Socializing Algorithms

JED R. BRUBAKER // UNIVERSITY OF COLORADO BOULDER

Matching people to former romantic partners. Categorizing photos of black people as gorillas. Painful reminders about the death of a loved one. These are just a few of the now familiar stories of algorithms behaving badly. Insensitive technology is nothing new. However, an explosion of social computing, personal data, and algorithmic curation and personalization has resulted in new aspects of social life around which algorithms can be insensitive. Take, for example, Eric Meyer's widely circulated experience with Facebook's Year in Review – a system that generates a video aggregating social media highlights from the previous year. While his friends shared compellations of vacations, parties, and birthdays, all accompanied by joyful and upbeat music, the celebratory video waiting for Mr. Meyer prominently featured the death of his four-year-old daughter. For Meyer, the encounter was nothing short of "cruel" [11].

In principle, the algorithms that power recommendation and personalization systems should serve user's needs. They should be sensitive to the benefits and risks of the content they present. But what kinds of social information should they consider? And while Meyer's experience was unquestionably cruel, in what cases might encountering photos of a deceased loved one be appreciated, nostalgic, or encourage social support?

As algorithms extend into our social lives, we must confront a pragmatic reality: algorithms do not understand our lives very well. This is for two related reasons:

First, data on which algorithms rely is necessarily partial. Take, for example, the "user." In order to make humans and human activity amenable to computation, system designers create partial "representational schemas" consisting of "a small vocabulary of discrete elements" [1] that we call user accounts, profiles, and timelines. These representations are always partial, raising questions about what kinds of user entities exist (or even can), what social and technological functions the entities serve, what breakdowns occur between representations and practice, and, finally, what social complications arise because of mismatches between practice and representation. These questions are especially prescient when algorithms aim for human-centeredness. The types of "human" we can center on is often predetermined and constrained by the ways our infrastructure has operationalized the user.

Second, algorithms do not fully understand the meaning, origin, and context of data. What algorithms understand are narrow slices of our social world, often pre-ordained by the interactions and data afforded by a given platform. The social lives they understand are bound by the availability of data and narrow scope of trace data afforded by a platform and data ecology. As a result, "social" is often understood in terms of system and business functions. "Likes" can easily be measured, but are rarely fully understood. We have socialized our algorithms to be sensitive to financial practices for credit scores, web navigation patterns for online advertising, and network density and tie-strength for our social feeds.

If we want algorithms to be worthy of our trust, they need to be socialized to understand the context in which they operate. To this end, extreme examples can be instructive. Recently, sensitive interactions and life experiences have emerged as active research areas in the CSCW and HCI communities (e.g., [6,11–13,17]). The existing research typically focuses on sensitive experiences and

broad design implications, but has not focused on unique design challenges when considering the design of algorithms.

Sensitive interactions are a useful case to think with as they require nuanced understandings of data and sensitivity to the circumstances under which it is presented. In my research group, we have begun focusing on what we are calling **sensitive algorithmic encounters**, focusing our attention on how people navigate encounters with algorithmic content related to the death of friends and family. Death is a challenge for algorithms in social media. Until just recently, computational representations of the "user" did not include mortality. And while big players such as Facebook and Google now have means of identifying deceased users, signals for death are typically not available and are never uniformly applied. Meanwhile, the death of a family member or close friend is considered one the most stressful of life events [9], and while funerals and memorials are communal events [2,15], experiences of grief and mourning are highly-individualized [10,13] during which the bereaved may have shifting needs [14,15]. It is for these reasons that people strive to be sensitive surrounding death. But how might an algorithm practice a similar sensitivity?

In our work to date, we have documented ways in which people encounter death as a result of algorithmic recommendations and prompts [3,4]. Whether through Facebook's News Feed, birthday reminders, an old email, or even a seemingly innocuous advertisement for Mother's Day flowers, these "unexpected encounters" [4] often appear out of context. Technically, the solution may seem straightforward: Yet, while participants report that these encounters can be upsetting, this assessment is not universal [4]. What some describe as "creepy", others find a thoughtful prompt to reminisce about a loved one.

Sensitive algorithmic encounters highlight the broader social context to which algorithms should be accountable. However, to date, it is unclear what sociotechnical factors shape how people experience these sensitive encounters and it is unclear how sensitive algorithms should be designed. This prompts four questions:

1. How do people experience and navigate encounters with algorithmically curated sensitive content?
2. What impact do algorithmic encounters with sensitive data have on people?
3. What social context and properties of data do algorithms fail to understand and accommodate?
4. What design requirements and strategies can be identified to support improving algorithmic curation of content and interaction design around sensitive social data?

If we are to build algorithms we can trust with our social lives, we need better ways of understanding how people experience algorithmically generated content and interactions. We need to identify what types of data and features are important for the designers of algorithms to include. And we need to develop new design practices that sit user experience and interaction designers with the ML/AI professionals whose work increasingly shapes these experiences and interactions.

REFERENCES

1.   Philip E. Agre. 1997. Beyond the mirror world: Privacy and the representational practices of computing. In *Technology and privacy: The new landscape*, P. E. Agre and M. Rotenberg (eds.). MIT Press, Cambridge, MA, 29–62.

2.   Philippe Ariès. 1975. *Western Attitudes toward Death: From the Middle Ages to the Present*. The Johns Hopkins University Press, Baltimore.

3.   Jed R. Brubaker, Lynn S. Dombrowski, Anita M. Gilbert, Nafiri Kusumakaulika, and Gillian R. Hayes. 2014.

Stewarding a Legacy: Responsibilities and Relationships in the Management of Post-mortem Data. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, ACM Press, 4157–4166.

4.   Jed R. Brubaker, Gillian R. Hayes, and Paul Dourish. 2013. Beyond the Grave: Facebook as a Site for the Expansion of Death and Mourning. *The Information Society* 29, 3: 152–163.

5.   Jed R. Brubaker, Funda Kivran-Swaine, Lee Taber, and Gillian R. Hayes. 2012. Grief-Stricken in a Crowd: The Language of Bereavement and Distress in Social Media. *ICWSM 2012 - Proceedings of the 6th International AAAI Conference on Weblogs and Social Media*, The AAAI Press, 42–49.

6.   Oliver L. Haimson, Jed R. Brubaker, Lynn Dombrowski, and Gillian R. Hayes. 2016. Digital Footprints and Changing Networks During Online Identity Transitions. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, ACM Press, 2895–2907.

7.   Oliver L Haimson, Jed R. Brubaker, Lynn Dombrowski, and Gillian R. Hayes. 2015. Disclosure, Stress, and Support During Gender Transition on Facebook. *CSCW '15: 18th ACM Conference on Computer Supported Cooperative Work & Social Computing Proceedings*.

8.   Daniel Herron, Nazanin Andalibi, Oliver Haimson, Wendy Moncur, and Elise van den Hoven. 2016. HCI and Sensitive Life Experiences. *Proceedings of the 9th Nordic Conference on Human-Computer Interaction - NordiCHI '16*, ACM Press, 1–3.

9.   Thomas H. Holmes and Richard H. Rahe. 1967. The Social Readjustment Rating Scale. *Journal of Psychosomatic Research*.

10. Elisabeth Kübler-Ross and David Kessler. 2005. *On grief and grieving: finding the meaning of grief through the five stages of grief*. Simon and Schuster, New York.

11. Eric Meyer. 2014. Inadvertent Algorithmic Cruelty. *meyerweb.com*. Retrieved June 2, 2016 from http://meyerweb.com/eric/thoughts/2014/12/24/inadvertent-algorithmic-cruelty/

12. Wendy Moncur, Lorna Gibson, and Daniel Herron. 2016. The Role of Digital Technologies During Relationship Breakdowns. *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing - CSCW '16*, ACM Press, 370–381.

13. Therese A. Rando. 1993. *Treatment of complicated mourning*. Research Press, Champaign, IL.

14. Margaret Stroebe and Henk Schut. 1999. The Dual Process Model of Coping With Bereavement: Rationale and Description. *Death Studies* 23, 3: 197–224.

15. Tony Walter. 1999. *On Bereavement: The Culture of Grief*. Taylor & Francis Group.