"I Wanted Them to Think That I Wrote That": AI-Generated Self-Presentation on Dating Apps and Implications of Non-Disclosure on Informed Consent

Meryem Barkallah
University of Michigan-Flint
Flint, Michigan, USA
meryemba@umich.edu

Douglas Zytko
University of Michigan-Flint
Flint, Michigan, USA
dzytko@umich.edu

Abstract

Generative artificial intelligence (AI) adds unprecedented scale to capabilities for self-presentation online that may diverge from one's physical-world identity, thus potentially misinforming consent to intimate interactions, such as in online dating. Yet there is little empirical understanding of AI-generated self-presentation and (non-)disclosure to interaction partners. We present a qualitative survey of 113 online daters who used AI-generated content in their profiles or messages seen by in-person meeting partners. Findings show that generative AI is often used to fabricate attractive dating personalities through profile text and bios, with no relevance to one's actual identity, and is seldom disclosed to meeting partners to avoid romantic rejection. Because sexual assault is defined by mis- or under-informed consent, the study positions generative AI as a potentially significant sexual assault risk factor through its use for presentation of non-physical traits that are influential to dating outcomes yet not readily identified as AI-generated upon meeting face-to-face.**Content warning: this paper discusses forms of sexual violence including rape by deception.**

CCS Concepts

• **Human-centered computing** → **Human computer interaction (HCI)**; *Collaborative and social computing*; *Empirical studies in HCI*; • **Computing methodologies** → Artificial intelligence.

Keywords

online dating, generative AI, disclosure, informed consent, AI-mediated communication, AI-mediated self-presentation, reflexive thematic analysis

ACM Reference Format:
Meryem Barkallah and Douglas Zytko. 2026. "I Wanted Them to Think That I Wrote That": AI-Generated Self-Presentation on Dating Apps and Implications of Non-Disclosure on Informed Consent. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3772318.3791593

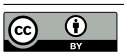

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3791593

1 Introduction

Chris is excited for their first date with Sam, whom they discovered on a dating app. Over days of messaging on the dating app, Chris becomes drawn to their personality, the way they make them laugh, and the sense of connection they express to each other. When the in-person date feels different, with awkward silences and jokes that no longer land, Chris rationalizes it away as nerves and anxiety that anyone feels on a date. They have sex. What Chris does not realize is that most of the messages they received on the dating app–the crux of their impression of a loveable personality and desire for meeting in-person–were AI-generated.

Consider that sexual assault is defined by the Rape, Abuse & Incest National Network (RAINN) as sexual contact that occurs "without clear, voluntary, and informed consent" [92]. In the scenario above, was Chris the victim of sexual assault due to misinformed consent from undisclosed use of generative AI? Could their consent to meeting face-to-face, more generally, be considered adequately informed?

Consent has received growing attention in the Human-Computer Interaction (HCI) literature, from personal data collection [65, 67, 81] to interpersonal behavior in Computer-Mediated spaces such as social Virtual Reality (VR) [59, 95, 122, 132], social networking platforms [60, 129], and dating apps [123, 124]. Regarding interpersonal behavior, consent has emerged as a contemporary lens for studying computer-mediated harm and unwanted behavior [95, 124, 132] due in part to the lens's acknowledgment that such harm may occur unintentionally due to misperceptions of agreement or the right to decline the respective behavior [9, 82]. HCI research has explored how social technologies can and do shape the ways that people give and (perceive to) receive consent to interpersonal behavior [29, 95, 124], and how technology may predispose individuals to commit unwanted acts due to misunderstandings of consent. To the latter effect, research has found that platforms like dating apps can misinform understanding of consent [124] through seemingly innocuous signals received online (e.g., interpreting a bikini picture on a dating app profile as a definitive indication of interest in sex).

In this paper, we focus not on technology's shaping of how consent itself is communicated, but on how consent decisions to interpersonal behavior are (mis)informed by emerging technology. In this sense, we connect the relatively newer HCI lens of consent with the more established frame of impression management: the use of technology to mediate one's self-presentation [12, 46, 48, 55, 120] and strategic self-disclosure online [36, 89].

Deceptive self-presentation has pervaded social platforms for decades [34, 55, 105]. Yet generative AI is an emerging technology that adds unprecedented scale to capabilities for realistic self-presentation that may diverge from one's physical-world self, thus potentially misinforming consent to incredibly intimate interactions as the online dating context exemplifies. The implications of generative AI on consent are only starting to be identified and explored, such as the nonconsensual creation and distribution of AI-generated pornography [106]. Beyond speculation [6, 54], there remains little empirical understanding of how generative AI is used in self-presentation on social platforms, as well as if and how use of generative AI is voluntarily disclosed to individuals whose consent to subsequent interaction may be predicated on AI-generated content.

We explore this gap in knowledge with a qualitative, open-ended survey of 113 online daters from 15 countries who used AI-generated content in their self-presentation on dating apps, most of whom met at least one user face-to-face who had seen or received their AI-generated content. The study sought to answer:

- RQ1. How is AI-generated content incorporated into self-presentation on dating apps?
- RQ2. How do online daters make decisions to disclose, or not disclose, use of generative AI to dating partners?
- RQ3. What are the implications of generative AI, and associated (non-)disclosure decisions, on informed consent to interaction between online daters?

Findings show that generative AI was most often used to fabricate an attractive online dating personality regardless of its relevance to one's own identity (**RQ1**). We call this AI-as-self because dating partners were effectively building a connection with a generative AI tool's idea of an attractive dater rather than the person using the generative AI tool. In some cases, these AI-generated personalities were tailored to specific dating partners by copy-and-pasting content from said partners' profiles and messages into a generative AI tool, but without permission or awareness of those users. Generative AI was predominantly used to create text-based content for profile bios or messages, and not pictures of one's physical appearance; thus the use of AI-generated content was not inevitably discovered by meeting other daters face-to-face.

Participants typically did not disclose their use of generative AI to face-to-face meeting partners (**RQ2**). Rationale for non-disclosure was primarily to limit risk to one's own dating success, with some acknowledging their AI-generated content as intentionally deceptive. Others justified non-disclosure by downplaying the impact of generative AI on the decision-making of their dating partners, with some indicating that generative AI needs to be disclosed only when used to modify physical appearance.

The study ultimately shows that generative AI poses considerable risk of misinforming consent to interpersonal behavior between online daters because of its use for presentation of non-physical traits that are influential to dating outcomes, yet not readily identified as AI-generated upon meeting face-to-face (**RQ3**). We conclude the paper by discussing how AI-generated self-presentation could shape future HCI research of computer-mediated harm, as well as design decisions around automated disclosure of generative AI.

2 Related Work

2.1 Computer-Mediated Consent to Interpersonal Behavior

Consent has steadily grown as a research focus in HCI, particularly consent to data collection [94, 109] given legal regulations around data [77, 107], and newfound uses for AI model training [66, 121]. Informed consent to data collection (i.e., providing people with information necessary to make an educated decision about their data) has been a special emphasis in light of challenges that ubiquitous computing poses when data can potentially be collected at all times through any number of hidden devices [49, 74].

Consent to interpersonal behavior, such as through social platforms, has also experienced mounting interest in HCI [129], in part because consent (or the lack thereof) is the defining quality of many forms of interpersonal harm, including sexual violence, intimate partner violence, and child abuse [7, 16, 19]. A benefit of a consent lens is that it sidesteps historical challenges to consistently defining harm, which has been especially obstructive to AI- and content moderation-based solutions when moderators/data annotators disagree over what is considered harm [64, 93]. Per Zytko and Chan, under a consent lens *"the notion of objectively harmful behavior is rejected, with harm instead being qualified on a user-by-user basis as unwanted interaction"* [122]. We similarly use a consent lens to qualify a behavior as unwanted according to the person experiencing the behavior, rather than by the judgment of an external entity.

Prior HCI research has used consent as a lens to study unwanted behavior on social media platforms [60], social VR environments [95], and dating apps [29, 124]. To the latter, online dating research of consent has found that unwanted behavior may occur unintentionally through misinterpretation of consent (agreement) to sex through signals received in dating app interfaces, such as physically revealed profile pictures and emojis in messages [124]. This speaks to a broader pattern in the literature of social platform design inadvertently facilitating unwanted acts through fueling misinterpretations of consent or failing to provide explicit mechanisms for consent exchange [60, 124, 130].

In considering how social technologies could be deliberately designed to scaffold consent exchange [84], most HCI research has advocated for affirmative consent as the basis for how consent to behavior should be exchanged [60, 68, 84, 100, 122, 123]. Distilled in the slogan "yes means yes," affirmative consent refocuses responsibility onto the initiator of a sexual or other social act to ask for and receive explicit permission [39]. Hallmarks of affirmative consent, as captured in the FRIES acronym [88], are freely given, reversible, **informed**, enthusiastic, and specific consent.

Informed consent speaks to an individual basing their consent on accurate and complete information. Even when sex is initially voluntary, the experience may later be realized as rape if it is learned that the consent decision relied on false or misleading information. In prior scholarship, this is often discussed as *rape by deception*. For example, an individual may consider their consent to have been misinformed if their partner secretly removed a condom during thought-to-be protected sex [28], or if their partner intentionally withheld or lied about their human immunodeficiency virus (HIV) status [10].

Informed consent has been a complicated construct even before computer-mediated communication and AI. Legal scholarship and court rulings have debated for some time the extent to which consent can be vitiated or invalidated on the basis of being misinformed, particularly on the basis of condom-use, as well as the standard for proving misinformed consent [3, 50, 51, 117]. In computer-mediated communication, informed consent is potentially "not, in most cases, even a realistic option" [100], given the freedom of self-presentation online and limited awareness of platform-specific dynamics. For example, Im and colleagues point to challenges with informing users that their sexual content is being shared online (revenge porn) [60], and Zytko and colleagues speculate on inconsistent sharing of sexual health information on dating apps like sexually transmitted infection (STI) status that are necessary for informed consent to "safe sex" [132]. Social VR environments add further complexities to informed consent through endless options for avatar design that may distort exactly "who" one is giving consent to, and to what particular acts [17, 59]. Challenges to informed consent have motivated the ideation of new consent models and practices for computer-mediated communication [100, 122, 123] . These include dedicated discussion of sexual expectations on dating apps before meeting a sexual partner face-to-face [123], deliberately designing social VR environments to inform users of normative behaviors before entering [132], and taking cues from sex-themed video games about making consent a more elaborate "conversation" [84].

Despite attention to how informed consent is or could be mediated by emerging technologies like social robots [26] and VR [59, 122], generative AI represents a relatively new complicating factor because of its potential to (mis)inform consent at scale.

2.2 The Implications of AI-Mediated Communication on Informed Consent

Generative AI tools, such as ChatGPT, Google Gemini, and Claude, have significantly altered capabilities for self-expression online by adding scale, speed, and accessibility. While the potential for misinformation through AI-generated news is well known [25, 85, 97], AI-generated content also has the potential to misinform consent to interpersonal behavior if it mischaracterizes the person it is intended to represent.

As of this writing, the implications of generative AI on interpersonal consent are largely speculative. Furlo et al. used participatory design to consider how generative AI could scaffold sexual consent exchange in positive ways by steering online communication towards discussion of explicit sexual interests and expectations [42]. Other literature that directly speaks to the intersection of consent and generative AI mostly pertains to creating sexual imagery of other people (deepfake porn) [87, 112], and the broader landscape of nonconsensual image abuse [8, 43, 44, 91]. Rather than studying implications on informed consent to interaction, such work looks at and problematizes the absence of consent to data, such as taking a person's social media images without permission to give to a generative AI tool, and distributing AI-generated images depicting another person without their consent [87].

The implications of AI-generated content about oneself on informed consent have not been directly studied, to our knowledge. However, research on AI-mediated communication provides some indirect indication of consent challenges, although most of this research is based on experiments with controlled conditions. AI-mediated communication refers to "interpersonal communication that is not simply transmitted by technology, but modified, augmented, or even generated by a computational agent to achieve communication goals" [54]. Multiple experiments on AI-mediated communication, including with online dating [61, 72, 118], find reduced trust [61] and increased uncertainty [72] about people perceived to have used AI-generated content, especially when that use is not disclosed [90]. One notable exception is Hohenstein and Jung's experiment [58] of messaging applications showing that AI-generated "smart replies" increase perceived trust by allowing blame for miscommunication or coercion to be placed on AI rather than the human messaging partner. Other experiments have found that receiving AI-generated content from others (without knowing that the content is AI-generated) does influence impressions of the sender [71, 78] and even the behavior of the recipient in turn, such as the content and pace of responses [41, 79].

Disclosing that content is AI-generated could rectify potentially misinformed consent; however, HCI research emphasizes that empirical knowledge on generative AI disclosure is in formative stages [30, 75, 98, 113]. This includes technical approaches to detection of AI-generated content [18, 27, 37, 69], and growing work on AI-use labels and content provenance. For example, 'Content Credentials' based on the Coalition for Content Provenance and Authenticity (C2PA) standard attach tamper-evident metadata so people can see how media was made and edited [21].Platforms like TikTok now label AI content, including items created off-platform, by reading and writing these credentials [102]. However such metadata can be stripped from C2PA manifests [83] and many large platforms also remove embedded metadata on upload, rendering AI labels far from foolproof, especially in dating apps where exchanges are often short-lived text chats. It is thus not surprising that deliberation is ongoing about design decisions or requirements for what information to include in disclosure and how to communicate that information to users consuming one's AI-generated content [33, 57, 63, 98, 116]. There is an early indication that people do not voluntarily disclose use of generative AI, including an empirical study of generative AI-use [5] by students (and their subsequent non-disclosure to teachers) as well as a relatively larger-scale experiment of AI "ghostwriting" [32]. In both studies, the justification for non-disclosure was to avoid negative judgment and reactions.

Ultimately, much of the research that is at least tangentially indicative of the impact of generative AI on informed consent is speculative at this point. More empirical insight is needed into how generative AI is used and disclosed in ecologically valid settings.

2.3 (Mis)Informing Consent in Online Dating through Strategic Self-Presentation and Disclosure

Online dating has been singled out as a context in which generative AI is "likely to impact both how people present themselves online and how they evaluate others" [54]. Strategic self-presentation long pre-dates generative AI, having been a hallmark of online dating research over at least two decades, starting with Ellison's 2006 seminal work [34, 46] and continuing in contemporary CHI

literature [6, 24, 86, 96, 120]. Self-presentation or Impression management, originating from Goffman [48], has been widely adopted as a lens to research how online daters curate the impression that other users form of them [46, 114, 125]. Early research found impression management to largely be motivated by goals to increase attractiveness to the broader userbase, which often involved deception or at least mild exaggeration of one's content on dating apps [38, 46, 53, 55, 56, 103, 104].

More recent HCI research on online dating has put emphasis on marginalized and at-risk user populations, learning how strategic self-presentation is used for personal safety, rather than only boosting attractiveness. Most safety-oriented self-presentation strategies involve strategic self-disclosure of personal details about one's identity that may be subject to harassment or other harms. For instance, users of apps for men-seeking-men sometimes disguise their identity because of safety implications of being openly interested in men [11, 12, 23], as well as their HIV status [110] and use of pre-exposure prophylaxis (PrEP) [111]. Transgender users [36] as well as users with disabilities [89] similarly withhold these characteristics to avoid identity-based harm. Interest in casual sex is also sometimes disguised to avoid stigma associated with open interest in sex [12, 126]

The flipside of self-presentation is impression formation, to, for example, decide whether to consent to further interaction with an online dater and meet them face-to-face. Online daters have historically had difficulty forming impressions of "experiential traits" such as personality [125, 128] that must be signaled indirectly through content online, and frustrations with making adequately informed in-person date decisions have been common [40, 125], leading to design efforts to improve impression formation capabilities [131]. Some impression formation strategies have been linked to safety, such as searching for more information about meeting partners on search engines and avoiding contact with profiles that do not have pictures [12, 47].

The historical dichotomy of self-presentation motivations in online dating would suggest that use and (non)disclosure of generative AI may be used for both deceptive and safety-oriented reasons, both of which may misinform consent decisions of other users who consume said content. Our study, presented in this paper, is the first to provide an empirical account of this role.

3 Method

We conducted an Institutional Review Board (IRB)-approved qualitative, open-ended survey study to investigate how online daters incorporate AI-generated content into self-presentation and how they make disclosure decisions about their generative AI-use to dating partners. We opted for a survey over other qualitative method choices for two reasons. One, in light of growing attention in the CHI and broader HCI literature to online dating in multicultural contexts [11], we opted for a survey to facilitate breadth of representation across diverse geographic regions and cultures, including non-Western and Global South participants. Two, because use of generative AI may carry a negative connotation according to prior literature [5, 30, 32], we considered that participants may be more comfortable with discussing potentially deceptive or public-disapproving uses of generative AI through asynchronous data collection methods in which the participant does not have to see and directly interact with a human researcher.

3.1 Recruitment and Participants

The final sample of 113 participants was recruited via a reputable third-party human subjects research recruitment platform commonly used in HCI studies. Eligibility criteria required participants to: (1) be at least 18 years old, (2) have current or prior experience using dating apps, and (3) have used generative AI tools for creating, modifying, or enhancing content on dating apps. Participants were compensated $10 USD for survey completion. We collected 151 total survey responses, however, 38 were removed for failing to meet the eligibility criteria.

While the survey was administered only in English, it amassed representation from 15 different countries across multiple regions. These included 46 from North America (the United States (40), Canada (4), and Mexico (2)); 41 from Africa (South Africa (32) and Kenya (9)); 22 from Europe (the United Kingdom (14), Poland (3), Portugal (2), Germany (1), Italy (1), and Greece (1)); 3 from Australasia (Indonesia (1), Australia (1), and New Zealand (1)); and 1 from South America (Chile).

Fifty-two participants identified as women and 61 as men, with 3 participants also identifying as transgender. Eighty-seven identified as heterosexual, 4 as homosexual, and 16 as bisexual; the remainder identified as pansexual (3), asexual (1), aromantic (1), or demisexual (1). Regarding ethnicity, 60 identified as Black or African American, 38 as White, 7 as Asian, 5 as Hispanic or Latino, 1 as Middle Eastern or North African, 1 as "other," and 1 as mixed (Asian + White). The mean age was 32.9 years, with a range of 18–83 years old.

Most participants (82) had used dating apps for over a year, and all but one used dating apps for at least a month at the time of the survey. Most participants used dating apps multiple times a week (77), and all but two used dating apps at least a "few times a month." Almost all participants used multiple dating apps, with the most popular being Tinder (102), Facebook Dating (54), Bumble (59), Grindr (18), Hinge (35), Match (21), OkCupid (33), and Plenty of Fish (24).

All participants reported having messaging interactions on dating apps with at least three other online daters, with the mean being 48.5 different messaging partners and the max being 1000. All but 4 participants met at least one online dater in-person, with a mean of 6.6 in-person dating partners and a max of 60. Most participants (93) had met an online dater in-person who had seen their AI-generated content before meeting. Generative AI tools used by participants to create content used in online dating included ChatGPT and common AI photo editors. Most participants (98) used generative AI to create multiple types of content, with the most common being private messages to other online daters (93), profile bios/"about me" sections (84), photos on one's profile or sent to an online dater in private messaging (55), and videos sent to online daters through private messaging (19).

3.2 Participant Care and Safety

We designed the study to minimize risk and center participant autonomy. **Discomfort alleviation:** Given the potentially sensitive nature of this study's subject matter, and associated discomfort that

participants may feel, all open-ended questions in the survey were skippable. Participants could also discontinue the survey at any time without penalty to their financial compensation.

Re-identification mitigation: We did not collect any direct identifiers. Our recruitment platform of choice handled participant compensation on our behalf and supported anonymous review of survey response quality by assigning a code to each respondent that only they could connect to a particular user on their recruitment platform. To further limit re-identification from this paper, we report demographic traits only in aggregate and do not attribute quotes to individual participant IDs. Survey data were stored on access-controlled institutional storage with team-only access; any indirect identifiers (e.g., references to specific dating apps in survey answers) were omitted during data cleaning.

Survey wording: The survey deliberately omitted the word "consent" from all question wording, based on guidance from three nursing and/or psychology professors with expertise in consent that we consulted for this study. The term may carry preconceptions or learned understandings of how consent "should" be exchanged, which could lead to self-censored answers and emotional duress if questions explicitly about "consent" incite participants to reflect on whether they committed a nonconsensual act or misinformed someone else's consent through their use of generative AI. We more broadly used question wording that sought to avoid connotations of "good" and "bad" use of generative AI so that participants would not perceive the survey as casting judgment on their generative AI strategies (this is also partly why we opted against interview or focus group methods). This also led us to ask about concepts in multiple ways, such as inquiring about (non)disclosure of AI-generated content through one's personal disclosure decisions and in conjunction with general opinions on when online daters should need to disclose.

Financial compensation: We ran multiple pre-launch survey pilots to estimate a completion time of 30 minutes, for which compensation of $10 USD was intended to pay 2.75x the USA-based federal minimum wage. The study advertisement informed participants that the expected completion time was 30 minutes and that they could skip any question and end the survey early with no adverse consequence to their compensation. The eventual median and mean completion times from our full sample were 48 and 50 minutes respectively, which were longer than our piloted completion times but still paid above the USA federal minimum wage of $7.25 USD per hour. As to why the median and mean completion times were longer than our pilot estimates, we suspect two reasons. One, the recommended compensation on our chosen recruitment platform was $12/hour, putting $10 USD for a median 48-minute survey in line with standard compensation rates. Participants may have thus more willing to give additional time. Two, the survey topic may have been of intrinsic importance to participants, leading to a voluntary decision to take more time.

3.3 Data Collection

We had anticipated a survey completion time of 30 minutes (and determined the $10 USD payment accordingly), however the mean completion time for our participants turned out to be 50.44 minutes (median: 47.54), with a minimum completion time of 15 minutes and maximum of 102 minutes.

The survey consisted predominantly of open-ended questions where participants could provide answers without any text limit. After eligibility and demographic questions, the survey first inquired about what types of online dating content participants used AI tools to generate, followed by motivations for using generative AI with the question "Why did you use AI tools for content creation, modification, or enhancement? Please be as specific as possible." They were then asked which generative AI tools they used, followed by a prompt to provide a detailed step-by-step account of how they used those generative AI tools to create content for online dating. The wording of that question was "How did you use these AI tools? Describe the steps or prompts as if you were teaching someone how to do it."

The next portion of the survey was about disclosure of generative AI-use to other online daters. Participants were first asked if they ever "met anyone face-to-face after they saw your AI-generated, modified, or enhanced content on the dating app?" For those that did, subsequent questions then asked if they disclosed their generative AI-use to any face-to-face meeting partners, followed by an open-ended prompt to explain why/not. For participants who had not met another user face-to-face who saw their AI-generated content, they were asked if they had disclosed their generative AI-use to "any online dater" and why/why not. The final portion of the survey inquired about personal opinions on when AI-generated content does and does not need to be disclosed on dating apps; this was intended to capture their broader mental models about generative AI disclosure norms and expectations.

3.4 Data Analysis

We used reflexive thematic analysis (RTA) [15] to examine the open-ended survey responses. RTA is one of three common approaches to thematic analysis (TA) in CHI literature; the other two being coding reliability TA and codebook TA [13]. Unlike these other thematic analysis approaches that use predefined codebooks and interrater reliability towards achieving an unbiased and objective analysis of data, reflexive thematic analysis (RTA) views knowledge as unavoidably *"situated and contextual"* where *"researcher subjectivity is conceptualised as a resource for knowledge production"* [14]. We found RTA especially useful for our study because of its accommodation of inductive and deductive coding, which allowed the foregrounding of the researchers' professional experience and literature surrounding consent, AI-mediated communication, and self-presentation to inform analysis.

The analysis team consisted of two researchers with prior publication and empirical research experience pertaining to computer-mediated consent, self-presentation, and impression formation. Due to the theoretical flexibility of RTA [15], the team's past work and professional experience contributed to the overall analysis through a preconceived deductive orientation to the data that centered these concepts, combined with inductive coding grounded in participants' language.

RTA has the following steps: 1) data familiarization; 2) initial coding; 3) initial themes; 4) development of themes; 5) refining themes; and 6) writing results. The authors familiarized themselves

with the data (step 1) through cleaning of the survey responses and prepping them for analysis by transporting them into both spreadsheets and a Miro board [80] (a virtual whiteboard web platform where each participant's quote is represented as a virtual notecard). We conducted initial inductive coding (step 2) in Miro by assigning semantic labels to participant responses. Initial themes (step 3) consisted of broad categories around content that was AI-generated, motivations for using generative AI in online dating, and decisions to (not) disclose use of generative AI. For instance, the use of generative AI to boost attractiveness was one amalgamation category at this stage. These initially-broad categories were further refined and broken down (steps 4–5) by repeatedly reorganizing quotes into visual sections in Miro; at this point, the differences in generative-AI use gradually became clearer (e.g., distinguishing AI-generated content that was intended to be broadly attractive from AI-generated content intended to be attractive to one specific dating partner). During the writing of this paper's findings (step 6), we produced tables to solidify our thematic structure around insights for AI-generated self-presentation strategies and disclosure strategies.

4 Findings

Through data analysis we organized themes into two halves: AI-generated self-presentation strategies (how AI-generated content was incorporated into dating app profiles and messages) and AI-generated content disclosure strategies (if, how, and when online daters disclosed their use of generative AI to dating partners).

Regarding self-presentation, most participants regardless of their gender or geographic location used AI-generated content to increase attention and attraction from other online daters. For these participants, AI-generated content had little or no personal relevance, but was rather crafted in accordance with how their AI tool of choice understood general romantic attractiveness. In contrast, a smaller set of participants spanning demographics used generative AI to improve the clarity of content intended as an authentic self-presentation. See Table 1 for a summary of four AI-generated self-presentation strategies found in our analysis, which are discussed in detail in Section 4.1.

Most participants did not voluntarily disclose use of generative AI to their face-to-face meeting partners, regardless of their intent for using AI-generated content. Rationale for non-disclosure revolved around avoiding romantic rejection that could be incurred if one's partner learned of their AI-use, acknowledging that generative AI carries a negative connotation to most participants. Those who did disclose use of generative AI cited transparency with their partner or an interest in discussing AI as a hobby. See Table 2 for a summary of reasons for disclosing and not disclosing generative AI-use. These are discussed in detail in Section 4.2. Passages in *italics* below are direct quotes from participants. We do not include participant identifiers to prevent re-identification.

4.1 Strategies for AI-Generated Self-Presentation on Dating Apps

Strategies of our participants for using AI-generated content in self-presentation were organized into four categories that vary based on the extent that which AI-generated content was intended to be relevant to their real identity. These included: AI as fabricated personality (no relevance to one's own identity), AI as personalized ideal dating partner (no relevance to one's own identity), AI for enhancement (partial relevance to one's own identity), and AI for authenticity (intended to be fully representative of one's identity). We discuss each strategy below; they are summarized in Table 1.

4.1.1 AI as Fabricated Personality. Several participants adopted a self-presentation strategy in which generative AI tools were prompted to create text content for profile bios and private messages that broadly aligned with personality traits and other non-physical qualities that were considered generally attractive. It was irrelevant as to whether the user believed they personally possessed these traits, and in many cases, participants acknowledged not having the traits that they wanted generative AI tools to create content about. Terms that participants commonly used to describe the personality they wanted to convey included *"bubbly," "intelligent," "witty," "creative,"* and *"funny."* Others had more abstract descriptions of the qualities they wanted to present, encapsulated with terms like *"catchy"* and *"attractive."* Participants enacting this self-presentation strategy dismissed the importance of romantic compatibility in online dating, and instead espoused a view that dating success is dictated by universally desirable traits. One participant exemplified this view by approaching generative AI as a *"dating coach"* that could provide them with content that would make them attractive to essentially every user on a dating app. They described their step-by-step use of generative AI like this: *"open the app. type the prompt 'act as a dating coach. help me make my profile on a dating app more inviting. I want to attract people of all genders' [...] i might then copy and paste what it suggested or give a prompt like 'sound younger' or 'make it funnier' or 'try again'."*

This view on universally desirable traits being key to dating success, and the associated expectation that generative AI tools could provide indiscriminately *"attractive answers,"* as one account put it, was also aligned with participants' stated goals for online dating. Participants freely acknowledged that their motivation was increasing matches on the dating app rather than cultivating genuine compatibility. For example, a participant wrote *"I use [generative AI tools] to make myself look unique and attract more people to match with me,"* and another indicated, *"I used an AI tool to generate a more attractive and attention trap bio to get more likes and attention from ladies on dating apps."*

Implied in several participant responses, such as the reference to *"attention trap"* profile content, was the belief that generative AI tools have an almost magical quality to create content that instigates sudden, strong attraction from any online dater. One online dater even aimed to make matches *"obsessed"* with them by asking ChatGPT, *"what are most interesting topics to chat about with a stranger to make them obsessed with you, then I select one topic and then I ask chatgpt,' would you craft a good looking message on this topic that I will send to this stranger, make the massage [sic] more appealing so that they would love to talk about this topic."*

While these mental models on generative AI capabilities and the mechanisms of romantic attraction may appear overly simplistic or naive, they become more understandable in light of participants' stated insecurities and self-doubt about dating. Several participants self-identified with dating challenges that they blamed on their own

Table 1: Strategies for incorporating AI-generated content into self-presentation on dating apps. The table summarizes four strategy types, what each involves, why participants used it, and which parts of their profiles or messages were generated with AI.

AI-generated self-presentation strategy	Description	Motivation	What content is AI-generated?
AI as a fabricated personality	AI used to fabricate a generally "attractive" personality, regardless of its relevance to the user's own identity	Users are insecure about their innate qualities as a dating partner; they value increased attention on dating apps rather than a genuine connection	Profile bio (open-ended text content) and private messages
AI as a personalized ideal dating partner	AI used to generate attractive messages that are personalized to a specific user by providing that user's profile and message content to AI, regardless of relevance to the sender's own identity	"Stand out" from the competition by appealing to supposed romantic compatibility with potential partners	Private messages
AI for enhancement	AI used to "enhance" profile pictures, ranging from minor lighting modification to intentionally misleading content such as changing background objects and body	Improve the physical attractiveness and general appeal of pictures	Pictures
AI for authenticity	AI used for clarity of authentic self-expression by editing grammar or generating content based on accurate information about the user	Avoiding misunderstandings by potential dating partners; elucidating genuine romantic compatibility	Profile bio (open-ended text content) and private messages

personality deficiencies or other unattractive non-physical qualities. In effect, they deemed their personalities inherently undatable or unattractive, thus needing generative AI to craft a new persona that would have better dating prospects. For example, a participant described being *"not creative enough"* and *"awkward."* Another recalled: *"I typically struggle with conversation with new people, so I decided to look to chatbots such as ChatGPT."*

It is important to note that participants' self-criticisms pertained exclusively to non-physical traits. No participant utilizing this self-presentation strategy made any reference to their physical appearance, and generative AI was seldom used to create pictures. Rather, AI-generated content was predominantly text-based, such as profile bios (what one online dater called the *"about me"* section) and private messages to other daters. Participants described prompting generative AI specifically for content to put in their dating app profile, e.g., *"a witty bio [. . .] catchy and creative."* For message content, generative AI was typically prompted to create generic conversation starters or *"pick up lines."* Participants provided specific examples of their prompts for such pickup lines that speak to the perception that a single conversation topic could spark universal attraction, e.g., *"the best pick up line that might get a girl's attention,"* or *"Write a flirty message that makes a great first impression on a dating app."*

4.1.2 AI as Personalized Ideal Dating Partner. Participants who used generative AI to create what we coded as a personalized ideal dating partner had similar practices as those in the previous theme about fabricated personality: both used generative AI to fabricate content that had little or no relevance to their actual identities. However, a key difference is that participants here used generative AI to increase attraction by fabricating unique romantic chemistry with individual online daters rather than a universally attractive persona. The strategy pertained exclusively to AI-generated messages that were tailored to a specific dating app user, and had two distinct phases: generating a first-message conversation starter specific to the other dater's profile, and automating responses throughout the messaging conversation with said dater until meeting face-to-face.

For opening messages, participants provided details of their match's profile into a generative AI tool along with a prompt to create a message tailored to that person. In one participant's words: *"what should I say or ask to make me sound like I'm interested in them romantically."* Some participants copy-and-pasted content from their match's profile verbatim into the generative AI tool, without that person's awareness, while others paraphrased the profile content. To the latter, an example of a prompt to generative AI was: *"I matched with someone who loves hiking and dogs. Create a fun and engaging opening message related to that."* Prompting by other participants became increasingly elaborate, in some cases by

asking the generative AI tool to come up with multiple variations of a message that the user could choose from.

Once messaging conversations were ongoing, participants used generative AI to automate most of their responses. In effect, these participants' dating partners were having messaging interactions entirely with a generative AI tool, with the participant's role being a data liaison: providing profile and message content from the dating partner to the AI tool, and relaying content created by AI back into the dating app chat interface as their own messages. One respondent provided a systematic breakdown: *"Login to the AI tool. Copy the response [received from the messaging partner on] the dating app. Paste it, with the command to suggest a response. Read through it, if satisfied, copy and paste it to the dating app and send the response."* In some cases, participants indicated specific personality traits they wanted to showcase in their responses, for example: *"[I would ask] it to respond to the messages that people would send to me in a witty, engaging, and contextually relevant manner."*

Participants further relied on AI for translation and interpretation, both to bridge language gaps and to decode intent (*"what do they mean by this?"*). This let them sustain interactions with partners who spoke other languages or used unfamiliar terms, but it also made them reliant on generative AI for meaning-making. As one participant noted, they would paste a confusing message into their generative AI tool and prompt it to *"Explain it to me."*

4.1.3 AI for Enhancement. Another strategy for generative AI is centered not on textual self-presentation but on optimizing profile photos. Participants described using generative AI tools to improve the visual quality of their images, with practices that ranged from subtle adjustments in lighting to transformative modifications of bodies, faces, and backgrounds. Unlike the strategies that fabricated dating personas entirely unrelated to one's personal identity, this approach was rooted in the optimization of content that, at least originally, was representative of the user.

For some participants in this theme, generative AI's appeal lay in minor corrections that made photos clearer and more professional without fundamentally altering their physical appearance. These adjustments included steps that would have previously been done in photo editing software, such as sharpening blurry images or *"adjustments to the lighting"* and *"exposure."* These seemingly minor changes were sometimes embedded in elaborate editing routines that layered multiple AI tools to refine appearance. One participant described a workflow that combined several applications: *"First, I used Lumii to adjust the lighting in my profile photos [. . .] I also applied minor High Dynamic Range effects to make the details pop without over-editing. With Fotor AI Photo Editor, I refined my images further [. . .] to create a warm and inviting look."*

Several participants emphasized restraint with use of generative AI on pictures so as not to cross the line into deception. One noted, *"I used the FaceApp to help edit my profile picture [. . .] I used this app with caution to avoid over-editing my profile picture to ensure it remained authentic."* Similarly, another participant described experimenting with features like *"Improve"* or *"Smooth"* but *"always keeping the intensity very low to achieve a natural enhancement [. . .] making minimal adjustments to look like a slightly better version of myself."* Such statements suggest participants recognized a fine line between minor touchups to their pictures and misrepresentation.

For others, however, enhancement extended beyond minor touchups into significant changes to their appearance and even location in photos. Examples of the latter were removing objects from pictures or changing the background of an image to an exotic location to imply having visited (e.g., a background implying a *"safari to Africa"*). Some participants justified these major photo changes in terms of *"confidence"*: improving photos was not just about aesthetics but about overcoming shyness or insecurity in dating. Per one participant: *"To be honest, I wasn't that much confident about my messages or bios, I wanted to enhance it, to spice it up a bit and etc, 'cause I am a little bit shy at the beginning."*

The intent behind these more extreme AI-driven photo enhancements was to improve physical appearance in order to attract *"high valuer"* dating partners, as one participant framed it, even if the resulting pictures had little resemblance to one's physical-world self. Words like *"beautiful"* and *"perfect"* were typical in responses describing this practice: *"It just made me look more beautiful,"* one participant wrote, and another commented, *"I used their beautifier or magic button [...] to make me look perfect."* While these AI alterations typically pertained to one's face, some described using generative AI to change their entire body shape to affect their perceived age. As one participant reflected, *"I also used a picture to enhance beauty and edit my body making myself young and [it] came out great."*

4.1.4 AI for Authenticity. In contrast to strategies that fabricated new dating personas with no relevance to user of the generative tool, a smaller group of participants described using generative AI to represent themselves more clearly and authentically. For these individuals, AI was not an attractive mask, but a way to overcome awkward phrasing, language limitations, or anxiety about honest self-expression. Their emphasis was on presenting who they *"really"* were without distortion, ensuring that potential partners could accurately gauge romantic compatibility. As a participant explained, *"I take online dating seriously, so if I am going to talk to someone, I want to get an honest shot at it so I don't want me not representing myself in the best way possible to hinder me from meeting the love of my life."*

Participants commonly framed generative AI as a tool for more clearly expressing their own thoughts in messages to other users, rather than inventing content to boost attraction. Use of the word *"my"* was common across responses in this theme, in emphasizing that generative AI was intended to help the expression of participants' own ideas. Per one participant: *"I sometimes struggle with putting my thoughts into words."* Several alluded to issues with grammar, wording choice, and *"sentence construction"* as motivations for seeking help from generative AI. Another respondent added: *"I wanted to express my thoughts in a better way. I chose to use AI because it is good in sentence construction."*

This pursuit of authenticity extended to dating profiles. Several participants used AI to generate bios that accurately described their hobbies, values, and relationship goals, but in a way that read more coherently than what they thought they could write on their own. Rather than fabricating traits, they saw AI as helping them assemble fragments of their identities into a more readable whole. One participant wrote, *"I mainly used it for making my Bio nice and that makes people know me better, I use it that way because sometimes it is hard for me to [write] a presentation about*

me." Another participant described feeding AI keywords such as *"hiking, cooking, live music, dogs, travel, witty, adventurous, craft beer"* to receive suggestions for how to write about these qualities in a profile bio. The intent was not to mislead but to ensure that the qualities they valued most about themselves did not get lost in poor wording.

Participants varied in how they enacted this authenticity strategy. Some relied on AI for revision and polish of already-written sentences, such as asking AI to make their content sound *"less harsh or more kind"* or using tools like Grammarly to ensure the content was free of grammatical errors. Others used AI in a more generative capacity, or what one participant called the *"fill in the blanks"* approach, by providing keywords or short phrases about their thoughts or qualities and letting generative AI tools craft full sentences based on the content. A survey participant reflected, *"I used ChatGPT to feed into my personal traits and strengths since I found it challenging to provide a clear description of myself."* Some participants spent considerable time crafting prompts to generative AI that described their personal qualities in detail, in hopes of further improving clarity in their self-descriptions. Another participant elaborated on their *"fill in the blanks"* approach when prompting AI with their interest in reading: *"The more specific you are, the better the AI can help. For example, instead of just 'reading,' you could say 'reading science fiction and historical biographies. "'*

4.2 Strategies for (Non-)Disclosure of AI-Generated Content to Online Dating Partners

Strategies for (not) disclosing use of AI-generated content to online dating partners were organized into six categories mapping to rationale provided for (non-)disclosure; see Table 2 for a summary. Most participants did not disclose their use of generative AI to any dating partners. Reasons for non-disclosure included fear of rejection, a perception that AI-generated content did not significantly affect dating outcomes, and a view that disclosure should happen, albeit later in the romantic relationship. For the participants that did disclose their use of AI-generated content, they did so because they either wanted to talk about AI as a hobby with their dating partner, they anticipated a worse impact on their dating success if their partner discovered their use of AI on their own, or they generally valued transparency in their dating interactions. Importantly, decisions to (not) disclose and their accompanying rationales were not tied to any particular self-presentation strategy from Section 4.1. In other words, most participants opted not to disclose their AI-use regardless if they used AI for authenticity, enhancement, or complete fabrication.

4.2.1 Non-disclosure of AI-Generated Content to Avoid Romantic Rejection. The dominant rationale for non-disclosure was a belief that honesty about AI-use would directly compromise romantic prospects. Participants weighed disclosure against the risks of rejection, embarrassment, or stigma, and overwhelmingly concluded that concealing their use of AI was the safer choice. For participants coded under this theme, the potential to mislead their dating partners through non-disclosure was either not mentioned or deemed acceptable as long as their own *"good first impression"* was maintained. In rare cases, participants admitted to intentionally trying to deceive their partners, for which disclosure of generative AI would be counterproductive. For example: *"I wanted them to think that I wrote that with my mind."*

Several participants described AI as a hidden advantage that allowed them to pass through early dating *"filters,"* fearing that acknowledgement of its use would undo the progress towards face-to-face encounters or subsequent romantic or sexual experiences. One participant argued, *"I feel like [disclosing generative AI-use] would produce a sense of falsehood which might undermine my goals with someone, whether it might be long term or short term."* Multiple participants made reference to being *"judged"* negatively by their dating partners if they disclosed their use of generative AI, which, in itself, was positioned as a reason not to disclose. As another participant explained, *"I fear that they might judge me and may choose [not] to meet me again."* A different participant echoed, *"Some people [...] might judge others harshly for using it [generative AI] in something as intimate as dating."*

In explaining why they feared rejection after AI disclosure, several participants referenced a concern of being perceived as *"fake"* or *"not real."* For instance, one participant admitted being *"afraid"* that their meeting partner would think they were *"different from the person they were chatting with."* Another worried about being seen as a *"fake person."* This viewpoint conveyed a passive acknowledgment that AI-generated content is–or could at least be seen as–deceptive, yet rather than use this as motivation to disclose their AI-use and remedy any such association, participants intentionally hid it for their own gain. As one respondent added, they *"didn't feel comfortable doing it [disclosing AI-use]"* because they thought the meeting partner *"would feel tricked or violated or lied to."*

Relatedly, other participants thought generative AI-use would be unfairly associated with unattractive personal qualities, and so non-disclosure was a way to avoid these unappealing and inaccurate impressions. Participants provided an extensive list of such qualities, including *"weird," "boring," "dumb," "lazy,"* and even *"psycho."* As one of the participants put it: *"the person may misunderstand me and assume that I'm a lazy person and one that lacks effort."*A different participant shared a similar concern: *"I was afraid that they would think I'm lazy and not smart enough to create my own attractive profile."*

4.2.2 Non-Disclosure of AI-Generated Content Due to Downplayed Impact on Dating Outcomes. Participants in the previous theme feared that disclosure of AI-generated content would undue its essential role in attracting other online daters. In contrast, participants under this theme minimized the impact of generative AI on their dating success altogether. For these individuals, AI was framed as an *"irrelevant thing,'* that polished or simplified communication but did not meaningfully alter who they were or why their dating partners would want to meet them. Because they saw its influence as negligible, they concluded that disclosure of generative AI was unnecessary.

Minimizing disclosure often reflected a broader belief that attraction is determined by *"real"* interpersonal chemistry, not technological intervention. Some participants outright denied that generative

Table 2: Rationale for (not) disclosing use of AI-generated content to dating partners. The table groups participants' explanations into themes and subcategories; most participants opted not to disclose their AI-use to any dating partners they met face-to-face.

Group	Theme	Subcategories
Reasons for not disclosing	1. Fear of rejection if partner learns of their AI-use	1.A Partner may think they are a fake person/scammer 1.B AI-use associated with undesirable traits 1.C AI was used to intentionally deceive the partner
	2. Downplayed impact of AI-generated content on dating outcomes	2.A AI-generated content conveyed the "real" me 2.B Perception that everyone uses AI 2.C Partner's responsibility to ask about AI-use
	3. Disclosure should happen later in the relationship formation	—
Reasons for disclosing	1. Disclosure as opportunity to talk about interest in AI	—
	2. Preemptive disclosure to avoid a negative impact on dating outcomes if AI-use is discovered	—
	3. Transparency is valued by the user of AI-generated content	—

AI had any influence on their dating partners. Another participant explained, *"it was irrelevant because I don't believe the prompts are what got the person engaged to meet with me in person."* Others dismissed it as *"just simple text"* or *"not important."*

For some, minimization of AI's role in their dating outcomes was tied to the view that their AI-generated content did not misrepresent their real identity. In participant responses, this sometimes took the form of minimizing the amount of content that was AI-generated. The word *"only"* appeared across multiple participants here, e.g., *"I only used AI to help me correct the description text and create initial conversation topics."* In another case a participant explained that *"everything the AI generated was all true,"* referring not to the amount of AI-generated content they used, but its veracity as a depiction of their physical-world identity. As one participant elaborated: *"...it just helped me present myself in the best way possible. My photos still looked like me. My profile description [...] the personality and interests were entirely mine. Even in messaging, the responses reflected my real thoughts and feelings. Since AI didn't alter my identity, I didn't feel the need to bring it up, especially once we met in person and connected naturally."*

Other participants took a different course by arguing that generative AI is so *"normal"* and *"widespread"* that any effect of AI-generated content on dating outcomes would be shared evenly across the userbase. Beyond the online dating context, some participants referred to AI's pervasive use in *"academics"* and *"writing"* more generally as a way to dismiss the necessity of discussing their own content. One participant explained: *"I believe there was no reason to go out of my way and inform the other person that I had edited photographs. I believe that everyone uses tools to make themselves look more appealing than they usually ar [sic], and we can all expect to see small differences in person compared to online profiles. Just like makeup, we can use tools to fix small errors on photographs that we like, such as better hair, beard, cleaner background, and better resolution. On the first meeting, we generally look like the best versions of ourselves anyway."*

Relatedly, some participants shifted the responsibility for generative AI disclosure onto their partners, arguing that they do not need to proactively disclose its use, but rather only if their partner directly asks about it. They argued that if AI-use truly mattered to other users, it was up to others to inquire: *"I thought it was not necessary to tell them because they didn't ask anything concerning my photos,"* one participant remarked. Another echoed, *"I did not feel it was necessary, especially since the person did not ask about it."* Implied in some of the responses in this theme is a form of conditional disclosure, based not only on a partner's verbal inquiries about generative AI, but also on the partner's ability to even suspect its use. For instance, a participant explained that the absence of partner suspicion validated nondisclosure: *"I felt it wasn't important at the time and the fact that they didn't notice and actually we even met I felt it was really a non-issue."*

4.2.3 Disclosure of AI-Generated Content as Hobby, Act of Transparency, and Step in Relationship Formation. Disclosure of AI-generated content to online dating partners was less common than nondisclosure, and participants gave a variety of reasons for doing so. Perhaps the most surprising was to discuss generative AI as a hobby. Some participants described wanting to have conversations

specifically about generative AI with their dating partners because AI was a general interest of theirs and disclosure was a natural way to steer conversation towards the topic. As one of the participants put it, *"it's a bit of fun to talk about how you come up with whats on your profile so as part of that chat I'd say I used ChatGPT to jazz it up a bit. also gives an insight to what the other person thinks about AI, tech, current state of affairs."* Some participants also encouraged their partners to download generative AI tools themselves and create their own content. For instance, one participant offered to give a tutorial to their dating partner.

Other participants disclosed their AI-use as a matter of principle: they simply *"wanted to be transparent"* and viewed AI disclosure as an opportunity to *"be truthful and free with the person."* Words like *"true"* and *"real"* were typical, with participants referring to transparency as a general *"value"* in their dating lives. There was an element of pragmatism in this transparency for a few participants, particularly those who used AI-generated pictures on their profiles. They had anticipated that dating partners may find their physical appearance to differ from their pictures, so they preemptively disclosed their use of generative AI before meeting face-to-face. As one participant noted, *"i may not [look] exactly as my AI-generated photo, so do not be alarmed."*

Another consideration raised by participants who did and did not disclose their generative AI use was *"when"* disclosure should happen. Some participants viewed AI disclosure as an important milestone in romantic relationship development, and one that should happen only after the relationship has had time to progress. The first date was sometimes singled out as a point that is *"too soon"* to disclose AI. As one participant said: *"I don't feel its relevant when meeting someone for the first time."* While a specific point in time when disclosure was appropriate was not clearly identified by participants in this category, at least one alluded to longer-term plans contingent on relationship success: *"My plan was to disclose the information once it worked out, but none have so far."*

5 Limitations

Our study has several limitations that shape how the findings should be interpreted. Because the survey was administered in English, participation was self-selected for English speakers, which may partly explain the relatively minimal representation from South America and Asia. Recruitment also explicitly framed the study as being about generative AI-use, which likely attracted individuals comfortable reflecting on that practice and may have excluded those for whom AI was a less significant part of their dating experiences. Relatedly, our sample may under-represent marginalized groups for whom safety-oriented self-presentation is most consequential and for whom AI-generated content could serve fundamentally different purposes. As with any self-report method, we cannot independently verify the accuracy of participant accounts. At the same time, participants' candid acknowledgement of dating struggles, including intentional deception through AI-generated content, should temper concern about false answers. Finally, our study focused on the producers of AI-generated content; we did not include the recipients of that content, and thus cannot determine whether partners themselves considered their consent to subsequent interaction to be misinformed.

6 Discussion

We conducted a qualitative survey study with 113 online daters about their use of AI-generated content in self-presentation on dating apps, as well as their decisions to (not) disclose such content as AI-generated to their dating partners. The study found that AI-generated content was often intended to boost one's attractiveness to other dating partners, even as an intentional misrepresentation of one's real identity and personal qualities **(RQ1)**. Furthermore, use of AI-generated content was seldom disclosed to dating partners, including those they met face-to-face for dates, to protect one's own dating prospects and avoid romantic rejection **(RQ2)**. The findings point to negative implications of AI-generated content on informed consent to interaction in online dating **(RQ3)**, which may go undetected even during face-to-face dates because generative AI was most often used for text content reflecting intangible personal attributes.

In this section, we first reflect on the implications of our findings on consent to interpersonal behavior, particularly whether undisclosed use of generative AI misinforms consent to the extent that subsequent interactions may be considered nonconsensual. We then discuss implications on the AI-mediated communication literature by comparing our findings to the largely-experimental results of past work about the influence of AI-generated content on communication outcomes, and considering how our findings may inform decisions about automated disclosure of AI-generated content on social platforms. Lastly, we compare our findings to "pre-generative AI" knowledge of self-presentation strategies on dating apps, through which we note the conspicuous absence of safety-oriented uses of generative AI in our study and future work that may explore that motivation for use.

6.1 Implications for Computer-Mediated Consent to Interpersonal Behavior

HCI literature has consistently shown that Computer-Mediated communication complicates consent to interpersonal behavior [129]. Consent itself can be misinterpreted through online signals like emojis and profile pictures on dating apps [124], and avatars in social VR [95]. This can potentially lead to unwanted sexual activity and other intimate behaviors simply because a partner's agreement to the act was misunderstood through prior interaction online [123] or not reconfirmed [29]. Making adequately informed consent decisions is also a challenge in Computer-Mediated Communication [60, 100]. Given the bevy of information potentially available about a person online, dating app and social VR users have openly questioned when consent becomes "too" informed [17]. Likewise, self-presentation capabilities in novel virtual contexts like VR enable the creation of online personas that diverge wildly from one's physical appearance and behaviors, thus obfuscating who a person is giving consent to [59].

Our findings portray generative AI as the latest complicating factor for informed consent to interpersonal behavior. Most of our participants used AI-generated content for self-presentation of traits and mannerisms that had no relevance to their actual identities; in many cases, this content was intended to diverge from their real self because of insecurities and doubts about being able to procure a date on their own. This leads us to conclude

that AI-generated content in online dating is indeed misinforming consent to interaction, including face-to-face dates and any intimate behaviors that occur therein.

Our study did have some participants who used generative AI for help with presenting an authentic version of self, which departs markedly from intentionally manipulative uses of AI. While less common in our analysis, this type of use suggests that generative AI can be helpful for individuals with introverted personalities or who otherwise struggle to express themselves in written form. Yet this type of AI-use could still be problematic through a consent lens because such users cannot presumably continue using AI for assistance in expression in face-to-face encounters, and they may be understating the importance of the way AI uniquely articulates their thoughts on partners' attraction and consent decisions. Ultimately, even well-intentioned uses of generative AI can still have unintended influence on informed consent decisions of interaction partners. This speaks to perhaps the most concerning implication on informed consent from our study: that AI-generated content was predominantly text-based (e.g., profile bios, private messages), intended to convey non-physical traits and compatibility. Prior online dating research indicates that such traits are important to attraction [125], but difficult to reliably communicate [31] and immediately identify as inaccurate during a face-to-face date relative to differences in physical appearance. As such, misinformed consent decisions may not be self-corrected simply by meeting face-to-face.

These findings lay the groundwork for several lines of future inquiry. For one, our study provides a baseline typology of AI-generated self-presentation strategies that can now be subjected to subsequent quantitative assessment for relative frequency of these strategies and how their adoption may vary across user groups at scale. For example, future work could explore how adoption of the generative AI practices discovered in our study vary on the basis of gender given misogynistic "pickup" culture that some men prescribe to [127] versus more safety-oriented self-presentation strategies adopted by other gendered user groups [36]. Studying variations across dating apps could also be fruitful given some have connotations for casual sex [101].

Further research can also look into what exactly "informed" means in consent decisions, and at what point consent can be considered mis- or under-informed in computer-mediated communication. There is opportunity for collaboration with legal scholars on this front, who have contended with the complexities of (mis)informed consent in regard to non-digital comparators such as condom-use [50, 117] and HIV status disclosure [1]. Questions about specificity of the "informed" element of consent to interpersonal behavior are starting to be raised in HCI literature [60], in part to problematize affirmative consent [17, 95] and persistent abstract descriptions of what qualifies as "informed." Yet progress on this question pales in comparison to research into informed consent to personal data collection [20, 119], which has experienced rich discussion and ideation about the concept of informed consent in emerging computing contexts [115] and system design for procuring informed consent to personal data [45, 73, 107].

This is an important line of inquiry because misinformed consent to certain behavior, such as sexual activity, can be qualified as rape by deception [2, 4]. Even though AI-generated content is created and distributed online, it can have implications on sexual acts across online and in-person modalities. Rape by deception has been well studied outside of HCI [99], and has garnered some attention in HCI in the context of Internet-enabled sex toys when the teleoperator of one's sex toy is different from the person they intended [99]. Computer-unmediated examples of rape by deception include undisclosed HIV status [22] and removing a condom without one's partner knowing [28, 52]. Some of our participants did use AI-generated content intentionally to deceive a partner into romantic attraction, and several more did not disclose their AI-use specifically to avoid accusations of deception from their dating partners. These uses, combined with non-disclosure of AI, may qualify as rape by deception if such interactions culminated in sexual activity and influenced their partners' willingness to have sex.

Clearly more research is needed into how generative AI may be repurposing dating apps into devices for rape by deception (dating apps have already been shown to be a significant sexual violence risk factor [108]). Future work could involve individuals who realized or suspected generative AI-use by their partner after sex, or participatory design of clearer, prescriptive rules around when and what uses of generative AI qualify as misinforming consent if not disclosed.

6.2 Implications for AI-Mediated Communication

AI-mediated communication is a relatively new research topic in HCI [54], and one largely comprised of experiments [41, 58, 61, 62, 72, 78, 79, 118] to assess how AI-generated content in controlled social scenarios impacts trust, uncertainty, and other relational dynamics. Our study provides an opportunity to extend that work with insight from ecologically valid situations of AI-generated content "in the wild."

Prior studies have shown that undisclosed AI-use can reduce trust and increase uncertainty in the sender of the content [61, 72, 118]. While we did not study the recipients of AI-generated content, our participants certainly thought their use of generative AI would reduce trust. It was this perception that fueled decisions of non-disclosure for many participants. In practice, disclosure of generative AI-use to dating partners was rare in our study. The primary reason participants withheld disclosure was fear of romantic rejection, reflecting the strong negative connotation they believed partners would attach to AI-use. This contrasts with prior work suggesting that known use of AI-generated "smart replies" can increase trust by serving as source of blame for miscommunication between human partners [58]. Of course, our participants' anticipation of reduced trust after disclosure may not be indicative of how their partners would actually react to such disclosure. Yet their use of AI to intentionally curate an attractive (and often misleading) dating persona is a far different use of AI from simply auto-reply suggestions, which renders their anticipation of a distrustful reaction from partners at least understandable, if not likely.

Research into the disclosure of generative AI-use is also starting to emerge. For example, research of students using generative AI for coursework [5] has found nondisclosure to teachers to be the norm due to fear of judgment or penalty. This echoes findings from a larger scale experiment in which AI authorship was not declared

due to anticipated negative reactions from readers [32]. Our findings from the online dating context show the same justification for nondisclosure: fear of negative repercussions to their impression and failure to achieve their goals (in our case, going on dates).

If this growing corpus of research is an indication that use of AI-generated content will typically go undisclosed, it begs the question of if or when disclosure should be automated in light of ever-improving technical advances in AI detection [30]. Emerging AI-use labels and content provenance standards such as C2PA "Content Credentials," and their deployment on social media apps like TikTok, show that tamper-evident metadata can in principle support disclosure at scale [21, 102]. However, these systems are optimized for relatively persistent images and videos and often rely on embedded metadata that is stripped or broken when users screenshot, copy and paste, or re-upload content across apps, which makes them a poor fit for the short-lived, cross-app text chats that shape impressions in online dating [83]. This question has been considered in the context of news misinformation [70], but contexts of human-to-human social interaction can be more complicated precisely because of the impacts on trust that prior experiments have shown [61, 72, 118] and which our own participants anticipated. Forced disclosure of AI-use in content that does not ultimately affect a partner's consent to subsequent interaction may unduly spark distrust and communication breakdown. Yet our study's findings also show how leaving disclosure as a completely voluntary decision can lead to misinformed consent to incredibly intimate interactions at scale. Future work could apply a consent lens to elicit stakeholder opinions on when use of AI should and should not be automatically disclosed in social computing contexts.

6.3 Implications for Self-Presentation in Online Dating

Two decades of online dating research have established self-presentation in dating apps to be a strategic and often deceptive practice, motivated by the need to appear attractive to potential dating partners [38, 46, 53, 55, 56, 103, 104]. At the same time, self-presentation has been found to be a practice for safety by marginalized and at-risk user groups, defined by strategic concealment and disclosure of stigmatized traits such as trans identity [36], disability [89], sexual orientation [11, 12, 23], and HIV status [110, 111] to avoid harassment or violence. Generative AI has added scale to capabilities for self-presentation in online dating, and our findings show both continuity and divergence from these "pre-generative AI" impression management tactics.

Motivations of most of our participants to boost their attractiveness to other daters through AI-generated content aligned strongly with self-presentation motivations in prior work. In particular, practices around photo "enhancement" align with prior findings around exaggerated self-presentation like using outdated photos on one's profile [34, 56, 105]. Yet there are notable differences in AI-mediated self-presentation strategies for attractiveness relative to prior work. For one, the extent of deception appears to be much more extreme than in prior literature, such as in Ellison and colleagues' "profile as promise" framework [35] that describes exaggeration as being tempered by the prospect of in-person meetings (in effect, a "promise" that one's profile will not deviate too significantly from their physical world self). Our participants, in contrast, deliberately created entirely new personalities that intentionally espoused attractive personality qualities that they did not personally possess.

Furthermore, while prior work on deception in online dating primarily focuses on self-presentation in profiles, AI-generated content was perhaps most crucial in private messaging for our participants. The role of messaging in self-presentation on dating apps [23, 46, 47, 114] has been discussed only occasionally in prior work [76, 131]. This may be a consequence of the sheer labor needed to maintain an impression over multiple messages, and with individual users, hence dating coaches and pickup artists selling one-size-fits-all messages that one could send indiscriminately to every dating partner [128]. Generative AI not only adds scale to message writing, but tailored self-presentation to each individual dating partner with little effort beyond copy-and-pasting content from each partner into a generative AI tool–one participant in our study even used the term "dating coach" when describing their AI-generated self-presentation.

In light of prior work, it is surprising that our study elucidated no safety-oriented strategies for AI-generated self-presentation. Unlike earlier work documenting how marginalized users hide or alter aspects of identity for self-protection, our participants rarely described using AI tools to shield themselves from harassment or identity-based harm. Some speculative examples would include AI-modified profile content to strategically describe personal interests and qualities without implying stigmatized personality traits, or AI-generated profile pictures that better align with one's gender identity. The absence of such uses in our data may be indicative of access barriers to generative AI, unknown motivations for avoiding AI-use, or an under-representation of user groups that would most benefit from safety-oriented uses of generative AI. Future work could include dedicated study of marginalized user groups, or the recruitment of individuals who self-identify safety as a motivation for using AI-generated content.

7 Conclusion

Generative AI adds unprecedented scale to the ability to craft dating personas that can diverge substantially from one's physical-world self. In doing so, it complicates whether partners' consent to interaction and intimacy can be considered misinformed if AI-use is not disclosed to them. Yet there is little empirical understanding of AI-generated self-presentation and disclosure decisions. Through a qualitative survey of 113 online daters, we found that generative AI was most often used to fabricate appealing personality traits in profile bios and messages that were not necessarily reflective of the user's actual identity. Non-physical traits are known to be important to impression formation and decision-making in dating contexts, yet unlike edited photos, divergences between AI-generated and actual personality traits cannot always be readily identified in face-to-face meetings. Disclosure of AI-use to partners was rare, with participants citing fear of rejection as the primary reason for concealment.

These findings carry two important implications. First, they point to a consent-relevant risk: partners may agree to meet or become intimate under impressions authored by an algorithm rather than

their partner. Second, they highlight a disclosure dilemma: while voluntary disclosure was seen as socially costly, the absence of disclosure poses a question as to whether dating partners could make fully informed choices. We conclude that generative AI has the potential to become a significant sexual assault risk factor by creating conditions of mis- or under-informed consent at scale. For HCI, this work underscores the urgency of reexamining consent frameworks, moving AI-mediated communication studies beyond laboratory settings, and expanding design conversations about automated disclosure of AI-generated content.

Acknowledgments

This work was supported in part by the U.S. National Science Foundation under Grant No. 2211896, 2401775, and 2339431, and by the University of Michigan-Flint through the CIT Research Seed Grant and the IST Large Grant Award. We thank the participants who shared their experiences for this study, and the reviewers and associate chair for their constructive feedback.

References

[1] 1998. R. v. Cuerrier. File No.: 25738.
[2] 2014. No Way around Consent: A Reply to Rubenfield on "Rape-by-Deception". *Yale Law Journal Forum* 123 (2014), 321. 2013–2014.
[3] 2014. R. v. Hutchinson. File No.: 35176.
[4] 2022. How to Expand Rape by Deception and Protect Consent. *University of Massachusetts Law Review* 17 (2022), 112. 2021–2022.
[5] Rudaiba Adnin, Atharva Pandkar, Bingsheng Yao, Dakuo Wang, and Maitraye Das. 2025. Examining Student and Teacher Perspectives on Undisclosed Use of Generative AI in Academic Work. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1071, 17 pages. doi:10.1145/3706598.3713393
[6] Meryem Barkallah, Yosr Aissa, and Douglas Zytko. 2025. Transparent Hearts: Balancing Privacy and Trust in AI-Generated Self-Presentation for Online Dating. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 883, 8 pages. doi:10.1145/3706599.3720311
[7] Kathleen C. Basile, Sharon G. Smith, Matthew J. Breiding, Michele C. Black, and Reshma R. Mahendra. 2014. *Sexual Violence Surveillance: Uniform Definitions and Recommended Data Elements* (version 2.0 ed.). Technical Report. Centers for Disease Control and Prevention.
[8] Amna Batool, Mustafa Naseem, and Kentaro Toyama. 2024. Expanding Concepts of Non-Consensual Image-Disclosure Abuse: A Study of NCIDA in Pakistan. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 398, 17 pages. doi:10.1145/3613904.3642871
[9] Laina Y Bay-Cheng and Rebecca K Eliseo-Arras. 2008. The making of unwanted sex: Gendered and neoliberal norms in college women's unwanted sexual experiences. *Journal of sex research* 45, 4 (2008), 386–397.
[10] Eric G. Benotsch, Vivian M. Rodríguez, Kristina Hood, Shannon P. Lance, Marisa Green, Aaron M. Martin, and Mark Thrun. 2012. Misleading sexual partners about HIV status among persons living with HIV/AIDS. *Journal of Community Health* 37, 5 (oct 2012), 1049–1057. doi:10.1007/s10900-011-9529-5
[11] Jeremy Birnholtz, Shruta Rawat, Richa Vashista, Dicky Baruah, Alpana Dange, and Anne-Marie Boyer. 2020. Layers of Marginality: An Exploration of Visibility, Impressions, and Cultural Context on Geospatial Apps for Men Who Have Sex With Men in Mumbai, India. *Social Media + Society* 6, 2 (2020), 2056305120913995. doi:10.1177/2056305120913995
[12] Courtney Blackwell, Jeremy Birnholtz, and Charles Abbott. 2014. Seeing and being seen: Co-situation and impression formation using Grindr, a location-aware gay dating app. *New Media & Society* (2014). doi:10.1177/1461444814521595
[13] Robert Bowman, Camille Nadal, Kellie Morrissey, Anja Thieme, and Gavin Doherty. 2023. Using Thematic Analysis in Healthcare HCI at CHI: A Scoping Review. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 491, 18 pages. doi:10.1145/3544548.3581203
[14] Virginia Braun and Victoria Clarke. 2021. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative Research in Psychology* 18, 3 (2021), 328–352. doi:10.1080/14780887.2020.1769238
[15] Virginia Braun and Victoria Clarke. 2021. *Thematic Analysis: A Practical Guide*. SAGE Publications Ltd, London.
[16] Matthew J. Breiding, Kathleen C. Basile, Sharon G. Smith, Michele C. Black, and Reshma R. Mahendra. 2015. *Intimate Partner Violence Surveillance: Uniform Definitions and Recommended Data Elements* (version 2.0 ed.). Technical Report. Centers for Disease Control and Prevention.
[17] Braeden Burger, Devin Tebbe, Emma Walquist, Toby Kind, and Douglas Zytko. 2025. Saying No to "Yes Means Yes": Limitations of Affirmative Consent for Mitigating Unwanted Behavior Online According to Women and LGBTQIA+ Stakeholders. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–17. doi:10.1145/3706598.3713236
[18] Deng Cai, Yan Wang, Lemao Liu, and Shuming Shi. 2022. Recent Advances in Retrieval-Augmented Text Generation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 3417–3419.
[19] Centers for Disease Control and Prevention. 2022. Fast Facts: Preventing Child Sexual Abuse.
[20] Yi-Shyuan Chiang, Omar Khan, Adam Bates, and Camille Cobb. 2024. More than just informed: The importance of consent facets in smart homes. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 849, 21 pages. doi:10.1145/3613904.3642288
[21] Coalition for Content Provenance and Authenticity. 2025. C2PA: Coalition for Content Provenance and Authenticity. https://c2pa.org/.
[22] Susan D. Cochran and Vickie M. Mays. 1990. Sex, lies, and HIV. *New England Journal of Medicine* 322, 11 (1990), 774–775. doi:10.1056/NEJM199003153221111
[23] Elena Francesca Corriero and Stephanie Tom Tong. 2016. Managing uncertainty in mobile dating applications: Goals, concerns of use, and information seeking in Grindr. *Mobile Media & Communication* 4, 1 (2016), 121–141.
[24] Yichao Cui, Naomi Yamashita, Mingjie Liu, and Yi-Chieh Lee. 2022. "So Close, yet So Far": Exploring Sexual-minority Women's Relationship-building via Online Dating in China. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. Association for Computing Machinery, New York, NY, USA, Article 394, 15 pages. doi:10.1145/3491102.3517624
[25] Valdemar Danry, Pat Pataranutaporn, Matthew Groh, Ziv Epstein, and Pattie Maes. 2024. Deceptive AI systems that give explanations are more convincing than honest AI systems and can amplify belief in misinformation. arXiv:2408.00024 [cs.AI]
[26] Isha Datey, Hunter Soper, Khadeejah Hossain, Wing-Yue Geoffrey Louie, and Douglas Zytko. 2023. Ethical Participatory Design of Social Robots Through Co-Construction of Participatory Design Protocols. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. 2454–2461. doi:10.1109/RO-MAN57019.2023.10309539
[27] Sumanth Dathathri, Abigail See, Sumedh Ghaisas, Po-Sen Huang, Rob McAdam, Johannes Welbl, Vandana Bachani, Alex Kaskasoli, Robert Stanforth, Tatiana Matejovicova, et al. 2024. Scalable watermarking for identifying large language model outputs. *Nature* 634, 8035 (2024), 818–823. doi:10.1038/s41586-024-08025-4
[28] Kelly Cue Davis, Julia F. Hammett, and Mitchell Kirwan. 2023. A Scoping Review of Nonconsensual Condom Removal ("Stealthing") Research. *Trauma, Violence, & Abuse* 25, 1 (2023). doi:10.1177/15248380221146802
[29] Christopher Dietzel. 2024. Clickable Consent: How Men Who Have Sex with Men Understand and Practice Sexual Consent on Dating Apps and in Person. *The Journal of Sex Research* 61, 3 (2024), 481–494. doi:10.1080/00224499.2023.2235584 PMID: 37526356.
[30] Hyo Jin Do, Molly Q Feldman, Jessica He, Angel Hsing-Chi Hwang, and Seyun Kim. 2025. Navigating Generative AI Disclosure, Ownership, and Accountability in Co-Creative Domains. In *Adjunct Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 2, 4 pages. doi:10.1145/3707640.3729212
[31] Judith Donath. 2007. Signals in Social Supernets. *Journal of Computer-Mediated Communication* 13, 1 (2007), 231–251. doi:10.1111/j.1083-6101.2007.00394.x
[32] Fiona Draxler, Anna Werner, Florian Lehmann, Matthias Hoppe, Albrecht Schmidt, Daniel Buschek, and Robin Welsch. 2024. The AI Ghostwriter Effect: When Users do not Perceive Ownership of AI-Generated Text but Self-Declare as Authors. *ACM Trans. Comput.-Hum. Interact.* 31, 2, Article 25 (Feb. 2024), 40 pages. doi:10.1145/3637875
[33] Abdallah El Ali, Karthikeya Puttur Venkatraj, Sophie Morosoli, Laurens Naudts, Natali Helberger, and Pablo Cesar. 2024. Transparent AI Disclosure Obligations: Who, What, When, Where, Why, How. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 342, 11 pages. doi:10.1145/3613905.3650750
[34] Nicole Ellison, Rebecca Heino, and Jennifer Gibbs. 2006. Managing Impressions Online: Self-Presentation Processes in the Online Dating Environment. *Journal of Computer-Mediated Communication* 11, 2 (2006), 415–441. doi:10.1111/j.1083-

6101.2006.00020.x

[35] Nicole B. Ellison, Jeffrey T. Hancock, and Catalina L. Toma. 2011. Profile as promise: A framework for conceptualizing veracity in online dating self-presentations. *New Media & Society* 14, 1 (2011), 45–62. doi:10.1177/1461444811410395

[36] Julia R. Fernandez and Jeremy Birnholtz. 2019. "I Don't Want Them to Not Know": Investigating Decisions to Disclose Transgender Identity on Dating Platforms. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–21.

[37] Pierre Fernandez, Guillaume Couairon, Hervé Jégou, Matthijs Douze, and Teddy Furon. 2023. The Stable Signature: Rooting Watermarks in Latent Diffusion Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, 22466–22477. doi:10.48550/arXiv.2303.15435

[38] Andrew T. Fiore and Judith S. Donath. 2004. Online Personals: An Overview. In *CHI'04 Extended Abstracts on Human Factors in Computing Systems*. 1395–1398. doi:10.1145/985921.986073

[39] Jaclyn Friedman and Jessica Valenti. 2019. *Yes Means Yes!: Visions of Female Sexual Power and a World Without Rape*. Seal Press.

[40] Jeana H. Frost, Zoe Chance, Michael I. Norton, and Dan Ariely. 2008. People are experience goods: Improving online dating with virtual dates. *Journal of Interactive Marketing* 22, 1 (2008), 51–61.

[41] Liye Fu, Benjamin Newman, Maurice Jakesch, and Sarah Kreps. 2023. Comparing Sentence-Level Suggestions to Message-Level Suggestions in AI-Mediated Communication. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3544548.3581351

[42] Nicholas Furlo, Jacob Gleason, Karen Feun, and Douglas Zytko. 2021. Rethinking Dating Apps as Sexual Consent Apps: A New Use Case for AI-Mediated Communication. In *Companion Publication of the 2021 Conference on Computer Supported Cooperative Work and Social Computing (CSCW '21 Companion)*. Association for Computing Machinery, New York, NY, USA, 53–56. doi:10.1145/3462204.3481770

[43] Chris Geeng, Natalie Chen, Kieron Ivy Turk, Jevan Hutson, and Damon McCoy. 2024. "Say I'm in public...I don't want my nudes to pop up.": user threat models for using vault applications. In *Proceedings of the Twentieth USENIX Conference on Usable Privacy and Security (SOUPS '24)*. USENIX Association, USA, 433–451.

[44] Christine Geeng, Jevan Hutson, and Franziska Roesner. 2020. Usable sexurity: Studying {People's} concerns and strategies when sexting. In *Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020)*. 127–144. https://www.usenix.org/conference/soups2020/presentation/geeng

[45] Nina Gerber, Alina Stöver, Justin Peschke, and Verena Zimmermann. 2023. Don't Accept All and Continue: Exploring Nudges for More Deliberate Interaction with Tracking Consent Notices. *ACM Trans. Comput.-Hum. Interact.* 31, 1, Article 1 (Nov. 2023), 36 pages. doi:10.1145/3617363

[46] Jennifer L. Gibbs, Nicole B. Ellison, and Rebecca D. Heino. 2006. Self-presentation in online personals: The role of anticipated future interaction, self-disclosure, and perceived success in internet dating. *Communication Research* 33, 2 (2006), 152–177. doi:10.1177/0093650205285368

[47] Jennifer L. Gibbs, Nicole B. Ellison, and Chih-Hui Lai. 2011. First comes love, then comes Google: An investigation of uncertainty reduction strategies and self-disclosure in online dating. *Communication Research* 38, 1 (2011), 70–100. doi:10.1177/0093650210377091

[48] Erving Goffman. 1978. *The Presentation of Self in Everyday Life*. Harmondsworth, New York, NY.

[49] Richard Gomer, m.c. schraefel, and Enrico Gerding. 2014. Consenting Agents: Semi-Autonomous Interactions for Ubiquitous Consent. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication (UbiComp '14 Adjunct)*. Association for Computing Machinery, New York, NY, USA, 653–658. doi:10.1145/2638728.2641682

[50] Lise Gotell. 2020. Thinly Construing the Nature of the Act Legally Consented to: The Corrosive Impact of R v Hutchinson on the Law of Consent. *UBCL Rev.* 53 (2020), 53.

[51] Isabel Grant and Lise Gotell. 2020. Does" No, Not without a Condom" Mean" Yes, Even without a Condom"?: The Fallout from R v Hutchinson. *Dalhousie LJ* 43 (2020), 767.

[52] Esperanza L. Gómez-Durán and Carles Martin-Fumadó. 2024. Non-consensual Condom-Use Deception: An Empirically Based Conceptualization of Stealthing. *Trauma, Violence, & Abuse* 25, 1 (2024), 87–101. arXiv:https://doi.org/10.1177/15248380221141731 doi:10.1177/15248380221141731 PMID: 36565266.

[53] Jeffrey A. Hall, Namkee Park, Hayeon Song, and Michael J. Cody. 2010. Strategic misrepresentation in online dating: The effects of gender, self-monitoring, and personality traits. *Journal of Social and Personal Relationships* 27, 1 (2010), 117–135. doi:10.1177/0265407509349633

[54] Jeffrey T Hancock, Mor Naaman, and Karen Levy. 2020. AI-Mediated Communication: Definition, Research Agenda, and Ethical Considerations. *Journal of Computer-Mediated Communication* 25, 1 (March 2020), 89–100. doi:10.1093/jcmc/zmz022

[55] Jeffrey T. Hancock, Catalina Toma, and Nicole Ellison. 2007. The truth about lying in online dating profiles. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 449–452. doi:10.1145/1240624.1240697

[56] Jeffrey T. Hancock and Catalina L. Toma. 2009. Putting your best face forward: The accuracy of online dating photographs. *Journal of Communication* 59, 2 (2009), 367–386. doi:10.1111/j.1460-2466.2009.01420.x

[57] Jessica He, Stephanie Houde, and Justin D. Weisz. 2025. Which Contributions Deserve Credit? Perceptions of Attribution in Human-AI Co-Creation. arXiv:2502.18357 [cs.HC] https://arxiv.org/abs/2502.18357

[58] Jess Hohenstein and Malte Jung. 2020. AI as a moral crumple zone: The effects of AI-mediated communication on attribution and trust. *Computers in Human Behavior* 106 (May 2020), 106190. doi:10.1016/j.chb.2019.106190

[59] Connor Homayouni and Douglas Zytko. 2023. Consensual XR: A Consent-Based Design Framework for Mitigating Harassment and Harm Against Marginalized Users in Social VR and AR. In *2023 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct)*. 360–364. doi:10.1109/ISMAR-Adjunct60411.2023.00077

[60] Jane Im, Jill Dimond, Melody Berton, Una Lee, Katherine Mustelier, Mark S. Ackerman, and Eric Gilbert. 2021. Yes: Affirmative Consent as a Theoretical Framework for Understanding and Imagining Social Platforms. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–18. doi:10.1145/3411764.3445778

[61] Maurice Jakesch, Megan French, Xiao Ma, Jeffrey T. Hancock, and Mor Naaman. 2019. AI-Mediated Communication: How the Perception that Profile Text was Written by AI Affects Trustworthiness. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3290605.3300469

[62] Kowe Kadoma, Marianne Aubin Le Quere, Xiyu Jenny Fu, Christin Munsch, Danaë Metaxa, and Mor Naaman. 2024. The Role of Inclusion, Control, and Ownership in Workplace AI-Mediated Communication. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 1016, 10 pages. doi:10.1145/3613904.3642650

[63] Donal Khosrowi, Finola Finn, and Elinor Clark. 2023. Diffusing the Creator: Attributing Credit for Generative AI Outputs. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (Montréal, QC, Canada) *(AIES '23)*. Association for Computing Machinery, New York, NY, USA, 890–900. doi:10.1145/3600211.3604716

[64] Seunghyun Kim, Afsaneh Razi, Gianluca Stringhini, Pamela J. Wisniewski, and Munmun De Choudhury. 2021. A Human-Centered Systematic Literature Review of Cyberbullying Detection Algorithms. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 325 (Oct. 2021), 34 pages. doi:10.1145/3476066

[65] Lin Kyi. 2024. Reimagining Online Consent for More Responsible and Human-Centred Data Collection. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, Article 453, 3 pages. doi:10.1145/3613905.3651129

[66] Lin Kyi, Amruta Mahuli, M. Six Silberman, Reuben Binns, Jun Zhao, and Asia J. Biega. 2025. Governance of Generative AI in Creative Work: Consent, Credit, Compensation, and Beyond. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 197, 16 pages. doi:10.1145/3706598.3713799

[67] Lin Kyi, Abraham Mhaidli, Cristiana Teixeira Santos, Franziska Roesner, and Asia J. Biega. 2024. "It doesn't tell me anything about how my data is used": User Perceptions of Data Collection Purposes. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 984, 12 pages. doi:10.1145/3613904.3642260

[68] Una Lee and Dann Toliver. 2017. Building Consentful Tech. Zine/primer, Consentful Tech Project.

[69] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (Vancouver, BC, Canada) *(NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, Article 793, 16 pages.

[70] Fan Li and Ya Yang. 2024. Impact of Artificial Intelligence–Generated Content Labels On Perceived Accuracy, Message Credibility, and Sharing Intentions for Misinformation: Web-Based, Randomized, Controlled Experiment. *JMIR Form Res* 8 (24 Dec 2024), e60024. doi:10.2196/60024

[71] Jiahao Li, Yang Chu, and Jie Xu. 2023. Impression transference from AI to human: The impact of AI's fairness on interpersonal perception in AI-Mediated communication. *International Journal of Human-Computer Studies* 179 (2023), 103119. doi:10.1016/j.ijhcs.2023.103119

[72] Bingjie Liu, Jin Kang, and Lewen Wei. 2024. Artificial intelligence and perceived effort in relationship maintenance: Effects on relationship satisfaction and uncertainty. *Journal of Social and Personal Relationships* 41, 5 (2024), 1232–1252. arXiv:https://doi.org/10.1177/02654075231189899 doi:10.1177/02654075231189899

[73] Juniper L. Lovato, Antoine Allard, Randall Harp, Jeremiah Onaolapo, and Laurent Hébert-Dufresne. 2022. Limits of Individual Consent and Models of Distributed Consent in Online Social Networks. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) *(FAccT '22)*. Association for Computing Machinery, New York, NY, USA, 2251–2262. doi:10.1145/3531146.3534640

[74] Ewa Luger and Tom Rodden. 2013. An Informed View on Consent for UbiComp. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. Association for Computing Machinery, New York, NY, USA, 529–538. doi:10.1145/2493432.2493446

[75] Atil Mansour, Ofra Amir, and Liat Levontin. 2024. The Potentially Negative Effects of AI Writing Assistants on Self-Disclosure. In *Proceedings of the Third Workshop on Intelligent and Interactive Writing Assistants* (Honolulu, HI, USA) *(In2Writing '24)*. Association for Computing Machinery, New York, NY, USA, 1–3. doi:10.1145/3690712.3690713

[76] Christina Masden and W. Keith Edwards. [n. d.]. Understanding the Role of Community in Online Dating. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA. doi:10.1145/2702123.2702417

[77] Celestin Matte, Nataliia Bielova, and Cristiana Santos. 2020. Do Cookie Banners Respect my Choice? : Measuring Legal Compliance of Banners from IAB Europe's Transparency and Consent Framework. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, USA, 791–809. doi:10.1109/SP40000.2020.00076

[78] Jingbo Meng, Renwen Zhang, Jiaqi Qin, Yu-Jen Lee, and Yi-Chieh Lee. 2025. AI-mediated social support: the prospect of human–AI collaboration. *Journal of Computer-Mediated Communication* 30, 4 (July 2025), zmaf013. doi:10.1093/jcmc/zmaf013

[79] Hannah Mieczkowski, Jeffrey T. Hancock, Mor Naaman, Malte Jung, and Jess Hohenstein. 2021. AI-Mediated Communication: Language Use and Interpersonal Effects in a Referential Communication Task. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 17 (April 2021), 14 pages. doi:10.1145/3449091

[80] Miro. [n. d.]. Miro: Online Whiteboard for Visual Collaboration. https://miro.com/.

[81] Alistair Morrison, Donald McMillan, and Matthew Chalmers. 2014. Improving consent in large scale mobile HCI through personalised representations of data. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational* (Helsinki, Finland) *(NordiCHI '14)*. Association for Computing Machinery, New York, NY, USA, 471–480. doi:10.1145/2639189.2639239

[82] Charlene L. Muehlenhard, Terry P. Humphreys, Kristen N. Jozkowski, and Zoë D. Peterson. 2016. The Complexities of Sexual Consent Among College Students: A Conceptual and Empirical Review. 53, 4 (2016), 457–487. doi:10.1080/00224499.2016.1146651

[83] neon2048. 2024. c2pastrip: Simple Tool to Remove C2PA Manifests from Files. https://github.com/neon2048/c2pastrip.

[84] Josef Nguyen and Bonnie Ruberg. 2020. Challenges of Designing Consent: Consent Mechanics in Video Games as Models for Interactive User Agency. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3313831.3376827

[85] Hyanghee Park and Daehwan Ahn. 2024. The Promise and Peril of ChatGPT in Higher Education: Opportunities, Challenges, and Design Implications. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 271, 21 pages. doi:10.1145/3613904.3642785

[86] Seora Park, Hajin Lim, and Joonhwan Lee. 2024. "Some Hope, Many Despair": Experiences of the Normalization within Online Dating among Queer Women in a Closeted Society. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 869, 15 pages. doi:10.1145/3613904.3642893

[87] Jessica Pater, Brandon T. McDaniel, Fayika Farhat Nova, Michelle Drouin, Kimberly O'Connor, and Douglas Zytko. 2025. A commentary on sexting, sextortion, and generative AI: Risks, deception, and digital vulnerability. *Family Relations* 74, 3 (2025), 1109–1120. doi:10.1111/fare.13152 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/fare.13152.

[88] Planned Parenthood. [n. d.]. Sexual Consent. https://www.plannedparenthood.org/learn/relationships/sexual-consent

[89] John R. Porter, Kiley Sobel, Sarah E. Fox, Cynthia L. Bennett, and Julie A. Kientz. 2017. Filtered out: Disability disclosure practices in online dating communities. *Proceedings of the ACM on Human-Computer Interaction* 1, CSCW (2017), 87.

[90] Zoe A. Purcell, Mengchen Dong, Anne-Marie Nussberger, Nils Köbis, and Maurice Jakesch. 2024. People have different expectations for their own versus others' use of AI-mediated communication tools. *British Journal of Psychology* (Sept. 2024). doi:10.1111/bjop.12727

[91] Lucy Qin, Vaughn Hamilton, Sharon Wang, Yigit Aydinalp, Marin Scarlett, and Elissa M. Redmiles. 2024. "Did they F***ing consent to that?": safer digital intimacy via proactive protection against image-based sexual abuse. In *Proceedings of the 33rd USENIX Conference on Security Symposium* (Philadelphia, PA, USA) *(SEC '24)*. USENIX Association, USA, Article 4, 18 pages.

[92] Rape, Abuse & Incest National Network. 2025. Get the Facts About Sexual Assault & Rape. https://rainn.org/get-the-facts-about-sexual-assault-rape.

[93] Afsaneh Razi, Seunghyun Kim, Ashwaq Alsoubai, Gianluca Stringhini, Thamar Solorio, Munmun De Choudhury, and Pamela J. Wisniewski. 2021. A Human-Centered Systematic Literature Review of the Computational Approaches for Online Sexual Risk Detection. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 465 (Oct. 2021), 38 pages. doi:10.1145/3479609

[94] Florian Schaub, Rebecca Balebako, Adam L. Durity, and Lorrie Faith Cranor. 2015. A Design Space for Effective Privacy Notices. In *Proceedings of the Eleventh Symposium on Usable Privacy and Security (SOUPS 2015)*. 1–17.

[95] Kelsea Schulenberg, Lingyuan Li, Caitlin Lancaster, Douglas Zytko, and Guo Freeman. 2023. "We Don't Want a Bird Cage, We Want Guardrails": Understanding & Designing for Preventing Interpersonal Harm in Social VR through the Lens of Consent. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2 (Oct. 2023), 323:1–323:30. doi:10.1145/3610172

[96] Chenxinran Shen, Yan Xu, Ray Lc, and Zhicong Lu. 2024. Seeking Soulmate via Voice: Understanding Promises and Challenges of Online Synchronized Voice-Based Mobile Dating. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 921, 14 pages. doi:10.1145/3613904.3642860

[97] Mohamed R. Shoaib, Zefan Wang, Milad Taleby Ahvanooey, and Jun Zhao. 2023. Deepfakes, Misinformation, and Disinformation in the Era of Frontier AI, Generative AI, and Large AI Models. arXiv:2311.17394 [cs.CR] https://arxiv.org/abs/2311.17394

[98] Alejo José G. Sison, Marco Tulio Daza, Roberto Gozalo-Brizuela, and Eduardo C. Garrido-Merchán. 2024. ChatGPT: More Than a "Weapon of Mass Deception" Ethical Challenges and Responses from the Human-Centered Artificial Intelligence (HCAI) Perspective. *International Journal of Human–Computer Interaction* 40, 17 (Sept. 2024), 4853–4872. doi:10.1080/10447318.2023.2225931 Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/10447318.2023.2225931.

[99] Robert Sparrow and Lauren Karas. 2020. Teledildonics and rape by deception. *Law, Innovation and Technology* 12, 1 (2020), 175–204. arXiv:https://doi.org/10.1080/17579961.2020.1727097 doi:10.1080/17579961.2020.1727097

[100] Yolande Strengers, Jathan Sadowski, Zhuying Li, Anna Shimshak, and Florian 'Floyd' Mueller. 2021. What Can HCI Learn from Sexual Consent? A Feminist Process of Embodied Consent for Interactions with Emerging Technologies. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (CHI '21)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3411764.3445107

[101] Sindy R. Sumter, Laura Vandenbosch, and Loes Ligtenberg. 2017. Love me Tinder: Untangling emerging adults' motivations for using the dating application Tinder. *Telematics and Informatics* 34, 1 (2017), 67–78. doi:10.1016/j.tele.2016.04.009

[102] TikTok. 2025. TikTok. https://www.tiktok.com/.

[103] Catalina L. Toma and Jeffrey T. Hancock. 2010. Looks and lies: The role of physical attractiveness in online dating self-presentation and deception. *Communication Research* 37, 3 (2010), 335–351. doi:10.1177/0093650209356437

[104] Catalina L. Toma and Jeffrey T. Hancock. 2010. Reading between the lines: Linguistic cues to deception in online dating profiles. In *Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work*. 5–8.

[105] Catalina L. Toma, Jeffrey T. Hancock, and Nicole B. Ellison. 2008. Separating Fact From Fiction: An Examination of Deceptive Self-Presentation in Online Dating Profiles. *Personality and Social Psychology Bulletin* 34, 8 (2008), 1023–1036. doi:10.1177/0146167208318067

[106] Rebecca Umbach, Nicola Henry, Gemma Faye Beard, and Colleen M. Berryessa. 2024. Non-Consensual Synthetic Intimate Imagery: Prevalence, Attitudes, and Knowledge in 10 Countries. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 779, 20 pages. doi:10.1145/3613904.3642382

[107] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. 2019. (Un)informed Consent: Studying GDPR Consent Notices in the Field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (London, United Kingdom) *(CCS '19)*. Association for Computing Machinery, New York, NY, USA, 973–990. doi:10.1145/3319535.3354212

[108] Julie L. Valentine, Leslie W. Miles, Kristen Mella Hamblin, and Aubrey Worthen Gibbons. 2023. Dating App Facilitated Sexual Assault: A Retrospective Review of Sexual Assault Medical Forensic Examination Charts. *Journal of Interpersonal Violence* 38, 9-10

(2023), 6298–6322. arXiv:https://doi.org/10.1177/08862605221130390 doi:10.1177/08862605221130390 PMID: 36310506.

[109] Emma Walquist, Isha Datey, Wenqi Zheng, Xiangyu Zhou, Kelly Berishaj, Melissa McDonald, Michele Parkhill, Dongxiao Zhu, and Douglas Zytko. 2025. Collective Consent: Who Needs to Consent to the Donation of Data Representing Multiple People? *Proceedings of the ACM on Human-Computer Interaction* (2025). Issue CSCW. doi:10.1145/3757567

[110] Mark Warner, Andreas Gutmann, M. Angela Sasse, and Ann Blandford. 2018. Privacy unraveling around explicit HIV status disclosure fields in the online geosocial hookup app Grindr. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–22.

[111] Mark Warner, Juan F. Maestre, Jo Gibbs, Chia-Fang Chung, and Ann Blandford. 2019. Signal Appropriation of Explicit HIV Status Disclosure Fields in Sex-Social Apps used by Gay and Bisexual Men. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–15.

[112] Miranda Wei, Christina Yeung, Franziska Roesner, and Tadayoshi Kohno. 2025. "We're utterly ill-prepared to deal with something like this": Teachers' Perspectives on Student Generation of Synthetic Nonconsensual Explicit Imagery. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3706598.3713226

[113] Justin D. Weisz, Shraddha Vijay Kumar, Michael Muller, Karen-Ellen Browne, Arielle Goldberg, Katrin Ellice Heintze, and Shagun Bajpai. 2025. Examining the Use and Impact of an AI Code Assistant on Developer Productivity and Experience in the Enterprise. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 673, 13 pages. doi:10.1145/3706599.3706670

[114] Monica T. Whitty. 2008. Revealing the "real" me, searching for the "actual" you: Presentations of self on an internet dating site. *Computers in Human Behavior* 24, 4 (2008), 1707–1723. doi:10.1016/j.chb.2007.07.002

[115] Maximiliane Windl, Petra Zsofia Laboda, and Sven Mayer. 2025. Designing Effective Consent Mechanisms for Spontaneous Interactions in Augmented Reality. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1225, 18 pages. doi:10.1145/3706598.3713519

[116] Chloe Wittenberg, Ziv Epstein, Gabriel Péloquin-Skulski, Adam J. Berinsky, and David G. Rand. 2025. Labeling AI-generated media online. *PNAS Nexus* 4, 6 (May 2025), pgaf170. doi:10.1093/pnasnexus/pgaf170

[117] JJ Wright, Alexandra M. Zidenberg, Stephanie Varsanyi, Kaelyn Macaulay, and Sara Sanabria. 2025. The Murky Circumstances of Consent in Canada: A Commentary on R. v Kirkpatrick. *Canadian Journal of Women and the Law* 36, 1 (2025), 48–61. arXiv:https://doi.org/10.3138/cjwl-04-x doi:10.3138/cjwl-04-x

[118] Yihan Wu and Ryan M. Kelly. 2021. Online Dating Meets Artificial Intelligence: How the Perception of Algorithmically Generated Profile Text Impacts Attractiveness and Trust. In *Proceedings of the 32nd Australian Conference on Human-Computer Interaction (OzCHI '20)*. Association for Computing Machinery, New York, NY, USA, 444–453. doi:10.1145/3441000.3441074

[119] Ziang Xiao, Tiffany Wenting Li, Karrie Karahalios, and Hari Sundaram. 2023. Inform the Uninformed: Improving Online Informed Consent Reading with an AI-Powered Chatbot. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 112, 17 pages. doi:10.1145/3544548.3581252

[120] Fan Zhang, Yun Chen, Xiaoke Zeng, Tianqi Wang, Long Ling, and Ray Lc. 2025. "An Image of Ourselves in Our Minds": How College-educated Online Dating Users Construct Profiles for Effective Self Presentation. *Proc. ACM Hum.-Comput. Interact.* 9, 2, Article CSCW031 (May 2025), 30 pages. doi:10.1145/3710929

[121] Wenqi Zheng, Emma Walquist, Isha Datey, Xiangyu Zhou, Kelly Berishaj, Melissa McDonald, Michele Parkhill, Dongxiao Zhu, and Douglas Zytko. 2024. "It's Not What We Were Trying to Get At, but I Think Maybe It Should Be": Learning How to Do Trauma-Informed Design with a Data Donation Platform for Online Dating Sexual Violence. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 743, 15 pages. doi:10.1145/3613904.3642045

[122] Douglas Zytko and Jonathan Chan. 2023. The Dating Metaverse: Why We Need to Design for Consent in Social VR. *IEEE Transactions on Visualization and Computer Graphics* 29, 5 (2023), 2489–2498. doi:10.1109/TVCG.2023.3247065

[123] Douglas Zytko and Nicholas Furlo. 2023. Online Dating as Context to Design Sexual Consent Technology with Women and LGBTQ+ Stakeholders. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 339, 17 pages. doi:10.1145/3544548.3580911

[124] Douglas Zytko, Nicholas Furlo, Bailey Carlin, and Matthew Archer. 2021. Computer-Mediated Consent to Sex: The Context of Tinder. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 189 (April 2021), 26 pages. doi:10.1145/3449288

[125] Douglas Zytko, Sukeshini A. Grandhi, and Quentin Jones. 2014. Impression management struggles in online dating. In *Proceedings of the 18th International Conference on Supporting Group Work*. 53–62. doi:10.1145/2660398.2660410

[126] Douglas Zytko, Sukeshini A. Grandhi, and Quentin Jones. 2015. Frustrations with Pursuing Casual Encounters through Online Dating. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems*. 1935–1940.

[127] Doug Zytko, Sukeshini A Grandhi, and Quentin Jones. 2016. The coaches said...what?: Analysis of online dating strategies recommended by dating coaches. In *Proceedings of the 19th International Conference on Supporting Group Work*. 385–397. doi:10.1145/2957276.2957287

[128] Douglas Zytko, Sukeshini A. Grandhi, and Quentin Jones. 2016. Online dating coaches' user evaluation strategies. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. doi:10.1145/2851581.2892482

[129] Douglas Zytko, Jane Im, and Jonathan Zong. 2022. Consent: A Research and Design Lens for Human-Computer Interaction. In *Companion Publication of the 2022 Conference on Computer Supported Cooperative Work and Social Computing* (Virtual Event, Taiwan) *(CSCW'22 Companion)*. Association for Computing Machinery, New York, NY, USA, 205–208. doi:10.1145/3500868.3561201

[130] Douglas Zytko, Nicholas Mullins, Shelnesha Taylor, and Richard H. Holler. 2022. Dating Apps Are Used for More Than Dating: How Users Disclose and Detect (Non-)Sexual Interest in People-Nearby Applications. *Proc. ACM Hum.-Comput. Interact.* 6, GROUP (Jan. 2022), 30:1–30:14. doi:10.1145/3492849

[131] Douglas Zytko, Victor Regalado, Nicholas Furlo, Sukeshini A. Grandhi, and Quentin Jones. 2020. Supporting Women in Online Dating with a Messaging Interface that Improves their Face-to-Face Meeting Decisions. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2, Article 137 (Oct. 2020), 30 pages. doi:10.1145/3415208

[132] Douglas Zytko, Devin Tebbe, Braeden Burger, Emma Walquist, and Meryem Barkallah. 2025. Beyond Affirmative Consent: Creating Alternative Models for Consent to Computer-Mediated Sexual Activity Through Social VR for Online Dating. *Proceedings of the ACM on Human-Computer Interaction* 9, CSCW, Article 377 (Nov. 2025), 34 pages. doi:10.1145/3757558

8 Appendix

Table 3: Participant demographics (N = 113). Values are counts (n) and percentages of the sample within each category.

Category	n	%
Age (years)		
18–24	24	21.2
25–34	49	43.4
35–44	27	23.9
45+	13	11.5
Gender		
Man	61	54.0
Woman	52	46.0
Transgender		
Yes	3	2.7
No	110	97.3
Country of residence		
United States	40	35.4
South Africa	32	28.3
United Kingdom	14	12.4
Kenya	9	8.0
Other	18	15.9
Ethnicity		
Black or African American	60	53.1
White	38	33.6
Asian	7	6.2
Hispanic or Latino	5	4.4
Other / multiracial	3	2.7
Sexual orientation		
Heterosexual	87	77.0
Bisexual	16	14.2
Homosexual	4	3.5
Pansexual	3	2.7
Other (e.g., asexual, aromantic, demisexual)	3	2.7