Community, Impact and Credit: Where should I submit my papers?

Aaron Halfaker

GroupLens Research University of Minnesota halfaker@cs.umn.edu

Cliff Lampe

School of Information Michigan State University lampecli@msu.edu

Amy Bruckman

College of Computing Georgia Inst. of Technology asb@cc.gatech.edu

Aniket Kittur

HCI Institute Carnegie Mellon University nkittur@cs.cmu.edu

R. Stuart Geiger

School of Information University of CA, Berkeley stuart@stuartgeiger.com

Loren Terveen GroupLens Research University of Minnesota terveen@cs.umn.edu

Brian Keegan Northeastern University bkeegan@northeastern.edu

Geraldine Fitzpatrick

Vienna University of Technology geraldine.fitzpatrick@tuwien.ac.at

Abstract

We (the authors of CSCWs program) have finite time and energy that can be invested into our publications and the research communities we value. While we want our work to have the most impact possible, we also want to grow and support productive research communities within which to have this impact. This panel discussion explores the costs and benefits of submitting papers to various tiers of conferences and journals surrounding CSCW and reflects on the value of investing hours into building up a research community.

Author Keywords

community; credit; impact; publishing; peer review

ACM Classification Keywords

H.5.0. [Information Interfaces and Presentation (e.g. HCI)]: General

Introduction

We (the authors of CSCWs program) have finite time and energy that can be invested into our publications and the research communities we value. In order to allow our work to have an impact, we must also grow and maintain productive research communities within which to share our work. This panel discussion explores the costs and benefits of submitting papers to various tiers of

Copyright is held by the author/owner(s).

CSCW '13 Companion, Feb. 23–27, 2013, San Antonio, Texas, USA. ACM 978-1-4503-1332-2/13/02.

Format:

- The panel presentation will start with a short introduction by the moderators.
- Each panelist will follow with a presentation of their perspective on an aspect of the panels themes.
- The moderators will guide a discussion between the audience and panelists.

conferences and journals surrounding CSCW and reflects on the value of investing hours into building up a research community. By comparing CSCW to restrictive, top tier but general conferences like CHI and less restrictive, specialized conferences like WikiSym and RecSys, the panelists will frame a discussion around the value of participation in different types of research communities and how decisions on where to invest ones time affect metrics used to evaluate our worth.

This discussion is valuable because it provides an opportunity for experienced researchers to advise young researchers on which communities to invest their time. It will also provide an opportunity to address several trends relevant to the CSCW community:

- The merging role of top conferences and journals
- The value of small, specialized conferences
- Inconsistencies between impact and impact metrics

To focus the discussion, panelists will address three themes of publication: community, impact and credit. Each of these themes will be discussed both in terms of their current state and where each panelist thinks they ought to be.

Themes

Community

How does the structure of research publication affect the kinds of communities we form?

Conferences are designed to bring people to the same physical space. This allows members of a research community to meet face-to-face to discuss their work, and theoretically, decreases the activation energy necessary for forming new collaborations. But the way that conferences accept and reject papers can have an effect on the type of communities that are formed.

The first-round rejection process for most conference proceedings encourages rejected submissions to be resubmitted to a different conference – if not, we must wait almost a year before resubmission. A common but often unacknowledged practice is to therefore to resubmit rejected papers to slightly or completely different conferences. Substantive comments about the quality of the paper obviously ought to be addressed in revisions, but how much do we really think about revising the same paper for a different conference?

If our conferences are, as is commonly claimed, journals that meet once a year, what implications does the debate over revise and resubmit versus first-round rejection have for the boundaries and bridges between research communities? If we are more like a journal and encourage revisions to be resubmitted, there is an argument that this would make authors target their submissions towards a particular venue. Yet is this a good thing or a bad thing? Does first-round rejection encourage authors to form communities which are not coextensive with conferences and are instead more ad-hoc around other areas of commonality? Or does this effectively silo similar kinds of research into areas that do not interact with others in the same conference (Wikipedia research at CSCW, for example).

On the one hand, we can think of communities as existing around a single conference that extends outward (e.g. the CSCW, CHI, RecSys, or ICTD community), but we can also think of communities as existing around other common areas, such as common methodologies, theoretical approaches, and specific applications?

- How does the structure of publication affect the kinds of communities that we form?
- How are the boundaries between research communities be drawn?
- Does the community have a conference or does a conference form a community?

Impact

What is impact and what types of impact can researchers have?

When we talk about impact, we usually mean moving the field forward via publications of important research. However, the impact of a publication is not solely dependant on the quality of the research and manuscript, but also the audience to whom it is presented. In the ecosphere of conferences and journals, a line might be drawn between small conferences with a limited focus and large conferences with a broader focus.

Submissions to smaller conferences may have fewer potential readers, they also offer direct address to a tight community which is more likely to respond productively to the submission. Conversely, papers published to large conferences and journals will usually see a larger audience of potential readers but a smaller proportion of readers are likely to find the research useful.

Further, not all impact happens via paper publications. Service work like reviewing papers and being part of the program committee is necessary for a functioning peer review system. Processes like revise and resubmit both increase the required investment required by reviewers but also allow for reviewers to have a more substantial impact on the quality of a program.

- Where should researchers direct their effort?
 - ... for the benefit of their work?
 - ... for the good of science?

Credit

How do we measure the worth of a researcher or a publication venue?

For better or worse, hiring and tenure committees often use impact metrics to judge the worthiness of applicants. The lack of an impressive citation count can dramatically affect a researchers career. In this way, metrics dont just describe the pattern of publishing, but direct it as well.

The quality of a conferences are also commonly evaluated by a simple metric: acceptance rate. While acceptance rate tends be an effective measure of the restrictiveness of conferences and journals that use similar review patterns, the metric ceases to be useful when comparing a conference with a revise and resubmit pattern to one with first round rejections only. Assuming that a revise and resubmit process is beneficial to the research community, how does one weigh such benefits against the cost of a higher acceptance rate?

But these measurements of impact do not capture all meaningful contributions. As was mentioned above, not all ways that a researcher can have impact in a field result in a paper with a lot of citations. When the measurement does not capture the value of a contribution, it is likely to misdirect our efforts towards less worthy end. In this case, are we obligated to ignore the metric.

What are some contributions that we make, but do not get credit for? Is it important that reviewers receive credit for reviewing papers? How might we change our measurements of impact to direct researchers better?

Panelists

Loren Terveen

Bio Loren Terveen is a professor in the Department of Computer Science and Enginereing at the University of Minnesota. He specializes in human-computer interaction and computer-mediated communication. His academic record includes about 50 refereed journal and conference papers, one book, four book chapters, nine U.S. Patents, in addition to numerous other publications.

Position Over the past 20 years, Computer Science researchers fought successfully to convince hiring and promotion committees that publications in selective conferences should be valued as highly as journal publications. The most selective conferences, like CSCW, CHI, and UIST are considered first rate, while conferences with acceptance rates above 30% or so are considered less prestigious, and publications there are generally not valued as highly. Thus, authors generally try to publish their work in the most prestigious conferences.

Of course, there are clear problems with this approach, as anyone who has submitted a few papers, served on a few program committees, and attended a few conferences knows. For example: reviewers often disagree significantly about the quality and value of submissions; partly due to this, any conference review process ends up with a large pool of submissions "somewhere in the middle", and the final decision to accept some and reject others feels arbitrary; authors sometimes find their favorite papers do not make it in to the top conferences, but when published at less prestigious conferences receive an enthusiastic response. Attending less selective conferences may expose one to more relevant and interesting work and offer much greater networking and community building opportunities, and high rejection rates can lead individuals and whole sub-communities to stop participating in a conference.

In short, lots of people think the current publication culture is broken. Within the past few years, a number of conferences have experimented with alternative models, like CSCW's 2-phase-with-a-full-revision-cycle approach. However, these alternatives raise challenges of their own, both logistical (do they scale?) and evaluative (if the acceptance rate increases, how do we convince people that our conference is still high quality?). I will talk briefly about some of the alternative models being considered and their pros and cons.

Cliff Lampe

Bio Cliff Lampe is an Assistant Professor in the School of Information at the University of Michigan. He researches the social processes that underly collective action in social media and social computing systems, and how features of technical systems can affect those underlying processes. In that research, he has studied sites like Facebook, Wikipedia, and Slashdot, among others.

Position In our field, we examine socio-technical systems that are a complex combination of social processes and technical systems, all across a highly heterogeneous set of topics. We employ theories from multiple fields and methods ranging from building systems, to qualitative, contextual methods, to quantitative methods like

experiments and surveys. This makes it very difficult to adequately assess the quality of work in the field, as any program committee requires a mix of different theoretical and methodological experts.

While we adopt quality measures like acceptance and citation rates as end results of success for our conferences, we limit the scope and type of work we can incorporate. When substandard work is accepted because the review system wasn't able to critically evaluate it, other disciplines where those methods or theories are more standard form negative impressions of social science work being done in CSCW and ACM conferences, marginalizing our work and limiting impact on broader fields.

Amy Bruckman

Bio Amy Bruckman is a professor in the School of Interactive Computing at the Georgia Institute of Technology. She does research on how people collaborate to create content online. She received her PhD from the MIT Media Lab in 1997, where she studied educational applications of Internet technology.

Position Conferences serve two purposes: advancing the state of the field (primary), and researcher assessment (secondary). Individuals rely on publications in highly selective conferences to help advance their careers, and institutions use those publications to decide who to reward with fellowships and jobs. The CSCW revise and resubmit process helps us to better fill our primary purpose, advancing the state of the field. Through revision, good papers get better and papers that represent good underlying research but that need rewriting get published. The resultant rise in acceptance rates arguably makes conference papers somewhat less useful for assessment. However, I will argue that conference paper

reviews weren't a particularly fair or effective way to assess researchers anyway, and new metrics can be invented. Regardless, the primary goal of advancing the field should take priority.

Geraldine Fitzpatrick

Bio Geraldine Fitzpatrick is a professor of technology design and assessment in the Faculty of Informatics at the Technical University of Vienna. Her research is at the intersection of social and computer sciences looking at how we design pervasive, tangible and Web 2.0 technologies to fit in with everyday contexts of work and daily life. She has a particular interest in health and well-being, older people, and supporting social interaction and collaboration.

Position CSCW and many of our main conferences are located predominantly in the US, with ECSCW in Europe another key venue. I originally come from outside both the US and Europe and am very aware that cultural and geographical contexts play a large part of where and why we publish. For many people, particularly those who need to travel internationally, their only chance to access travel support is through publication. Yet even under a revise and resubmit model, more than half of those wanting to participate are not accepted to do so. Other 'home' opportunities for international peer contact can be rare. And then there are our smaller conferences. These are not just about specific sub-topic areas but can also be more regional versions of the larger conference, Examples are OzCHI, Nordichi, SouthCHI, and British HCI, many of which make their papers accessible through the ACM digital library. All have their own evolving cultures and discussions, are more accessible to local people in terms of time, travel and costs, and are important places where

critical local relationships and support networks are formed and nurtured, and where local policies can be influenced. How can we, at the same time, promote high quality content and broad participation? How do we value both global and local contributions?

Aniket Kittur

Bio Aniket Kittur is an Assistant Professor in the Human-Computer Interaction Institute at Carnegie Mellon University. His research on *collaborative cognition* focuses on harnessing the efforts of many individuals to make sense of information together in ways that exceed their individual cognitive capacities in domains ranging from Wikipedia to crowdsourcing markets to scientific collaboration. Dr. Kittur has received an NSF CAREER award, Google Research Awards, a Microsoft Research Award, and his work has been reported in venues including Nature, The Economist, The Wall Street Journal, NPR, Slashdot, and the Chronicle of Higher Education.

Position Early career researchers face a difficult challenge in navigating the ecology of conference and publication venues and deciding where to invest their resources. Viewing this as an resource optimization and prediction problem provides a lens that may be useful for understanding the factors involved and ways researchers and program committees might provide better navigation signals within the ecology.

Brian Keegan

Bio Brian Keegan is a post-doctoral research fellow in computational social science at Northeastern University. He draws on methods in network analysis to understand the structure and dynamics of information sharing in social media. His work focuses on the responses of socio-technical systems like Wikipedia and Twitter to bursts of activity and unexpected events. His Twitter followers can beat up your Twitter followers.

Position In the social sciences, journals absurdly delay publication for years to better evaluate the impact of a submission. In the computer and information sciences, conferences absurdly invite a mass of submissions for a single annual deadline and euthanize the majority of them. Both camps criticize the other for their absurdities, but they quarter no reflexivity for their own shortcomings. Furthermore, while appeals to "interdisciplinarity" may be effective in conning funders of their dollars and applicants of other opportunities, scholars are evaluated against criteria for traditional disciplines like "computer science" or "communication". New models like PLoS ONE and CSCW offer alternatives for interdisciplinary scholars, but they face bootstrapping costs and institutional inertia. If we're to eat our own dog food about the value of interdisciplinarity, scholars betwixt disciplines should (1) acknowledge that knowledge comes from diverse networks of sources and (2) reward outlets that align with our values.