Considering Social Factors in New Technologies

New technologies that alter how we interact with other people come and go, creating new opportunities but also upending social norms. How should builders of new technologies consider the social implications of their systems?

By Amy X. Zhang

DOI: 10.1145/3301479

New consumer technologies are coming out at a dizzying pace. From alleged "email killers" like Slack to numerous social VR applications, the number of tools we can use to communicate and collaborate seems to be growing at an increasing rate. But a proliferation of tools doesn't necessarily translate into an increase in lasting innovations. Indeed, many novel systems that are around today likely won't be around in five or ten years' time. Looking back, some tools were a flash in the pan—exciting in their novelty but ultimately cast aside partly because of their violation of social norms. In the case of Google Glass, not enough attention was placed on the privacy norms that were violated when one could continuously record others and with little awareness by those being recorded. In contrast to short-lived technologies, there are tools that have been around for decades and are still widely used. Despite plenty of op-eds declaring that email is dead, it persists.

Separating the innovations that fizzle out from the ones with long-lasting appeal is tricky because of the many factors that can affect adoption, which have little to do with the inherent idea behind a new tool [1]. Maybe the marketing didn't hit the right demographic, or the account creation workflow was too cumbersome. Or maybe difficulties with scaling led to a botched launch. As product designers and developers, some of these factors are in our control but many are not. Often implementation details are manageable—these are the elements of execution that some product will get right someday, given the same underlying idea.

As builders of social technologies, how then should we evaluate the promise of a new idea? While it may be tempting, it would be shortsighted to throw up our hands and say we should focus only on the technical aspects of a new system. If history is our guide, the novel technology behind a new application needs to be considered alongside the social considerations.

However, considering the social implications of a new tool can be a lot harder than evaluating a tool's novelty or technical contribution. Sometimes a change in degree technically can lead to a similar change socially. For instance, the gradual shrinking of personal computing devices from desktops, to laptops, to phones, to watches involved many technical innovations that we can measure in terms of size, weight, and speed. But along the way, each of these incremental changes fundamentally shifted the nature of social interaction using technology, making it easier to interact with others outside the home and on the go. Other times, an idea doesn't fit into the broad social norms of a particular time and place but then becomes popular in a different context. An example of this is the Bluetooth headset, which was

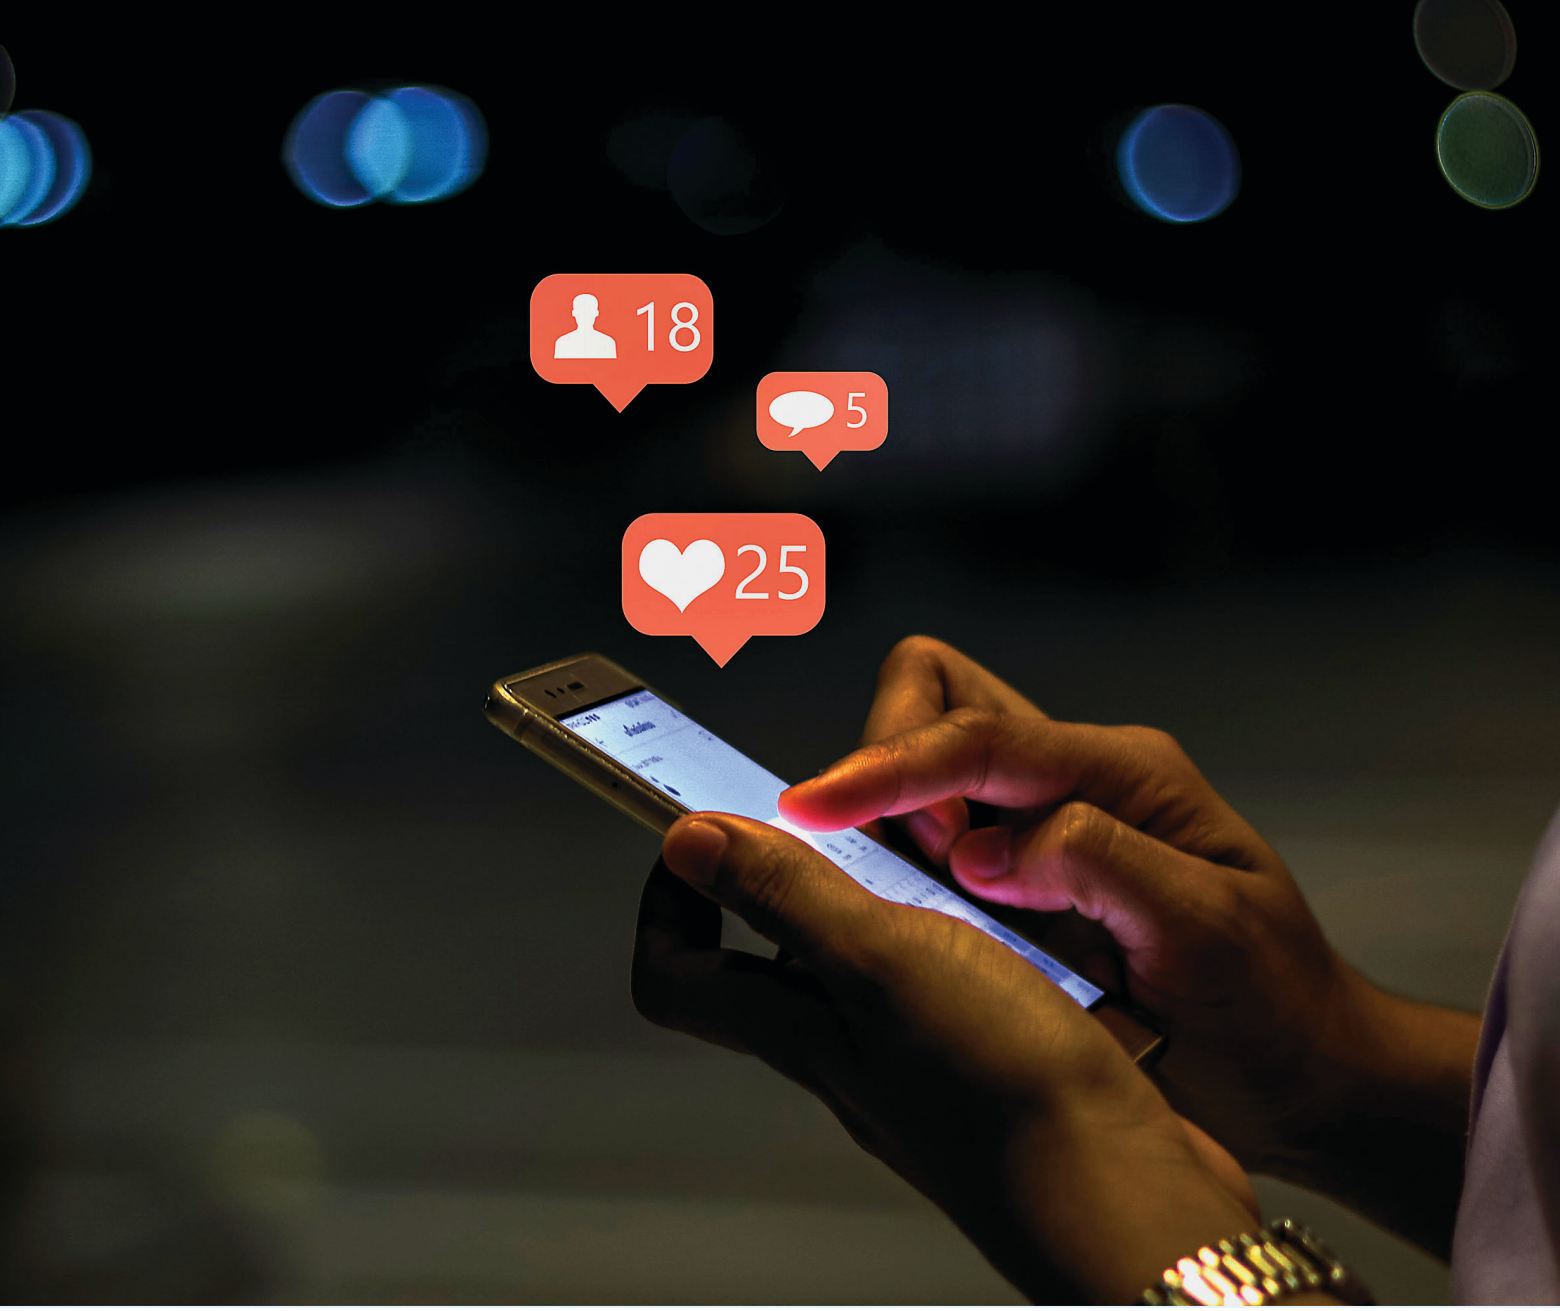

Image by Phunkod / Shutterstock.com

derided in the mid-2000s for making the wearer seem self-important and unfashionable. Fast forward to today, Apple airpods are commonly worn and socially acceptable in public.

Given these difficulties, there are a number of strategies that can be used to consider social implications when coming up with or developing a new technology.

LEARNING FROM LASTING TOOLS

When brainstorming new ideas to old problems, take lessons from existing, long-lasting solutions that haven't changed with the times. Why are these technologies still around after all the years of boom-bust cycles? And despite what they fundamentally get right, what problems do they still have that users have learned to tolerate or that have been exacerbated with time and changing contexts?

These were the exact questions I asked when I first set out to innovate tools for online group discussion. A look around at the existing systems quickly revealed a mighty dinosaur—the mailing list [2]. Here was a tool that has seen almost no innovation in the 40-odd years it has been around. Yet almost every person I knew was a member of several lists, each with upwards of thousands of members. Even while they were using newer systems that could have substituted, like Facebook Groups, why did they still use mailing lists?

It turned out the things people like about mailing lists are related to the reasons they still prefer email over social media for certain activities. Social norms around email still revolve around work, while social media is more associated with procrastination. Email feels more private, while social media feels more public. Additionally, people feel more confident that emails will be seen and in a timely fashion compared to other methods. For example, posting to a Facebook Group,where algorithms mediate distribution.

Interestingly, I found these conceptions held even when they weren't technically true or had not kept up with changing behaviors. Several of the mailing lists I studied were publicly archived unbeknownst to members. Also, due to practices like Gmail auto-filtering mailing lists into a sep-

arate tab or folder, many people no longer see all posts from their mailing lists. This means they access their mailing list emails in a "pull" fashion not unlike how people browse social media feeds, ignoring content or reading it when they feel like it. This is a shift from the "push" behavior of accessing traditional email or SMS, where recipients receive all messages and in real-time. Changes in email delivery can impact how people feel about a list and how they then act as senders. People who automatically filter their emails may assume others are doing the same and thus send more emails. Conversely, people who read all emails from their main inbox may be hesitant to post for fear of "spamming" the list.

From this exploration into a widely-used tool, we can understand what enduring qualities people expect from their online discussion tools and where tools are failing to keep up with those expectations. These findings can help design new tools that fit into existing social norms by resolving the discrepancies between how we think a tool should behave and how it actually behaves. In the case of mailing lists, features could be added that provide more transparency around who has access to the list and more controls to signal what messages are desired. I've developed these ideas into a tool called Murmur (murmur.csail.mit.edu), a new mailing list system where users can designate what kinds of messages they want to receive, as well as how and to whom their messages should be sent.

EXAMINING SOCIAL MOTIVATIONS AGAINST RISKS

However, as tool builders, we don't always want to build something that completely conforms to existing social expectations. Sometimes, we seek to build a tool that nudges existing behavior toward encouraging new ones or brings social expectations more in line with what people actually want or how they actually behave. As seen in the case of mailing lists, when people think other people are acting in a certain way that they're actually not, these gaps can lead to tensions from clashing behaviors. But how can we know whether the new features in the tool we're building will be embraced?

One way to think about this question is to consider the social motivations of people using the tool. While there are often personal motivations for using a tool, examining social motivations can help guide the design of systems that involve novel social interactions. Think about all the behaviors we are comfortable performing socially today that would have been completely foreign only a few years ago.

As one example, recall when the behavior of publicly broadcasting "check-ins" to physical locations—pioneered by social networking services like Dodgeball and Foursquare and now built into most social media—was introduced in the early 2000s. At first, the practice was questioned and sometimes even derided as oversharing but then was eventually normalized. This was in part due to personal and financial benefits, namely keeping track of places or earning discounts. But it also broadly conferred social benefits, such as crafting a particular public image or facilitating serendipitous meetups.

When it comes to benefits that involve greater self-disclosure, however, users must weigh them against the risks, including loss of privacy or the collapse of different social contexts into one space [3]. Given the calculations users must make, tool builders can improve their chances of social adoption by providing more social benefits, while mitigating potential risks by providing easy-to-use tools to allow users to negotiate the terms of their self-disclosure.

This was one approach I took when developing a tool that aims to bring social awareness and interaction to web browsing [4]. While social features exist on some specific sites, the general concept of social web browsing has never been mainstream, despite some web browsers incorporating social elements in the mid-2000s, such as Flock or RockMelt.

The consideration of social implications early in the conception, design, and development process can help avoid some major social missteps.

Instead, most browsing today is done in private with little awareness of others, despite the potential benefits to users of social interaction while traversing the web. Part of the reason there is no ecosystem of consumer tools around making use of browsing activity is that almost all browser activity tracking happens in the dark. Private companies primarily reap the benefits of tracking and collecting browsing activity. In contrast, the users who create that data get little utility from having it collected, and usually do not provide consent for data collection or even have awareness that it is happening. Our tool, called Eyebrowse (eyebrowse.csail.mit.edu), changes this equation by offering users the ability to share only certain aspects of their browsing activity publicly, thus mitigating potential risks, and providing social benefits in the places on the web they consider public.

When studying the potential benefits of web activity sharing, we found several social benefits, some mirroring the benefits of physical check-ins. These included self-presentation toward impression management, ambient awareness of friends' activities, serendipitous connections, and the ability to have conversations within the shared context of a webpage. These benefits in turn suggested features that are built into Eyebrowse, such as indications when a user "bumps" into a friend on a page or the ability to highlight a particular passage on a page to have a discussion "in the margins." In terms of mitigating risks, we examined what concerns people have regarding sharing aspects of their web activity. These included misrepresentation, inadvertently revealing information, and concern about being judged. The privacy risks described led us to build opt-in whitelisting at the domain level in Eyebrowse to allow users to manage what sites they felt comfortable sharing.

CONSIDERING THE LONG TAIL

Sometimes when we think about the social implications of a new technology, we focus on what will happen to the majority of users or what the mainstream reaction to the tool will be. This makes sense as a way to maximize the likelihood of adoption or minimize potential harms. However, focusing on a subset of users, even if they are the largest subset, can lead to blind spots regarding how a new technology can have adverse social consequences. Though the tool in question may work as intended in the general case, some users in the long tail may use the tool to achieve disproportionate impact on other users. In addition, there may be other groups of users in the long tail, oftentimes people who are marginalized in other ways, who are doomed to be severely impacted without active consideration. These concerns are exacerbated in technologies where social interactions can happen at scale.

Consider the existing social features that live in many of the tools we use. In many cases, an action by a stranger somewhere around the world leads to a notification on a user's phone or a message in their inbox. For the majority of people, this is a good thing. The feature works as intended. But for a subset of users, social features with no safeguards can leave them vulnerable to deception or attack.

A small example of a feature like this is the way users were added to repositories in Github. In the past, anyone was able to automatically add someone to their repository without explicit permission. This feature was useful for quickly establishing collaborations on projects. Unfortunately, this also meant users could create bogus repositories with offensive names and associate them with people they were aiming to harass. Github rectified the problem when they added an explicit invitation step, where the recipient of the invitation must agree in order to be added to the repository. They also have the option of blocking further invitations from the sender. This change seems like a subtle thing but it likely drastically improved the social experience of Github for a subset of users in the long tail.

Separating the innovations that fizzle out from the ones with long-lasting appeal is tricky because of the many factors that can affect adoption.

Twitter has also been grappling with these issues when it comes to being added to Twitter Lists, which suffers from similar problems as Github repository membership. One solution that was implemented was to suppress notifications when getting added to a list on Twitter, so that users would not have to see the harassment embedded within the list name. This change was quickly rolled back because some users wanted to be notified so they could block the offending user. It turns out that the right solution may not be the same for everyone.

In my research around designing online harassment interventions, I similarly found people have different ideas about how they would like to handle harassment, and even have different definitions about what constitutes harassment [5]. Some people prefer to be oblivious and block all harassing content, while others want the ability to review their harassment in order to plan follow-up actions, such as to alert their family when they've been "doxxed" (had their personal information published without their consent). Some people want the ability to respond to their harassers while others prefer the harasser be left in the dark about what happened to their correspondence.

As a result, the tool we ended up building has a heavy emphasis on customization. The tool is called Squadbox (squadbox.org), and it aims to help people with email harassment. In addition to being customizable, a central element of the tool is the use of friend-moderators or trusted individuals who triage potentially harassing messages according to the specifications of the person being targeted. This was because we found harassment is often contextual as well as highly personalized to the recipient.

Looking around online, there are unprotected social features like the ones described herein. Left unattended, users who are harassed will leave, harassers will become emboldened, and the general social norms around the use of a tool will degrade. This is why it is imperative to consider the social impacts of our systems not just to the average user but to the long tail.

At the end of the day, we all want to build tools that do right by users. Tools that help users with what they want to get done, don't violate social norms, provide social benefits to users while reducing risk, and consider how all users will be affected. As tool builders, we want our creations to succeed, and while nothing can guarantee adoption, the consideration of social implications early in the conception, design, and development process can help avoid some major social missteps.

References

[1] Bernstein, M. S., Ackerman, M. S, Chi, E. H. and Miller, R. C. The trouble with social computing systems research. In *Extended Abstracts on Human Factors in Computing Systems* (CHI'11). ACM, New York, 2011, 389-398.

[2] Zhang, A. X., Ackerman, M. S. and Karger, D. R. Mailing Lists: Why are they still here, what's wrong with them, and how can we fix them?. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI'15). ACM, New York, 2015, 4009-4018.

[3] Ma, X., Hancock, J., and Naaman, m. Anonymity, intimacy and self-disclosure in social media. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI'16). ACM, New York, 2016, 3857-3869.

[4] Zhang, A. X., Blum, J., and Karger, D. R. Opportunities and challenges around a tool for social and public web activity tracking. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing* (CSCW'16). ACM, New York, 2016, 913-925.

[5] Mahar, K., Zhang, A. X. and Karger, D. R. Squadbox: A tool to combat email harassment using friendsourced moderation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (CHI'18), 586. ACM, New York, 2018.

Biography

Amy X. Zhang is a Ph.D. candidate in the Computer Science and Artificial Intelligence Lab at Massachusetts Institute of Technology. She builds new online discussion systems to improve how we conduct discourse online.