

On the Intersection of Human-Centered Computing and Web Science

Alejandro Jaimes*

Division of Web Science and Technology, Korea Advanced Institute of Science and Technology, Daejeon, Korea
ajaimes@caa.columbia.edu

Abstract

Human-centered computing (HCC) as a field was defined several years ago. Since then, the Web has grown significantly, and so has its impact in society at every level, creating the grounds for a new Web science. In this paper, we examine the definitions of HCC and Web science, and discuss how they might be related. On the one hand, HCC can be viewed as a set of methodologies, and on the other hand, Web science can be viewed as a platform or repository, from which behavioral data can be drawn. We examine the relationship between the two, and summarize three different research approaches that in many ways show this intersection.

Category: Human computing

Keywords: Computer science and engineering; Human-computer interaction

I. INTRODUCTION

Computing is leading a revolution—everything we do is changing, at a pace never experienced before in human history. Whether we talk about the pervasive, ubiquitous, mobile, grid, or even social computing revolution, we can be sure that computing is impacting the way we interact with each other, the way we design and build our homes and cities, the way we learn, the way we communicate, the way we play, and the way we work. Simply put, computing technologies are increasingly affecting and transforming almost every aspect of our daily lives [1].

In [2], we defined human-centered computing (HCC) as an emerging field that aims at bridging the existing gaps between the various disciplines involved in the design and implementation of computing systems that support people's activities. HCC aims at tightly integrating human sciences (e.g., social and cognitive) and computer science

(e.g., human-computer interaction (HCI), signal processing, machine learning, and ubiquitous computing), for the design of computing systems with a human focus from beginning to end. This focus should consider the personal, social, and cultural contexts in which such systems are deployed [3]. Beyond being a meeting place for existing disciplines, HCC also aims at radically changing computing with new methodologies, to design and build systems that support and enrich people's lives [2].

One important shift since the publication of these papers [1, 2] on HCC is the growth, both in the Web, and in terms of the devices and ecosystems built around them. In particular, the expansion in the use of the tablet and other mobile devices is significantly changing the way we use the Web. At the same time, applications built on the Web, for social networking, photo sharing, news consumption, and in many other domains, continue to proliferate. Search, in spite of being fairly "old" as a technology

Open Access <http://dx.doi.org/10.5626/JCSE.2013.7.4.211>

<http://jcse.kiise.org>

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/3.0/>) which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

Received 7 June 2013, Accepted 1 July 2013

*Corresponding Author

to access the Web, has continued to grow, and to play a crucial role in people’s daily activities.

At the same time, the amounts of data generated and collected by a plethora of Web services, have in part facilitated new research in Web science, “a term that refers to processing the information available on the Web in similar terms to those applied to the natural environment” [4].

Key questions include the relationship between HCC and Web science, what the research process might look like, and what types of applications could emerge from the intersection of HCC and Web science—if they are to be considered as separate disciplines. Web science is fairly new, boosted mainly not just by significant growth in the use of the Web, but also due to the availability of technologies to efficiently process massive quantities of data. One could argue that while Web science is all about embracing big data, HCC is still in its infancy, in terms of how large quantities of data are used.

In this paper, we first examine the definitions of HCC, and how they relate to Web science. Then we focus on three examples of published research, which potentially exemplify work at the intersection of Web science and HCC. Finally, we discuss what the HCC-Web science research process might consist of, and what the challenges are, moving forward.

The rest of the paper is organized as follows. Sections II and III focus on definitions of HCC and Web science. In Section IV, we present three research examples, focusing on data analysis, algorithms, and an application. We conclude with a section of observations, and discuss future work.

II. HUMAN-CENTERED COMPUTING: DEFINITIONS

In the last few years, many definitions of HCC have emerged. In general, the term HCC is used as an umbrella to embrace a number of definitions that were intended to express a particular focus or perspective [1]. In 1997, the United States National Science Foundation supported a workshop on human-centered systems [5], which included position papers from 51 researchers from various disciplines, and application areas that included electronics, psychology, medicine, and the military. Participants proposed various definitions for HCC, including the following [2]:

- HCC is “a philosophical-humanistic position regarding the ethics and aesthetics of the workplace.”
- An HCC system is “any system that enhances human performance.”
- An HCC system is “any system that plays any kind of role in mediating human interactions.”
- HCC is “a software design process that results in interfaces that are really user-friendly.”

- HCC is “a description of what makes for a good tool—the computer does all the adapting.”
- HCC is “an emerging inter-discipline requiring institutionalization and special training programs.”

Other definitions of HCC have also appeared in the literature [1]:

- According to [6], HCC is “the science of designing computations and computational artifacts in support of human endeavors.”
- For Canny [3], HCC is “a vision for computing research that integrates technical studies with the broad implications of computing in a task-directed way. HCC spans computer science and several engineering disciplines, cognitive science, economics and social sciences.”

The field has not necessarily changed much in recent years, although arguably a lot of research has taken on a stronger “HCC” flavor. On the one hand, recent years have seen the success of various Web products and apps that excel at being simple, intuitive, easy to use, and that result in explosive growth. This has been largely spearheaded by the availability of tablet devices and smart phones.

The Web, on the other hand, has grown to be used as a serious and credible research platform in its own right, with data providing insights into innumerable aspects of human behavior, and as a reflection of society.

In the next section we briefly discuss how HCC and Web science might relate.

III. HCC AND WEB SCIENCE

One possible view of Web science is that it is strongly based on the web as a “platform”, on which research can be performed on many human-related issues. As shown




Fig. 1. An illustration of the components of Web science (Creative Commons attribution, Steffen Staab; <http://west.uni-koblenz.de>).




Fig. 2. Human-centered computing can be viewed as a set of methodologies that falls in the intersection of data analysis, human-computer interaction (HCI), and human issues.




Fig. 3. The human-centered computing-Web science research process can have various starting points, but must have human abilities and needs at its center, and consist of one or more “loops”, as indicated in the figure.

in Fig. 1, the Web can thus be interpreted as a large repository of observations on human behavior, both individual and collective.

Since HCC, on the other hand, is defined as a set of methodologies, one could argue that both can easily go hand-in-hand. HCC can be viewed as a set of methodologies that is applicable specifically *on* the web, to inform and guide Web science research. At the same time, HCC can be viewed as illustrated in Fig. 2, as that set of methodologies which focus on the intersection of data analysis (i.e., “big data”), HCI, and what can be roughly referred to as “Human Issues”, which include cognitive, psychological, social, and other aspects.

In terms of the research process itself, one could argue that HCC, as well as Web science, should be “human” focused. That is, all of the research, development, and deployment should center on human needs and abilities. A key question is, however, what should the research process be? Where should the research process start, and what should the necessary steps be?

In the next section we present three examples of research in HCC that fall into this paradigm. First, we present work on the analysis of a large-scale data set.

Then, we present a study of various ways of ranking, which roughly corresponds to starting the work with a set of given hypotheses. We then present work that focuses on design aspects (in Flickr) (Fig. 3).

IV. EXAMPLES

A. Data Analysis in Web Search

Web search is a multi-billion dollar industry; and yet one of the biggest challenges is understanding the general characteristics of Web searchers, in terms of their demographic profiles, what they search for, and how they search [7]. Unlike in traditional businesses, where customers come face-to-face with vendors at some point of the production and sales process, the Web is unique, because it can be used anonymously, and search portals in particular act like gateways to a plethora of information and services across the globe.

As described in [7], analysis of user logs has therefore become an important research area, because it forms the basis for many business decisions. The results can be used for search engine optimization strategies, for improving user experience, for advertising, and for several other purposes. One area of particular interest has been the analysis of large-scale search query logs, in order to understand user intent, topics being searched for, and the quality of search results, among others.

In this section, we present the main findings of [7], as an example of an HCC process that starts with data analysis.

1) Data Set

The main data source for the study reported in [7] was a large sample of the Web search query log of the Yahoo! search engine between 2008 and 2009. Queries were cast to lower case, whitespaces at either end were removed, and all other consecutive sequences of whitespaces were replaced by a single “space” character. No other normalization, such as the removal of stop words or punctuation, was performed. Our sample only included queries of registered U.S. Yahoo! users with an identifiable cookie, who had provided demographic information (gender, birth year, country, and ZIP code) upon registration. Furthermore, we only used the log for the U.S. Yahoo! site, and only for users whose country in the self-provided information was the United States. We automatically removed users who claimed to have ZIP codes that did not exist, or who claimed to live in an area with 0 population (as of the United States Census 2000). Finally, our sample only included “active” users, that is, those who had issued at least 100 queries over the sample period. To filter out additional bots, on top of a proprietary bot filtering algorithm, we also removed users with more than 100,000 queries, and users who, on average, had clicked on organic results for fewer than 1/100 of their queries, or

more than 100 times per query. Our final sample thus contained queries for 2.3 million registered users, for which no personal information was used. All analysis was anonymous, and performed in aggregate.

2) Analysis

The results reported in [7] consist of analyzing the behavior of Web searchers along the following three orthogonal dimensions:

a) Query topics (“what?”): What are the topics that the user issues queries on?

In our study, we used the Yahoo! Directory classification as a definition of a “topic”, but other topic definitions could be used; or, instead of looking at topics, one could consider the fraction of queries having a local focus, or containing product names.

b) User demographics (“who?”): What is the demographic profile of the user?

In our study, we used a mix of user-provided information (age and gender) and information derived from the user’s ZIP code (expected income, expected educational level, expected political party affiliation). If available, other features, such as their actual educational level, their marital status, or even their body mass index, could be included.

c) Session characteristics (“how?”): Does a user have many/few, short/long, and navigational/informational sessions?

For the study in [7], we limited ourselves to basic measures, such as session length, number of queries per session or the fraction of sessions with suggested/guided queries. We also looked at the fraction of queries with a very high/low click entropy. If available, one could include features about the session’s success, or the user’s frustration level.

3) Main Findings

The analysis reported in [7] consists of analyzing the behavior of web searchers, and in this sub-section we summarize the main findings of [7] on a few selected clusters, and describe them.

Cluster names are deliberately oversimplified, to emphasize the relative differences between the clusters.

a) Baby boomers (Cluster ID 1): Users in this cluster tend to be older than the typical Web searcher, with an average age of 50 years. Their predominant topic of interest is finance, and a large fraction of their sessions consist of simple navigational queries related to online banking.

b) Adult content seekers (Cluster ID 4): A large fraction of these mostly male users’ queries are of a transactional kind, involving adult content. They are slightly older than the average user, and are often “satisfied” with a single click in a session.

c) Liberal females (Cluster ID 8): The typical user in this group is female, and more likely to have voted Democrat in the 2008 elections. The biggest single query topic

is shopping, and sessions are comparatively long, hinting at possible browsing and comparison behavior.

d) White conservatives (Cluster ID 12): Users in this group are more likely to be white, and live in areas that voted Republican in the 2008 elections. Mostly male, they often search for automotive related topics, business pages, and, compared to the average, relatively often for home and garden information.

e) Challenged youth (Cluster ID 14): These users are comparatively young, with an average age of 34. They tend to live in low-income neighborhoods with a low level of education. Their searches are centered around music, and their sessions are often of a navigational kind.

In [7] we presented an analysis of a large-scale search query log of anonymous registered users from a major web search engine. We jointly studied on-line behavior in terms of *who* these users might be (demographics), *what* they searched for (query topics), and *how* they searched (session analysis). We used information provided by users during registration (age, sex, zip code), along with data obtained from the United States Census, and classified queries into topic categories, analyzed basic session statistics, and clustered users based on the topics of the queries they issued. Our analysis showed that by examining the resulting clusters, it is possible to identify distinct patterns of behavior along different demographic features. We were also able to classify query topics into informational, navigational, and transactional categories based on click entropy, and upon examining the clusters obtained, we were able to suggest different types of searchers (e.g., baby boomers, white conservatives, etc.), in terms of the topics they search for, and their behavior.

Overall, the findings are “stereotypical” [7], with men more interested in adult content than women, with health topics attracting more of a research behavior (rather than navigational behavior), and with people with higher educational levels issuing fewer navigational queries (which are arguably more efficiently typed in the browser’s address bar directly), etc.

B. Hypotheses

Many social media platforms function as somewhat independent ecosystems, in which users carry out a number of social activities, and in this section we briefly summarize the work presented in [8].

In Flickr, in particular, users can share content, and participate in multiple communities, by submitting their photos to groups, by joining groups, and by performing several types of actions over Flickr content (e.g., comment, add notes, favorite, etc.). The result of this is that the way the content is consumed is strongly influenced by all of the different social navigation paths that lead to it: a photo on Flickr, for example, can be linked to from a user’s favorite photo collection, from several groups, gal-

leries, and via other mechanisms, including the “external” Web (i.e., URLs outside of the Flickr domain).

As more social media platforms emerge, one of the key questions is whether traditional ranking algorithms, which may not take into account the subtleties of navigation patterns driven by social connections, can be successful within those ecosystems. In particular, the problem we are interested in addressing is the general ranking of images in Flickr (i.e., as in PageRank, we would like to rank all of the images, or a subset of them, in order of importance). Such ranking can have many applications, including retrieval, and information discovery, among others.

The importance of images in Flickr, or of “nodes” in similar social media platforms might depend on a number of internal and external factors. For example, an image that is very popular in a group that has a cult following, may have been marked by many users as a favorite image. The image has a large number of favorites, because people within the Flickr community, and in particular, in the specific cult, view and favorite the image. In contrast, an image of an important real world event (e.g., the British Royal Wedding) may get a high number of views, not by belonging to groups, but instead because it is linked to by multiple external (i.e., outside of the Flickr domain) media outlets, and get comparatively few favorite marks. One of the key questions is thus, what is the impact of those external and internal factors on the ranking and selection of content?

1) Data Set

For the purpose of the study reported in [8], we took a sample of the pageviews of more than 10 million anonymous users, from approximately two months of Flickr user log data, from August to October 2011. The pageviews are represented as plain text files that contain a line for each HTTP request satisfied by the Web server. For each pageview, our dataset contains the following fields: *UserId*, *Time*, *ReferrerURL*, *CurrentURL*, *User Agent*. The *UserId* is a unique anonymized identifier computed from the Flickr *userId*, in the case of logged-in users, and from a browser cookie, otherwise. *CurrentURL* and *ReferrerURL* represent the current page the user is visiting, and the page the user visited before arriving at the destination page. The *User Agent* identifies the browser in use, and the timestamp indicates when the page was visited. All of the data processing was anonymous, and performed in aggregate. Flickr allows users to set specific pages to “private”, but in our analysis, we considered only public pages.

2) Analysis

When comparing different picture sets, image quality is just one of the parameters. In particular, when images are embedded in dynamic social environments, the interest people have in particular photos can be determined (or

influenced) more by the social dynamics of a community (e.g., a group in Flickr), than by the inherent quality of the photos themselves. Similarly, interest can originate externally (i.e., many photos in Flickr are linked to from outside of Flickr), and thus be important, independently of their aesthetic qualities (e.g., photos of important events).

Given that several factors can be taken into account in considering a ranking of images, we identified four importance notions, and we list some quantitative features for each of them. All of the features were then used as evaluation parameters, to compare the rankings.

a) Internal popularity: Popularity of a photo inside the Flickr community. Popularity does not necessarily imply quality, but directly expresses the interest of users in a particular item. Features describing photo popularity are the number of comments the picture receives, and the number of internal Flickr groups in which it appears.

b) External popularity: We considered measures of external popularity: the number of search results obtained from a Google search using the photo page URL as a query, the Google PageRank of that URL, and the number of browsing sessions originating from an external URL that visit the photo page as the first Flickr page.

c) Collective attention: Users not logged into Flickr, as well as Flickr users who do not actively give feedback on photos, implicitly express their interest in specific photos by visiting the pages that contain them, and by spending time on them. Therefore, we use the total number of views of a photo, and the cumulative time spent on the photo as an aggregate measure of attention that a generic Web user, whether or not logged into Flickr, devotes to that image.

d) Diversity: One of the applications of ranking a large set of photos might be to display the most interesting ones. In this case, a very homogeneous set of pictures may result in appealing to some user categories, but are less likely to attract a wide public. Assuming that photos belonging to the same user are on average more homogeneous than pictures taken from different users, diversity can be estimated by the number of different photo owners. Additionally, an analysis on the diversity of the corpus of tags of the photos can be a measure of the variety of concepts represented.

In [8], we used the number of views and the view time as both ranking metrics and evaluation parameters, to draw a more complete analysis of other rankings. We could have done the same for the number of comments, but we omitted it as a ranking metric, because its performance was very similar to the favorites. Since many applications need much shorter ranked lists, we reported that the relative position of the different curves is nearly unchanged for the top 20 and top 100 photos, for all the metrics considered.

Results reveal that most favorites have good internal

popularity, being the top metric in both number of groups and comments, but behave worse than any other metric, in terms of external popularity and collective attention. In contrast, photos with top BrowseRank scores are less popular internally (even though their scores are comparable to favorites up to the top 100), but they attract relatively more collective attention, and position above any other metric, when counting external relevance and owner diversity. PageRank behaves worse than BrowseRank, except for collective attention.

Finally, Pageviews and Time perform reasonably well for external popularity, and by definition in collective attention, but surprisingly, the ranked photos have relatively little popularity in groups, and receive few comments.

3) Main Findings

The overall scenario emerging from the comparison shows that different metrics promote different types of photos. Rankings based on explicit feedback, namely favorites, boost pictures that are well spread across Flickr groups, and that receive attention from active Flickr users, but that may not have great impact outside of Flickr. Top rated images tend to belong to a small set of owners, and convey a lower semantic variety, than the pictures from centrality-based rankings.

Artistic photos made by professionals are prevalent. BrowseRank and PageRank, instead, overshadow a bit the very popular content inside Flickr, to provide images with a higher semantic variety, and with apparently stronger interest from a broader part of the Web (outside of Flickr). This includes popular photos on trendy social events, or pictures about popular fun facts, or peculiar subjects. A positive side effect of this is that photos that are related to popular memes just inside Flickr (e.g., horror house pictures) are downgraded, and tend to disappear from the top ranking.

Moreover, being based on the data from the navigation log only, centrality rankings are fully implicit. They do not need an active user base commenting or voting on the images.

This means that BrowseRank and PageRank are effectively more able to pick up diverse image collections, and produce more balanced lists, by considering external links to the photos.

Such algorithms can be profitably parallelized, making efficient their computation, even for big social media sites like Flickr.

C. Design

Social media platforms such as Flickr provide a wide range of functionalities, and different ways to share and view content. Given the sequential nature of browsing photographs, it is common, however, for people to share and view images in sequences, whether the photos are

arranged in galleries, slide shows, or in groups. In this section we summarize the work presented in [9] around this concept.

In Flickr, in particular, photos uploaded by a user to his or her account are placed in a “photostream”, which is, in essence, a sequence of photos. Although there are many ways to reach individual photographs, such sequences constitute a fundamental part of the interaction.

In [9] we referred to such sequences as photostreams (or simply *streams*). Furthermore, navigation across sequential units of content is present in other fields of social media, e.g., social networks, music streaming and microblogging platforms. In popular social networks photos are organized into albums, and can be viewed sequentially. Songs in music streaming services can be listened to one after another, usually as a part of an album, or a playlist. Posts in microblogging platforms are chronologically organized in independent blogs. Therefore, methods developed for photostreams could be adapted to other social media as well.

A key question for social media platforms, then, is how users navigate inside, and between various photostreams. In particular, such photostreams may be considered not just as collections of images, but rather as fundamental units of content. On the one hand, understanding how users navigate between specific photostreams is crucial in designing interfaces and algorithms that improve user experience, by providing the right content in the right places. On the other hand, analyzing the semantic categories of such streams can also provide important insights on general topics of interest. In addition, investigating how users transition between photostreams allows us to understand how topics may be related.

In [9], we treated photostreams as content units, and analyzed a large sample of navigation logs, to gain insights into how users navigate between different photostreams. More specifically, we examined user navigation logs containing several million pageviews, in order to create a photostream transition graph, to analyze frequent topic transitions (e.g., users often view “train”, followed by “fire truck” photostreams).

As described in [9], we implemented two photostream recommender systems: a collaborative filtering approach, based on transitions between photostreams, and a content-based method, using tag-similarity of photostreams. Finally, we reported the results of a user study involving 40 participants, to explore the fundamentals for creation of an effective recommender system in a large social media platform.

1) Data Set

For the purpose of the study reported in [9], we took a sample of the pageviews of more than 10 million anonymous users from 2011. Since Flickr allows users to set specific pages to be private, in our analysis we considered only public pages. All of the data processing was

anonymous, and performed in aggregate.

2) Analysis

Photos in Flickr are organized into photostreams. Each photo in Flickr belongs to a photostream of the owner, but it can belong to other streams of photos as well: groups, sets, galleries, and favorites. Apart from favorites, all of these photostreams are either chosen, or created by the owner of the photo. Users always view and browse photos in the context of a particular photostream.

There are two main ways of viewing photostreams: (a) grid view, i.e., grid of photos from the stream and (b) photo-focused view, i.e., a single zoomed-in photo with a possibility of browsing neighboring photos.

Although Flickr allows different variations of grid views, they share a common feature, namely that they show several pictures from the browsed stream at a glance. The photo focused view is the same for all the streams: it shows a large selected photo, and on the right side of it, thumbnails of 4 neighboring photos from the stream are presented, which the user can switch to by clicking on them. This way, one can change the focus from the current photo, to another one from the currently browsed stream. Below the thumbnails, a list of all photostreams that the photo belongs to is shown in the form of hyperlinks.

One can expect that users first enter the grid view of a photostream, and then select one of the photos they like, and see it in a photo-focused view. Then, they can continue on, browsing other photos from this photostream. The grid view may be used for purposes which seem less involving to the user, e.g., quick browsing many photos from a stream, having an overview of a stream, or seeking interesting content.

Photo-focused views give the user options of performing many different actions in reference to the photo, e.g., he or she can comment on the photo, favorite it, download it, see it at different sizes, or in a light-box setting. For the purpose of the study, we define a stream-browsing sequence as an uninterrupted chronological sequence of pageviews that contains at least one photo-focused view, and an indefinite number of grid views of one particular photostream. Each browsing session can consist of a number of stream browsing sequences. The Flickr logs in our dataset contain a total of 264 million pageviews, out of which a considerable part form stream-browsing sequences. Among the sequence of pageviews, there are photo-focused views, and grid views of the photostream. Distributions of both the number of distinct streams viewed per session, and the number of photo-focused views per stream, have a heavy-tail, showing high variability in user browsing patterns.

3) Main Findings

Since clusters contain streams of similar topic, an interesting question to ask is: which clusters do people

switch between most often? This can be answered by creating a node in place of every cluster of streams, and aggregating edges of all streams belonging to this cluster. In this manner, in [9] we obtained a directed and weighted network of transitions between clusters from the network of transitions between streams. After the conversion, there are self-loops in such a network, which we removed. This network is dominated, however, by the connections between large nodes. In order to account for the size effect of the nodes, and to extract meaningful information about relations between clusters, we took the following approach. If connections between clusters were spread randomly between nodes of known degrees, then $l_{rand,ij}$ would be expected to be the number of edges between particular nodes. To see which connections between clusters are the furthest from random configuration, we calculated the actual number of connections lij and expected $l_{rand,ij}$. We called aij the abundance ratio. If this ratio is larger than 1, then transitions from stream i to j are over represented, while if it is lower than 1, then they are underrepresented.

We picked the parts of the network formed by edges with abundance ratio aij higher than 10, and actual weight lij also higher than 10. Here we provide a short description of each of the examples presented in [9] (details can be found in [9]):

a) Clusters of fans of cars and machinery: the first cluster seems to be on the verge of cars and photography, while the next one is more narrowly about cars, especially classic ones. Users from this cluster tend to switch between both, to see photos of trains and railroads, as well as fire trucks.

b) Event-orientated clusters: photography of live music shows is related to the cluster of journaling, blogging and fisheye photography.

c) Household-centered clusters: clusters of cakes and vintage style, which incorporate elements from previous eras into modern fashion and style, are related to the cluster of sewing and fabrics, to dolls, and then handmade, are related. Dolls and Disney/Disneyland are also related.

d) Toys and military: photography of Lego constructions, mostly of Star Wars, is related to army and military photography. This is quite interesting, and shows an interest from toys and plastic soldiers, to real ones. The military cluster is related to natural disasters, in which the army and powerful natural forces are often involved.

It is also possible to find underrepresented connections between clusters, and it would certainly be interesting to examine more clusters in detail.

V. DISCUSSION

In this paper we discussed the definitions of HCC and Web science, and examined how the two might be inter-related.

In particular, we examined one possible definition of the HCC process, in terms of research that parts either from data, from hypotheses, or from design, and “loops” around the user within a particular context (Fig. 3). We provided three examples of such research.

The first one focused on large-scale data analysis of search logs, in order to gain insights into what people search for, and how they search for it. The work occurs within a given context, as exemplified by the uses of census data, which is very specific to a particular country or location. In this example, thus, we show how starting from large-scale data, we move back and forth between data analysis and hypotheses, guided by a particular context given by the Census data.

In the second example, we focused our starting point on hypotheses, in particular, on the meaning of different ranking algorithms in Flickr. From there we “jumped” to data analysis, and examined how different ranking algorithms may impact what users see, and jumped back to hypotheses, in examining what those rankings (and actions) might mean.

Finally, we started the last research example by analyzing the design of most social photo sharing platforms, in which the paradigm for photo viewing consists of a stream. In that paper, the work went further, in not only performing a large-scale data analysis, but in also designing a new interface, and testing hypotheses about the implications of such interface. Although the work discussed here focused on discussing viewing patterns, and some sample clusters were presented, we note that the work parted from the observation that in the Flickr design, where the work was carried out, photos were viewed in sequences (“streams”).

V. CONCLUSIONS AND FUTURE WORK

This paper presents a first unified view of HCC and Web science, using three examples of published research, to explore how they might relate to the research process. In all three cases, it is possible to identify the starting points, and in general terms, the research process that was carried out, in which a “human-centered” view was adopted.

Although these works can be used as examples, a lot more has to be done, to formalize the process. On the one hand, the definitions need to be narrowed, and further detailed. On the other hand, the process itself needs to be further specified, and in the case of these examples, the work would ideally be followed up on, to result in either actionable insights, in the case of data, or in actual

human-centered applications. One possible way to view this is to think of the research process as a meta-process: can we abstract different research steps in each case, and identify commonalities and differences?

ACKNOWLEDGMENTS

The author would like to thank the various co-authors and collaborators who have contributed to the definition of the field of human-centered computing, and the co-authors of the papers described in this article. The author was supported in part by the World Class University program, under the National Research Foundation of Korea, and funded by the Ministry of Education, Science and Technology of Korea (Project No. R31-30007).

REFERENCES

1. A. Jaimes, D. Gatica-Perez, N. Sebe, and T. S. Huang, “Human-centered computing: toward a human revolution,” *Computer*, vol. 40, no. 5, pp. 30-34, 2007.
2. A. Jaimes, N. Sebe, and D. Gatica-Perez, “Human-centered computing: a multimedia perspective,” in *Proceedings of the 14th Annual ACM International Conference on Multimedia*, Santa Barbara, CA, 2006, pp. 855-864.
3. J. Canny, “Human-center computing retreat overview,” presented at the fifth US Berkeley HCC Retreat, Berkeley, CA, Jul 5-7, 2001.
4. B. Shneiderman, “Web science: a provocative invitation to computer science,” *Communications of the ACM*, vol. 50, no. 6, pp. 25-27, 2007.
5. R. R. Hoffman, P. J. Feltoovich, K. M. Ford, D. D. Woods, G. Klein, A. Feltoovich, “A rose by any other name ... would probably be given an acronym,” *IEEE Intelligent Systems*, vol. 17, no. 4, pp. 72-80, 2002.
6. Georgia Institute of Technology, “HCC educational digital library,” <http://hccedl.cc.gatech.edu/>.
7. I. Weber and A. Jaimes, “Who uses Web search for what? and how?” in *Proceedings of the 4th ACM International Conference on Web Search and Data Mining*, Hong Kong, China, 2011, pp. 15-24.
8. M. Trevisiol, L. Chiarandini, L. M. Aiello, and A. Jaimes, “Image ranking based on user browsing behavior,” in *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Portland, OR, 2012, pp. 445-454.
9. L. Chiarandini, P. A. Grabowicz, M. Trevisiol, A. Jaimes, “Leveraging browsing patterns for topic discovery and photostream recommendation,” in *Proceedings of the 7th International AAAI Conference on Weblogs and Social Media*, Cambridge, MA, 2013, pp. 71-80.



Alejandro Jaimes

Dr. Alejandro Jaimes was a visiting professor at KAIST under the WCU program until August 31st, 2013, and is Director of Research at Yahoo! Labs, where he is leading new initiatives at the intersection of Web-scale data analysis and user understanding (user engagement and improving user experience). He manages the Social Media Engagement (SOMER) and Learning for Multimedia and Vision (LMV) groups in Barcelona and Bangalore, respectively. Dr. Jaimes is the general chair of ACM Multimedia 2013, Practice and Experience track chair for WWW 2013, the founder of the ACM Multimedia Interactive Art program, Industry Track chair for ACM RecSys 2010 and UMAP 2009, and panels chair for KDD 2009, among others. His work has led to over 80 technical publications in international conferences and journals, and to numerous contributions to MPEG-7. He has been an invited speaker at Practitioner Web Analytics 2010, CIVR 2010, ECML-PKDD 2010 and KDD 2009 and (Industry tracks), ACM Recommender Systems 2008 (panel), DAGM 2008 (keynote), 2007 ICCV Workshop on HCI, and several others. Before joining Yahoo!, Dr. Jaimes was a visiting professor at Universidad Carlos III de Madrid, and founded and managed the User Modeling and Data Mining group at Telefónica Research. Prior to that Dr. Jaimes was Scientific Manager at IDIAP-EPFL (Switzerland), and was previously at Fuji Xerox (Japan), IBM TJ Watson (USA), IBM Tokyo Research Laboratory (Japan), Siemens Corporate Research (USA), and AT&T Bell Laboratories (USA). Dr. Jaimes received his PhD in Electrical Engineering (2003), and his MS in Computer Science from Columbia University (1997) in New York City.