

Designing for the Bittersweet: Improving Sensitive Experiences with Recommender Systems

Caitlin Lustig
celustig@uw.edu
University of Washington
Seattle, Washington, USA

Artie Konrad
akonrad@fb.com
Meta
Menlo Park, California, USA

Jed R. Brubaker
jed.brubaker@colorado.edu
CU Boulder
Boulder, Colorado, USA

ABSTRACT

It is difficult to design systems that honor the complex and often contradictory emotions that can be surfaced by sensitive encounters with recommender systems. To explore the design and ethical considerations in this space, we interviewed 20 people who had recently seen sensitive content through Facebook’s Memories feature. Interviewees typically described how (1) expectedness, (2) context of viewing, and (3) what we describe as “affective sense-making” were important factors for how they perceived “bittersweet” content, a sensitizing concept from our interviews that we expand upon. To address these user needs, we pose provocations to support critical work in this area and we suggest that researchers and designers: (1) draw inspiration from no/low-technology artifacts, (2) use empirical research to identify contextual features that have negative impacts on users, and (3) conduct user studies on affective sense-making.

CAUTION: This paper discusses difficult subject matter related to death and relationships.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**;
Empirical studies in collaborative and social computing.

KEYWORDS

technology-mediated reflection, social media, death, breakup

ACM Reference Format:

Caitlin Lustig, Artie Konrad, and Jed R. Brubaker. 2022. Designing for the Bittersweet: Improving Sensitive Experiences with Recommender Systems. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*, April 29-May 5, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3491102.3502049>

1 INTRODUCTION

CAUTION: Throughout the paper there are mentions of difficult subject matter, including murder and suicide. A fatal overdose and an abusive relationship are mentioned in Section 5.3.2.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '22, April 29-May 5, 2022, New Orleans, LA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9157-3/22/04...\$15.00
<https://doi.org/10.1145/3491102.3502049>

Jeremy loves being a father, but would be the first to admit that parenthood has been harder than he expected. As a stay-at-home dad in Louisiana, Jeremy described some reoccurring challenges with his son that often left him feeling defeated. It was during a low moment when he opened up Facebook and found a picture of his father and his uncle hunting at the top of his News Feed, an old post resurfaced by Facebook’s Memories feature. Many technology-mediated reflection (TMR) systems, such as Memories, are supported by recommender systems that curate past content and are designed to give users an opportunity to reflect on the past. For Jeremy, “it felt like a sign.”

Jeremy’s father passed away 18 years ago from a cardiac arrest, but three years ago Jeremy’s sister scanned dozens of old photos, which Jeremy then shared on Facebook. Encountering one of these photos brought back the pain of losing his father, amplified by the confusion and frustration he now felt as a parent. But he also felt like his dad was looking down on him, saying, “don’t give up; just keep going.”

When he reached out to his sister, she told him that she had gotten the same Memory. They talked about their father, and how much they still missed him. That night Jeremy found himself staring at the photo, crying as he went to sleep. And while in many ways this experience was painful for Jeremy, he also found it beautiful and uplifting.

Revisiting our digital pasts can be powerful. These experiences can be meaningful in ways that are complicated, messy, joyous, and sad—often at the same time. TMR systems are designed to make that experience ultimately positive for a user. Facebook’s Memories feature, for example, is one of the most loved features on the platform with, as of 2018, over 90 million people using it daily.¹ Research has shown that TMRs are generally successful in having a positive impact on users’ well-being [33, 42]. Even though the recommender systems used in TMR systems sometimes show users sensitive content that may be distressing—such as pictures of a deceased loved one—studies have shown that TMRs provide well-being benefits for both positive *and* negative memories [33], although they may initially negatively impact a person’s current mood before the well-being effects emerge. An individual may need to reflect further before they experience positive impacts [43]. However, not all memories have the same impacts on users, and designing recommender systems to accommodate the wide range of users’ needs is an open challenge.

In this work we turn to encounters with sensitive content (i.e., content that has the potential to evoke a strong emotional response)

¹<https://about.fb.com/news/2018/06/all-of-your-facebook-memories-are-now-in-one-place/>

to ask how recommender systems can and should account for positive and negative experiences, as well as those experiences that evoke complex, bittersweet emotions. Users may not want to go through negative feelings even if they would ultimately lead to greater happiness, and conversely, some users do not want to only see content that makes them happy—they may want to see content that is meaningful to them even if it hurts. We need to understand the benefits and harms of systems that curate sensitive content, but currently it is difficult to do so in a nuanced way through the mechanisms used to assess recommender systems (e.g., engagement, share-rates, and click-through-rates). To that end, we interviewed 20 Facebook users who had recently seen a post via Facebook’s Memories feature related to a sensitive life event.

This paper makes five main contributions to social computing research on the sociotechnical impacts of recommender systems that curate content related to people’s personal social lives:

- (1) we uncover the complex emotions and experiences that people have with these systems,
- (2) we introduce and explore the concept of “bittersweet” content, a sensitizing concept from our interviews,
- (3) we introduce and explore the concept of “affective sense-making”,
- (4) we identify open challenges for creating more compassionate and sensitive recommender systems that curate personal content, and
- (5) we develop provocations and suggested practices for researchers and designers to address these challenges.

While we focus on TMRs in this paper, this work has wider implications for recommender systems that curate social and personal content. We can imagine sensitive encounters with other kinds of recommender systems, such as receiving an advertisement for Mother’s Day flowers after one’s mother has passed away; recommender system generated playlists curating music that reminds one of an ex; and recommender system generated photo albums of pictures taken in a hospital. These systems have the capacity to both delight and distress users—and we hope that this paper will help researchers and designers to reflect on the impacts of these systems and their ethical responsibilities to users.

2 RELATED WORK: ALGORITHMICALLY CURATED BITTERSWEET MEMORIES

We discuss the potential benefits of technology-mediated reflection systems (TMRs), but also discuss the shortcomings of the recommender systems they employ, including the technical limitations of incorporating user feedback. These systems are particularly challenged by the difficulty of identifying sensitive content that is not “good” or “bad”—in other words, the bittersweet content which exists in a liminal space that resists algorithmic categorization or quantification. We turn to HCI and critical algorithm studies literature to discuss users’ qualitative experiences with recommender systems.

2.1 Technology-Mediated Reflection

Technology-mediated reflection systems (TMRs) (e.g., Timehop, Google’s Rediscover This Day, Apple’s Memories feature in Photos, and Facebook Memories) are applications or features that present

users with these technology-mediated recordings of their lives and implicitly or explicitly prompt users to reflect on them. These systems may have structured prompts (e.g., a question for users to answer) or present the user with an unstructured space to reflect (e.g., Facebook allows users to repost Memories with added commentary but does not prompt people to reflect on any particular topic) [42].

In this context, reflection is the act of “actively reviewing” a memory [43]. One motivation behind the early design of TMRs was to enhance the well-being benefits of unmediated reflection techniques [42, 43]—such as through better accessibility (i.e., a user can access their data on their mobile phone) and more complete recordings of memories. (The benefits of reflection through non-technological means is discussed further in Section 7.1.) Reflection on both positive and negative experiences has been shown to typically have well-being benefits [52]. Conversely, some people tend to ruminate (i.e., they perseverate on distressing events that normally would be “edited” or forgotten) [50], which can have negative health outcomes [51], and TMR systems may act in a similar way if they prevent users from forgetting events. However, Konrad, Isaacs, and Whittaker argue that they also hold promise because providing detailed information to users about their pasts may prevent rumination by helping them identify solutions to problems that they faced [42]. For example, if a user had a bad experience with public speaking, reflecting on the experience might help them identify reasons that it did not go well, such as not eating before giving the speech. These kinds of reflections can imbue experiences with different or additional meanings—Facebook posts that are a record of “banal and everyday experiences of users” may be made into “intimate and significant” memories when users view and reflect on them [54].

TMRs often present users with content that has been algorithmically curated, such as the algorithmically generated albums of Google Photos that use location data and facial recognition to put together albums around a common theme [46]. Through this curation, people are able to view their pasts in new ways that are not easy for them to curate on their own or access through their organic memory; however, some have argued that curating content in this way may have some troubling implications—recommender systems classify content to select what is “worthy” of being remembered [48] and filter out what “should” be made invisible and forgotten [34]. However, they do not do so without human input—TMRs rely on human feedback loops in order to perform these functions as “self-learning machines” that adapt to observable and quantifiable human actions [46].

2.2 Curation in practice

Although TMRs are shown to be emotionally helpful for many users, it is difficult for many recommender systems to meaningfully incorporate user feedback about content that evokes complex emotions; while recommender systems attempt to personalize their results to what specific users want, ways for users to provide feedback to recommender systems (so that they can be personalized further) are somewhat limited. In this section, we discuss the approaches to gathering and using implicit and explicit user feedback, the limitations to these approaches, and what HCI research has uncovered

about both the beneficial and harmful impacts of these approaches on users.

Recommender systems primarily use implicit and explicit user feedback in order to personalize their results. Explicit feedback mechanisms include rating questionnaires or indicating whether content is liked/disliked [36, 39]. Implicit feedback is much more bountiful and involves detecting user preferences through user behavior and how they interact with the recommender system, such as user purchase history, what webpage elements a user clicks on, browsing history and how long a user spends on each webpage, and search history [31], as well as contextual factors, such as a user's location [4].

Both implicit and explicit feedback suffer from performance challenges due to noisy data. Explicit feedback has mixed success even in typical settings because users may not provide these data frequently [39] and because the data they do provide may be biased or noisy [1, 2, 64]. Having users re-rate items may help mitigate issues with accuracy but may not be a realistic solution [2]. Nonetheless, explicit feedback is considered to be more accurate than implicit feedback [1]; however, it is far less plentiful than implicit feedback [31].

Explicit feedback can present a different kind of noise than implicit feedback because it is more “subjective”: some people are more likely to give explicit feedback than others [8], some people provide more expert feedback [9], and some people are more likely to give feedback in certain contexts (e.g., they have strong positive or negative feelings [30] or perceive that a purchase has a high transactional risk [40]). But explicit feedback can be more meaningful than implicit feedback because it uses scales to indicate whether a user liked or disliked content (e.g., a user gives a song thumbs up or thumbs down); whereas, implicit feedback typically only gathers positive feedback, in other words, a user's “degree of preference” (e.g., because a user listened to a song multiple times, they likely want to hear it again more than they want to hear a song they have listened to once) [31, 35].

Both kinds of feedback have privacy concerns, such as of data breaches by employees or third parties and re-identification of anonymized data, and, as we explore in Section 6.1, implicit data is especially troubling in terms of user consent [28, 37, 49].

2.2.1 Sensitive social media content in HCI research. Sensitive content presents recommender systems with additional challenges related to feedback. Recommender systems are typically designed to give people a “good match”/“more like this”—however, this is not how we experience the world when it comes to sensitive content. As we discuss in this paper, users may also find that the content is neither entirely negative or positive, which makes it difficult to decide whether to provide feedback on whether they liked or disliked it. Furthermore, research has shown that users have a wide range of experiences with sensitive content.

Recommender systems are sometimes designed to provide users with serendipity by showing them unexpected or new content that evokes a positive emotional response [44] (i.e., an unexpected moment can brighten a user's day), but emerging work in HCI has also documented the negative experiences that can occur when

unexpected encounters with sensitive content occur, such as unexpectedly seeing content related to an ex-partner [53] or a deceased person [11, 38].

In contrast to some of the negative effects that users sometimes (but not always) experience when unexpectedly encountering algorithmically curated sensitive content, users may experience greater well-being when they intentionally create and share sensitive content. Research suggests that users may find they receive social support when they make sensitive self-disclosures on social media [13], such as sharing experiences with depression on Instagram [3] or coming out as trans on Facebook [29]. Researchers have identified that a variety of contextual factors, including a user's location/setting, influence how users use digital tools to interact with social support systems (e.g., users are more likely to engage with their online support systems when they are physically located in places that are private and comfortable) [12]. In Section 5.2, we also identify that location and other contextual factors are important when *viewing* sensitive content.

Much HCI research on mental health on social media has been devoted to detecting emotional distress due to mental health issues through automated systems (e.g., [18, 19, 59]). However, researchers have argued for a more nuanced analysis of multiple viewpoints and the myriad ways that users share and interpret content related to mental health and illness on social media, and they point out the ways that classification systems fail to take into account these nuances [25]. We aim to contribute to this conversation by recognizing that users have a multiplicity of experiences with how social media interacts with their well-being. In this paper, we also explore how automatically detecting whether content is bittersweet is both a technical challenge and an ethical dilemma.

2.3 Bittersweet emotions and nostalgia

In this section, we turn to the psychology research as a starting point for thinking about bittersweet as a concept, although these conceptions have limits that we discuss in this section. Throughout this paper, we bring this literature in conversation with HCI and critical algorithm studies research, and we discuss how mixed emotions create challenges for recommender systems.

The concept of “bittersweet” emotions has been explored in the psychology discipline as “mixed valence emotions” [61] or “ambivalence” [58], which are described as feeling both “happy” and “sad” simultaneously. There has been some debate among emotion theorists about whether it is truly possible to feel more than one emotion at once² or if people simply vacillate between emotions [45]. In this literature, bittersweet emotions are also tightly coupled to the concept of “nostalgia” which has taken on new meanings over the years [63]—where nostalgia was once considered to be “bitter” and was pathologized, it now is often described as a “bittersweet emotion” [6].

As we describe in our findings in Section 5.1, the process of reflecting on the past can change the narratives people have about their life experiences—reflection can bring about both redemption sequences (in which a person's “narrative progresses from a negative life scene to a positive or triumphant one” [63]) or contamination sequences (i.e. “the narrative moves from an affectively

²Our participants generally held this viewpoint.

positive life scene to an affectively negative one” [63]). Nostalgia typically brings about redemption sequences rather than contamination sequences [58, 63], which can have well-being benefits. Nostalgic feelings are thought to often be triggered by disruptive life events or negative affect (especially loneliness [65]), but they generally help buffer people from these negative experiences and ultimately have a more “pleasant” or “positive” effect because they allow people to reflect on their perceptions of their “true selves” [5], and they have social and existential benefits [58]. However, for people who tend to worry, nostalgia may be more bitter due in part to the dissonance between how a person feels in the present and their past feelings [60].

The approaches to understanding bittersweet emotions that we discussed above require emotions to be framed as things that can be measured and operationalized. Like other critiques around how algorithms classify sensitive content—of algorithms’ lack of nuance and context [25], and of how the discourses used to describe these systems can be dehumanizing [14]—we think there are nuanced ways to discuss bittersweet emotions in our context that honor the complexity of human emotions. We think of bittersweet emotions as highly contextual and evolving and beyond the dichotomies of bitter/sad and sweet/happy. This concept of bittersweet resists quantification or categorization; conceptualizing bittersweet content in this way allows us to look at a wider range of human experiences than we might in the design of a recommender system or in a study in which emotions are quantified. We use Facebook Memories as a case study of the complexity of the ways that human emotions and recommender systems interact.

3 BACKGROUND: FACEBOOK MEMORIES

Facebook Memories (see Figure 1) provides a compelling site for studying TMRs in part because Facebook has such a large and diverse user base and a long history of using TMRs (and thus, users are familiar with these features). Furthermore, Facebook has a rich repository of content to present regularly to these users—thus we were easily able to find users that had experiences with sensitive content. A Facebook Memory is typically a past post from a user that Facebook resurfaces. The most well-known of the Memories features is “On This Day” (OTD), which uses a recommender system to select and show a user a post they made on Facebook on that same day in a prior year. People have the option of sharing the post with others and adding commentary. Our study focused on this feature. Other types of Memories features include or have included Father’s Day and Mother’s Day memories, Year in Review videos, anniversary memories with a friend, and various personalized videos like Friend’s Day, Say Thanks, and Lookback.

People do not always see Memories when they log into Facebook, but when they do, they are displayed at the top of their News Feed. People also have the option of going to the Memories Home page that shows them all of their Memories for that day — including Memories that were not selected for display on the News Feed. Settings on Memories Home also allow people to select specific users and time frames to exclude from their Memories, and well as controls for Memories related notifications.

4 METHODS

We recruited participants via a screening survey on Facebook and then interviewed participants about their experiences with seeing sensitive content via the Memories feature. We then analyzed the interviews with thematic analysis and subsequently identified “bittersweet” as a sensitizing concept. We used our analysis to explore the contextual factors that make content bittersweet and what participants believed about why they were recommended bittersweet content.

4.1 Recruitment and interviews

In 2019, we interviewed 20 Facebook users from the United States who had seen a post with sensitive content through a Facebook Memory. To recruit interview participants, a standard survey invitation (displayed at the top of a user’s News Feed) was broadcast to Facebook users who had seen a Facebook Memory in the past day on their News Feed or on the Memories Home page. The survey invitation asked if users would like to share their opinion. Clicking on the invitation took participants to an in-platform survey screen. The survey asked participants to describe what the Memory was about in an open-ended text box and then indicate how they felt about it using the options: “Good”, “Bad”, “Both good and bad”, or “I did not feel anything.” Participants were then asked to explain their response and were given the options: “the memory was painful”, “the memory was enjoyable”, “the memory was uninteresting”, and “I don’t like looking at memories on Facebook.” Participants could select multiple options. Finally, participants were given the option of leaving their name and email address if they would like to be contacted for an interview. The survey data were only used to recruit participants and inform us about what issues might come up before each interview.

Of those users who had said they would like to participate and had shared their email address, we followed up with people who reported viewing Memories about common but sensitive life events, told them about the study, and asked if they would like to participate in an interview. We had originally planned to only recruit participants who had seen a Memory related to a deceased loved one, but expanded the scope of our study to focus on Memories related sensitive life transitions. This decision was in part because of the types of responses we received through the screening survey and in part because it allowed us to extend HCI literature that has focused on sensitive content curation related to a single type of life event (e.g., breakups [53]). Two of the authors discussed this new study design decision extensively and selected participants together—all of whom had experienced a life transition typically associated with grief (described in Table 1). We scheduled interviews as soon after their encounter with the Memory as possible, with the majority being interviewed within 48 hours. The first four interviews were conducted by two of the authors, and we iteratively refined our protocol based on those initial interviews. The rest were conducted by the first author.

Our 20 interview participants’ ages ranged from 21–69. Participants were located in the United States, predominantly from the South and Midwest. When asked about their occupation, participants said they worked in marketing, sales, medicine, education,

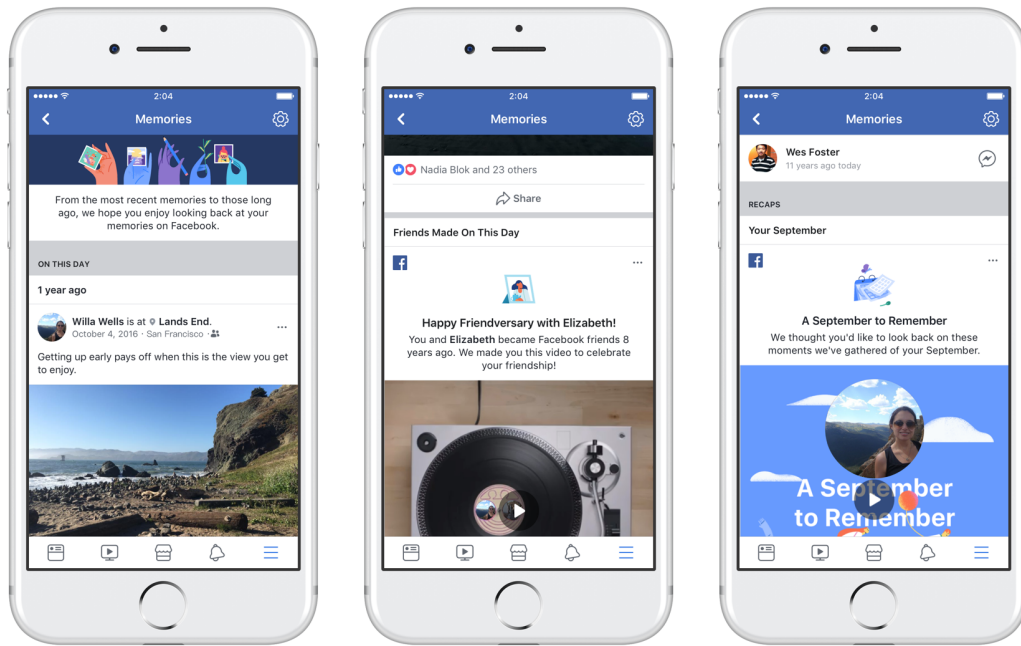


Figure 1: Screenshots from 2018 of Facebook Memories. (Courtesy of Meta.)

were stay-at-home parents or partners, and/or retired. Four participants also said they were disabled. See Table 1 for additional demographic information including age, gender identity, and details about their life transition and encounter.

We spent approximately one hour with each participant over video chat or phone and asked them questions about the content of the post they saw, details about the life transitions they experienced, the context in which they saw the post, how they felt about the post, their thoughts about why they had been shown the post, and what could improve their experience.

4.2 Research ethics

Prior to conducting this work, we developed a protocol for halting the interview if participants showed distress based on the recommendations in [21] and [41], and the last author had extensive experience with conducting interviews on sensitive topics. Interviewees were also informed that the interview might be emotionally difficult and that their participation was voluntary and they could choose to stop participating at any time. Interviewees gave us permission to share their stories and quotes. We encouraged participants to speak about whatever they found most salient as opposed to adhering strictly to our interview protocol. We did not need to stop any interviews, but we did send follow up emails to thank all participants and also did an additional check in with one participant after an

interview that was particularly emotionally intense. Participants were given a \$75 Amazon gift card as compensation for their time.

We have pseudonymized all participant names, and we strove to include enough details to tell their stories in a way that honors the complexity of the participants' experiences without presenting them in a sensationalist manner.

We had a sensitivity reader review our story about Jeremy at the top of this paper, and we followed their feedback to omit some details. Per the sensitivity reader's request, we made a donation to a relevant organization to thank them for their time and effort. We also consulted with HCI researchers of sensitive topics to learn about how to conduct researcher self-care after difficult interviews.

4.3 Analysis

We adopted an inductive approach to our analysis drawing on constructivist grounded theory [16]. Our analysis involved iterative rounds of open coding [55] and discussion to identify and characterize the themes presented in our results [10]. When we began the study, we were initially examining how users were negatively impacted by encounters with sensitive content; however, in early analysis we noticed that a large set of experiences with Memories were both negative and positive—what our participants referred to as “bittersweet”. We chose to explore this type of experience further in our analysis. The codes we developed could be grouped into: types of Memories, what action the participant took after seeing the

Pseudonym	Gender	Age	Life transition (LT)	LT occurred	Post made	Reaction
Roger	Male	41	Difficult living situation	4 years ago	4 years ago	Bad
Nic	Non-binary	21	Estranged family	5 years ago	5 years ago	Good and bad
Scarlett	Female	33	Divorce	1 year ago	8 years ago	Good and bad
Beth	Female	42	Death of friend (murder)	0.5 years ago	Data unavailable	Good
Nelly	Female	48	Death of husband (cancer, sudden)	2 years ago	7 years ago	Good
Susan	Female	46	Death of mother	3.5 years ago	3 year ago	Good and bad
Trish	Female	53	Breakup	7 years ago	7 years ago	Good and bad
Mark	Male	48	Death of wife (suicide)	1 year ago	6 years ago	Good and bad
Stella	Female	38	Death of dog (accident)	1 year ago	2 years ago	Good and bad
Charlene	Female	23	Death of cousin (suicide)	1 year ago	3 years ago	Good and bad
Bridgett	Female	30	Divorce	0.5 years ago	6 years ago	Good and bad
Jeremy	Male	36	Death of father (cardiac arrest)	18 years ago	3 years ago	Good
Marianne	Female	37	Life-threatening illness	3 years ago	3 years ago	Good and bad
Eunice	Female	69	Family members' terminal cancer	Ongoing	2 years ago	Good and bad
Dana	Female	34	Death of friend (cancer)	1 year ago	7 years ago	Good
Cass	Female	47	Divorce (sudden)	0.5 years ago	6 years ago	Good and bad
Hannah	Female	55	Death of sister (cancer)	1 year ago	1 year ago	Good and bad
Linda	Female	64	Death of daughter (overdose)	7 years ago	7 years ago	Good
Shauna	Female	39	Life-threatening illness	2 years ago	2 years ago	Good and bad
Anna	Female	53	Divorce (sudden)	4 years ago	4 years ago	Good and bad

Table 1: Participants' demographics and information about their life transition and reaction to the Memory. All participants were from the United States, with most residing in the South and Midwest.

Memory, contextual factors (e.g., time of day of viewing), the participant's suggestions for how to improve encounters with sensitive content, and the participant's theories about how the Memories recommender system chooses posts. We then wrote memos, and we discussed which issues cut across the memos. We identified that participants' expectations, contexts in which they viewed the content, and how they made sense of the Memories recommender system were key components of how they navigated encounters and how the encounters impacted them.

5 FINDINGS

We found that participants often described the content they saw as "bittersweet", which we found to be an important sensitizing concept that allowed us to explore the breadth and complexity of people's experiences with TMRs. We first describe what that concept meant for our participants and how it could evoke positive and negative emotions. Even when participants felt negative emotions, they typically wanted to see the bittersweet content—but whether an encounter was (un)wanted was dependent on multiple factors, including whether it was an (un)expected encounter with the content, whether the content was viewed in an appropriate context,

and whether the participant's understanding of how the recommender system worked aligned with how it behaved. We discuss the actions that participants took based on whether the content was more wanted or more unwanted. In Section 6, we reflect on known challenges in recommender systems that are further complicated by people's often contradictory and complicated reactions.

5.1 What makes content bittersweet

We found our participants' use of the word "bittersweet" to be powerful because the concept allows us to think beyond dichotomies of bitter and sweet. We do not envision bitter and sweet as a Likert scale or a spectrum, but two spectra (bitter/not bitter and sweet/not sweet). We think of bittersweet as just one type of "sensitive content", which we described earlier as content that has the potential to evoke a strong emotional response. Often content that participants described as bittersweet tended to be both intensely bitter and intensely sweet.

"Bittersweet" was described in diverse ways. Participants spoke of content that was simultaneously happy and sad, content that made them feel a happiness twinged with sadness, and content that made them feel a sadness twinged with happiness. Neither was bittersweet always a static or inherent quality of the content:

Feelings about content can shift over time and in different contexts—content could be bittersweet in one moment and not another. These varied experiences are similar to the multiple, sometimes contradictory ways that bittersweet emotions and nostalgia were theorized in Section 2.3. Although how participants understood and experienced bittersweet emotions differed, it generally was associated with movement through emotions: content was not bittersweet when it was originally posted but became so over time.

Content was often bittersweet when it had changed emotional valence after it had been posted (e.g., something that was sweet at the time was now also somewhat bitter or something that was bitter at the time was now also somewhat sweet). Encounters with bittersweet content afforded participants the opportunity for reflection about a difficult life transition, and often that reflection prompted the change in emotional valence. TMR systems support modes of reflection that would otherwise be difficult—they provide serendipity and unexpectedness because users have the opportunity to see content that they would otherwise have to deliberately seek out. For example, our interview with Stella demonstrated how the seeing the Memory post led her to remember how different her life was six years ago when she was regularly in chronic pain.

Stella: Sometimes it's like, "wow, what was going on at that time?" I hadn't had any bouts in probably, I would say, close to six years now. And it amazes me when I see the fibromyalgia posts, how much I was hurting. And it's hard for me to remember how well I'm doing now compared to then. It's so different now. I'm not married anymore. I don't have all the weight that I had on me. I don't have all the stress that I had. I'm not hurting like I was. I was thinking about how it was then, and it's kind of powerful.

Reflections on difficult life transitions, like Stella's, often help people to reframe negative events as positive events that led to growth (e.g., what the psychology literature refers to as "redemption sequences" [63]), and participants often expressed that they were thankful for the support they had received during the life transition and grateful that they were no longer in that place emotionally or physically. However, the opposite ("contamination sequences" [63]) were not infrequent—initially positive experiences, such as a photograph of a vacation with an ex, now caused participants to have mixed feelings because they were reflecting on it with new information (e.g., that the ex was cheating):

Cass (about an ex-husband who cheated on her): It just brings up stuff, what you thought life was going to be like now, and then it's just so different than what you thought it was going to be. [...] It makes the memory bittersweet. It doesn't make it horrible, but it just makes it sad now because it reminds you of what you no longer have.

Content that evoked contamination sequences, like Cass' encounter with a picture of her ex-husband on their vacation, could be difficult for participants. While it was a bittersweet experience for Cass, a similar experience would have been just bitter for some of our other participants. Even though we saw that participants generally did want to see sometimes bittersweet content regardless of whether it was difficult, participants generally did not want

to see content about exes or estranged people, often because the memories were significantly tainted by a sense of hurt or betrayal that other kinds of bittersweet content were not.

While encounters with bittersweet content were often experienced as more positive (redemption sequences) or as less positive (contamination sequences) than a participant's original experiences, a common factor across all the bittersweet content was that it was transformed by a participant's reflections and knowledge that they had gained since their original experience.

As we discuss later, participants often discussed bittersweet content with others to process it, and they were able to reflect on this content with us in the interviews as well. While sometimes a participant's encounter with sensitive content evoked bittersweet emotions right away, the encounter was sometimes only became bittersweet after they reflected on the experience of encountering it. Thus, we see bittersweet content as not just content that is bittersweet in the moment of viewing, but bittersweet over time. We later discuss how participants tried to anticipate when they might see bittersweet content in order to ameliorate some of the more difficult aspects of these encounters.

5.2 Factors that influence whether an encounter with bittersweet content is (un)wanted

Participants had a range of feelings about whether they wanted to have encounters with bittersweet content based on three main factors: (1) whether they expected to encounter the content, (2) the context of their environment when viewing the content (e.g., what time of day they saw the content), and (3) their beliefs about how the recommender system decided to curate the content, which were often based on how well they thought the recommender system could understand contextual factors and on their perceptions of the affect of the recommender system.

5.2.1 (Un)expectedness of the encounter. Most participants knew that there were algorithms that curated their News Feeds and that this curation would create unexpected encounters. However, they often went to their News Feeds to see their friends and groups' posts, not with the intent to reflect or reminisce. Therefore, they found some encounters with Memories to be unexpected because the encounters did not align with their intent when they opened Facebook.

Expectedness often made a difference in how participants reacted when encountering bittersweet content: many participants reported that they would prefer not to be surprised by a sensitive Memory at the top of their News Feed.³ Participants often did want to see—and even expected to see—sensitive content on dates that were significant to them, such as seeing content about deceased person on the person's birthday, anniversary of their death, other anniversaries, and holidays. For example, Dana looked forward to seeing a video of her grandfather every year on Facebook.

Dana: [It] goes back to my faith system and believing that I will see [my grandfather] again [...] I learned so much more about him [through the video shown at his

³While almost all of our participants saw Memories pop up at the top of their News Feed, four of our 20 participants actively sought out Memories by going to the Memories Home page as part of their daily routine, and a fifth participant occasionally looked at the Memories Home page.

memorial] and then just getting to see the conglomeration of all the different pictures and it's a celebration of his life. I mean I cry of course, but sometimes those are good tears, those are happy tears, but yeah, I just, to take the time to remember my grandpa by watching that video it feels very special and painful to me.

For Dana, it was almost a ritual to take a certain time of year to remember her grandfather, but she did not want to spend other times of year doing so because "he wouldn't want us to be sad frequently throughout the year." Dana would prefer not to see content about her grandfather other times because it would cause her to feel complicated and intense feelings at times that she had not intentionally set aside for them.

5.2.2 Context of viewing: time of day, location, and mood/headspace. The context of viewing had an important effect on the reception of the content itself and often influenced whether the content was experienced as bittersweet. As discussed earlier, there is an aspect of nonlinear movement into/through bittersweet emotions, and changing contexts can influence its direction. We thus see content and context as distinct, but intertwined aspects of bittersweet emotions. In the following, we describe what we mean by "context" in further detail.

The time of day was a major factor in how participants felt about posts. Memories became part of their routine and an integral part of some participants' experiences with Facebook. Participants usually viewed Facebook on their phone, and they usually saw the post in the morning as part of their routine.

Cass (about her ex-husband): I'm usually sadder in the morning. I don't know why. So, if it's first thing in the morning and that's the first thing I see when I'm looking at Facebook while I'm drinking my coffee or eating breakfast, it might have the tendency to set the tone for the day, and just made me sad for the day. Especially if it was a really great, great memory, and I'm like, "God why did you [the ex] fuck this up?"

Some locations, such as a workplace, were less ideal for viewing sensitive content:

Susan: [The picture of my deceased mother] just kind of caught me, and I sat there and thought about different things for a little bit. Then I had to get back to work. But I was distracted the whole time.

Another contextual factor was the mental health, mood, or headspace of the participant at the time that they saw the content:

Beth: So I have depression and anxiety. I think maybe if I was in the middle of an episode or something like that, you know what I mean? I think it would be a perfect storm of really horrible stuff happening.

Some participants said that they consistently reacted more strongly to sensitive content than most might because they had behavioral health conditions (e.g., anxiety disorders, personality disorders), which are associated with rumination [47]. As discussed earlier, people who ruminate focus repetitively on negative aspects of a memory, which causes them to feel worse.

5.2.3 Understandings of algorithmic "affect" and the limitations of algorithms. How participants understood behavior of the Memories recommender system impacted their expectations about when they would see bittersweet content. They often tried to determine what emotion the recommender system was trying to convey—which we refer to as "affective sense-making".

Affective sense-making occurred when participants ascribed affect to the recommender system—both in terms of how a recommender system might make sense of a participant's emotions and in terms of how it presented the content to them—and they used their theories about the recommender system's affect to make sense of its behavior. They felt that it did not make space for their emotional responses in the ways that a close friend or therapist would. In human interactions, we have a bidirectional outlet to express our feelings—the person we are speaking with understands the context of the conversation and we can have dialogue that can help with sense-making [17]. Furthermore, the other person's goal is not necessarily to make us happy but to make space for us to process our feelings. However, some of the participants theorized that the recommender system was simply designed to show them content that would make them happy.

They felt negatively when the information was presented to them in a way that suggested a recommender system had misinterpreted bittersweet content as "happy" content. Some participants had beliefs about the limitations of the recommender system which softened the blow when there was a mismatch in what they wanted the recommender system to (not) show them and what they were shown. Their reactions were complicated, messy, and often contradictory—leading to behavior that could be difficult to interpret using feedback mechanisms.

At times participants were upset when Memories showed them content they did not expect, and they believed that it did so because the platform was designed in such a way that it made incorrect assumptions about what they would want to see.

Mark (about his wife who died by suicide): A handful of times in the past year [there] has been a photo that has either a picture of her or an association to something we did. You might be able to imagine, each and every time that happens it's sort of a bittersweet kind of situation where, on one hand you're like, "Oh, here's when we did this thing" or "I'd even forgotten about this. Oh, isn't that pleasant? Oh, by the way, here's a reminder that your wife died a year ago."

Mark later suggested that perhaps others find this suite of features to be "entertaining". Along with many other participants, he hypothesized that the Memories recommender systems were designed to only select happy and/or significant Memories. Nic, for example, shared an experience where Memories shared a photo of them playing a board game with estranged family members. As a student/web designer, Nic had put a significant amount of thought into how Facebook's algorithms might go about ensuring Memories were positive. And indeed, the Memory in that photo had been positive, initially, and Nic was quick to reason that that was why the photos had been presented. Participants like Mark and Nic demonstrate how reactions to encounters with bittersweet content

are the result of both the content and one's understanding of why the content (even erroneously) was presented.

Thus they felt the "tone" was inappropriate, even if the content was presented in a neutral way—just the act of showing it to them, rather than any particular aspect of the interface, was perceived as being cheerful. These participants were subsequently surprised when a Memories post had a negative association or was a reminder of things that they found inconsequential or unimportant to them.

Nic: Sometimes Facebook will try to pull up the positive memories with someone who you no longer want to associate with, which I like the focus on the positivity but sometimes the person themselves is the problem.

Participants did not feel that much could be done about seeing unwanted Memories because they felt the algorithms were just not "smart" enough to detect contextual factors that were highly individualized to them. While Facebook does allow users to list people and date ranges that should be excluded from Memories, our participants also discussed more subtle and contextual factors that cannot be easily articulated to a recommender system. However, the participants also acknowledged how complex and difficult those factors would be to identify. Participants like Nic felt that it was impossible for Facebook to know the context of a post (e.g., whether the picture was taken in a place that they had bad associations with) because it was not technically feasible. And participants repeatedly told us that there was almost no way that Facebook could possibly know that a particular Memory was bittersweet (e.g., in the case where a breakup was not reported on Facebook). Furthermore, sometimes content is sensitive for very indirect reasons:

Dana: [Making the algorithms more sensitive] would be almost impossible because you would never know what triggers those things. Some of the big things that I think kind of define my grandpa were skiing and honeybees. I mean how are you supposed to know that an ad for honey might be a trigger, you know?

In many cases, participants seemed more accepting of sensitive content because they found the perceived limitations of recommender systems to be understandable—in other words, they may have been more forgiving when the recommender system showed them upsetting content because they thought they knew why it had happened. For example, Eunice told us about a picture of her father-in-law with her grandson that she found to be bittersweet, but she did not fault the recommender system because "*How [is Facebook] supposed to know that my father-in-law, that was the last picture of him with the baby, you know?*" Not only would it be difficult to determine that it was the last time that her father-in-law saw his great-grandson, but it would be difficult to determine how Eunice felt about it. She generally did not feel sad about seeing sensitive content on Facebook, but she told us that she found this particular Memory both "happy and sad" because she wished that her father-in-law had seen her grandson grow older. Participants also felt that Facebook's recommender systems could not know the contextual factors that we described earlier, such as a user's mood:

Stella: There's no way Facebook can prevent [me from seeing something when I'm feeling really down], because Facebook, it's an online app. It's not you in

life. It's virtual. It's something that you go to. It's not something that you're going through.

Stella's comment also brings up salient ethical questions about the obligations that designers of a platform have to their users (e.g., should a platform be trying to "prevent" users from seeing upsetting content?), which we further explore in Sections 6 and 7. Additionally, Stella raises important questions about how the role of a platform in a user's life impacts their reactions to bittersweet content. While Stella thought of Facebook as something separate from her "real life", for some participants there was not a clear separation. How participants thought of the role of Facebook in their life may have impacted their reactions to misaligned expectations about the behavior of the recommender system.

5.3 Actions taken depending on whether an encounter with bittersweet content was (un)wanted

How participants navigated bittersweet encounters was influenced by whether it was a wanted or an unwanted encounter. We found that when participants felt that they did not want to see bittersweet content, they would often delete posts or they would take breaks from Facebook. However, more often than not, we saw that participants felt positively about the encounter with bittersweet content. When participants did want to see the bittersweet content, they often shared the Memory with a select group of people because they wanted to experience the content in a more private setting with people who would understand its significance and could provide emotional support. Sometimes the ways in which people shared content was akin to how they would share a physical artifact (e.g., photo album), and we further explore the implications of these practices in Section 7.1.

5.3.1 Actions taken for more unwanted content often involved disengaging or venting. Usually participants scrolled past unwanted content and went on with their day, but sometimes participants deleted the post if it was about a topic they never wanted to see again (usually an ex—we had no participants mention deleting content about a deceased person), and one person mentioned that they used it as a kind of curation tool to identify which posts needed to be deleted.⁴ However, in general, participants did not delete posts because they felt that they might want to view them at a later date even if they did not want to see them currently. For example, some mentioned that they do not delete pictures of exes because they want to share them with their children (sometimes saving them for when their children are older).

When participants did share posts about things they did not want to see, they did so out of frustration or amusement that the recommender systems would do such a poor job of predicting what they would like to see. As discussed in 5.2.3, they felt that the algorithm was not "smart enough" to avoid these mistakes. For example, Scarlett, who saw a post about her ex-husband, said, "*I thought it was just ironic and I was like, 'this is stupid', so I just*

⁴Hiding posts was another option, but only one participant mentioned hiding posts. It is unclear whether participants were unaware that they had the option to hide posts, whether they thought that deleting and hiding were the same thing, or whether hiding was simply an uncommon action. It is possible that people find scrolling by more effective than dismissing a post.

posted it” and that “it was more of an annoyance [towards Facebook]. I didn’t really feel anything towards the relationship. It was just an annoyance, like “it’s 8:00 in the morning and I don’t want to see this.”” Scarlett’s friends also responded by saying they were “cracking up” and responded with “haha” reacts. For participants like Scarlett, one reason for sharing was to express annoyance (and perhaps even some excitement about finding a “bug” in the system) about Memories’ recommender system’s poor choice in selecting a post due to how they perceived the goal of the recommender system: to show them content that they *wanted* to see.

While some participants coped with unexpected sensitive content through humor, others shared the post out of frustration and sadness. For example, Hannah shared a Memory about her deceased sister because she wanted to gain the support of people who understood why it was painful to see the post.

Some participants even reported that people they knew would stay off Facebook for periods of time, such as a month they associated with a suicide attempt or the death of a loved one, in order to avoid these kinds of encounters.

Anna: I have a friend whose dad died in August two years ago... and when it comes close to the anniversary of his death, she just gets off Facebook for a couple weeks because she doesn’t want to relive the memories... I have a friend whose son died in September six years ago, and she just stays off Facebook for the month of September. So, there are people who do stay off of it because of bad personal memories. I respect that.

Instead of avoiding Facebook during times of the year associated with difficult events, some participants would only avoid going on Facebook on specific days in order to avoid encounters with sensitive content (e.g., the deathday of a loved one). In this case, being able to expect the encounter gave the participant the ability to make a decision about how they would like to engage with past content.

Mark: I might have been on very briefly [on the anniversary of my wife’s death], but there was a conscious choice to like, “I’m not going to hang around here. I’m just checking the direct messages. That sort of thing.” [...] The Memory feature was part of [the reason I stayed off Facebook]. [Also] it was not a day where, “Oh, I wonder who’s saying what and what are they sharing? What pithy comments are being made today?” That sort of thing.

In extreme cases, when there were simply too many posts to delete that appeared throughout the year (such as pictures that were taken with an ex that a person had been with for many years), participants said that some people they knew deleted their accounts and started new ones.

Hannah: I actually did share [a picture of my sister who was seriously declining in health a week before her death] because I just sort of said, “Oh wow. Thanks Facebook for reminding me of this sad week.” That’s when my other friend popped up and said, “Now do you see why I deleted my old account and got a new

one?” Because her husband, I think, died of brain cancer, and she just didn’t like the constant reminders.

Hannah knew people from through brain cancer advocacy groups who could not handle the constant bombardment of posts about deceased loved ones. In this case, the content tipped over from bittersweet to simply bitter.

5.3.2 Actions taken for more wanted content often involved discussing it with a select audience. Some participants moved on with their day after the encounter and did not continue to think about the bittersweet content, but others talked about it with loved ones through mediums that allowed them to have more private conversations, such as instant message, Facebook groups, or face-to-face. When participants discussed them with someone in person, sometimes it was because the other person was present when they saw the Memory and sometimes because they wanted to discuss them in person because it felt more appropriate to do so.

Dana: Social media and texting is not always the best way to communicate the way you feel about something. There’s something to be said for an in person or face-to-face discussion [...] Sometimes you know you share something and you want to talk about it a little bit more, but texting and Facebook isn’t always the best avenue for it, sometimes it’s just best done face-to-face.

Dana had taken a screenshot of a post that her deceased friend had written and sent it to one of her friend’s two sons, whom she was close to. (Dana took a screenshot of the post because the friend’s son was not active on Facebook.) She chose that particular son because she knew she would see him later to “check in on him” in person. She was not concerned about his reaction, but wanted to share because “it would be something that he would appreciate getting since he wouldn’t necessarily have access to some of the words that she had spoken”, and she wanted to share in that joy in person.

Linda engaged in a similar practice. Before her daughter had overdosed and died, she had defriended some of her longtime friends. After she died, the friends wanted to see the content on Linda’s daughter’s account to remember her, but were unable to because they had been defriended. They would come over to Linda’s house sometimes and she would pull up her daughter’s account and they would view it together.

Others shared the content with select groups of people because they wanted to limit the audience of their post or for privacy reasons. In the case of Anna, this decision was related to her safety. She still wanted to engage with the content, but could not do so publicly without fear of backlash.

Anna: [I share Memories of pictures with my kids over instant message because] I’ve learned that if I share it on my wall, there are people who can’t stay out of the middle and automatically text [my abusive ex-husband whom I have a restraining order against], call him, or forward it to him, and then I get nasty phone calls and stuff [from him]; so I just like to keep it private, between the kids and I.

And for other participants Memories were shared with select groups of people because the participant felt that they should only be shared with people who would really appreciate and understand the Memory. Furthermore, they worried that other people might judge why the participant was sharing it (e.g., they might think that the person was not “moving on” or engaging with the content in a “healthy” way).

Beth: Even though [my deceased local celebrity friend] was super popular in [large Midwest city], she still only had a very tiny, small group of friends that she was tight with. It's a group that we've had in Messenger since we were planning her funeral... So we just always share our little stories or memories in there. I don't want to make [the general public] sad all the time, you know what I mean?

Beth shared the video of her best friend singing happy birthday to the city she lived in a private chat group with others who had been close to her friend. Her best friend's murder made the news, and people inundated her friend's Facebook page with comments shortly after her death, some of whom Beth felt were insincere. Thus, she wanted to share the video with people who were close to her friend and would respond appropriately. She did not want to share it to make a “spectacle” of it. Some participants had support groups (both in-person and on Facebook) that they were part of and felt that they could talk with people from those groups about the Memories they saw.

Hannah: I think it's kind of because sometimes you can't talk to people in your ordinary life about things, so you have an audience that you already know, that, “Wow I know about 50 people in the brain tumor community, they're totally going to get this. And they're going to be able to relate to it.”

Another reason to share posts was that they conveyed something important to the participant that may have been difficult to write about again—in other words, they do not have to create new content about something difficult.

Hannah: You know, I knew [my sister's death day], and that was more expected and I was going to post a picture of her anyways, and so it was kind of nice because you actually have something you've already posted and you can just go, “Oh yeah.” You know? Thanks everyone and remember my sister and stuff like that. I didn't have to go recreate anything because I could find a picture very easily.

As described in this section, sometimes people with different emotional reactions engaged in similar behavior (e.g., sharing a post out of frustration and sharing a post out of joy), highlighting one the challenges of detecting bittersweet content and obtaining feedback on it. We further explore the design implications of sharing with a select audience in Section 7.1.

6 TWO CHALLENGES FOR THE RECOMMENDATION OF BITTERSWEET CONTENT

Bittersweet content presents sociotechnical challenges for detection and feedback. These challenges are the result of the diversity of ways that people experience, engage with, and even conceive of this content. It breaks the assumptions of systems that presume that content is desirable or not, and that make decisions accordingly. Bittersweet content does not fit cleanly into the thumbs up/thumbs down model of personalization and recommender systems, and the features needed for a recommender system to make choices are not always clear, nor are the data always available.

In this section we detail the challenges identified in our analysis. Then, to address these challenges, in Section 7 we then adopt a more holistic perspective around the design of such systems, focusing on the issues identified in Section 5.2: expectations, context, and sense-making.

6.1 Challenge 1: Detection of Bittersweet Content

In HCI literature, we tend to focus on designing and improving upon designs—but it is difficult to make design recommendations when the core of users' experiences with a system is dynamic, such as with curated content. Not only is the behavior of machine learning technology dynamic, but the user's contexts, desires, and situation can change. What use case scenarios are designers able to test for when building curation features like Instagram suggested posts or Amazon purchasing recommendations? How can designers make sure that they are testing for all the scenarios that they need to be? Bittersweet content is by its nature hard to account for. Because few participants said that they did *not* want to see bittersweet content ever, designers are left with a challenge of identifying what content might be sensitive, and accounting for contextual factors about the user. We found that users had different needs at different times of day, when viewing the content in different locations, and when they were in different headspaces. They wished there were ways for TMRs to take these contextual factors into account. However, their preferences were highly individualized. While we can imagine various signals that could enable recommender systems to account for such factors, many are not technically, socially, or ethically viable. Take for example headspace: There are no sensors for headspace like there are for location. Likewise, relying on users to constantly provide such information through direct-report would be cumbersome, even if we imagine it would be effective. Moreover, inferring headspace or mental health based on other signals might present privacy concerns or violate users' assumptions about how their data are used [15].

A potentially even greater challenge is posed by the associative nature of memories. Content may trigger unexpected memories about sensitive topics that are seemingly unrelated. For example, participants found that the location where a photograph was taken may have negative associations even if someone who used to be part of the participant's life was not in the photo (e.g., the place may have been a frequent location of dates with an ex). Can systems be designed to take into account that sensitive content might be sensitive only because it is indirectly related to another memory

(i.e., it triggers that memory)? How can a recommender system take that indirectness into account, and how many “hops” in association should it take into account? While there are no easy answers to these questions, we offer them as a way of highlighting the complexity of this space—as well as the limitations when recommender systems are tasked with handling nuanced content.

6.2 Challenge 2: Understanding Feedback on Bittersweet Content

Relying on user feedback to tailor curation becomes challenging when we consider ways that content might be emotionally impactful. Sensitive content as a whole presents a challenge here: It is not reasonable for a system to show someone a piece of upsetting content and then expect them to reliably provide feedback that they have been hurt. Even in the case of bittersweet content, where participants described momentary breaks from the platform and long periods of reflection on the content, it is unclear when it might be most appropriate to solicit feedback.

Asking users for feedback on recommendations is further complicated given that what users want can be layered, quickly evolving, and even contradictory. As a result, disconnects between system behavior and user preferences is likely. Prior work has noted these disconnects in other contexts. For example, a person might miss content from a friend who needs emotional support because an algorithm “assumed” they were not close [24]. A person may feel worse if told they have a behavioral health disorder based on their behavior online [25, 56]). And people with historically marginalized identities can be harmed by algorithmic attempts to debias users’ preferences around dating [32]. With bittersweet content, our data show that disconnects can happen in both directions: users can be negatively impacted by content the system thinks they want to see—and conversely, when recommender systems are overzealous in filtering out bittersweet content, users may not have wanted encounters with content that is meaningful to them.

Furthermore, research has shown that folk theories about how recommender systems work influence user behavior [20, 22], and if users are concerned about these dissonances, they may provide biased feedback in an attempt to mitigate the dissonance (e.g., they may be concerned that if they say that they do not like content, they might miss content they might like to see in the future). Research has shown that negative reactions to the dissonance between what users want and how a recommender system behaves can be mitigated somewhat by explanations, which can help users understand the limitations of recommender systems [23], and similar strategies are worth considering with bittersweet content.

Providing feedback will not be straight-forward. Our current work does not suggest clear ways to improve feedback mechanisms that will result in new data, signals, or features that recommender systems can leverage in better determining *when* to show bittersweet content. However, we do see ways that interaction and interface design can address *how* to show bittersweet content. We turn to this topic in the next section.

7 PROVOCATIONS FOR DESIGN PRACTICE AND IMPLICATIONS FOR DESIGN

In this section, we discuss the three factors that influence whether bittersweet content is (un)wanted (as introduced in Section 5.2): expectedness, context of viewing, and sense-making. These factors are reframed as user needs that can guide the (re)design of recommender systems that curate bittersweet content. For each of these needs, we discuss (1) the supporting findings that we used to identify these needs, (2) suggested design practices for addressing the needs, (3) specific design recommendations for our context that we derived using those practices, and lastly, (4) broader provocations for design, inspired by Baumer and Silberman [7] and generalized from the needs that we found. In addition to using the suggested design practices to develop recommendations for a specific context, researchers and designers can use the provocations to reflect on broader challenges with recommender systems that curate bittersweet content (such as the ones we discussed in Section 6). An overview of this section is provided in Table 2.

Baumer and Silberman [7] urge HCI researchers to “engage in a critical, reflective dialog about how and why [HCI systems] are built.” They ask a series of questions to technologists to “offer one set of techniques for engaging in such dialog”, which we find relevant here, including questions about the viability of non-technical solutions, the difficulty of relying on technology, and how designs often replace human problems with computationally-tractable alternatives. Algorithmic systems present new challenges as the experiences they enable were not always considered by designers. In big data algorithmic systems it is not just the designers’ intentionality—but also training data, user feedback, and commercial third parties—that shape the behavior of a system.

As we noted earlier, it is difficult to fully predict the actions of a recommender system because of this diffusion of intentionality. (In our context, it is difficult to know when and how a recommender system will interact with content that is bittersweet.) This diffusion requires that designers account for a broader set of possibilities that may emerge from their systems.⁵ We believe that the three additional questions that we pose here—adapted from the questions forwarded by Baumer and Silberman [7]—can serve as provocations for designers of recommender systems to reflect on their responsibilities to users. These questions are not just salient when deciding whether to design a technology, but can also guide the (re)design of systems as well:

- (1) How can designs draw inspiration from familiar low-tech and non-tech artifacts where sensitive content is common?
- (2) How can we identify effective solutions that ameliorate the negative aspects of encountering bittersweet content when bittersweet content cannot always be readily identified or negative outcomes predicted?
- (3) How can a technology be designed to match how humans understand a problem rather than a computationally tractable version of the problem?

It is our hope that researchers and designers can use these provocations to find their own design implications for their use cases.

⁵See [62] for summaries of long-standing debates about accountability in algorithmic systems.

Topic	Designing for expectedness: drawing inspiration from non-technological artifacts	Designing for contextual factors: examining the relationship between context and sensitive content	Designing for how humans understand problems: affective sense-making and computational tractability
Finding	People want to be able to expect when they will see bittersweet content	The context in which someone views sensitive content shapes whether it is upsetting	People engage in affective sense-making when they encounter curated bittersweet content
Design practice	Examining the benefits and harms of expectedness and discovery	Moving from simple signals to holistic and contextual feature sets	Using studies on affective sense-making to inform metrics
Design implications	Enabling expectations	Detecting contextual factors	Enabling affective sense-making
Provocation	How can designs draw inspiration from familiar low-tech and non-tech artifacts where sensitive content is common?	How can we ameliorate negative aspects of encountering bittersweet content when bittersweet content cannot always be readily identified or negative outcomes predicted?	How can a technology be designed to match how humans understand a problem rather than a computationally tractable version of the problem?

Table 2: Structure of this section: findings which led to general design practice recommendations, our design recommendations using those findings and practices, and wider provocations.

7.1 Designing for expectedness: drawing inspiration from non-technological artifacts

In this section, we explain how referencing other sites of bittersweet content can inspire design approaches. Following our findings, we focus on increasing expectedness and user agency.

7.1.1 Finding: People want to be able to expect when they will see bittersweet content. Our participants found it upsetting when they did not have control over whether they could expect sensitive content. Research on sensitive content demonstrates that algorithmic harms can result from unexpectedly engaging with our own pasts (e.g., [53]). Drawing insights about expectedness from other contexts—especially non-technological ones—where bittersweet content is common presents a promising avenue. Physical artifacts have different affordances and uses than algorithmically curated content, and we do not want to assume that they are a replacement for recommender systems; however, there are affordances that we can incorporate. Consider scrapbooks: While filled with memories from the past, people make intentional choices to pull the scrapbook off the shelf. They are often seen with loved ones sitting on a couch with the goal to reminisce. And the organization of the photos allows people the ability to anticipate painful photos on the next page, and even skip over a section if they need.

Non-technological artifacts may achieve similar goals as TMRs, however, today people are increasingly documenting their lives through technology, rather than scrapbooks (indeed, Facebook has become a digital archive of co-created memories [54]). As a result, encounters with bittersweet Memories often happen as people browse Facebook idly and see Memories on the top of their feeds. The practice is often an individual one, and when content is interjected into the News Feed there is less ability to contextualize the content or anticipate what is coming.

We saw people emulating the affordances of non-technological systems in myriad ways. Most common was making the *intentional choice* to go to the Memories Home page to browse photos, in some cases with friends or family, huddled around a computer. Others talked about the importance of sharing bittersweet Memories—typically via more private channels such as Messenger, a phone call, or a face-to-face conversation.

7.1.2 Design practice: examining the benefits and harms of expectedness and discovery. One reason recommender systems are compelling is that they help users to discover new content that they might otherwise not [44] — in our case, reminders of things forgotten in the past. Yet, as participants described to us, this unexpectedness can create emotionally upsetting situations in which a user is reminded of something they would rather not confront at that moment. For some participants, intentionally visiting the Memories section would have resulted in a more positive experience than having an old post displayed unexpectedly at the top of their feed. However, participants also found that unexpectedly seeing old content could also be delightful and uplifting, especially during difficult times, such as Jeremy’s experience with seeing a picture of his father.

In order to evaluate the trade-offs between expectedness and discovery, designers will have to consider the benefits and harms of algorithmically curated systems (which enable discovery) relative to their non-digital counterparts (which behave in more predictable ways). Designers must also consider the *differences* between computational systems and non-digital artifacts when conducting these investigations—for example, recommender systems typically have a much larger reach than non-digital artifacts (e.g., only a few people might ever look at a photo album; whereas, there are often many users of a recommender system), and benefits and harms will impact people on different scales. Part of these considerations

will require designers to investigate how people are not only using technological systems but also using non-technological artifacts.

7.1.3 Design implications: enabling expectations. In these design recommendations, we draw inspiration from physical photo albums and scrapbooks while still seeking to give users the option of engaging with curated content.

(1) *Recommender systems could put sensitive content into its own special place for people to view when they are ready.* Sensitive content could be treated more like a photo album that people view intentionally (rather than content that people see without warning). Participants noted that they did not always want to see bittersweet content unexpectedly, but that they did not want to delete the content either—instead, they wanted it to be saved somewhere separate from the rest of their content. Similarly, Sas and Whittaker [57] have proposed capturing sensitive materials that might be upsetting into a “Pandora’s box” that would be a “technology for self-control”. The Pandora’s box would prevent people from impulsively deleting content but would also restrict access so that people could not obsessively keep reviewing painful content. This recommendation is in line with what we heard from our participants, but returns us to the challenges around detection of sensitive content we discussed earlier.

(2) *Provide users with opportunities to initiate expected encounters with sensitive curated content.* Users do not always come to applications with the intent to engage with sensitive content; however, they may want these encounters when they are expected. As opposed to creating recommender systems that give users encounters at unexpected times, we could imagine recommender systems similar to Google’s “I’m Feeling Lucky” button, which would allow users to initiate encounters with curated or random content.

(3) *Provide more channels (and more frictionless sharing) to allow spaces for private groups to come together over the content.* Users may also benefit from more private ways of engaging with sensitive content as a group (which would also support the creation of contexts where bittersweet content was expected). For example, people could use a shared scrapbook page or, like Beth, use group messages to privately reminisce around the bittersweet content.

7.1.4 Provocation: How can designs draw inspiration from familiar low-tech and non-tech artifacts where sensitive content is common? We encourage designers to turn to non-algorithmic systems for inspiration when considering how to address bittersweet content. After all, uncomfortable encounters are not unique to algorithmic systems, or even technology. When considering our previous examples of scrapbooks and photo albums, it is clear how these artifacts allow users more agency in ways that designers can learn from. We do not mean to suggest that there is an abundance of agency in non-algorithmic systems, but simply that engaging with these systems may produce different experiences, some of which may be less upsetting for some people. It is for these reasons that we find turning to familiar alternatives as a useful point of comparison.

7.2 Designing for contextual factors: examining the relationship between context and sensitive content

We argue that understanding a user’s contextual factors and personal definitions of “sensitive” will help researchers and designers to create better recommender systems.

7.2.1 Finding: the context in which someone views sensitive content shapes whether it is upsetting. As we have discussed, our findings do not suggest that researchers and designers should try to avoid creating applications that produce encounters with bittersweet content. Participants said they *did* want to see bittersweet content, and research has shown that TMRs can help people to re-appraise their pasts in ways that are beneficial to their well-being [33]. Participants generally did not want to be shielded from bittersweet content and said they did not want to delete such content. Therefore, there may be significant negative emotional impacts if these features or content are removed.

We found that the context in which participants viewed bittersweet content could contribute to the pain people experienced in the encounter. Specifically, people found it to be a particularly difficult experience when they viewed Facebook in contexts that they felt were inappropriate for viewing emotionally intense content (e.g., while in bed in the morning, during a lull at work, or while parked in the car when waiting to pick someone up). While we acknowledge detecting context raises a number of technical and social issues, context was one of the most clear factors that impacted sensitive encounters.

7.2.2 Design practice: Moving from simple signals to holistic and contextual feature sets. We do not feel it is possible to build recommender systems that do not produce encounters with bittersweet content that are more negative because it is difficult to predict what a user will find bittersweet (as we discuss in Section 6.1). Therefore, we are left with some formidable design challenges. But we believe understanding the relationship between context and sensitive content will begin to address these challenges.

Empirical research on user feedback and interactions can help identify which contextual features are important to consider. We were able to identify that some contextual factors make encounters with bittersweet content negative. We urge researchers and designers to work closely and sensitively with users to identify the contexts in which they have harmful experiences with content. There is a wealth of prior work in HCI that has produced methodologies that can help researchers and designers to interrogate their own assumptions about what constitutes a harm (e.g., reflexive design [26, 27]). Extending this work to consider algorithmic systems is a promising direction of future research.

7.2.3 Design implications: Detecting contextual factors. We recommend ways of mitigating painful encounters by using a user’s context and implicit user feedback to make decisions about whether to display content. These implications are largely based off of participants’ suggestions and their ways of understanding encounters with bittersweet content.

(1) *Make it clear to users that there are options to customize recommender systems to exclude certain content.* In the case of Facebook

Memories, all participants were unaware that they had the option to exclude date ranges and users from their Memories. Like the “Take a Break” option on Facebook,⁶ which is suggested to users when they change their relationship status to indicate they are no longer in a relationship, recommender systems could inform users of such features if they can detect that a user likely did not want to see the content (e.g., they delete the content after it shows up in their Memories).

(2) *Change the time of day or location that sensitive content is shown.* Participants said that they did not want to see sensitive content first thing in the morning—they were often still half awake and it sometimes set the tone for their day. Some participants did not like seeing sensitive content at nighttime. Location may also play into user experiences—for example, some users may not want to see sensitive content while they are at work. As we discuss in Section 8, this design implication may not be appropriate for all contexts—it may be the case that avoiding negative encounters in the morning might also prevent users from having positive experiences with more sweet content in the morning.

7.2.4 *Provocation: How can we ameliorate negative aspects of encountering bittersweet content when bittersweet content cannot always be readily identified or negative outcomes predicted?* Designers may want to normalize the expectation that content may not all be sweet and that users may find experiences with bittersweet content to be positive in their own ways—not ways that a recommender system might “expect”, but in very human, complicated ways. Instead of presenting users with content in a cheerful light, which *can be harmful* to users when there is emotional dissonance, there may be ways in interfaces to acknowledge the breadth and complexity of experiences with curated content—without presuming to be able to know a user’s feelings.

7.3 Designing for how humans understand problems: affective sense-making and computational tractability

We examine how users try to make sense of the intent and affect of the curated content that they see, and how to design to support user sense-making. Understanding users’ behaviors around sense-making will also provide recommender systems with information about how humans understand problems.

7.3.1 *Finding: people engage in affective sense-making when they encounter curated bittersweet content.* We saw that even with a relatively (seemingly) straightforward recommender system, such as On This Day, participants still engaged in complex sense-making and they attributed intentions and affect to the recommender system, as well as reasoning about its limitations. Participants perceived affective dimensions of a recommender system’s decision-making process (i.e., what we refer to as “affective sense-making”). They felt that recommender systems were unintelligent and that recommender systems often “thought” that content was happy when it was in fact more bittersweet. Some participants found those incongruities upsetting, but they felt there was nothing they could do to stop that content from being shown to them because they did not have a way to provide feedback on the selection criteria used

to curate content. We see a relationship between affective sense-making, agency, and expectedness. If users are able to accurately understand the recommender systems of the platforms that they use, they may not be caught off guard when the system behaves in a way that is incongruous with their feelings about the content.

7.3.2 *Design practice: using studies on affective sense-making to inform metrics.* Understanding affective sense-making will require researchers to engage in qualitative research on users’ affective sense-making in order to create human-centered metrics that can be used to evaluate the benefits and harms of such systems—instead of relying on convenient metrics that are computationally tractable. To do so, we will need to broaden what we consider “tractable” in order to include human understandings of problems. Likewise, we will need to carefully consider what is a “problem.” Baumer and Silberman [7] argue that computational systems are often designed with the assumption that there *is a problem* that needs to be fixed. In our current context, the term “problem” highlights the complexity of bittersweet content and how it can result in users feeling difficult (although not necessarily negative) emotions at inopportune times or in inappropriate contexts. As researchers, we in turn are asked to consider what the “problem” is. Even in our own research, we initially approached sensitive content as a category of content that we needed to better understand such that we could make recommendations for how to prevent it from being curated. However, as our analysis clearly demonstrates, our initial objectives were not in line with our participants’ preferences.

Our (the authors’) initial instinct to detect and remove sensitive content is a good example of the tractability problem. We use “computationally tractable” to refer to taking a problem and simplifying its scope and complexity (in terms of societal and personal impacts) such that feasible solutions to the problem will have quantifiable metrics for success. In recommender systems, computational tractability often means framing problems so that their solutions can be evaluated by metrics related to user behavior (e.g., whether a change to a website increases traffic, whether people are purchasing more) in order to determine whether a technology achieves the utility it was created for.

Bittersweet content, however, requires we take a more nuanced approach to metrics. If nothing else, our analysis demonstrates that bittersweet content is not content that sits at the mid-point *between bitter and sweet*—in other words, bitter and sweet are not a dichotomy, and content can be both very bitter and very sweet at once. As researchers, we need to focus on developing metrics that capture the complexity experienced by users. Doing so will require new ways of soliciting feedback from users and comparing it with inferred signals based on deeper qualitative research.

7.3.3 *Design implications: enabling affective sense-making.* We developed design implications based on the affective sense-making that our participants described to us. Their affective sense-making was not necessarily grounded in empirical data about how algorithms work, but about their own perceptions. We think that providing users with more data on why they see content would improve their experiences with sensitive content.

(1) *Users should see information about why bittersweet content was curated.* Participants sometimes assumed that the recommender system only selects posts that it “thinks” are “happy”, and sometimes

⁶<https://www.facebook.com/help/1638212473101795>

participants were distressed when the bittersweet content was presented in a neutral or even cheerful way. In order to give users more clarity about why content was curated, we could imagine features similar to Facebook’s “Why am I seeing this ad?” (which gives users information about why a user was selected for targeted advertising), as well as features that would allow users to say that they want to see less content that is similar to what they were shown.

(2) *Ask users for feedback about whether those selection criteria were accurate.* When people are shown information about why content was selected, a platform could also ask whether the criteria used to select the content were relevant to the user. This method could help to determine which criteria are most relevant to each user (e.g., popular content, content posted on an important date, positive content). In Section 6.2 we acknowledged the difficulty of soliciting feedback after an upsetting encounter, but we believe that gaining feedback on the criteria rather than on the content itself could be more beneficial for users.

7.3.4 Provocation: How can a technology be designed to match how humans understand a problem rather than a computationally tractable version of the problem? Designing for supporting affective sense-making requires us to also aid users in understanding a system and provide users with more opportunities to give feedback. More user feedback could help recommender systems learn to be more responsive to user needs—in other words, user feedback could expand computational tractability. Thus, we see a feedback loop between enabling affective sense-making and expanding computational tractability. However, researchers and designers will have to consider how to encourage users to provide information about their affective sense-making process (and when/if it is appropriate to do so if a user may be distressed), and they will have to consider how to best represent algorithmic intentionality when, as we have stated earlier, recommender systems are influenced by a variety of factors (e.g., training data, third party interests, designers’ intents).

Lastly, although we framed this section in terms of affect and intent, these concepts come from how our participants talked about recommender systems. Researchers and designers will also have to consider their own stances on *whether recommender systems have intent and/or affect* and how their stance relates to their users’ understandings, and they must also reflect on how these stances impact their designs.

8 LIMITATIONS

Because of the content of our interviews, this paper was primarily focused on our participants’ more negative encounters with bittersweet content, especially bittersweet content that was particularly bitter—and our design recommendations erred towards suggesting ways to avoid these kinds of encounters. However, recommender systems that curate personal social content need to be designed to support a wide range of user needs—including the need to have positive experiences with bittersweet content. Furthermore, designing to avoid unwanted encounters with particularly bitter content may foreclose opportunities for wanted encounters with content that is particularly sweet. For instance, we recommended avoiding showing users sensitive content in the morning; however, some encounters with sensitive content can be positive and inspire the

start to a good day. Given that these trade offs will differ in each research and design context, we recognize that our findings and design recommendations may not generalize to all applications, and we encourage researchers and designers to use the practice recommendations and provocations as a starting point for reflection and analysis.

9 CONCLUSION

Ultimately, this paper is a story about how people find bittersweet content meaningful. People want to engage with things that they find to be bittersweet, and this paper points at ways to prevent bittersweet from becoming just bitter. We argued that recommender systems often use binary mechanisms (e.g., “thumbs up, thumbs down”) to determine what a user wants to see more or less of, but these mechanisms are inadequate when there are complex emotions at stake. They cannot necessarily take into account aspects of a user’s experience that take place off of the platform, such as whether a user expected to see sensitive content, the context in which a user views sensitive content, and how a user interprets a recommender system’s affect and intent.

We provided some design methods/practices, implications for design, and provocations to guide reflection on researchers and designers’ responsibilities to users with regard to bittersweet content:

- (1) how to enable expectedness by using non-digital artifacts as inspiration,
- (2) how to use a user’s context to understand which encounters are upsetting, and
- (3) how to enable affective sense-making to expand what is computationally tractable.

Ultimately, we believe that researchers need to conduct studies on these off-platform aspects of user experience. We also argue that researchers and designers need to be reflexive when it comes to examining their own beliefs about experiences with bittersweet content.

These suggestions and our solutions partially address issues around upsetting encounters, but also there is still more to reflect upon and learn. We want to re-emphasize that most participants wanted to engage with bittersweet content, so we are not advocating for never developing these systems or that they must be perfect—but that their design must be tempered with considerations about user expectations, about the context in which a user views content, and about ways to support sense-making and solicit feedback.

Our hope is that further research can help us understand the contextual factors that influence how people experience sensitive content. We want researchers and designers to create systems that do not make sensitive content any more bitter than it needs to be, with the hope that they can represent users’ desires for both the bitter and the sweet.

ACKNOWLEDGMENTS

We wish to thank Meta for helping us recruit the participants who were interviewed in this study (which was separately run by CU Boulder). Meta was not involved in the data collection or analysis for this study and no data was shared with Meta. CU Boulder’s research was supported by NSF #1756028, and approved by CU Boulder’s IRB #17-0727. We thank the TAT Lab at the University

of Washington and the Identity Lab at CU Boulder for their invaluable feedback, as well as Daniela Rosner, Jennifer Turns, Mark Zachry, Sean Munson, and Mar Drouhard. We thank the reviewers for thoughtful suggestions that improved this paper. We also have the deepest appreciation for our participants, who trusted us to represent their stories. And lastly, we are indebted to all of the recommender systems that have, over the years, given us an appreciation for the joys and heartbreaks of encountering bittersweet content.

REFERENCES

- [1] Xavier Amatriain, Josep M. Pujol, and Nuria Oliver. 2009. I Like It... I Like It Not: Evaluating User Ratings Noise in Recommender Systems. In *User Modeling, Adaptation, and Personalization (Lecture Notes in Computer Science)*, Geert-Jan Houben, Gord McCalla, Fabio Pianesi, and Massimo Zancanaro (Eds.). Springer, Berlin, Heidelberg, 247–258. https://doi.org/10.1007/978-3-642-02247-0_24
- [2] Xavier Amatriain, Josep M. Pujol, Nava Tintarev, and Nuria Oliver. 2009. Rate it again: increasing recommendation accuracy by user re-rating. In *Proceedings of the third ACM conference on Recommender systems (RecSys '09)*. Association for Computing Machinery, New York, NY, USA, 173–180. <https://doi.org/10.1145/1639714.1639744>
- [3] Nazanin Andalibi, Pinar Ozturk, and Andrea Forte. 2017. Sensitive Self-disclosures, Responses, and Social Support on Instagram: The Case of #Depression. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. Association for Computing Machinery, New York, NY, USA, 1485–1500. <https://doi.org/10.1145/2998181.2998243>
- [4] E. Ashley-Dejo, S. Ngwira, and T. Zuva. 2015. A survey of Context-aware Recommender System and services. In *2015 International Conference on Computing, Communication and Security (ICCCS)*, 1–6. <https://doi.org/10.1109/ICCCS.2015.7374144>
- [5] Matthew Baldwin, Monica Biernat, and Mark J. Landau. 2015. Remembering the real me: Nostalgia offers a window to the intrinsic self. *Journal of Personality and Social Psychology* 108, 1 (2015), 128–147. <https://doi.org/10.1037/a0038033>
- [6] Krystine Irene Batcho. 2013. Nostalgia: The bittersweet history of a psychological concept. *History of Psychology* 16, 3 (2013), 165–176. <https://doi.org/10.1037/a0032427>
- [7] Eric PS Baumer and M. Six Silberman. 2011. When the implication is not to design (technology). In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2271–2274.
- [8] Bharat Bhole and Brid Hanna. 2017. The effectiveness of online reviews in the presence of self-selection bias. *Simulation Modelling Practice and Theory* 77 (Sept. 2017), 108–123. <https://doi.org/10.1016/j.simpat.2017.05.005>
- [9] Kyoungsoo Bok, Jongtae Lim, Heetae Yang, and Jaesoo Yoo. 2016. Social group recommendation based on dynamic profiles and collaborative filtering. *Neuro-computing* 209 (Oct. 2016), 3–13. <https://doi.org/10.1016/j.neucom.2015.11.120>
- [10] Virginia Braun and Victoria Clarke. 2012. Thematic Analysis. In *APA handbook of research methods in psychology, Vol 2: Research designs: Quantitative, qualitative, neuropsychological, and biological*, Harris Ed Cooper, Paul M. Camic, Debra L. Long, A. T. Panter, David Ed Rindskopf, and Kenneth J. Sher (Eds.).
- [11] 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 (May 2013), 152–163. <https://doi.org/10.1080/01972243.2013.777300>
- [12] Eleanor R. Burgess, Kathryn E. Ringland, Jennifer Nicholas, Ashley A. Knapp, Jordan Eschler, David C. Mohr, and Madhu C. Reddy. 2019. "I think people are powerful": The Sociality of Individuals Managing Depression. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 41:1–41:29. <https://doi.org/10.1145/3359143>
- [13] Moira Burke and Mike Develin. 2016. Once More with Feeling: Supportive Responses to Social Sharing on Facebook. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. Association for Computing Machinery, New York, NY, USA, 1462–1474. <https://doi.org/10.1145/2818048.2835199>
- [14] Stevie Chancellor, Eric P. S. Baumer, and Munmun De Choudhury. 2019. Who is the "Human" in Human-Centered Machine Learning: The Case of Predicting Mental Health from Social Media. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 147:1–147:32. <https://doi.org/10.1145/3359249>
- [15] Stevie Chancellor, Michael L. Birnbaum, Eric D. Caine, Vincent M. B. Silenzio, and Munmun De Choudhury. 2019. A Taxonomy of Ethical Tensions in Inferring Mental Health States from Social Media. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*. Association for Computing Machinery, New York, NY, USA, 79–88. <https://doi.org/10.1145/3287560.3287587>
- [16] Kathy Charmaz. 2014. *Constructing Grounded Theory*. SAGE.
- [17] Herbert H. Clark and Susan E. Brennan. 2004. Grounding in communication. In *Perspectives on socially shared cognition*. American Psychological Association, 127–149. <https://doi.org/10.1037/10096-006>
- [18] Munmun De Choudhury, Scott Counts, Eric J. Horvitz, and Aaron Hoff. 2014. Characterizing and predicting postpartum depression from shared facebook data. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing (CSCW '14)*. Association for Computing Machinery, New York, NY, USA, 626–638. <https://doi.org/10.1145/2531602.2531675>
- [19] Munmun De Choudhury, Sanket S. Sharma, Tomaz Logar, Wouter Eekhout, and René Clausen Nielsen. 2017. Gender and Cross-Cultural Differences in Social Media Disclosures of Mental Illness. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. Association for Computing Machinery, New York, NY, USA, 353–369. <https://doi.org/10.1145/2998181.2998220>
- [20] Michael Ann DeVito, Darren Gergle, and Jeremy Birnholtz. 2017. "Algorithms ruin everything": #RIPTwitter, Folk Theories, and Resistance to Algorithmic Change in Social Media. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. Association for Computing Machinery, New York, NY, USA, 3163–3174. <https://doi.org/10.1145/3025453.3025659>
- [21] Claire Burke Draucker, Donna S. Martsoff, and Candice Poole. 2009. Developing Distress Protocols for Research on Sensitive Topics. *Archives of Psychiatric Nursing* 23, 5 (Oct. 2009), 343–350. <https://doi.org/10.1016/j.apnu.2008.10.008>
- [22] Motahhare Eslami, Karrie Karahalios, Christian Sandvig, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. 2016. First I "like" it, then I hide it: Folk Theories of Social Feeds. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. Association for Computing Machinery, New York, NY, USA, 2371–2382. <https://doi.org/10.1145/2858036.2858494>
- [23] Motahhare Eslami, Sneha R. Krishna Kumaran, Christian Sandvig, and Karrie Karahalios. 2018. Communicating Algorithmic Process in Online Behavioral Advertising. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3174006>
- [24] Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. "I Always Assumed That I Wasn't Really That Close to [Her]": Reasoning About Invisible Algorithms in News Feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 153–162.
- [25] Jessica L. Feuston and Anne Marie Piper. 2018. Beyond the Coded Gaze: Analyzing Expression of Mental Health and Illness on Instagram. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (Nov. 2018), 51:1–51:21. <https://doi.org/10.1145/3274320>
- [26] Benjamin Fish and Luke Stark. 2020. Reflexive Design for Fairness and Other Human Values in Formal Models. *arXiv:2010.05084 [cs]* (Oct. 2020). <http://arxiv.org/abs/2010.05084>
- [27] Barton Friedland and Yutaka Yamauchi. 2011. Reflexive design thinking: putting more human in human-centered practices. *Interactions* 18, 2 (March 2011), 66–71. <https://doi.org/10.1145/1925820.1925835>
- [28] Arik Friedman, Bart P. Knijnenburg, Kris Vanhecke, Luc Martens, and Shlomo Berkovsky. 2015. Privacy Aspects of Recommender Systems. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer US, Boston, MA, 649–688. https://doi.org/10.1007/978-1-4899-7637-6_19
- [29] Oliver L. Haimson, Jed R. Brubaker, Lynn Dombrowski, and Gillian R. Hayes. 2015. Disclosure, Stress, and Support During Gender Transition on Facebook. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. Association for Computing Machinery, New York, NY, USA, 1176–1190. <https://doi.org/10.1145/2675133.2675152>
- [30] Nan Hu, Paul A. Pavlou, and Jennifer Zhang. 2006. Can online reviews reveal a product's true quality? Empirical findings and analytical modeling of online word-of-mouth communication. In *Proceedings of the 7th ACM conference on Electronic commerce (EC '06)*. Association for Computing Machinery, New York, NY, USA, 324–330. <https://doi.org/10.1145/1134707.1134743>
- [31] Y. Hu, Y. Koren, and C. Volinsky. 2008. Collaborative Filtering for Implicit Feedback Datasets. In *2008 Eighth IEEE International Conference on Data Mining*. 263–272. <https://doi.org/10.1109/ICDM.2008.22>
- [32] Jevan A. Hutson, Jessie G. Taft, Solon Barocas, and Karen Levy. 2018. Debiasing Desire: Addressing Bias & Discrimination on Intimate Platforms. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (Nov. 2018), 73:1–73:18. <https://doi.org/10.1145/3274342>
- [33] Ellen Isaacs, Artie Konrad, Alan Walendowski, Thomas Lennig, Victoria Hollis, and Steve Whittaker. 2013. Echoes from the past: how technology mediated reflection improves well-being. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1071–1080.
- [34] Benjamin N. Jacobsen. 2020. Sculpting digital voids: The politics of forgetting on Facebook. *Convergence* (Feb. 2020), 1354856520907390. <https://doi.org/10.1177/1354856520907390>
- [35] Gawesh Jawaheer, Martin Szomszor, and Patty Kostkova. 2010. Comparison of implicit and explicit feedback from an online music recommendation service. In *Proceedings of the 1st International Workshop on Information Heterogeneity and Fusion in Recommender Systems (HetRec '10)*. Association for Computing

- Machinery, New York, NY, USA, 47–51. <https://doi.org/10.1145/1869446.1869453>
- [36] Gawesh Jawaheer, Peter Weller, and Patty Kostkova. 2014. Modeling User Preferences in Recommender Systems: A Classification Framework for Explicit and Implicit User Feedback. *ACM Transactions on Interactive Intelligent Systems* 4, 2 (June 2014), 8:1–8:26. <https://doi.org/10.1145/2512208>
- [37] Arjan J. P. Jeckmans, Michael Beye, Zekeriya Erkin, Pieter Hartel, Reginald L. Lagendijk, and Qiang Tang. 2013. Privacy in Recommender Systems. In *Social Media Retrieval*, Naeem Ramzan, Roelof van Zwol, Jong-Seok Lee, Kai Clüver, and Xian-Sheng Hua (Eds.). Springer, London, 263–281. https://doi.org/10.1007/978-1-4471-4555-4_12
- [38] Jialun “Aaron” Jiang and Jed R. Brubaker. 2018. Tending Unmarked Graves: Classification of Post-mortem Content on Social Media. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (Nov. 2018), 81:1–81:19. <https://doi.org/10.1145/3274350>
- [39] Michael Jugovac and Dietmar Jannach. 2017. Interacting with Recommenders—Overview and Research Directions. *ACM Transactions on Interactive Intelligent Systems* 7, 3 (Sept. 2017), 10:1–10:46. <https://doi.org/10.1145/3001837>
- [40] Radu Jurca, Florent Garcin, Arjun Talwar, and Boi Faltings. 2010. Reporting incentives and biases in online review forums. *ACM Transactions on the Web* 4, 2 (April 2010), 5:1–5:27. <https://doi.org/10.1145/1734200.1734202>
- [41] Elaine Kasket. 2009. Protocol for responding to participant distress, adapted version for telephone/skype. *London Metropolitan University* (2009).
- [42] Artie Konrad, Ellen Isaacs, and Steve Whittaker. 2016. Technology-mediated memory: is technology altering our memories and interfering with well-being? *ACM Transactions on Computer-Human Interaction (TOCHI)* 23, 4 (2016), 1–29.
- [43] Artie Konrad, Simon Tucker, John Crane, and Steve Whittaker. 2016. Technology and Reflection: Mood and Memory Mechanisms for Well-Being. *Psychology of Well-Being* 6, 1 (June 2016), 5. <https://doi.org/10.1186/s13612-016-0045-3>
- [44] Denis Kotkov, Shuaiqiang Wang, and Jari Veijalainen. 2016. A survey of serendipity in recommender systems. *Knowledge-Based Systems* 111 (Nov. 2016), 180–192. <https://doi.org/10.1016/j.knsys.2016.08.014>
- [45] Jeff T. Larsen and A. Peter McGraw. 2011. Further evidence for mixed emotions. *Journal of Personality and Social Psychology* 100, 6 (2011), 1095–1110. <https://doi.org/10.1037/a0021846>
- [46] JeongHyun Lee. 2020. Algorithmic Uses of Cybernetic Memory: Google Photos and a Genealogy of Algorithmically Generated “Memory”. *Social Media + Society* 6, 4 (Oct. 2020), 2056305120978968. <https://doi.org/10.1177/2056305120978968>
- [47] Katie A. McLaughlin and Susan Nolen-Hoeksema. 2011. Rumination as a transdiagnostic factor in depression and anxiety. *Behaviour Research and Therapy* 49, 3 (March 2011), 186–193. <https://doi.org/10.1016/j.brat.2010.12.006>
- [48] Ana Lúcia Migowski and Willian Fernandes Araújo. 2019. “Looking back” at personal memories on Facebook: Co-constitutive agencies in contemporary remembrance practices. *Journal of Aesthetics & Culture* 11, sup1 (Nov. 2019), 1644130. <https://doi.org/10.1080/20004214.2019.1644130>
- [49] Itishree Mohallick and Özlem Özgöbek. 2017. Exploring privacy concerns in news recommender systems. In *Proceedings of the International Conference on Web Intelligence (WI '17)*. Association for Computing Machinery, New York, NY, USA, 1054–1061. <https://doi.org/10.1145/3106426.3109435>
- [50] Susan Nolen-Hoeksema. 1991. Responses to depression and their effects on the duration of depressive episodes. *Journal of Abnormal Psychology* 100, 4 (1991), 569–582. <https://doi.org/10.1037/0021-843X.100.4.569>
- [51] Susan Nolen-Hoeksema, Blair E. Wisco, and Sonja Lyubomirsky. 2008. Rethinking Rumination. *Perspectives on Psychological Science* 3, 5 (Sept. 2008), 400–424. <https://doi.org/10.1111/j.1745-6924.2008.00088.x>
- [52] James W. Pennebaker and Sandra K. Beall. 1986. Confronting a traumatic event: Toward an understanding of inhibition and disease. *Journal of Abnormal Psychology* 95, 3 (1986), 274–281. <https://doi.org/10.1037/0021-843X.95.3.274>
- [53] Anthony T. Pinter, Jialun Aaron Jiang, Katie Z. Gach, Melanie M. Sidwell, James E. Dykes, and Jed R. Brubaker. 2019. “Am I Never Going to Be Free of All This Crap?” Upsetting Encounters with Algorithmically Curated Content About Ex-Partners. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–23.
- [54] Brady Robards, Sian Lincoln, Benjamin C. Pinkard, and Jane Harris. 2018. Remembering Through Facebook: Mediated Memory and Intimate Digital Traces. In *Digital Intimate Publics and Social Media*, Amy Shields Dobson, Brady Robards, and Nicholas Carah (Eds.). Springer International Publishing, Cham, 75–91. https://doi.org/10.1007/978-3-319-97607-5_5
- [55] Johnny Saldana. 2012. *The Coding Manual for Qualitative Researchers*. SAGE.
- [56] Pedro Sanches, Axel Janson, Pavel Karpashevich, Camille Nadal, Chengcheng Qu, Claudia Daudén Roquet, Muhammad Umair, Charles Windlin, Gavin Doherty, Kristina Höök, and Corina Sas. 2019. HCI and Affective Health: Taking stock of a decade of studies and charting future research directions. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–17. <https://doi.org/10.1145/3290605.3300475>
- [57] Corina Sas and Steve Whittaker. 2013. Design for forgetting: disposing of digital possessions after a breakup. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1823–1832.
- [58] Constantine Sedikides and Tim Wildschut. 2016. A bittersweet emotion that confers psychological health benefits. In *The Wiley Handbook of Positive Clinical Psychology*, Alex M. Wood and Judith Johnson (Eds.). 125.
- [59] Sho Tsugawa, Yusuke Kikuchi, Fumio Kishino, Kosuke Nakajima, Yuichi Itoh, and Hiroyuki Ohsaki. 2015. Recognizing Depression from Twitter Activity. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 3187–3196. <https://doi.org/10.1145/2702123.2702280>
- [60] Bas Verplanken. 2012. When bittersweet turns sour: Adverse effects of nostalgia on habitual worriers. *European Journal of Social Psychology* 42, 3 (2012), 285–289. <https://doi.org/10.1002/ejsp.1852>
- [61] David Watson and Kasey Stanton. 2017. Emotion Blends and Mixed Emotions in the Hierarchical Structure of Affect. *Emotion Review* 9, 2 (April 2017), 99–104. <https://doi.org/10.1177/1754073916639659>
- [62] Maranke Wieringa. 2020. What to account for when accounting for algorithms: a systematic literature review on algorithmic accountability. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT* '20)*. Association for Computing Machinery, New York, NY, USA, 1–18. <https://doi.org/10.1145/3351095.3372833>
- [63] Tim Wildschut, Constantine Sedikides, Jamie Arndt, and Clay Routledge. 2006. Nostalgia: content, triggers, functions. *Journal of Personality and Social Psychology* 91, 5 (Nov. 2006), 975–993. <https://eprints.soton.ac.uk/40445/>
- [64] Hua Zhou and Kenneth Lange. 2009. Rating Movies and Rating the Raters Who Rate Them. *The American Statistician* 63, 4 (Nov. 2009), 297–307. <https://doi.org/10.1198/tast.2009.08278>
- [65] Xinyue Zhou, Constantine Sedikides, Tim Wildschut, and Ding-Guo Gao. 2008. Counteracting Loneliness: On the Restorative Function of Nostalgia. *Psychological Science* 19, 10 (Oct. 2008), 1023–1029. <https://doi.org/10.1111/j.1467-9280.2008.02194.x>