The Editor's Spotlight: TOCHI Issue 24:1

Prognostications. Well, 2016 has been consigned to the dustbin of history and 2017 is here whether you wanted it or not.

What exactly the new year will herald remains unclear beyond the nigh-certainties of rising global temperatures and a hothouse of political rhetoric, though I must admit to a certain reticence before I consult my Magic 8-Ball for its prognostications.

OUTLOOK UNCLEAR, it informs me, with an admonishment to ASK AGAIN LATER when I have the temerity to prod it again. I guess I shouldn't have asked.

But as a practicing seer of the technological near-future, one thing I'm certain of is constantly being surprised by the new insights, analyses, and breakthroughs that grace TOCHI's pages. I'm always excited to usher in the latest advances in the field of human-computer interaction, starting especially with the six sterling results you hold right now in your hands.

* * *

The sealed envelope. Originally, I had planned to announce the TOCHI 2016 Best Paper Award in this editorial.

But it belatedly dawned on me that the final issue of 2016 wouldn't be finalized with enough lead time to actually make that feasible. Apparently, our Associate Editors like to go home and see their families over the winter holidays, the nerve of them! (grin).

I also wanted to allow some time for TOCHI readers to nominate any article from 23:6 that may have caught their collective eyes.

But the Best Paper process is now in full swing and I expect to have the results for 2016 in hand when I sit down to write my next editorial two months hence. Be sure to keep an eye open for the April issue of TOCHI (in the eventuality that our internationally televised awards ceremony fails to materialize) so that you can join me, with my ever-thinning hair brilliantined and my tux fitting a bit more snugly than it did in my youth, as I open the sealed envelope for all to behold!

* * *

A turn of fortune. I found this month's allotment of TOCHI contributions even more interesting and diverse than usual, which I suppose says quite a bit about the depth and strength of this batch of papers. And the authors range from CHI Academy veterans to up-and-coming stars of the field.

Meanwhile, even as I type this, another acceptance has rolled in from the diligent authors and editors who are burnishing our long-awaited Special Issue on End-User Development for the Internet of Things. That, too, I expect to unveil two months hence.

And by the time you read this, our window for submissions to our Special Issue on Reimagining Participatory Design will have closed. Already, it looks like response to this call will be strong. Will a double-issue spectacular become necessary? It very well just might.

We also have some new Associate Editors coming onto our board, which will inject some much-needed youth and expand our frontiers. Several prospective editors have

already accepted my invitation, and I expect more to materialize, so once the dust settles I'll announce all our new associates in one go. My general rule is to only bring in people who have distinguished themselves with exceptional review work, and who are much smarter than myself. The first criterion is hard to find, but the latter fortunately leaves me with a rather generous pool of well-qualified individuals to consider.

And as your friend and humble editor, what more could I hope for, than to start out the New Year with such good fortune?

IN THE SPOTLIGHT, PART 1: PRODUCTIVE USE OF DIVERSE WAITING MOMENTS

The visual gestalt of modern society has changed dramatically.

As little as 15 to 20 years ago, one could ride on public transport, or walk around a city center, and not every single adult would have their eyes glued to the soothing blue-white glow of their mobile phone. And while it may be going too far to say that people now worship their devices, the *pull-to-refresh* gestures on mobile email clients, and Twitter feeds, certainly seem to offer the modern equivalent of nervously fingering one's rosary beads, or perhaps the spiritual quietude brought about by thumbing a prayer wheel.

Against this backdrop, the present article recognizes the shifting reality of modern "work," and its increasing fragmentation, across multiple locations, multiple contexts, and multiple devices. Yet even so, perhaps with clever design those moments of idleness amongst our ever-present companions—our mobile phones—could become agents of self-improvement, if not downright productive work.

WaitSuite is a set of applications that, through carefully considered design, present users with vocabulary words for learning a second language in a subdued manner. As the authors explain through a series of meticulously crafted design considerations, the key is to present such "wait-learning" opportunities at exactly the right moment, and in exactly the right context, such that users can take advantage of the learning opportunity (if desired) without interfering with their primary task. While certainly context sensing could play a role in such systems, the authors demonstrate a number of examples where software can present such opportunities at the right moment through the time-tested device known as "good design."

What emerges is a set of design insights and recommendations for a class of software systems that may help to overcome what is perhaps the most stubborn problem of human existence: the limitation of the clock to a mere 24 hours in each day. By harvesting these lurking fragments of inattention, WaitSuite convincingly illustrates a vision where we can co-opt our mobile devices to instill positive habits of learning something new, rather than having these elusive moments raft away and melt into nothingness on the great ocean of wasted time.

And while simple in conception, the idea raises interesting research questions for the field, and a whole new set of directions for interaction design.

Carrie J. Cai, Anji Ren, and Robert C. Miller. 2017. WaitSuite: Productive Use of Diverse Waiting Moments. ACM Trans. Comput.-Hum. Interact. 24, 1, Article 7 (February 2016), 37 pages.

DOI: <http://dx.doi.org/10.1145/3044534>

IN THE SPOTLIGHT, PART 2: HOW PEOPLE WRITE TOGETHER NOW

Every so often it hits me that we are still in the early days of coming to grips with how the internet—or "the cloud," if one prefers the more evocative image that has become

the buzzword *du jour*—will impact productivity tasks that were once taken for granted as single-user activities.

Certainly, collaboration in the cloud through tools such as Google Docs has made it easier for people to write together—both synchronously and asynchronously—than ever before. Yet, almost all previous papers discussing this topic report only generalized observations, or are based upon self-reports of what people *think* they do while engaged in such activities—as opposed to hard data that shows what people *actually* do.

As such, the field needs “in the wild” studies that capture and analyze how people collaborate with modern cloud-aware text editors. How often do documents actually exhibit instances of synchronous editing? What patterns of activity are manifest in how small teams produce documents together? Has the emergence of collaborative editing capabilities allowed new patterns of document production to emerge?

To answer these types of questions, this article analyzes keystroke-by-keystroke data for 96 documents collaboratively authored by advanced undergraduate students. The results provide the first detailed, quantitative investigation of collaborative writing based on hard data. While the keystroke-level data grounds the analysis in key observational details, the authors do a wonderful job of raising the level of discourse in the article to produce high-level insights. For example, a series of mesmerizing visualizations reveal interesting (and at times, previously undocumented) patterns in how small teams produce collaborative documents. The authors also observe that people use such documents as shared containers—a “place” for discussions, intermediate results, and even whimsical social exchanges—that do not end up as part of the final document.

While the present analysis is restricted to students, there are good reasons to believe that many of the insights should generalize. To this end, the paper concludes with a number of system design implications, as well as behavioral guidelines to better support people writing together.

Judith Olson, Dakuo Wang, Gary Olson, and Jingwen Zhang. 2016. How People Write Together Now: Beginning the Investigation with Advanced Undergraduates. ACM Trans. Comput.-Hum. Interact. 24, 1, Article 4 (February 2017), 48 pages.

DOI: <http://dx.doi.org/10.1145/3038919>

OVERVIEW OF VOLUME 24, ISSUE NUMBER 1:

FEELING THE ♥ FOR EMOJI, UNDERSTANDING THE SEMANTICS OF COLOR PALETTES, THE BENEFITS (?) OF MOOD SELF-TRACKING AT WORK, AND RECONCILING PRIVACY PREFERENCES IN MULTIUSER SHARING SCENARIOS

While I've highlighted just two contributions above, I found this issue of TOCHI to be thought-provoking on a wide range of topics, so as you will see from the tenor of my remarks below, I easily could have directed my spotlight on any and all of the other articles.

It's always a great feeling to be confronted by such an embarrassment of riches when it comes to publishing a journal that aspires to be the flagship of the field.

Beyond Just Text: Semantic Emoji Similarity Modeling to Support Expressive Communication 🌲 🍷 ☺. Texting on mobile phones has added not only to the lexicon to the English language (LOL, OMG, and the like), but also an entire new symbology—☺, 🙄, 🌟, 100, 🍷, and everything in between—known as the *emoji*. Indeed, at times I wonder how historians of some far-flung future will look back on these curious pictograms, wondering where the Rosetta stone lies that can unlock their mysterious messages.

Well, this TOCHI paper very well might be it.

Based on a corpus of 21 million collected tweets containing emoji—and in what may be the most delightfully illustrated article ever to appear in TOCHI—the authors carefully document emoji usage. The result is a semantic model of emoji similarity that paves the way for novel, more efficient, and more intuitive designs for emoji keyboards. Along the way, the authors provide many wonderful examples of how people use emoji to add expressiveness and emotional undertones to messages that would otherwise be ambiguous (e.g. “Susan’s coming over later 😊” vs. “Susan’s coming over later ☹”).

And the paper’s impressive depth of analysis represents a gift to the literature that should keep on giving for years to come. 📖

(DOI: <http://dx.doi.org/10.1145/3039685>).

Colors—Messengers of Concepts: Visual Design Mining for Learning Color Semantics. Anyone who’s played with color palettes knows that it can be fiendishly difficult to find a set of colors that complement one another while conveying a concept, or emotion, in just the right way.

Many ways of organizing color spaces have been devised, as have methods to intuitively map interactive controls to color selection, yet—until this article—there has been little traction on the fundamental problem of understanding the semantics of colors.

From a corpus of 2,654 magazine covers (spanning 71 distinct titles, in 12 genres), this work derives semantic concepts as a combination of color and word distributions. These resulting *color-word topics*, validated through crowd-sourcing, unlock new scenarios for interacting with colors and images.

For example, users can query using natural-language terms such as “science,” “gardens,” and “health” to discover color palettes that satisfy weighted combinations thereof. The same underlying color-word model allows the system to fetch design examples that satisfy the query, or to retrieve images based on color semantics. Likewise, one can recolor an image, or select certain colors from an image based on simple concepts such as “travel,” “beach,” and “trip.”

I also must note that the work is a visual feast: as the reader might expect given its subject matter, this article continues the trend of delightfully illustrated contributions that grace TOCHI Issue 24:1.

(DOI: <http://dx.doi.org/10.1145/3009924>).

Introducing Mood Self-Tracking at Work: Empirical Insights from Call Centers. Collecting data about one’s own behaviors, habits, and emotional states has become a popular way to gain quantitative insights into well-being, health, sustainable living, and other aspects of self-improvement.

Yet investigations of the “quantified self” have largely focused on private life, and as such the potential impact and implications of mood self-tracking in the workplace are still largely unknown.

By investigating the deployment of the “MoodMap App,” a mood self-assessment tool co-designed with call center workers and managers, this article takes a hard look at the benefits—as well as some of the moral ambiguities—raised by such an approach. The results provide empirical evidence that mood self-tracking can indeed improve work-related performance, and can help workers by making them aware of one another’s challenges and difficulties so that they can ask for help by a subtle new back-channel of emotional communication.

Even though the authors show that mood self-tracking—when tightly integrated with the existing workflow of tasks—can yield productivity advantages for workers, the authors also document that managers remained unwilling to cede even two extra minutes a day to allow workers to stop and reflect back on their actions. In the context

of a call center, such *reflection-on-action* in two-minute increments adds up to a lot of valuable time. But in a more enlightened view, two minutes seems a small price to pay for more productive workers who can better process the negative emotions evoked by a rude caller, for example, or who can take a short break to ponder how they could handle a difficult situation better the next time.

The results and insights shared by the authors are intriguing, while making apparent many unanticipated complexities raised by mood-tracking in work-related contexts.

(DOI: <http://dx.doi.org/10.1145/3014058>).

Sharing Policies in Multiuser Privacy Scenarios: Incorporating Context, Preferences, and Arguments. Imagine that someone you don't know very well has just posted a photo on a social network, taken at an event that you—or perhaps, your child—recently attended. They've made the post open to everyone. You (and perhaps your young one) are tagged in the picture, linking semiprivate information for anyone to access.

Do they have the right to post this?

Should you have the right to block them from sharing it?

What if the photo was taken at a public school performance by another parent, who therefore has a good reason to post it, versus one taken surreptitiously at a beach where your entire family was in swimsuits?

This article tackles difficult questions of this sort related to sharing content online. The authors argue that a simple policy that “respects each user's privacy” cannot reconcile the preferences of multiple users for such conflicts. And as such, the status-quo approach of aggregating user preferences and letting “majority rule” is not satisfactory either.

Indeed, the authors show that people's arguments—their reasons for sharing (or not sharing) content—are an essential ingredient that allows people to align conflicts, or agree to exceptions in well-justified circumstances.

These results and insights, in addition to a tool known as “Muppet” for assisting decisions involving multiuser privacy scenarios, makes a decisive contribution to this very tricky problem of maintaining control over one's personal information in the social-sharing reality of modern life.

(DOI: <http://dx.doi.org/10.1145/3038920>).

* * *

A realm governed by truth. Well, there you have it for TOCHI 24:1, yet another milestone in the annals of human-computer interaction.

But, I must admit that as I write this editorial for a publication grounded in design, invention, and the scientific method, it is quickly becoming apparent that a global chill has descended on scientific discourse. The large-scale kneecapping of the scientific method that we are presently witnessing undermines ongoing endeavors, of course, but the impact this could have on the training of future generations of scientists, technologists, and designers casts an even greater pall over my mind.

If we are not willing to come to grips with scientifically established facts—disconcerting though they may at times be—then our inability to make rational decisions in light of predictable consequences could lead us to a very dark future indeed.

And while political commentary should have no place in a scientific journal, by the same token if we cannot defend the truth (much less the carefully garnered datasets that could lead us to it) then science has no “place” in the future, either.

I suppose that's not much of a resistance, but nobody ever promised it would be easy to advance the frontiers of *all human knowledge*:

A realm that pointedly lacks great walls.

A realm, governed by truth, that allows unlimited freedom of thought and freedom of movement.

A realm that admits any and all the world's peoples, no matter their faith.

Ken Hinckley

Editor-in-Chief

Redmond, Washington

January 25, 2017