automatically.

ANALYSIS

We wanted to better understand CrowdFlow's ability to flexibly target a specific point in the speed-cost-quality tradeoff space. Running CrowdFlow repeatedly for every possible point in the space would consume time and money needlessly. Instead, we ran the HITs once and used simulation to explore the space. We chose two domains.

Human detection. In this task, the goal is to identify human subjects in a photograph and determine their bounding boxes. Turkers did this using a web interface. The machine ran the DetectorPLS human detection software [8], which implements a partial least squares algorithm. We found this to be 60% accurate on the data we used. The web interface was pre-populated with the machine's results so that if the machine was correct, the Turker could simply press a button to confirm that.

Turkers did 120 tasks in a total of 3 hours 42 minutes (1 minute 51 seconds per task) for $2.40 ($0.02 per task) with 90% accuracy, using standard cheating deterrence methods.

We can extrapolate to other scenarios. For example, had we needed to do 1000 such tasks under a time constraint of 10 hours, the Turkers would have done only $3.24 tasks, with the remainder done by the machine. That would yield a combined accuracy of 70% and a cost of $6.48.

Sentiment polarity of movie reviews. The goal of this task is to determine if a movie review (typically several paragraphs) is more positive or more negative. Turkers read the review and answered by selecting a radio button. The machine used a classifier based on the algorithm, described by Pang and Lee with the implementation provided by the publicly available LingPipe NLP toolkit. Accuracy on our data was 83.5%. Since the answer format was simple, the worker role was used.

Turkers did a total of 1083 movie reviews (in batches of 3) over a period of 8 hours 7 minutes for $18.05 with overall accuracy of 91%. Extrapolating, we can estimate that had the user needed only 90% accuracy (for example), 1083 movie reviews could have been done for $15.64 in 7 hours 1 minute. Although the cost and time savings here is modest, CrowdFlow lets the user control the tradeoff.

FUTURE WORK

The CrowdFlow toolkit demonstrates the ability to blend human and machine resources to achieve a desired tradeoff point. In the future, we hope to achieve even greater flexibility using more human roles (e.g. verifying correct or incorrect, appraising whether an incorrect answer would be worth fixing as opposed to doing it from scratch, etc.). Also, although our current experiments relied on some existing ground truth to measure accuracy, we hope to develop methods in the future that can estimate within some confidence interval even without the ground truth.

REFERENCES

1. Hoffmann, L. *Crowd control*. CACM 52, 3 (2009), 16-17.
2. Howe, Jeff, *The Rise of Crowdsourcing*. Wired Mag. June 2006.
3. Licklider, J. C. R. Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*, 1960.
4. Mechanical Turk. http://mturk.com.
5. Malone, T.W., Laubacher, R., and Dellarocas, C. Harnessing Crowds: Mapping the Genome of Collective Intelligence. (February 3, 2009). MIT Sloan Research Paper No. 4732-09.
6. Pang, B., Lee, L., and Vaithyanathan, S. 2002. Thumbs up?: sentiment classification using machine learning techniques. *ACL '02*
7. Quinn, A. J., Bederson, B. (2009) A Taxonomy Of Distributed Human Computation. Tech. Rpt, Univ. of Maryland, HCIL-2009-23.
8. Schwartz, W. R., Kembhavi, A., Harwood, D., Davis, L. S.. Human Detection Using Partial Least Squares Analysis. *ICCV* '09
9. von Ahn, L. and Dabbish, L. 2004. Labeling images with a computer game. *CHI '04*.
10. von Ahn, L. v. 2006. Games with a Purpose. *Computer* (Jun. 2006)

Notes

SESSION III: Search – 1:45pm
Session chair: *Michael Pack, Director, CATT Laboratory*

How Children Search Online at Home
Allison Druin, Beth Foss, Evan Golub, Leshell Hatley

Finding Temporal Patterns in Electronic Health Records
Krist Wongsuphasawat, David Wang, Catherine Plaisant, Ben Shneiderman

Analyzing Trends in Science & Technology Innovation
Ben Shneiderman, Ping Wang, Yan Qu, Cody Dunne

How Children Search Online at Home

Allison Druin, Elizabeth Foss, Hilary Hutchinson[1], Evan Golub, Leshell Hatley
College of Information Studies/Dept. of Computer Science, University of Maryland/Google[1]
Contact: allisond@umiacs.umd.edu

Traditionally, researchers have portrayed information-seeking as systematic, orderly, and procedural. But as this child shows (Figure 1), seeking information using a keyword search interface on the Internet can lead to uncertainty and confusion, with a search process that can be repetitive, complex and at times end in frustration.

This is a child we have come to call a *Developing Searcher*. He has challenges with spelling, typing, query formulation and results interpretation. Over the last year, our work with 83 children, ages 7, 9, and 11, have shown that these young people demonstrate seven distinctive *search roles*, sometimes with multiple roles present during any given information-seeking experience. To define these roles we examined their behavioral patterns by age and gender with particular interest in what triggered searching and what the barriers were.

STUDY METHODS

Between September 2008 and July 2009, we undertook a qualitative study to better understand how children search for information on the Internet. We realized that both the interfaces that children use to search and the content returned are subject to change on a daily basis, making quantitative analysis a challenge. We did not want to circumvent this challenge by constraining children only to fixed tasks, interfaces, and results, as this would subvert our goal of observing their natural behavior. Consequently, we chose to collect data by letting children search freely and then moving into more targeted tasks. We then used a rigorous qualitative approach, to structure our analysis.

Participants

We worked with 83 children from the Washington DC metro area: 42 boys and 41 girls, 28 age seven, 29 age nine, and 26 age 11. The children and at least one of their parents were self-selected to participate in the study. All the participants in our study had computers and an Internet connection in their homes. All but 12 of the children used computers regularly both at home and school.

Data collection methods

We conducted qualitative, in-home interviews. Each session began with the researcher interviewing the parent, captured with audio recording. We then interviewed the child, taking notes and video recording the child's keyboard and screen actions. We did not mention any search engines or tools by name in the interviews until either the child mentioned them or had searched on their own several times. We began with open-ended questions to see what search tools children used, and how they approached searching for their own pleasure: "Can you show me and explain how you usually look for something on the computer?" From this, we moved to a more task-oriented set of questions: "How would you look for information on dolphins? How would you look for information on what dolphins eat?" Next, we asked a more personal but targeted question: "If you were searching on Google for your own interest, what would you look for?" Finally, we asked a complex, multi-step question with only one right answer: "Which day of the week will the vice president's birthday be on next year?"

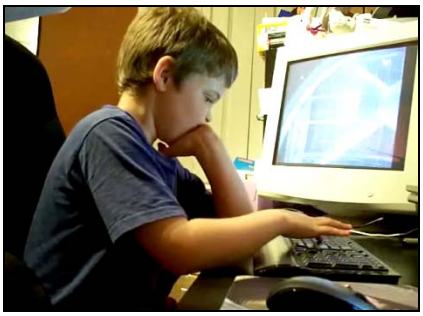

Figure 1 – A 7-year old child searching online

Data Analysis Methods

We used qualitative methods to understand the process and outcomes of the children's search experiences. We used the data coding research methods described by Strauss [1]. Specific categories for analysis were first developed using "open coding" methods of sorting, comparing, and categorizing data. We then used "axial coding" to further refine specific areas of the data. Using these techniques, we found that one useful way to consolidate the rich qualitative data to reveal the larger trends, strengths, and challenges of the search process was to describe the children as having *search roles.*

SEARCH ROLES DEFINED

Based on this data analysis process, seven search roles were ultimately defined. They are described below in order of how frequent these roles were found in the study, with the most frequent discussed first.

Developing Searcher

This role is the most common role children ages 7, 9, and 11 years old exhibit in this study. The defining behavior for the developing searcher is *willingness*, but not consistently successful, ability to search. *Developing Searchers* tend to search by using natural language syntax as opposed to keywords. Quite frequently they will by-pass a search engine and go directly to a website.

Domain-specific Searcher

Domain-specific Searchers are children who typically limit their searches to finding specific content of personal interest, which can include online games, sports scores, shopping, and videos. These searchers continually return to a small number of specific websites, and therefore, are limited in their knowledge of how to use a search engine to find new content. *Domain-specific Searchers* feel an ownership towards the content they search for and use.

Power Searcher

Power Searchers possess sophisticated searching skills. A defining characteristic of this group is their ability to understand and use keywords while searching. They are also reflective during the searching process, and can explain their searching strategies if asked. *Power Searchers* approach searching using tips or rules that are helpful to searching which they have learned from experience or from others.

Non-motivated Searcher

Unlike *Developing Searchers, Non-motivated Searchers* lack the *will* to search. *Non-motivated Searchers* are also not persistent when searching; they may attempt to find an answer briefly, and then give up or offer an alternative solution. They usually do not ask for help, admit their difficulty, or try something new.

Distracted Searcher

The role of *Distracted Searcher* is defined by children going off-task easily and wandering off on new search paths. They are difficult to get back on task, requiring multiple verbal prompts. Visual movement such as animation, blinking text, or videos within the searching interface or on linked websites is often distracting. In addition, the physical environment can also be a distraction for these searchers.

Visual Searcher

Visual Searchers are characterized by their desire to search within a visual context and have search results presented either as images or as videos. *Visual Searchers* do not simply click on an image or video result; they intentionally narrow their search results down in a visual format. These searchers frequently are able to effectively use search tools such as Google Images, YouTube, and Google Video.

Rule-bound Searcher

The least common, but clearly defined role for children in this study is that of the *Rule-bound Searcher*. As the name implies, these searchers seek information online according to an inflexible set of rules that they have learned through experience or other people, such as teachers or parents. These children are not able to adjust their rules to adapt to different types of searches. Yet, despite their frustration, they display persistence in their searching.

ANALYSIS USING SEARCH ROLES

We found in this study that most children exhibited multiple search roles. Most children in this study exhibited from 1 to 4 roles with an average of less than 2 roles per child. The 7-year olds exhibited the most roles per child, and the 11-year olds the fewest. This suggests that as these children get older; their search roles become more consistent.

Children had varying abilities to understand whether they had found what they were seeking, due to a wide variety of barriers. Not surprisingly, we also found very few children in this study were successful in formulating multi-step queries. We also found that among all of the search roles, children were more successful when they looked for information of personal interest, due to motivation, past experience, or both.

The girls at all ages in this study tended to offer their concerns when faced with challenges, while the males would suggest a way to fix the situation. For example, one 9-year old girl said, "Oh, oh, I'm looking for the wrong thing. I'm not good at math, I don't know." On the other hand, the boys would explain a new path forward. For example, a 9–year old boy suggested, "I could just get off the computer and look at a calendar." For a more complete understanding of this research see our full papers [1, 2].

ACKNOWLEDGEMENTS

We gratefully acknowledge the financial support of two Google Research Awards 2009 & 2010. We could not have accomplished this work without the parents and children who participated.

REFERENCES

1. Druin, A., Foss, E., Hutchinson, H., Golub, E., & Hatley, L. (2010). Children's roles using keyword search interfaces at home. In *Proceedings of ACM CHI2010*, Atlanta, Georgia.
2. Druin, A., Foss, E., Hatley, L., Golub, E., Guha, M. L., Fails, J., & Hutchinson, H. (2009). How children search the Internet with keyword interfaces. In *Proceedings of Interaction Design and Children (IDC 2009)*, Cuomo, Italy.

Notes

Finding Patterns in Temporal Data

Krist Wongsuphasawat[1,2], Taowei David Wang[1,2], Catherine Plaisant[2], Ben Shneiderman[1,2]
Department of Computer Science[1], Human-Computer Interaction Lab[2]
Contact: {kristw, tw7, plaisant, ben}@cs.umd.edu
www.cs.umd.edu/hcil/temporalviz

Temporal categorical data is a collection of records that contain *timestamped events*, such as (11/24/2009,Flu), (11/30/2009, Stroke), etc. Temporal categorical data is increasingly being collected by many organizations. Health organizations have Electronic Health Record (EHR) databases containing millions of records of patient histories. Each patient history contains hospital admissions, patients' symptoms and treatments. Transportation systems generate logs of incidents (notifications and arrivals of each unit on the scene). Temporal categorical data also includes student reports, web logs, financial histories, market baskets and data in many other domains.

These temporal categorical databases contain many patterns. For example, using EHR database that keeps track of patient transfers in the hospital, our physician partner wants to find "bounce backs" patients, which are patients in Intensive Care Unit (ICU) who seemed to get better so the physicians transferred them to normal rooms, or "Floor". However, they were transferred back to the ICU within 48 hours. These situations impact the quality of care and are undesired by the hospital. Therefore, the physician wants to find these patients to analyze what might have been the reasons that led to those transfers. There has been a long history of visualizing temporal categorical data, from visualizing intensive details of a single record to visualizing and searching for multiple records. Last year we presented two interactive search and visualization tools, LifeLines2 and Similan, and have been using these tools to support the physician in finding patients with specified patterns.

By working on these case studies, we discovered a need to explore the sequential patterns in the data. For example, a physician might ask, "What usually happened to the patient after they arrived the hospital?" Many visualizations do not support this task well due to a scalability issue. When there is a large number of records, users cannot see all records on

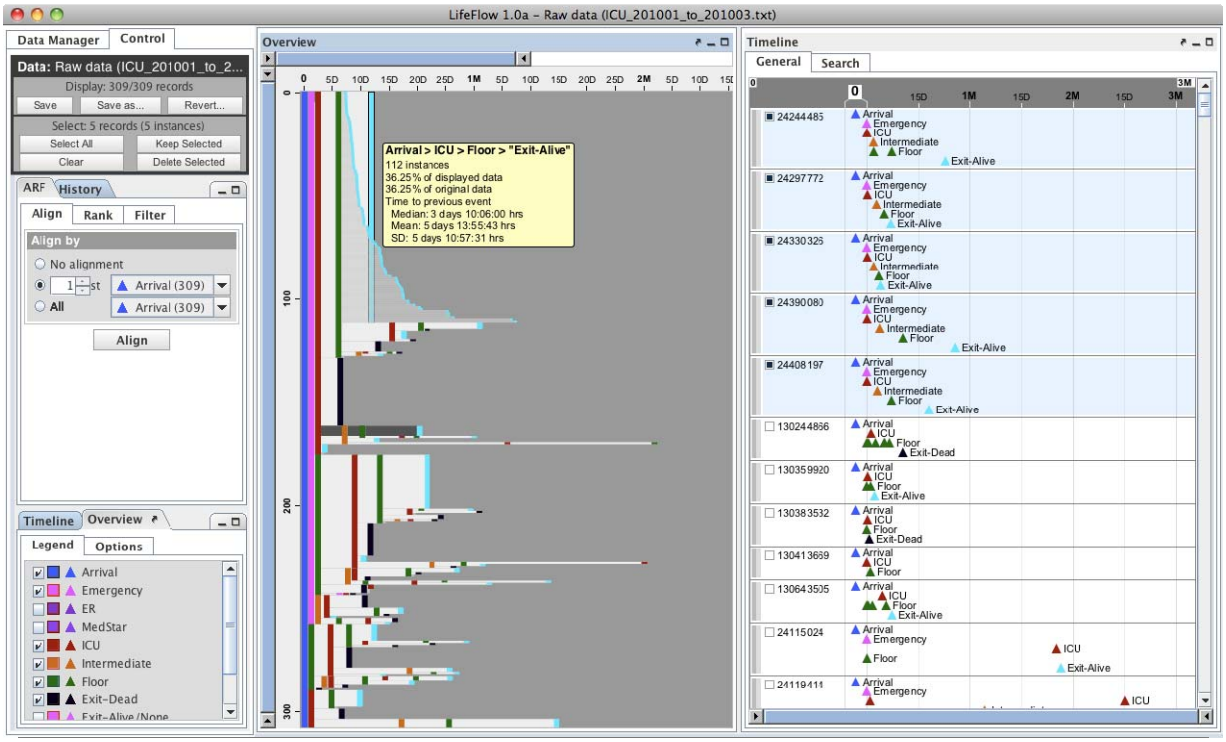

Figure 1. A working prototype, LifeFlow 1.0a –The control panel is on the left. In the middle is the LifeFlow visualization, showing patient transfer patterns from the time of arrival. On the right is a view inherited from LifeLines2, showing each patient in details. The selected path on LifeFlow is highlighted in dark gray and all records in that path is selected and highlighted in light blue on the LifeLines2 view.

the screen at once, which makes it difficult to recognize any pattern. A few scalable visualizations were later developed, but these techniques aggregate the records in a way that the sequential information is obscured. According to the information visualization mantra: "overview, zoom and filter, then details-on-demand", an overview of sequential patterns is needed.

LifeFlow: An Overview Visualization

Therefore, we have recently developed a new visualization called *LifeFlow* to provide a scalable overview of sequential patterns in temporal categorical data. Our technique aggregates multiple records into a tree, based on sequence of events, and visualizes this tree with the LifeFlow visualization.

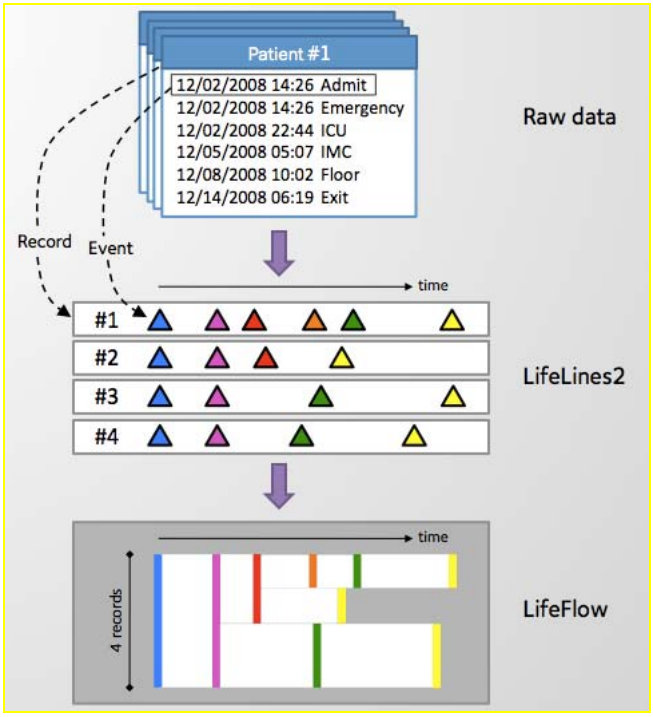

Figure 2. A diagram showing four records of raw data and how they are visualized as LifeLines2 and LifeFlow visualization.

LifeFlow uses a color-coded rectangle, or node, to represent an event. The height of each node represents the number of records. The first node (blue) and second node (pink) in Figure above has 100% height because all records fall into this node. The horizontal axis is used as a timeline. All nodes are placed on a horizontal position according to the representative time gaps between them and their previous nodes. (A mean is currently used by default.) An edge between each pair of nodes contains a distribution of the time gap between the two nodes, which will be overlaid on the nodes when users move their cursors over. A working prototype LifeFlow 1.0a inherited several features from our previous work on LifeLines2, including the Align, Rank and Filter (ARF) Framework and visualization of the records. By connecting these components together, user now can explore the data following the information visualization mantra: see an overview from the LifeFlow visualization; use the ARF Framework and other interactions to zoom and filter; then select any path on LifeFlow to see more details on demand.

Ongoing Case Study

We are conducting a multidimensional in-depth long-term case study with our physician partner at the Washington Hospital Center by helping him analyze patient transfers. An example can be seen from Figure 1. LifeFlow can easily show the common transfer pattern as well as point out the uncommon cases. It also shows the average time that the patients were transferred from one room to another. The preliminary result shows a promising opportunity to use LifeFlow to monitor the Emergency Department (ED) performance over time.

Conclusion and Future Work

We introduced a new visualization called LifeFlow that provides an overview of sequential patterns in temporal categorical data to support users' exploration, and reported an ongoing case study with the physicians. In the mean time, we are working on including additional features, such as displaying statistical information or comparison between two or more data sets. While the examples given and case studies described are medical use cases, we believe our technique can be applied to many other fields, where temporal categorical data is the main focus. We are eager to conduct new case studies with experts in other domains.

ACKNOWLEDGEMENT

We would like to thank the Washington Hospital Center, NIH grant CA147489 and Center for Advanced Transportation Technology (CATT Lab) for their support.

PAPERS

1. Wang, T., Plaisant, C., Quinn, A., Stanchak, R., Shneiderman, B., and Murphy, S., Aligning temporal data by sentinel events: Discovering patterns in electronic health records, *Proc. CHI 2008*, ACM, 457-466
2. Wang, T., Plaisant, C., Shneiderman, B., Spring, N., Roseman, D., Marchand, G., Mukherjee, V., Smith, M., Temporal Summaries: Supporting Temporal Categorical Searching, Aggregation and Comparison, *Proc. IEEE Information Visualization 2009*, IEEE, 1049-1056
3. Wongsuphasawat, K., Shneiderman, B., Finding Comparable Temporal Categorical Records: A Similarity Measure with an Interactive Visualization, Proc. IEEE Symp. Visual Analytics Science and Technology 2009, IEEE, 27-34

Notes

Analyzing Trends in Science & Technology Innovation

Ben Shneiderman[1], Ping Wang[2], Yan Qu[2], and Cody Dunne[1]
[1]Department of Computer Science & Human-Computer Interaction Lab
[2]College of Information Studies & Human-Computer Interaction Lab
Contact: ben@cs.umd.edu, pwang@umd.edu, yanqu@umd.edu, cdunne@cs.umd.edu
Web site: http://terpconnect.umd.edu/~pwang/STICK/

This project is part of the larger NSF-supported grant on Science and Technology Innovation Concept Knowledge-base (STICK). Previous research in trends of information technology innovations have found helpful patterns by studying the frequency of mention of key terms in trade and scholarly articles in bibliographic databases (Fig. 1).

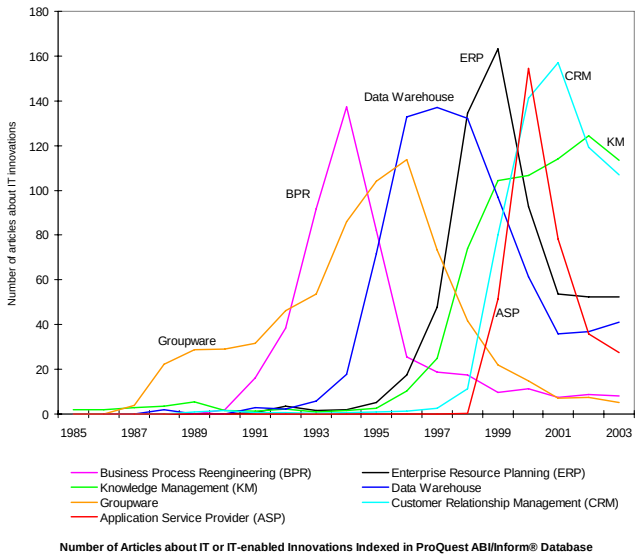

Figure 1: These trend lines show the number of articles mentioning seven key terms over two decades

The current efforts on this project extend this work to even more diverse sources of information and a richer model of technology innovation that includes academic papers, patents, trade press articles, news, blog mentions, product adoptions, organizations, people, and more careful topic analysis. Our work focuses on the development of ontologies to guide database schema design and two case studies of innovation trajectory analysis.

STICK Project Goals

The Science & Technology Innovation Concept Knowledge-base (STICK) project's goal is to overcome the bias in the Science of Science and Innovation Policy (SciSIP) towards popular or ultimately successful innovations by providing the much needed data and tools for analyzing innovations of all possible outcomes. This comprehensive endeavor enables SciSIP researchers to build and test theories that explain the differentiated trajectories of science and technology innovations and their associated communities. The project also spans disciplinary boundaries by bridging the artificial divide in SciSIP research between the production and the use of innovations, piecing together a holistic view of the dynamic supply and demand in the innovation ecosystem. Specifically, the project builds a large-scale, multi-source, longitudinal database, Science & Technology Innovation Concept Knowledge-base (STICK), and develops a set of visual analytic tools for monitoring and understanding the emergence and revolution/evolution of innovations in three exemplar science and technology fields: information technology, biotechnology, and nanotechnology.

The knowledge-base captures innovations, the individual and organizational actors associated with the innovations, and the relationships among the innovations and the actors through a hybrid approach that combines computational analysis of text (e.g., natural language processing) and social information processing (e.g., social tagging and collaborative writing by the users of the knowledge-base). State-of-the-art visualization tools are customized for SciSIP researchers and other innovation stakeholders to visualize innovation networks and analyze patterns and trends. The design of the knowledge-base and toolset is grounded in a demonstration study on the popularity of innovations. The study aims to address important questions concerning the complex relationships among innovations and the evolution of communities, with implications to the popularity and ultimate success of innovations.

Broader Impacts: STICK will be institutionalized at the University of Maryland at College Park as a free public service that offers web access to the data and tools developed in this project. This service also produces quarterly reports on the status of science and technology innovations, including the National Innovation Popularity Index, analogous to the Consumer Confidence Index for

the state of the economy. This research-based service will be designed as a tool for science and technology education. For most fields where specialization is the theme nowadays, students' and the public's interests increase with the capability to monitor and make sense of the fast-changing arenas where innovations emerge, converge, and diverge. For scientists and engineers, STICK's visual analytic toolset helps accelerate scientific discoveries and innovations by identifying and establishing collaborations within and across innovation communities. Finally, STICK will help science and technology policy makers monitor and understand the evolutionary paths of innovations, appraise the significance of innovations in rigorously charted terrains, and proactively foster, promote, and advance innovations with benefits to the society.

Tree Visulization Innovation Trajectories

One of our first case studies was a 20-year review of the innovation trajectories for three strategies for exploring tree-structured data: treemaps, cone trees, and hyperbolic trees (Fig. 2). These were introduced in academic papers in the early 1990s, which all became widely cited, although substantial follow-on work by academics grew only for treemaps starting after 2000. Patents were filed by Xerox PARC for cone trees and hyperbolic trees mainly in the late 1990s. The University of Maryland did not file patents for treemaps but refinement patents were filed by others. A flurry of trade press articles accompanied the commercial introduction of hyperbolic trees in the late 1990s, but subsided, while trade press articles on treemaps gained when The Hive Group (www.hivegroup.com) and others began marketing commercial versions.

These distinctive innovation trajectories shown by these three sources are a starting point for understanding the determinants of success for innovations. A few of the determinants of success are the entrenched alternatives, resistance to change, interactions with existing technologies, intellectual property rights, perceived usefulness and ease-of-use, availability of free versions, influence of entrepreneurial individuals, and responsiveness of the trade press. This complex mix of personalities, ideas, institutions, and economic constraints calls for a comprehensive theory that may better predict innovation outcomes so as to guide entrepreneurs and policy analysts.

Acknowledgement: We appreciate the support of the National Science Foundation, Award 0915645 under the Science of Science and Innovation Policy program. Thanks to Puneet Sharma for his data cleaning effort on patents.

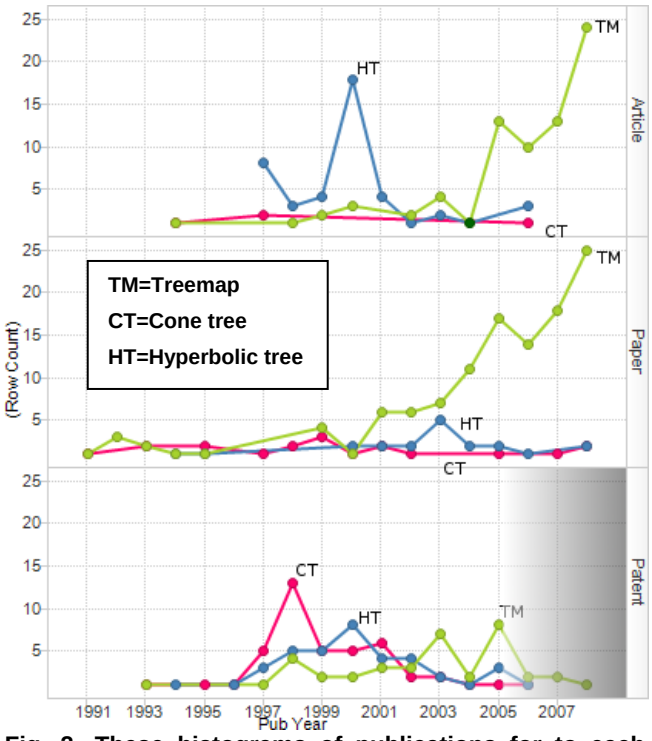

**Fig. 2: These histograms of publications for to each innovation show the number of trade press articles, academic papers, and patents published each year. **

PAPERS

1. Wang, P. "Chasing the Hottest IT: Effects of Information Technology Fashion on Organizations," *MIS Quarterly* (34:1), 2010, pp. 63-85
2. Tsui, C.-J., Wang, P., Fleischmann, K.R., Oard, D.W., and Sayeed, A.B. "Exploring the Relationships among ICTs: A Scalable Computational Approach Using Kl Divergence and Hierarchical Clustering," *43rd Hawai'i Int'l Conference on System Sciences*, Kauai, HI, 2010. http://doi.ieeecomputersociety.org/10.1109/HICSS.2010.203
3. Tsui, C.-J., Wang, P., Fleischmann, K. R., Sayeed, A. B., and Weinberg, A. "Building an IT Taxonomy with Co-Occurrence Analysis, Hierarchical Clustering, and Multidimensional Scaling," *iConference*, Champaign, IL., 2010. http://www.ideals.illinois.edu/handle/2142/14918
4. Shneiderman, B., Dunne, C., Sharma, P., and Wang, P., "Innovation Trajectories for Information Visualizations: Comparing Treemaps, Cone Trees, and Hyperbolic Trees", University of Maryland Human-Computer Interaction Lab Technical Report, 2010.

Notes

Technical Report Abstracts

2010, 2009

HCIL Papers and Technical Reports 2010, 2009

For a complete listing of our papers and technical reports, please refer to our website, www.cs.umd.edu/hcil. Our papers and technical reports are available on-line in several formats, including HTML, PDF, Postscript, and text only.

2010

Quinn, A., Bederson, B., Yeh, T., Lin, J. (May 2010)
CrowdFlow: Integrating Machine Learning with Mechanical Turk for Speed-Cost-Quality Flexibility
HCIL-2010-09

Humans and machines have competing strengths for tasks such as natural language processing and image understanding. Whereas humans do these things naturally with potentially high accuracy, machines offer greater speed and flexibility. CrowdFlow is our toolkit for a model for blending the two in order to attain tighter control over the inherent tradeoffs in speed, cost and quality. With CrowdFlow, humans and machines work together to do a set of tasks at a user-specified point in the tradeoff space. They work symbiotically, with the humans providing training data to the machine while the machine provides first cut results to the humans to save effort in cases where the machine's answer was already correct. The CrowdFlow toolkit can be considered as a generalization of our other domain-specific efforts aimed at enabling cloud computing services using a variety of computational resources to achieve various tradeoff points.

Yeh, T., White, B., Davis, L., Katz, B. (May 2010)
Searching the Web Using Screenshots
HCIL-2010-08

Many online articles contain useful know-how knowledge about GUI applications. Even though these articles tend to be richly illustrated by screenshots, no system has been designed to take advantage of these screenshots to visually search know-how articles effectively. In this paper, we present a novel system to index and search software knowhow articles that leverage the visual correspondences between screenshots. To retrieve articles about an application, users can take a screenshot of the application to query the system and retrieve a list of articles containing a matching screenshot. Useful snippets such as captions, references, and nearby text are automatically extracted from the retrieved articles and shown alongside with the thumbnails of the matching screenshots as excerpts for relevancy judgment. Retrieved articles are ranked by a comprehensive set of visual, textual, and site features, whose weights are learned by RankSVM. Our prototype system currently contains 150k articles that are classified into walkthrough, book, gallery, and general categories. We demonstrated the system's ability to retrieve matching screenshots for a wide variety of programs, across language boundaries, and provide subjectively more useful results than keyword-based web and image search engines.

Resnik, P., Hu, C., Buzek, O., Bederson, B. (May 2010)
Using Monolingual Human Computation to Improve Language Translation via Targeted Paraphrase
HCIL-2010-07

We introduce a new approach to the problem of obtaining cost-effective, reasonable quality translation, by exploiting simple and inexpensive human computations by monolingual speakers. The key insight behind the process is that it is possible to spot likely translation errors with only monolingual knowledge of the target language, and it is possible to generate new ways to say the same thing (i.e. paraphrases) with only monolingual knowledge of the source language. Initial evaluation demonstrates substantial improvements in translation quality.

Rios-Berrios, M., Sharma, P., Schwartz, R., Lee, T., Shneiderman, B. (May 2010)
TreeCovery : Coordinated Dual Treemap Visualization for Exploring the Recovery Act
HCIL-2010-06

The American Recovery and Reinvestment Act dedicated $787 billion to stimulate the US economy and mandated the release of the data describing the exact distribution of that money. The dataset is a large and complex one; one of its distinguishing features is its bi-hierarchical structure, arising from the distribution of money through agencies to specific projects and the natural aggregation of awards based on location. To offer a comprehensive overview of the data, a visualization must incorporate both these hierarchies. We present TreeCovery, a tool that accomplishes this through the use of two coordinated treemaps. The tool includes a number of innovative features, including coordinated zooming and filtering and a proportional highlighting technique across the two trees. TreeCovery was designed to facilitate data exploration, and initial user studies suggest that it will be helpful in insight generation. RATB(Recovery Accountability and Transparency Board) has tested TreeCovery and is considering including the concept into their visual analytics.

Gregory, M., Shneiderman, B. (May 2010)
Shape Identification in Temporal Data Sets
HCIL-2010-05

Shapes are a concise way to describe temporal variable behaviors. Some commonly used shapes are spikes, sinks, rises, and drops. A spike describes a set of variable values that rapidly increase, then immediately rapidly decrease. The variable may be the value of a stock or a persons blood sugar levels. Shapes are abstract. Details such as the height of spike or its rate increase, are lost in the abstraction. These hidden details make it difficult to define shapes and compare one to another. For example, what attributes of a spike determine its spikiness? The ability to define and compare shapes is important because it allows shapes to be identified and ranked, according to an attribute of interest. Work has been done in the area of shape identification through pattern matching and other data mining techniques, but ideas combining the identification and comparison of shapes have received less attention. This paper fills the gap by presenting a set of shapes and the attributes by which they can identified, compared, and ranked. Neither the set of shapes, nor their attributes presented in this paper are exhaustive, but it provides an example of how a shapes attributes can be used for identification and comparison. The intention of this paper is not to replace any particular mathematical method of identifying a particular behavior, but to provide a toolset for knowledge discovery and an intuitive method of data mining for novices. Spikes, sinks, rises, drops, lines, plateaus, valleys, and gaps are the shapes presented in this paper. Several attributes for each shape are defined. These attributes will be the basis for constructing definitions that allow the shapes to be identified and ranked. The second contribution is an information visualization tool, TimeSearcher: Shape Search Edition (SSE), which allows users to explore data sets using the identification and ranking ideas in this paper.

Sopan, A., Noh, A., Lee, G., Rosenfeld, P., Karol, S., Shneiderman, B. (May 2010)
Community Health Map: A Geospatial and Multivariate Data Visualization Tool for Public Health Datasets
Submitted to *VAST 2010*
HCIL-2010-04

Trillions of dollars are spent each year on health care. The Department of Health and Human Services keeps track of a variety of health care indicators across the country, resulting in a large geospatially multivariate data set. Current visualization tools for such data sets make it difficult to make multivariate comparisons and show the geographic distribution of the selected variables at the same time. Community Health Map is a web application that enables users to visualize health care data in multivariate space as well as geospatially. It is designed to aid exploration of this

huge data repository and deliver deep insights for policy makers, journalists, and academic researchers. Users can visualize the geospatial distribution of a given variable on an interactive map, and compare two or more variables using charts and tables. By employing dynamic query filters, visualizations can be narrowed down to specific ranges and regions. Our presentation to policy makers and pilot usability evaluation suggest that the Community Health Map provides a comprehensible and powerful interface for policy makers to visualize health care quality, public health outcomes, and access to care in an effort to help them to make informed decisions about improving health care.

Guha, M., Druin, A., Fails, J. (April 2010)
Investigating the Impact of Design Processes on Children
In press, *IDC 2010,* Barcelona, Spain
HCIL-2010-03

While there is a wealth of information about children's technology and the design processes used to create it, there is a dearth of information regarding how the children who participate in these design processes may be affected by their participation. In this paper, we motivate why studying this impact is important and look at the foundation provided by past research that touches on this topic. We conclude by briefly proposing methods appropriate for studying the impact of the design process on the children involved.

Fails, J., Druin, A., Guha, M. (April 2010)
Mobile Collaboration: Collaboratively Reading and Creating Children's Stories on Mobile Devices
In press, *IDC 2010,* Barcelona, Spain
HCIL-2010-02

This paper discusses design iterations of Mobile Stories - a mobile technology that empowers children to collaboratively read and create stories. We present the design and discuss the impact of different collocated collaborative configurations for mobile devices including: content splitting and space sharing. We share design experiences that illustrate how Mobile Stories supports collaboration and mobility, and identify how the collocated collaborative configurations are best suited for reading and sharing tasks. We also identify how creative tasks foster more mobility and dynamic interactions between collaborators.

Sopan, A., Freire, M., Plaisant, C., Golbeck, J., Shneiderman, B. (April 2010)
Distribution Column Overviews in Tabular Visualizations
HCIL-2010-01

While visual overviews of tables of numerical and categorical data have been proposed for tables with a single value per cell, it is difficult to summarize tables containing columns of distributions (e.g. the distributions of ages or ethnicity across counties, or the distribution of movie ratings and trust ratings in recommender systems). We present a novel way of displaying and manipulating overviews of distribution columns to reveal interesting patterns and outliers or find possible correlations between columns. We illustrate our approach with examples drawn from the analysis of telephone calls and movie recommendation systems.

2009

Fails, J. (August 2009)
MOBILE COLLABORATION FOR YOUNG CHILDREN: READING AND CREATING STORIES
HCIL-2009-34

Within the last decade, mobile devices have become an integral part of society, at home or work, in industrialized and developing countries. For children, these devices have primarily been geared towards communication, information consumption, or individual creative purposes. Prior research indicates social interaction and collaboration are essential to the social and cognitive development of young children. This dissertation research focuses on supporting collaboration among mobile users, specifically children ages 6 to 10 while collaboratively reading and creating stories. I developed Mobile Stories, a novel software system for the Windows Mobile platform that supports collaborative story experiences, with special attention to two collocated collaboration experiences: content splitting and space sharing. Content splitting is where interface parts (e.g. words, pictures) are split between two or more devices. Space sharing is where the same content (e.g. a document) is spread or shared across devices. These collocated collaborative configurations help address mobile devices primary limitation: a small screen.

The three research questions addressed are: how does Mobile Stories affect children's collaboration and mobility, what are some appropriate interfaces for collocated mobile collaboration with children, and when are the developed interfaces preferred and why. Mobile Stories was designed and develop using the Cooperative Inquiry design method. Formative studies furthered the design process, and gave insight as to how these collaborative interfaces might be used. A formal, mixed method study was conducted to investigate the relative advantages for each of the collocated collaborative interfaces, as well as to explore mobility and collaboration.

The results of the formal study show children were more mobile while creating stories than when reading and sharing them. As for task effectiveness, children read more pages when they were closer, and created more pages when they were further apart and more mobile. Children were closer together when they read using the content split configuration. While creating their stories, children rarely used the collocated collaborative configurations and used verbal collaboration instead. Several indicators pointed to relative advantages of the split content configuration over the share space configuration; however, the advantages of each are discussed.

Druin, A., Foss, E., Hutchinson, H., Golub, E., Hatley, L. (December 2009)
Children's Roles using Keyword Search Interfaces at Home
New York Times Article:
Published in *ACM CHI 2010 Conference on Human Factors in Computing Systems*
HCIL-2009-33

Children want to find information about their world, but there are barriers to finding what they seek. Young people have varying abilities to formulate complex queries and comprehend search results. Challenges in understanding where to type, confusion about what tools are available, and frustration with how to parse the results page all have led to a lack of perceived search success for children 7-11 years old. In this paper, we describe seven search roles children display as information seekers using Internet keyword interfaces, based on a home study of 83 children ages 7, 9, and 11. These roles are defined not only by the children's search actions, but also by who influences their searching, their perceived success, and trends in age and gender. These roles suggest a need for new interfaces that expand the notion of keywords, scaffold results, and develop a search culture among children.

Golbeck, J., Grimes, J., Rogers, A. (December 2009)
Twitter Use by the U.S. Congress
HCIL-2009-32

Twitter is a microblogging and social networking service with millions of members and growing at a tremendous rate. With the buzz surrounding the service have come claims of its ability to transform the way people interact and share information, and calls for public figures to start using the service. In this study, we are interested in the type of content that legislators are posting to the service, particularly by members of the United States Congress. We read and coded over 6,000 posts from all members of Congress using the site. Our analysis shows that Congresspeople are primarily using Twitter to post information, particularly links to news articles about themselves and to their blog posts, and to report on their daily activities. These tend not to provide new insights into government or the legislative process or to improve transparency; rather, they are vehicles for self-promotion. However, Twitter is also facilitating direct communication between Congresspeople and citizens, though this is a less popular activity. We report on our findings and analysis, and discuss other uses of Twitter for legislators.

Zalinger, J., Freier, N., Freire, M., Shneiderman, B.
Reading Ben Shneiderman's Email: Identifying Narrative Elements in Email Archives
HCIL-2009-31

This paper describes techniques for finding narrative elements in the archived email of a scholar. The goal is to test a narrative approach to searching using a 15-year email archive containing nearly 45,000 messages belonging to University of Maryland Professor Ben Shneiderman and ranging from 1984-1998. The goal is not to find complete narratives (although, many do exist) but to search for narrative elements, the building blocks that make up a narrative. Thus, narrative search is defined as both a set of search techniques and a way of thinking like a storyteller that allows designers and users to uncover narrative elements. We argue that narrative search is a promising strategy that can be productively applied to other email archives. This paper makes a contribution to HCI by showing that a narrative approach to search can be productive and compelling. By encouraging designers (and users) to think like storytellers, we can create robust interfaces that help users make narrative sense out of overwhelming amounts of messages.

Freire, M., Plaisant, C., Shneiderman, B., Golbeck, J. (September 2009)
ManyNets: An Interface for Multiple Network Analysis and Visualization
CHI 2010
HCIL-2009-30

Traditional network analysis tools support analysts in studying a single network. ManyNets offers these analysts a powerful new approach that enables them to work on multiple networks simultaneously. Several thousand networks can be presented as rows in a tabular visualization, and then inspected, sorted and filtered according to their attributes. The networks to be displayed can be obtained by subdivision of larger networks. Examples of meaningful subdivisions used by analysts include ego networks, community extraction, and time-based slices. Cell visualizations and interactive column overviews allow analysts to assess the distribution of attributes within particular sets of networks. Details, such as traditional node-link diagrams, are available on demand. We describe a case study analyzing a social network geared towards film recommendations by means of decomposition. A small usability study provides feedback on the use of the interface on a set of tasks issued from the case study.

Walsh, G., Druin, A., Guha, M., Foss, E., Golub, E., Hatley, L., Bonsignore, E., Franckel, S. (October 2009)
Layered Elaboration: A New Technique for Co-Design with Children
HCIL-2009-29

As technology for children becomes more mobile, social, and distributed, our design methods and techniques must evolve to better explore these new directions. This paper reports on Layered Elaboration, a co-design technique developed over the past year. Layered Elaboration allows design teams to generate ideas through an iterative process in which each version leaves prior ideas intact while extending concepts. Layered Elaboration is a useful technique as it enables co-design to take place asynchronously and does not require much space or many resources. Our intergenerational team used the technique to design a prototype of an instructional game about energy conservation

Bederson, B., Hu, C., Resnik, P. (October 2009)
Translation by Iterative Collaboration between Monolingual Users
HCIL-2009-28

In this paper we describe a new iterative translation process designed to leverage the massive number of online users who have minimal or no bilingual skill. The iterative process is supported by combining existing machine translation methods with monolingual human speakers. We have built a Web-based prototype that is capable of yielding high quality translations at much lower cost than traditional professional translators. Preliminary evaluation results of this prototype confirm the validity of the approach.

Golbeck, J., Hu, C. (October 2009)
Impact of Visualization Methods on Interaction with Search Results
HCIL-2009-27

There are many search and browsing tasks online where relevance scores are not particularly important to the user, but other scores like popularity or average rating can be very informative. If and how these scores are shown varies widely between systems. In this paper, we investigate different methods for visualizing these scores and how they affect user behavior. We conducted a con- trolled study with 21 subjects who each completed tasks with six different visualization methods. We found that there was no significant difference between the methods with respect to their impact on the user interaction with search results, but that there was a strong preference for having some sort of visualization. We discuss the experiment, results, and design implications that follow from this work.

Golbeck, J. (October 2009)
A Single Strong Disagreement Ruins a Recommender: Improving Recommendation Accuracy with a Simple Statistic
HCIL-2009-26

Research on the use of social trust relationships for collaborative altering has shown that trust-based recommendations can outperform traditional methods in certain cases. This, in turn, lead to insights that tie trust to certain more subtle types of similarity between users which is not captured in the overall similarity measures normally used for making recommendations. In this study, we investigate the use these trust-inspired nuanced similarity measures directly for making recommendations. After describing previous research that identified these similarity statistics, we present an experiment run on two data sets: FilmTrust and Movie- Lens. Our results show that using a simple measure - the single largest difference between users - as a weight produces significantly more accurate results than a traditional collaborative altering algorithm and in some cases also outperforms a model-based approach.

Walsh, G., Golbeck, J. (October 2009)
Curator: A Game with a Purpose for Collection Recommendation
HCIL-2009-25

Collection recommender systems suggest groups of items that work well as a whole. The interaction effects between items is an important consideration, but the vast space of possible collections makes it difficult to analyze. In this paper, we present a class of games with a purpose for building collections where users create collections and, using an output agreement model, they are awarded points based on the collections that match. The data from these games will help researchers develop guidelines for collection recommender systems among other applications. We conducted a pilot study of the game prototype which indicated that it was fun and challenging for users, and that the data obtained had the characteristics necessary to gain insights into the interaction elects among items. We present the game and these results followed by a discussion of the next steps necessary to bring games to bear on the problem of creating harmonious groups.

Golbeck, J. (October 2009)
The More People I Meet, The More I Like My Dog: A Study of Pet-Oriented Social Networks on the Web
HCIL-2009-24

There has been extensive research into friend-oriented and professional social networking websites, but relatively little work on passion-oriented sites designed to connect strangers around a shared passion. In this study, we extend the work on passion-oriented social networking through an examination of pet-oriented social networks. We address two questions. First, do people interacting in semi-anonymous passion-oriented social networks behave in observably different ways from users in friend-oriented networks? Second, do groups that, on the surface, appear quite similar (i.e. dog and cat owners) use the passion-oriented networks in significantly different ways? Our results show that passion- oriented networking behavior is significantly different from that on friend-oriented networks and that despite the apparent similarities, dog and cat owners use these sites quite differently. We discuss the implications these results have theoretically for understanding passion- oriented social networking and practically for supporting the human-animal bond in virtual environments.

Quinn, A., Bederson, B. (October 2009)
A Taxonomy of Distributed Human Computation
HCIL-2009-23

Distributed Human Computation (DHC) holds great promise for using computers and humans together to scaling up the kinds of tasks that only humans do well. Currently, the literature describing DHC efforts so far is segmented. Projects that stem from different perspectives frequently do not cite each other. This can be especially problematic for researchers trying to understand the current body of work in order to push forward with new ideas. Also, as DHC matures into a standard topic within human-computer interaction and computer science, educators will require a common vocabulary to teach from. As a starting point, we offer a taxonomy which classifies and compares DHC systems and ideas. We describe the key characteristics and compare and contrast the differing approaches.

Quinn, A., Bederson, B., Bonsignore, E., Druin, A. (October 2009)
StoryKit: Designing a Mobile Application for Story Creation By Children And Older Adults
HCIL-2009-22

As the capabilities of smartphones and similar mobile devices advance, opportunities increase to use them for meaningful creative tasks. Incorporating text, images, and sounds in documents is commonplace when using desktop office or graphics software. However, multimedia authoring interfaces for mobile devices remain undeveloped. Working with a participatory design group composed of children, older adults, and researchers, we developed StoryKit, an iPhone application for creating and sharing audio-visual stories on an iPhone. Our initial goal was to support children making stories together with older adults as a form of informal learning. To that end, it lets users create books on the touchscreen device by arranging their text, photos, drawings, and sounds on pages, and then sharing them via e-mail and the web. The design of StoryKit uncovered solutions to several general interface challenges that affect a wide range of mobile authoring applications. Thus, we think elements of the StoryKit interaction design may serve as a starting point for developers of mobile document authoring applications.

Bederson, B. (October 2009)
The Promise of Zoomable User Interfaces
HCIL-2009-21

Zoomable User Interfaces (ZUIs) have received a significant amount of attention in the 17 years since they were introduced. They have enjoyed some success, and elements of ZUIs are widely used in computers today, although the grand vision of a zoomable desktop has not materialized. This paper describes the premise and promise of ZUIs along with their challenges. It describes design guidelines, and offers a cautionary tale about research and innovation.

Wongsuphasawat, K., Plaisant, C., Shneiderman, B. (October 2009)
Event Sequence Queries by Example or by Filters: Design and Empirical Evaluation
HCIL-2009-20

Specifying event sequence queries is challenging even for skilled computer professionals familiar with SQL. Most graphical user interfaces for database search use a query-by-filters approach, which is often effective, but applies an exact match criteria. We describe a new query-by-example interface, in which users specify a pattern by simply placing events on a blank timeline, producing a similarity-ranked list of results. Users customize the similarity measure by four decision criteria, enabling them to reduce the impact of missing, extra, or swapped events or the impact of time shifts. We describe an example of use with electronic health records based on our ongoing collaboration with hospital physicians. Then we report on a controlled experiment with 18

participants that compared query-by-filters and query-by-example features. We report on the advantages and disadvantages of each approach and conclude with recommendations for the design of a hybrid approach combining both interfaces.

Bonsignore, E., Dunne, C., Rotman, D., Smith, M., Capone, T., Hansen, D., Shneiderman, B. (August 2009)
First steps to NetViz Nirvana: Evaluating social network analysis with NodeXL
In SIN '09: Proc. International Symposium on Social Intelligence and Networking. IEEE Computer Society Press.
HCIL-2009-19

Social Network Analysis (SNA) has evolved as a popular, standard method for modeling meaningful, often hidden structural relationships in communities. Existing SNA tools often involve extensive pre-processing or intensive programming skills that can challenge practitioners and students alike. NodeXL, an open-source template for Microsoft Excel, integrates a library of common network metrics and graph layout algorithms within the familiar spreadsheet format, offering a potentially low-barrier-to-entry framework for teaching and learning SNA. We present the preliminary findings of 2 user studies of 21 graduate students who engaged in SNA using NodeXL. The majority of students, while information professionals, had little technical background or experience with SNA techniques. Six of the participants had more technical backgrounds and were chosen specifically for their experience with graph drawing and information visualization. Our primary objectives were (1) to evaluate NodeXL as an SNA tool for a broad base of users and (2) to explore methods for teaching SNA. Our complementary dual case-study format demonstrates the usability of NodeXL for a diverse set of users, and significantly, the power of a tightly integrated metrics/visualization tool to spark insight and facilitate sensemaking for students of SNA.

Shneiderman, B. (July 2009)
iParticipate.gov: A National Initiative for Social Participation
Science 323 (March 13, 2009), 1426-1427.
HCIL-2009-18

Hansen, D., Rotman, D., Bonsignore, E., Milic-Frayling, N., Rodrigues, E., Smith, M., Shneiderman, B. (September 2009)
Do You Know the Way to SNA?: A Process Model for Analyzing and Visualizing Social Media Data
The paper has been submitted to CHI 2010.
HCIL-2009-17

Traces of activity left by social media users can shed light on individual behavior, social relationships, and community efficacy. Tools and processes to analyze social traces are essential for enabling practitioners to study and nurture meaningful and sustainable social interaction. Yet such tools and processes remain in their infancy. We conducted a study of 15 graduate students who were learning to apply Social Network Analysis (SNA) to data from online communities. Based on close observations of their emergent practices, we derived the Network Analysis and Visualization (NAV) process model and identified stages where intervention from peers, experts, and an SNA tool were most useful. We show how the NAV model informs the design of SNA tools and services, education practices, and support for social media practitioners.

HCIL In The Press

Human-Computer Interaction Laboratory
University of Maryland, College Park, MD 20742-3255
www.cs.umd.edu/hcil

HCIL In The Press

May 24, 2010. The new iPhone programming course being taught at the University of Maryland by Evan Golub of the HCIL, Adam Porter from Computer Science, and Chuck Pisula of Apple Inc. was featured in Fortune Magazine, *"Apps 101: The Hot Course on Campus."*

May 5, 2010. The new iPhone programming course being taught at the University of Maryland by Evan Golub of the HCIL, Adam Porter from Computer Science, and Chuck Pisula of Apple Inc. was featured on Fox 5 DC.

May 4, 2010, Ben Bederson was quoted in Technology Review*, "Redesigning the Web for Touch Screens."*

April 2010, Allison Druin and Ben Bederson win the *SIGCHI Social Impact Award*.

April 2010 FutureGov Magazine by Kelly Ng, *Singapore Description of ICDL for iPad application.*

April 17, 2010, The International Children's Digital Library, ABC 7 News, WLS-TV Chicago, "*Chicago Students Face Off in Battle of Books"*

April 16, 2010 Ben Shneiderman quoted in Washington Post on *"Library of Congress plan for Twitter: a big, permanent retweet"*

April 8, 2010, Ben Bederson was interviewed in Publisher's Weekly, *"The iPad Meets the Children's Book"*

April 3, 2010, The International Children's Digital Library was mentioned in The Independent, London*, "Chlidren's book iPad apps offer multilingual tales, long-distance bedtime stories"*

March 22, 2010, Ping Wang's research was featured in the ComputerWorld article *"Fashionable IT Investments Don't Help the Bottom Line."*

February 17, 2010, Ben Shneiderman elected to National Academy of Engineering.

February 9, 2010, Ben Shneiderman received an Honorary Doctorate from University of Castilla-LaMancha in Spain.

January 21, 2010, Tom Yeh's project, Sikuli, was featured on Slashdot.

January 28, 2010, Ben Bederson was quoted in the Baltimore Sun, Baltimore, MD, *"Apple unveils tablet computer"* by Gus G. Sentementes

January 7, 2010, Kent Norman was interviewed on the nationally syndicated show, "Parent's Perspective.*"The topic was "Should We 'Friend' Our Child?"*

December 25, 2009, Allison Druin's research on How Children Search was featured in the New York Times, *"Helping Children Find what they need on the Internet."*

December 2009 FutureGov Magazine by Kelly Ng, Singapore *"Inside the Teacher's Lounge"* covered Ben Bederson's keynote UNESCO talk in Hangzhou, China

November 22, 2009 Ben Shneiderman quoted in Washington Post on *"Augmented reality fuses the Web and the world around you."*

November 20, 2009, Ben Bederson and the International Children's Digital Library were featured in the World Bank Blog.

November 2, 2009, Ben Bederson was interviewed in Technology Review.com, *"First Test for Election Cryptography"*

October 22, 2009 Ben Shneiderman quoted in San Francisco Chronicle on *"An Experiment in 'Creativity of the Crowds'"*

October 21, 2009, Jenn Golbeck was interviewed on BBC Newshour about Twitter use by MPs.

Fall 2009 Ben Shneiderman and Catherine Plaisant's work is presented in the University of Maryland TERP magazine article (pages 22-25) on *"Health IT: An Rx for Health Care."*

September 18, 2009, Justin Grimes and Tony Rogers were interviewed about Jen Golbeck's research in The Washington Post: *Politicians' Tweets Are Mostly Self-Promotional,Researchers Say*. This was picked up by 19 other media outlets.

August 3, 2009 Ben Shneiderman's letter to President Obama is presented in *Government Computer News*

Directions to Demos 2:45pm – 4:00pm
@ HCIL Lab, 2117 Hornbake Bldg, South Wing

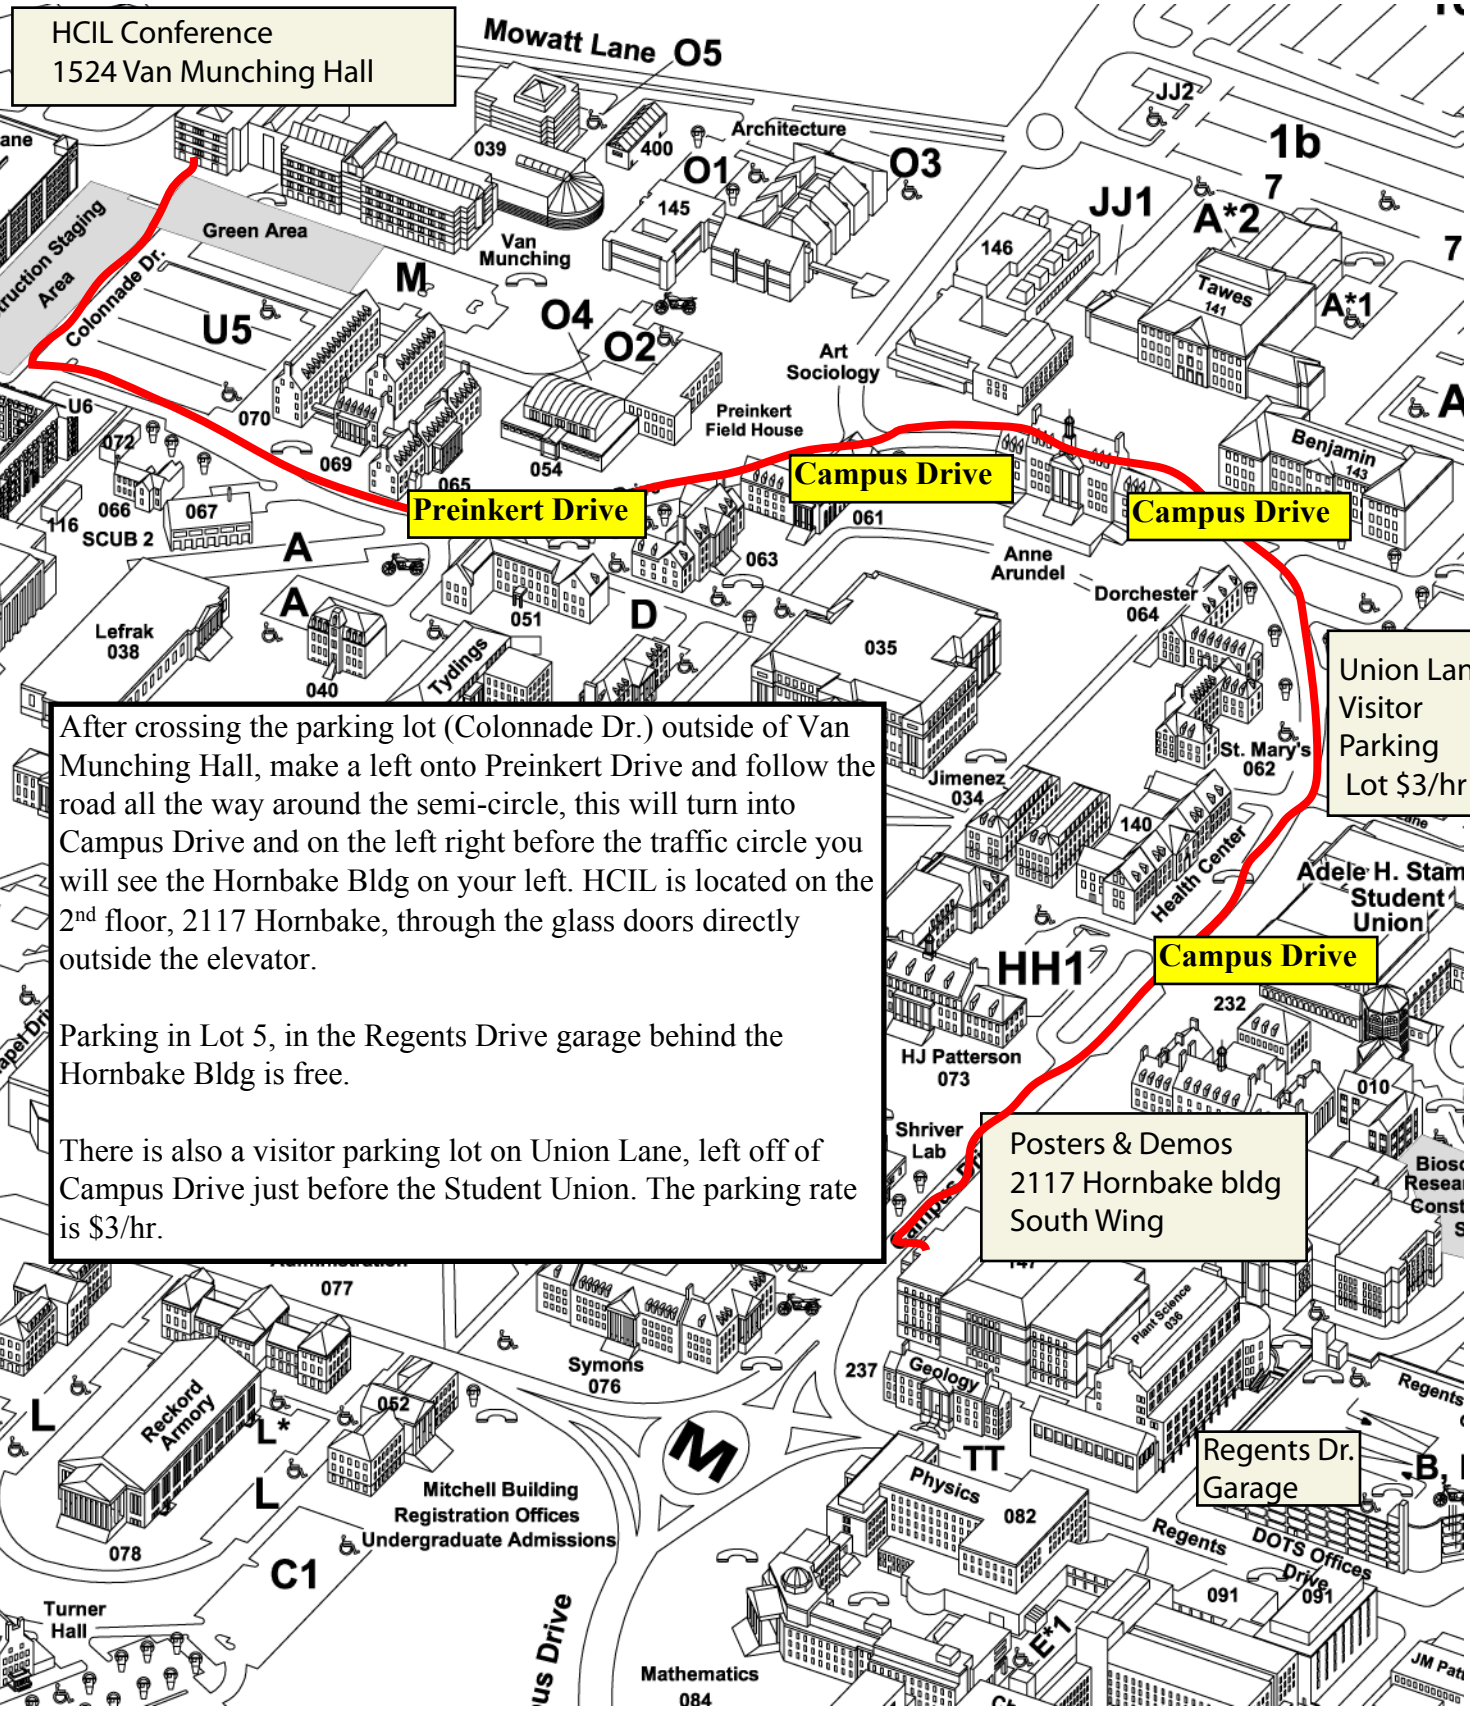

After crossing the parking lot (Colonnade Dr.) outside of Van Munching Hall, make a left onto Preinkert Drive and follow the road all the way around the semi-circle, this will turn into Campus Drive and on the left right before the traffic circle you will see the Hornbake Bldg on your left. HCIL is located on the 2nd floor, 2117 Hornbake, through the glass doors directly outside the elevator.

Parking in Lot 5, in the Regents Drive garage behind the Hornbake Bldg is free.

There is also a visitor parking lot on Union Lane, left off of Campus Drive just before the Student Union. The parking rate is $3/hr.