Catfished! Impacts of Strategic Misrepresentation in Online Dating

Oz Kilic
Carleton University
Ottawa, ON, Canada
OzKilic@cmail.carleton.ca

Alan Tsang
Carleton University
Ottawa, ON, Canada
Alan.Tsang@carleton.ca

ABSTRACT

Online dating is a multibillion-dollar global industry, and is increasingly becoming the go-to method for finding partners. Intricate dynamics mark its operation, influenced by varying user preferences, strategies, and traits, as well as by the underlying matchmaking algorithm. This complexity renders it a pertinent subject for multiagent systems research. Despite its relevance, an established simulation framework for online dating is lacking. This paper introduces a multiagent simulation framework for this domain. The framework is extensible and capable of modeling agents with diverse attributes and preferences, either reported or latent. It also supports varied strategies, outcomes, and types of matchmaking logic. Using this framework, we simulate an online dating platform based on real-world demographics to examine the effects of strategic misrepresentation, a notable concern in online dating. Surprisingly, the negative effect of strategic misrepresentation on users is marginal. Moreover, it disproportionately benefits female or honest agents more, enhances the overall welfare of the user population, and benefits attractive users — whether deceitful or not — over less attractive ones.

KEYWORDS

Online Dating; Agent-based Simulation; Preference; Truthfulness; Social Computing

ACM Reference Format:
Oz Kilic and Alan Tsang. 2024. Catfished! Impacts of Strategic Misrepresentation in Online Dating. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 8 pages.

1 INTRODUCTION

The online dating (OD) industry has generated $5.61 billion revenue in 2021 [34], increasing its popularity and slowly becoming the norm. The penetration rate of OD is especially high among young adults at 53% [49]. Currently, most heterosexual couples in the United States meet online and the stigma has waned [37].

This study proposes a multiagent OD simulation framework based on Tinder (currently the most popular OD platform in the West [1]), where agents with different attributes and preferences have limited like allowances for each day (round), and use their allowance to get matches that yield utility/happiness. Using our framework, we aimed to answer "How does strategic misrepresentation affect people in online dating?"

The contributions of this paper are as follows:

- We present the first multiagent online dating simulation framework in academic literature. Our open-source and extensible framework [30] can be used to simulate OD, professional networking, or other matchmaking problems where agents/entities can have varying attributes and strategies.
- This is the first multiagent simulation of OD. Existing studies either focus on capturing and analyzing user preferences in a laboratory setting or quantify misrepresentation without analyzing its effects. Through simulating user behavior and misrepresentation at a system level, we analyze how misrepresentation affects the whole population or specific groups.

2 BACKGROUND

When users sign up in modern OD platforms, they often fill out their profile with personal attributes (ex: gender, height, weight, ethnicity, hobbies, pictures) and preferences (ex: sexual orientation, preferences over personal attributes, and priority of those preferences). The platform uses this information to generate recommendations for the user. These recommendations are presented as a queue of profiles where users must "like" or "pass" a presented profile[1] before viewing the next one. If two users "like" each other, the platform declares a match and allows users to message each other. Crucially, this is not the end of the process; the platform continues to recommend other profiles and produce additional matches, giving the user the freedom to follow up with the match and/or continue swiping for additional matches[2]. To lure people, sometimes OD users create made-up profiles or strategically misrepresent themselves by using outdated/unclear pictures or providing incorrect information. This deception is called "catfishing."

In mathematical and social choice literature, OD is commonly modeled as a two-sided matching game [38] where agents have preferences over other agents, and a mechanism pairs off agents according to some solution concept. Two-sided matching can be solved using the Gale-Shapley algorithm [16] which also removes the incentive for one side to misreport their preferences. While applicable to many real-world problems, matching is not as suitable for OD: the function of OD platforms is to recommend *many*

[1] This interaction is colloquially called "swiping" by Tinder and became a common word in OD in general.
[2] Some platforms offer variations on this format – for example, some display the profiles of nearby users, or allow messaging of users not in the recommendation queue. Platforms may also relax restrictions when users subscribe to a higher tier of service. We focus on the "restricted queue" format as it is most common in modern platforms.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand.

matches to users, who have the final say on an offline meeting; indeed most matches do not result in offline meetings [43].

In this sense, OD is conceptually closer to a repeated hedonic game [26] with non-transferable utility [2] in a partially observable environment where the bidders are also the items. Bidders have valuations over each other and submit a bid ("like") for the recipient's consideration. The mechanism is online from the perspective of the users: Each user must "like" or "pass" on a recommendation from a queue without knowledge of subsequent items. A match is only revealed when both users like each other. However, a successful match may not necessarily result in the termination of the game, nor does not block the users from being matched to others. These practical considerations make the OD game an ideal topic for agent-based investigations.

3 MODEL

The simulation models an OD platform where agents receive recommendations each day (round) and they have a limited like allowance for each round. Each agent can like or pass an agent recommended by the system according to their strategy and reported attributes of the candidates. Two agents that like each other are matched, and derive utility from this match (see Section 3.5).

3.1 Agents

Each agent mainly has an ID, attractiveness, attributes and preferences, like allowance per round, and strategy. Agents are notified when a new round starts or they are matched with another agent.

3.1.1 Attractiveness. OkCupid data [39] suggests men have an overall evaluation that is almost normally distributed for women, with a mode close to 3 within a range of $[1, 5]$, while women found most men unattractive, with a mode close to 2. Following this, the attractiveness of an agent is a decimal value between 1 and 5, reflecting their average perceived physical attractiveness according to conventional beauty standards (5 being the most desirable). For simplicity, we assume an agent's attractiveness is perceived the same by all other agents, and agents can see this attractiveness value of their candidates. Based on Rudder [39], we sampled agent attractiveness continuously, using separate beta distributions for males ($\alpha = 2, \beta = 6$) and females ($\alpha = 4, \beta = 4$), scaled to 1–5, where females have a more normal distribution while males have a right-skewed distribution (the average male has below-3 attractiveness). As evidenced by [31, 39, 48], men rate women higher than women rate men, and women are far more likely to get overwhelming amounts of attention compared to men.

Supported by Greitemeyer [19], agents can only estimate their own attractiveness with a margin of error, which varies with their actual attractiveness. Based on the deviation maps presented by Greitemeyer [19], two linear functions were created to represent the upper and lower bounds of one's estimation. Within the bounds of these functions, for a given actual attractiveness g between $[1, 5]$, the estimated attractiveness g' is sampled from a truncated normal distribution that has a mean corresponding to the midpoint between these functions shown in Equations 1 and 2:

$$upperbound(g'|g) = 3g/8 + 25/8 \tag{1}$$

$$lowerbound(g'|g) = 3g/4 + 1/4 \tag{2}$$

These equations capture Greitemeyer's findings [19] that very low attractiveness agents tend to overestimate their attractiveness with greater margins compared to more attractive agents, while very attractive agents underestimate their attractiveness with a smaller margin. The mean of the distribution meets the actual attractiveness at around 3.857, which generalizes the plots from [19]. Figure 1 shows our sampling space of estimated attractiveness for a given attractiveness. This estimation is an important factor for agent strategy, as high self-esteem leads to more pickiness.

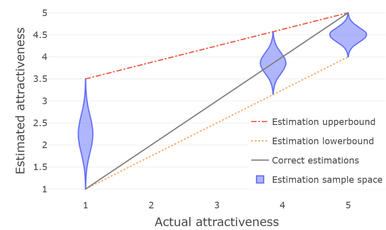

Figure 1: Estimation bounds and sampling spaces for attractiveness values of 1, 3.857, and 5. The self-estimate of attractiveness is unbiased when attractiveness is ~3.857.

3.2 Attributes and preferences

Agents have attributes that define their qualities (for example, their physique, personality, beliefs, and so on). However, not all attributes are necessarily reported. Reported attributes are the attributes shown to other agents and the matchmaker, used to decide if two agents may be compatible. Each attribute has a preference[3] that determines the attribute-specific compatibility.

Some attributes may not be reported or may be intentionally misreported, as frequently occurs in OD [7, 18, 21, 25]. When an attribute is not reported, it is not used by the mechanism or other agents for decision-making. However, once matched, these attributes will be revealed in subsequent interactions, and so are incorporated into the computation of utility that agents derive from the match. Through the attribute mechanism, agents can also lie about their preferences, as is often the case as well in OD [39]. Similarly, the system only uses their reported preferences for matchmaking, while the agent uses their actual preferences for candidate evaluation. While the framework allows both attribute misrepresentation and preference misrepresentation, this study focuses on the former where "deceitful" agents strategically misrepresent their attributes, while "honest" agents correctly disclose their attributes. The details of strategic misrepresentation are explained in Section 4.

We implemented four reported attributes for compatibility: gender, age, height, and weight. Our framework allows varying weights for each agent's preferences. A weighted average of attribute-specific

[3]Preferences can have a continuous range or set of discrete values. Scores for preferred and non-preferred values can be specified. For the preference ranges, it is possible to also pass a custom scoring function that maps the difference to a score using the difference and allowed value range. For a value that is outside the preferred range, its compatibility loss is calculated using the closest value from the preferred range and the custom function. Therefore, an agent can have binary or continuous attribute-specific compatibility scores with other agents.

compatibilities is used to calculate the overall compatibility between two agents, which can be used by the system for matchmaking purposes or by agents for decision-making, similar to OkCupid [39]. The framework allows misreporting preferences and preference weights, though we do not focus on them in this study.

3.2.1 Gender. We implemented gender as an attribute and sexual orientation as its preference. Due to simplicity and the referred studies in the literature having a heterosexual focus, this study is limited to two genders (male and female) and heterosexual agents. Therefore, in our study, a male agent is interested in female agents and vice versa. However, the framework allows creating non-binary gender identities and an orientation spectrum through its attribution and preference mechanism, and the framework is extensible. We sampled genders using a 72:28 ratio of males to females [41].

3.2.2 Age. We sampled agent ages using the 2022 BRFSS Survey [15] conducted in the US with over 445,000 respondents. We used ordinal age values corresponding to the age brackets provided by the survey dataset[4] and normalized their frequencies based on a study that analyzed OD prevalence in the U.S. in 2022[5] [49]. To sample age preferences, we used the 2017 U.S. Current Population Survey's opposite-sex married couples table [5]. Roughly mapping the age difference frequencies of all opposite-sex married couples to our case with ordinal age bracket values, for each agent, we first sampled an age difference, obtaining one extremity of the preferred range. Using an allowed difference of 3 (roughly corresponding to 15 years) from the sampled difference, we set the other extremity. If the sampled age difference was 0, we set the age preference as $[bracket-1, bracket+1]$ (roughly corresponding to 15 years again). We clamped the age preference range using (1) the minimum and maximum possible age values and (2) the agent's own age so that an agent who prefers to be the older party in a relationship would not seek someone older than themselves or the other way around. Apart from fostering matchmaking, we gave the agents some flexibility because characteristics in marriage data may be too strict for OD preferences. In line with existing studies [9, 22, 28], we ended up with populations in which females on average prefer slightly older males, while the reverse is true for males. We set a compatibility weight of 1 for age for both genders. Age-wise compatibility was continuously calculated based on the difference between the closest preferred age and the candidate's age, normalized between 0.25 and 1. For candidates that matched the preferred range, the compatibility score of 1.25 was used.

3.2.3 Height. We conditionally sampled agent heights (in cm) based on their sampled gender and age values, using the 2022 BRFSS Survey [15] again. We used kernel density estimation to handle the conditional height gaps in the data. According to Yancey and Emerson [52], females care more than males about height. 48.9% of the female respondents disclosed wanting to only date men who are taller than them, mostly explaining their preference through societal expectations and gender stereotypes. Meanwhile, 13.5% of the male respondents disclosed wanting to date only women who are shorter than them. Hitsch et al. [25] also suggest, relative to themselves, men prefer shorter women and women prefer taller men. Loosely following this narrative, we set the preferred range to be $[height - 30, height + 10]$ for males and $[height + 5, height + 50]$ for females. We used a continuous compatibility calculation, as with age. However, we used a compatibility weight of 1.5 for females and 1 for males, reflecting the idea that females care more than males.

3.2.4 Weight. As with height, we conditionally sampled agent weights based on their gender and age values from the 2022 BRFSS Survey [15]. Since weight is considered with respect to height for attractiveness, our attribute has ordinal values that correspond to the four body mass index (BMI) groups (1: Underweight, 2: Normal weight, 3: Overweight, 4: Obesity). Studies suggest that men prefer thinner partners [6, 25, 35]. Therefore, our male agents have a preference range of $[max(1, BMI - 1), BMI]$. Hitsch et al. [25] also suggest women prefer men who are bulkier than themselves and do not prefer underweight men. Hatoum and Belle [23] state that self-image and dating issues of underweight men and overweight women are very similar. However, Pawlowski and Koziel [35] state that weight was not an influential factor for men to get a higher response rate to their personal advertisement, while it was a factor for women. Fallon and Rozin [14] suggest men overestimate the body weight amount found ideal by women and women underestimate the body weight amount found ideal by men (which may be the reason they tend to understate their weight across the spectrum [46]). Considering these findings, we applied a more homophilic yet tolerating preference range for females. While underweight females have a range of $[1, 2]$, other females prefer $[max(2, BMI-1), min(BMI+1, 4)]$. We used the same compatibility score calculation for weight, but we assigned a compatibility weight (importance) of 1.5 for males and 1 for females, indicating weight compatibility is more important for males. Weight preferences are also shown as matrices in Figure 2.

Figure 2: Preferred (checked) and non-preferred (crossed) BMI groups for each BMI group and gender combination.

3.3 Matchmaking

The act of recommending one agent to the other is called "matchmaking," and the logic that decides on them is called a "matchmaker" or "matcher." For simplicity, when an agent likes or passes a recommended agent, they never encounter the same agent again, and matchmaking continues. After all likes and passes are processed by the matchmaker for the current round, it checks for new matches and informs the matched agents. Different matchmakers can use various criteria such as preferences, attractiveness, ranking, or a combination of them. While we created different matchmakers and the framework allows adding new ones, we focused on a preference-based matcher for this study.

[4] 1: 18–24, 2: 25–29, 3: 30–34, 4: 35–39, 5: 40–44, 6: 45–49, 7: 50–54, 8: 55–59, 9: 60–64, 10: 65–69, 11: 70–74, 12: 75–79, 13: 80+

[5] Both had compatible breakpoints for ages: [15] subdivides the age brackets from [49]. We applied the normalization from the larger age bracket across subdivided brackets.

Preferential matchmaker prioritizes agents based on the compatibility of reported attributes and preferences. Compatibilities can be calculated from the judging agent's perspective, the judged agent's perspective, or a weighted average of both. We used a weight of 0.99 for matchmaking, which almost entirely considers the evaluating agent's perspective for recommendations while the 0.01 weight ensures candidates who would never like them back due to a deal-breaker (having a negative attribute compatibility) are never recommended.

3.4 Strategy

While the framework allows having different strategies, we used the same threshold-based strategy for all agents. When the matchmaker presents a candidate to an agent, they calculate the overall compatibility using the candidate's reported attributes and preferences from their perspective. Then, they multiply the compatibility with the candidate's attractiveness and compare this value with their estimated attractiveness. If the value is greater than or equal to their estimated attractiveness, they like the candidate. This allows a less attractive agent with superior compatibility to compete with more attractive but less compatible agents. However, the evaluating agent's estimated attractiveness can play a significant part. An agent with a low attractiveness but a high estimated attractiveness can pass moderately compatible and attractive candidates (which is less likely for the top candidates). The liking strategy is loosely based on the preference for "better mates" observed in the literature [4, 19, 24, 25, 31–33].

In real life, agents may stop using the platform due to frustration or satisfying their needs. Agents in this simulation never stop using the platform because agents represent specific populations with a set of characteristics and a new agent with similar characteristics can replace a leaving one. Also, the simulation is not run long enough for agents to deplete their candidate pool.

3.5 Happiness

Agents' happiness (utility) from a match depends on multiple factors. Since existing literature does not quantify this utility, we designed our happiness function qualitatively based on their findings. Studies suggest that both males and females prefer an attractive partner [4, 32, 33], and more attractive people tend to prefer more attractive dates [31]. Therefore, an important component of our happiness function is the matched candidate's attractiveness.

While Tinder previously used user ratings for matchmaking [44], many OD platforms also consider compatibility. For example, OkCupid was solely focusing on questionnaires with thousands of questions[6], allowing users to specify preferred answers to those questions, and how much the answer matters to them [39]. Hinge explains their "most compatible" feature pairs people with each other using the Gale-Shapley algorithm [3, 36], similar to how two-sided matchings are treated in the stable marriage problem. However, it is actually a combination of Gale-Shapley algorithm and machine learning, allowing the system to learn users' preferences and make recommendations based on both parties' likelihood to like each other. Grindr, a dating and hookup application for queer people, reportedly simply filters people based on preference filters and online status, and ranks them by distance to the user [27]. The application sometimes adds a little randomness to keep the results "fresh." Considering these approaches, the second component of our happiness function was compatibility.

Combining both components, the match happiness in our study is calculated by multiplying the matched agent's attractiveness by their average compatibility factor between $-\infty$ and 1.25, which can enhance or diminish one's attractiveness to the point of not being recommended at all. However, we also considered some other factors. McNulty et al. [33] adds that both males and females may be happier in a relationship when the female is more attractive, explaining "less attractive wives may be less satisfied and behave more negatively in response to their more attractive husbands, whereas more attractive wives should be more satisfied and behave less negatively in response to their less attractive husbands." Meanwhile, males' preference is affected less by their own attractiveness [31]. Garcia and Khersonsky [17] support this phenomenon and state that participants predicted relationships in which the male is more attractive would be less satisfying. Another factor is that while OD encourages "the numbers game" mentality [24], the effect of getting a new match is not the same for someone who has no matches and someone who has 100 matches (matches have diminishing returns as visualized in Appendix B). For these reasons, agent i's utility from a matched agent j is formulated in Equation 3 as

$$u_i(j) = \begin{cases} \frac{((g_j \cdot c_{i,j}+2)^{0.9}-1)}{d_{i,j}} \cdot 0.999^{m_i} & \text{if } Gender_i = ♀ \\ ((g_j \cdot c_{i,j}+2)^{0.9}-1) \cdot 0.999^{m_i} & \text{if } Gender_i \neq ♀ \end{cases} \tag{3}$$

where g_i represents i's attractiveness, $c_{i,j}$ represents the overall compatibility for i and j, m_i represents the number of total matches i has, and $d_{i,j}$ represents the discomfort divisor for i when i is a female (♀). As agents get more matches, the marginal difference of getting a new match decreases due to being multiplied by 0.999^{m_i}. Since the discomfort females feel when their partner is significantly more attractive is not quantified [31, 33], a small penalty is used, specified below.

When i is a female with an estimated attraction $g'_i \in [1,5]$, for a male j with attraction $g_j \in [1,5]$, this function is formulated in Equation 4 as

$$d_{j,i} = \begin{cases} 1 & \text{if } g'_i \geq g_j \\ \frac{5-g'_i}{5 \cdot g'_i} + \sqrt[100]{g_j - g'_i} & \text{if } g'_i < g_j \end{cases} \tag{4}$$

where the discomfort only occurs when the female *thinks* (as estimated attractiveness may not correspond to the collectively perceived one) she is less attractive than their match, and it is only noticeable when the difference is significant. Importantly, because the agents are not explicitly aware of this discomfort factor, it is not incorporated into their liking strategy.

4 EXPERIMENTS

Our research question was "How does strategic misrepresentation affect people in online dating?" While we expected misrepresentation to improve agent utility (in line with existing studies [18, 25]), its impact on different groups under different settings was unclear.

[6]Before its sale to Tinder's parent company, Match Group, which made the platform more Tinder-like .

To explore these, we created a population of 8000 agents[7], sampling their attractiveness, estimated attractiveness, attributes, and preferences as explained in Section 3. Then, we set the daily like allowance to 200, the daily maximum recommendation limit to 400, and ran the simulation for 20 days (rounds). Therefore, agents could potentially see everyone but only like half of the population. All agents truthfully reported their attributes and preferences.

Next, we analyzed the total happiness of all agents after 20 rounds and we retrieved $X\%$ of the population from the bottom X percentile in happiness. These underperforming agents were assigned to be "deceitful" while the rest were assigned to be "honest." We reran a counterfactual experiment where the deceitful agents misrepresented themselves, and we analyzed the difference between these scenarios. In the literature, there are different views about strategic misrepresentation in OD. Ellison et al. [11] suggest people do not frequently lie in OD, and they lie because they perceive themselves differently, or they lie about small things such as a few pounds they can lose in a short time. Also, intentions and the possibility of an in-person interaction can decrease misrepresentation. Two publications based on another study [21, 46] suggest misrepresentation is common and intentional. Heino et al. [24] suggest people not only lie about themselves but also expect others to lie, and mentally adjust reported values accordingly (for example, expecting a reportedly 180 cm person to be 175 cm). Cornwell and Lundgren [7] suggest the percentage of people who lie about themselves in OD (27.5%) is significantly higher than in real life (12.5%). Meanwhile, Toma et al. [46] report 81.3% of the participants lied in at least one attribute (mostly, height for men and weight for women).

Since existing research is not definitive on the misrepresentation rate on dating platforms, we used two different deceitfulness percentages (27.5% and 81.3%) that span the full range of possibilities from the literature [7, 46], resulting in 2200 (Counterfactual A) and 6504 (Counterfactual B) deceitful agents. Due to the skewness and ties, Counterfactual A samples the bottom 27.5% of the population. To increase the reliability of our results, we repeated our experiment five times, obtaining five populations for each scenario (all-honest, low deceit, high deceit), yielding 15 different results in total.

In accordance with gender-specific importance of height and weight, the findings of [21, 46] suggest women consistently understate their weight and men consistently overstate their height. For these reasons, we decided to make deceitful agents misrepresent their height or weight depending on their gender. To see its effect better, while we followed Hancock et al. [21]'s findings, we went with more significant deviations from the truth. It was stated that the average height deviation for men was 1.45 cm (SD = 2.06) and the average weight deviation for women was -3.85 kg (SD = 4.02). We increased the reported height of the deceitful males by 8 cm by rounding the observed extreme of 7.62 cm. We decreased the reported weight group of the deceitful females by 1[8]. Considering the maximum observed understatement difference was −9.25 kg and a BMI group has a weight range of roughly 15 kg, our weight deviation amount was also relatively extreme yet within possibility.

After calculating the difference between the base happiness levels and the deceitful scenarios (Counterfactual A and Counterfactual B), we examined the Spearman rank correlations between variables and the outcomes, and descriptively analyzed them. To see if strategic misrepresentation yielded significantly different results (and to whom), we used Wilcoxon signed-rank and Mann-Whitney U tests, along with calculating their effect sizes.

5 RESULTS

For simplicity, we analyzed the correlations and conducted statistical tests across five simulations at once, combining their populations and treating them as a single population of 40,000 agents[9] for each scenario: the base scenario (*Base*), Counterfactual A (*CA*), and Counterfactual B (*CB*). All reported in-text correlations and statistical results are significant with a p-value lower than .001 (unless otherwise specified). Some additional plots, tables, and our experiment parameters (explained in Sections 3 and 4) are provided as supplementary materials [30].

34.8% of males and 8.8% of females were found deceitful in *CA*, while 87.6% of males and 65% of females were found deceitful in *CB*. Agents' final happiness in the honest condition (*Base*) is called "baseline." The difference between their happiness in *Base* and an counterfactual scenario is called "happiness difference."

For continuous variables and age[10], descriptive statistics across repetitions, grouped by gender, are given in Table 1. An average male was 177.9 cm tall and ~43 years old, while an average female was 163.2 cm tall and ~45 years old. As expected, the average male attractiveness was much lower (2.268 compared to 2.987). The maximum number of likes a male could get was 2289, which was strikingly less than females (5793). The difference was even larger for medians (58 vs. 3474). For these reasons, despite their average relative pickiness (15.7% vs. 55.2%), females were far more likely to match with the agents they liked (27.1% vs. 8% on average). Ultimately, females had higher happiness (max: 2420.308, mean: 419.168, median: 302.461) compared to males (max: 2392.018, mean: 168.798, median: 3.501). However, the difference in happiness did not match the extreme difference in likes because of the non-linear utility function. Gender-specific weight distributions, along with grouped happiness levels, are given in Appendix A. Both genders had very few underweight agents, males had a more unimodal distribution where the mode was overweight, and females had an almost uniform distribution excluding the underweight.

The overall match count and happiness for both genders increased with deceitfulness. *CB* yielded ~9.6% more matches than *Base*, yielding 12 more matches to females and 4.7 more matches to males on average. Using the Wilcoxon signed-rank test with agent-wise paired outcomes, we found that both counterfactual scenarios had significantly more matches and happiness ($p = .013$ for *CA*) levels compared to *Base*. Using Kerby's simple difference formula[11]

[7]This population size strikes a balance between computational costs and sampling stability. Apart from extremely rare attribute combinations, increasing population size should have little effect on our results, as we did not observe major differences with a smaller population of 2000 agents.

[8]It should be noted that, due to deceitful females always understating their weight, and some of them already being underweight, some effectively did not misrepresent. This affected 2–3% of deceitful females in counterfactual scenarios. Since we assigned deceitfulness purely based on underperformance, we did not interfere.

[9]Combined and individual/averaged results are compared in Appendix H. Apart from happiness-related tests in Counterfactual A, test results and averaged correlations with a magnitude higher than 0.10–0.15 were consistent with the combined versions.

[10]Age brackets are treated as continuous due to having more than 10 brackets.

[11]Calculated by subtracting the unfavorable evidence ratio from the favorable evidence ratio. Even the one-tail hypotheses use only non-zero differences.

[29], the effect size of the match difference was 1.0 for *CA* (due to not having unfavorable evidence) and 0.97 for *CB* (indicating strong positive association). The effect size of happiness diff. was 0.03 for *CA* (weak) and 0.88 for *CB* (strong). The Mann-Whitney U tests showed that females had significantly greater match count improvements in both *CA* and *CB*, with effect sizes[12] of 0.01 and 0.35 (weak to moderate). Happiness improvements were similarly significantly in favor of them in both *CA* and *CB*, with effect sizes of 0.03 and 0.36. The deceitful had significantly better match and happiness ($p = 0.004$) difference compared to the honest in *CA* with the effect sizes of 0.015, but it was the reverse in *CB* with the effect sizes of 0.328 and 0.348. The Mann-Whitney U test also showed that deceitful females benefited significantly more than everyone else from the changes in match counts and happiness with weak to moderate effect sizes (both 0.09 for *CA*, 0.34 and 0.36 for *CB*).

A scatter plot of attractiveness, estimated attractiveness, and baseline is given in Appendix C. Among the agents in the top 10% in happiness (> 240.45), female attractiveness spanned 2.5–4.76, and male attractiveness spanned 2.82–5. However, males had a more spread est. attractiveness/attractiveness ratio distribution (0.83–1.14) compared to females (0.82–1.1). Females under an attractiveness of 3 fared better than their male counterparts. Female happiness in the 25$^{\text{th}}$ percentile was 38.6 compared to males' 0. Agents with lower estimated than actual attractiveness performed better. However, for males and females, attractiveness had a higher significantly positive Spearman rank correlation with baseline ($r = .9$ and $.82$) compared to est. attractiveness ($r = .56$ and $.45$). The correlation between attractiveness and happiness diff. in *CB* was considerably lower for females ($r = .17$) than males ($r = .66$). Correlations between the attributes, baseline, and happiness changes in counterfactual scenarios are shown in Figure 3 for each gender. Gender-behavior-specific versions of the correlations in *CB* are given in Appendix D. For males, we found positive but negligible correlations with baseline and height ($r = .02$) or weight ($r = .05$). In *CA*, the correlation between height and happiness diff. was 0.05. However, in *CB*, the correlations between happiness diff. and both height ($r = -.12$) and weight ($r = -.15$) became negative. Age had negligible negative correlations with baseline ($r = -.01$, $p = .03$) and happiness differences ($r = -.01$, $p = .02, .01$). As expected, females had their weight ($r = -.09$) and age ($r = -.13$) negatively correlated with baseline. Honest females had a weak negative correlation between happiness diff. and weight ($r = -.09$) in *CB*, while deceitful females had a positive moderate one ($r = .51$) due to misrepresenting their weight. For males, the correlation between the deceitful's happiness difference and the misrepresented attribute was equal to the correlation between the honest's happiness difference and attractiveness ($r = -.17$). For females, the correlation between the deceitful's misrepresented attributes and the happiness difference was equal to the correlation between the honest's est. attractiveness and happiness difference ($r = .51$).

Happiness differences by gender and behavior are given in Appendix F for *CA* and Table 2 for *CB*. In *CA*, 17% of agents saw a net negative effect, but it was very negligible (at most −0.758). In *CB*, only 3.5% of agents saw a net negative effect, all groups were 6.4 to 33.5 times more likely to be affected positively than negatively, and the mean difference for positively affected agents was 40.059. The mean difference for negatively affected agents was still small (−1.54). Deceitful females were the most positively affected group (max. = 449.451, mean = 50.061) while deceitful males were the most negatively affected one (min.[13] = −69.676, mean = 9.764). Grouped by gender, behavior, and change direction, deceitful males had the most negative change (−3.753) and deceitful females had the most positive change (68.217) on average. The happiness distributions for each decile in *Base* and how they differ in the counterfactual scenarios are shown in Appendix E. Deciles that had the most significant change were in the 0.6–0.8 range.

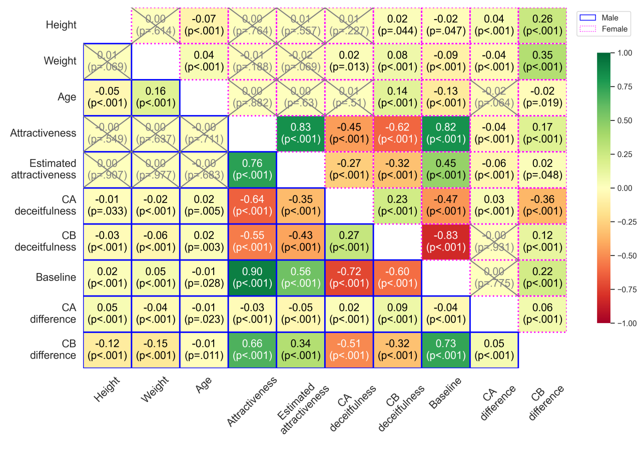

Figure 3: Spearman rank correlations between agent attributes/outcomes (insignificant ones are crossed out, males and females reported separately).

Heatmaps for happiness differences with respect to gender, behavior, height, and weight are given in Figure 4. In *CA*, the most obvious positive impact was on deceitful obese females, which was not surprising considering the importance of females' weight for men. The second group who were noticeably positively impacted, to a limited extent, was shorter deceitful males. *CB* had a clear difference between weight groups in honest males. While the normal-weight and overweight ones mostly saw an increase in happiness, the obese ones saw no difference, and there was almost no honest underweight male (they were mostly under the happiness threshold). Similar to *CA*, shorter deceitful males saw an overall improvement (due to closing the height gap through misrepresentation).

6 DISCUSSION

Here, we discuss the results of misrepresentation from different perspectives in Counterfactual A (*CA*) and Counterfactual B (*CB*).

Gender. In our study, gender played a distinct role: females were less common but more attractive, leading to higher match rates and greater satisfaction compared to males. Although most deceitful participants were male, females gained more in hypothetical

[12] For the Mann-Whitney U tests, effect sizes were calculated using the rank-biserial correlation formula [8].

[13] Focusing on the negatively affected ones, the mathematically minimum change corresponds to the most negative change.

Table 1: Descriptive statistics of the combined population over five repetitions, grouped by gender, for continuous variables

		Males					Females				
Scenario	Attribute/Outcome	Min.	Max.	Mean	SD	Median	Min.	Max.	Mean	SD	Median
All scenarios	Height (cm)	90	235	177.855	8.345	178.000	105	217	163.241	7.628	163.000
	Age (bracket)	1	13	5.550	3.547	5.000	1	13	6.016	3.564	6.000
	Attractiveness	1.000	5.000	2.268	0.751	2.158	1.000	5.000	2.987	0.777	2.993
	Estimated attractiveness	1.042	4.823	2.965	0.539	2.949	1.498	4.832	3.369	0.524	3.398
Base scenario	Happiness	0.000	2392.018	168.798	351.191	3.501	0.000	2420.308	419.168	425.972	302.461
	Got liked	0	2289	352.012	562.218	58.000	0	5793	3179.122	1968.559	3474.000
	Liking ratio	0.005	0.999	0.552	0.214	0.566	0.005	0.908	0.157	0.124	0.121
	Matched/Liked	0.000	1.000	0.080	0.191	0.001	0.000	1.000	0.271	0.304	0.132
Counterfactual A (*CA*)	Happiness	0.000	2391.788	168.829	351.188	3.558	0.000	2421.280	419.241	425.954	302.461
	Got liked	0	2289	352.371	562.005	60.000	0	5793	3184.328	1961.999	3474.000
	Liking ratio	0.005	0.999	0.553	0.215	0.567	0.005	0.903	0.157	0.124	0.121
	Matched/Liked	0.000	1.000	0.080	0.191	0.001	0.000	1.000	0.271	0.304	0.132
Counterfactual B (*CB*)	Happiness	0.000	2494.381	183.042	368.389	4.042	0.000	2425.322	457.146	427.932	385.546
	Got liked	0	2289	359.443	565.095	63.000	0	5793	3335.070	1935.842	3763.000
	Liking ratio	0.009	1.000	0.579	0.216	0.603	0.005	0.911	0.161	0.126	0.123
	Matched/Liked	0.000	1.000	0.081	0.191	0.001	0.000	1.000	0.288	0.307	0.167

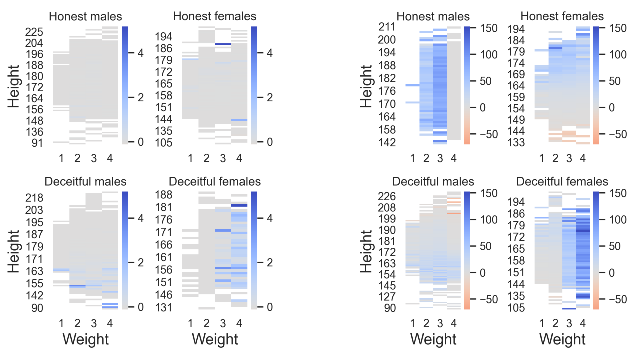

(a) Counterfactual A (b) Counterfactual B

Figure 4: Mean happiness difference heatmap with genders, honesty, height, and weight through counterfactual scenarios.

Table 2: Happiness difference by group in Counterfactual B

			Difference			
Gender	Behavior	Change	Min.	Max.	Mean	Median
Male (72.0%)	Honest (8.9%)	None (0.0%)	0	0	0	0
		Negative (1.2%)	−0.729	−0.001	−0.086	−0.049
		Positive (7.7%)	0.001	179.632	53.005	52.364
	Deceitful (63.1%)	None (39.8%)	0	0	0	0
		Negative (0.8%)	−69.676	−0.001	−3.753	−1.486
		Positive (22.5%)	0.001	279.604	27.495	13.275
Female (28.0%)	Honest (9.8%)	None (0.0%)	0	0	0	0
		Negative (1.1%)	−31.265	−0.001	−1.709	−0.069
		Positive (8.7%)	0.001	158.442	17.775	9.598
	Deceitful (18.2%)	None (4.5%)	0	0	0	0
		Negative (0.4%)	−24.332	−0.001	−1.087	−0.005
		Positive (13.4%)	0.001	449.451	68.217	33.999

scenarios. Specifically, under frequent misrepresentation, deceitful females derived much more utility than their male counterparts, being on opposite sides of the happiness difference spectrum. Our statistical tests showed that deceitful females benefited significantly more than everyone else. Our mean matching rate for males closely followed Tyson et al. [47] (0% vs. 0.6%), a study not referenced at the time of the simulation. While our matching rate for females was higher (27.1% vs. 10.5%) our median (13.2%) was comparable.

Attractiveness and happiness. Attractiveness was key to baseline happiness, more so for males (r = .9) than females (r = .82). For males, it had the highest correlation with the happiness difference in high-deceit environments, implying that attractiveness ensures male happiness. For females, attractiveness had a weak positive correlation with the happiness diff. in both settings. While women were more likely to be happier despite below-average attractiveness, reasonably attractive yet over-confident males fared better than their female counterparts. Extremely attractive females not being in the top 10% percentile confirms the findings of Rudder [39]. For honest males and females, estimated attractiveness was more correlated with happiness difference (r = −.3, −.51) than attractiveness (r = −.17, −.39), in the negative direction. For the deceitful males and females, attractiveness was more correlated with happiness diff. (r = .69, .57) than est. attractiveness (r = .29, .25). The direction difference is most likely tied to the fact that the deceitful are mostly from the lower percentile and the middle ranges benefit more from the change. The correlations between the (estimated) attractiveness and the misrepresented attributes with the happiness differences being equal raises the question of whether (1) low self-esteem and misrepresentation may have a similar effect on happiness and (2) misrepresentation could be a game-theoretical response to unrealistic expectations due to high self-esteem in dating.

Height. As expected, height was more influential for males than females and positively correlated with happiness. Its correlation with the happiness difference in *CB* was weakly positive for the honest (r = .06) but weakly negative for the deceitful (r = −.17), suggesting taller deceitful males sabotaged their success by going out of the females' height preference range. For females, the correlation between height and happiness difference in *CB* was .26 for all females and .69 for the honest, which may be due to taller females finding more males that are (seemingly) in the suitable height range. Extremely short honest females and extremely tall deceitful males saw a decrease due to losing their compatibility with each other.

Weight. Weight having a positive correlation for males and a negative correlation for females with baseline was expected. Overall, weight misrepresentation benefited obese females the most. The honest obese males were unaffected, while the deceitful experienced a positive or negative effect based on their heights.

Age. Due to the tendencies in the underlying age preference distribution, age had significant but weakly negative correlations with various outcomes (between −.07 and −.12), slightly more important for females and generally favoring younger agents. Due to agents not misrepresenting their age, its effect was relatively minimal.

Honesty. In *CA*, no deceitful agent saw a negative impact, and the deceitful benefited *more* from their deceitful behavior than the honest (albeit the happiness diff. having weak effect size and correlations). In *CB*, negative effects of strategic misrepresentation were marginal. All groups were more likely to be affected positively than negatively, and the deceitful benefited *less* from their deceitful behavior than the honest, which was counter-intuitive. These findings suggest that misrepresentation, even in very high density, can be tolerated and even benefit social welfare, especially the honest. However, this is likely the case due to agents not being resource-efficient (which may be the case in real life as well). Our findings about honesty yielding no difference or benefit in deceitful scenarios do not perfectly align with the view that honesty hurts success [18] in the larger scale.

6.1 Platform's perspective

Due to the lack of overall negative impact of strategic misrepresentation, the platform may not have an immediate incentive to crack down on misrepresentation. Considering the overall improvement may positively reinforce the platform use, misrepresentation may pose a conflict of interest for platforms.

6.2 Limitations

Our simulation has noteworthy simplifications and approximations. For example, in our study, perceived attractiveness is objective, any match has more utility than not having a match due to no agent crossing a deal-breaker through misrepresentation, and it is assumed that top-performing agents would not have an incentive to misrepresent themselves. Furthermore, certain aspects of the simulation either lack corresponding research/consensus or are too abstract to be definitively quantified. For this reason, a perfect parallel of reality is not possible. To tackle the difficulties of simulating online dating, we ground our study in existing research, as much as possible, either by directly using quantitative data or modeling it after the qualitative findings; these are referenced in the model explanation. Finally, populations with wildly different distributions may see different outcomes.

7 RELATED WORK

There is considerable research on dating and online dating. However, existing studies [4, 6, 7, 9, 11, 14, 17–25, 28, 31–33, 35, 39, 41, 46, 47, 49, 50, 52] mostly focus on capturing and analyzing preferences through surveys or lab experiments and/or quantifying the amount of misrepresentation without exploring its effects. In some rare cases, publications were produced in cooperation with the OD platforms [39, 50, 51]. The closest experiment we could find to our study was done by OkCupid through lying about candidate compatibilities, published in their blog [40]; their results suggest perceived partner compatibility generated longer conversations, even when the actual compatibility was very low.

Our experiments are simulation-based. The only relevant and publicly available simulator we could find was a NetLogo-based [45] model for an *offline* and simplistic dating environment [10]. To the best of our knowledge, ours is the first multiagent simulation framework for online dating.

Matchmaking algorithms are vital and closely guarded industry secrets for online dating platforms, with limited information available online. Platforms have used various methods such as an Elo-like rating system[14], compatibility questionnaires, preference elicitation, physical distance, and filters [3, 27, 36, 39, 42, 44].

8 CONCLUSION

In this study, we approached online dating from a multiagent simulation perspective, created an extensible online dating simulation framework, and used it to analyze how strategic misrepresentation in online dating affects people. Our simulation details were, as much as possible, based on existing studies. We found that strategic misrepresentation in our simulations increased overall happiness with a large effect size when misrepresentation was extremely frequent (81.3%). Moreover, its negative effect on individuals was marginal, and honest agents benefited more from deceitful agents' misrepresentation than the deceitful themselves. However, our results suggest the improvement was not distributed evenly over the population, favoring people who were already better off in happiness more. Although the deceitful were mostly male, we found that female agents benefited more from misrepresentation than males in both counterfactual scenarios. In a high-deceit environment, being attractive or fitter for males, misrepresenting weight for deceitful females, and being tall for honest females were the best indicators for improvement in happiness. Overstating height was harmful for already tall males. Our findings suggest misrepresentation can have complex and counter-intuitive outcomes for the population and specific groups, which may be one of the reasons behind misrepresentation in the current online dating scene. Our Python framework and supplementary materials are publicly available [30].

Future work can explore learning agents that adapt their strategies or self-image, incorporate some of the mentioned factors such as preference misrepresentation, examine the effects of the platform's revenue models, and design novel mechanisms to incentivize truthfulness while equitably improving the experience of all users.

ACKNOWLEDGMENTS

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), [RGPIN-2021-04196]; and we thank Kiki Zheng for her contribution to the paper title.

[14] The Elo rating was initially invented to rank chess players [12, 13], and it is still being used in chess and competitive video games.

REFERENCES

[1] Niloofar Abolfathi and Simone Santamaria. 2020. Dating disruption-how Tinder gamified an industry. *MIT Sloan Management Review* 61, 3 (2020), 7–11.

[2] Robert J Aumann. 1985. An axiomatization of the non-transferable utility value. *Econometrica: Journal of the Econometric Society* (1985), 599–612.

[3] Nana Baah. 2020. How does the Hinge algorithm work? https://www.vice.com/en/article/z3e3bw/how-does-the-hinge-algorithm-work

[4] Michele Belot and Marco Francesconi. 2006. Can anyone be'the'one? Evidence on mate selection from speed dating. (2006).

[5] United States Census Bureau and U.S. Bureau of Labor Statistics. 2017. Current Population Survey 2017: Table FG3 - All Families by Type, Age, and Presence and Age of Related Children: 2017. https://www2.census.gov/programs-surveys/demo/tables/families/2017/cps-2017/tabfg3-all.xls

[6] David M. Buss and David P. Schmitt. 1993. Sexual strategies theory: an evolutionary perspective on human mating. *Psychological review* 100, 2 (1993), 204.

[7] Benjamin Cornwell and David C Lundgren. 2001. Love on the Internet: Involvement and misrepresentation in romantic relationships in cyberspace vs. realspace. *Computers in Human Behavior* 17, 2 (2001), 197–211.

[8] Edward E Cureton. 1956. Rank-biserial correlation. *Psychometrika* 21, 3 (1956), 287–290.

[9] Charlotte De Backer, Johan Braeckman, and Lili Farinpour. 2008. Mating intelligence in personal ads. *Mating intelligence: Sex, relationships, and the mind's reproductive system* (2008), 77–101.

[10] Anuja Doshi. 2010. GSB Dating Scene. http://modelingcommons.org/browse/one_model/2800

[11] 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.

[12] Arpad Emmerich Elo. [n.d.]. The new U.S.C.F. rating system. *Chess Life* 16 ([n. d.]), 160–161.

[13] Arpad Emmerich Elo. 1978. *The rating of chessplayers, past and present.* Arco Pub.

[14] April E Fallon and Paul Rozin. 1985. Sex differences in perceptions of desirable body shape. *Journal of abnormal psychology* 94, 1 (1985), 102.

[15] Centers for Disease Control and Prevention. 2022. 2022 BRFSS Data. https://www.cdc.gov/brfss/annual_data/2022/files/LLCP2022ASC.zip

[16] David Gale and Lloyd S Shapley. 1962. College admissions and the stability of marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.

[17] Stella D Garcia and Diana Khersonsky. 1997. " They Are a Lovely Couple": Further Examination of Perceptions of Couple Attractiveness. *Journal of Social Behavior and personality* 12, 2 (1997), 367.

[18] 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.

[19] Tobias Greitemeyer. 2020. Unattractive people are unaware of their (un) attractiveness. *Scandinavian Journal of Psychology* 61, 4 (2020), 471–483.

[20] 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.

[21] 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.

[22] Martie G Haselton, David M Buss, Viktor Oubaid, and Alois Angleitner. 2005. Sex, lies, and strategic interference: The psychology of deception between the sexes. *Personality and Social Psychology Bulletin* 31, 1 (2005), 3–23.

[23] Ida Jodette Hatoum and Deborah Belle. 2004. Mags and abs: Media consumption and bodily concerns in men. *Sex roles* 51 (2004), 397–407.

[24] Rebecca D Heino, Nicole B Ellison, and Jennifer L Gibbs. 2010. Relationshopping: Investigating the market metaphor in online dating. *Journal of Social and Personal relationships* 27, 4 (2010), 427–447.

[25] Günter J Hitsch, Ali Hortaçsu, and Dan Ariely. 2010. What makes you click?—Mate preferences in online dating. *Quantitative marketing and Economics* 8 (2010), 393–427.

[26] Vincent Iehlé. 2007. The core-partition of a hedonic game. *Mathematical Social Sciences* 54, 2 (2007), 176–185.

[27] Anna Iovine. 2021. How do all the best dating app algorithms work? https://mashable.com/article/tinder-bumble-hinge-okcupid-grindr-dating-app-algorithms

[28] Douglas T. Kenrick, Richard C. Keefe, Cristina Gabrielidis, and Jeffrey S Cornelius. 1996. Adolescents' age preferences for dating partners: Support for an evolutionary model of life-history strategies. *Child development* 67, 4 (1996), 1499–1511.

[29] Dave S Kerby. 2014. The simple difference formula: An approach to teaching nonparametric correlation. *Comprehensive Psychology* 3 (2014), 11–IT.

[30] Oz Kilic and Alan Tsang. 2024. Supplementary Materials for "Catfished! Impacts of Strategic Misrepresentation in Online Dating". https://osf.io/bngr7/

[31] Leonard Lee, George Loewenstein, Dan Ariely, James Hong, and Jim Young. 2008. If I'm not hot, are you hot or not? Physical-attractiveness evaluations and dating preferences as a function of one's own attractiveness. *Psychological Science* 19, 7 (2008), 669–677.

[32] Juan-Ignacio Martínez-Pastor. 2017. How important is physical attractiveness in the marriage market? *Revista Española de Investigaciones Sociológicas (REIS)* 159, 159 (2017), 91–110.

[33] James K McNulty, Lisa A Neff, and Benjamin R Karney. 2008. Beyond initial attraction: Physical attractiveness in newlywed marriage. *Journal of family psychology* 22, 1 (2008), 135.

[34] Business of Apps. 2022. Dating app revenue and usage statistics (2022). https://www.businessofapps.com/data/dating-app-market/

[35] Boguslaw Pawlowski and Slawomir Koziel. 2002. The impact of traits offered in personal advertisements on response rates. *Evolution and Human Behavior* 23, 2 (2002), 139–149.

[36] Christy Piña. 2019. Here's how Hinge's algorithm works to make sure you're getting compatible matches. https://www.elitedaily.com/p/how-does-hinges-algorithm-work-the-app-ensures-youre-connecting-with-compatible-matches-15910445

[37] Michael J Rosenfeld, Reuben J Thomas, and Sonia Hausen. 2019. Disintermediating your friends: How online dating in the United States displaces other ways of meeting. *Proceedings of the National Academy of Sciences* 116, 36 (2019), 17753–17758.

[38] Alvin E Roth and Marilda Sotomayor. 1992. Two-sided matching. *Handbook of game theory with economic applications* 1 (1992), 485–541.

[39] Christian Rudder. 2014. *Dataclysm: Who we are (when we think no one's looking)*. Random House Canada.

[40] Christian Rudder. 2014. We Experiment On Human Beings! https://web.archive.org/web/20141001084951/http:/blog.okcupid.com/index.php/we-experiment-on-human-beings/

[41] Maria Stoicescu. 2020. Social impact of online dating platforms. A case study on Tinder. In *2020 19th RoEduNet Conference: Networking in Education and Research (RoEduNet)*. IEEE, 1–6.

[42] John Tierney. 2013. A match made in the code. https://www.nytimes.com/2013/02/12/science/skepticism-as-eharmony-defends-its-matchmaking-algorithm.html

[43] Elisabeth Timmermans and Cédric Courtois. 2018. From swiping to casual sex and/or committed relationships: Exploring the experiences of Tinder users. *The Information Society* 34, 2 (2018), 59–70.

[44] Tinder. 2019. Powering Tinder® - the method behind our matching. https://www.tinderpressroom.com/powering-tinder-r-the-method-behind-our-matching/

[45] Seth Tisue and Uri Wilensky. 2004. NetLogo: Design and implementation of a multi-agent modeling environment. In *Proceedings of agent*, Vol. 2004. 7–9.

[46] 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.

[47] Gareth Tyson, Vasile C Perta, Hamed Haddadi, and Michael C Seto. 2016. A first look at user activity on Tinder. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. IEEE, 461–466.

[48] Emily A. Vogels. 2020. 10 facts about Americans and online dating. https://www.pewresearch.org/fact-tank/2020/02/06/10-facts-about-americans-and-online-dating/

[49] Emily A. Vogels. 2023. Key findings about online dating in the U.S. https://www.pewresearch.org/short-reads/2023/02/02/key-findings-about-online-dating-in-the-u-s/

[50] Stephen Whyte and Benno Torgler. 2017. Preference versus choice in online dating. *Cyberpsychology, Behavior, and Social Networking* 20, 3 (2017), 150–156.

[51] Peng Xia, Bruno Ribeiro, Cindy Chen, Benyuan Liu, and Don Towsley. 2013. A study of user behavior on an online dating site. In *Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining*. 243–247.

[52] George Yancey and Michael O Emerson. 2016. Does height matter? An examination of height preferences in romantic coupling. *Journal of Family Issues* 37, 1 (2016), 53–73.