Learning online dating preferences from neuroimaging data

Mira Nencheva

I. INTRODUCTION

A. Why study online dating?

What is it about a person that makes you want to approach them? With the rise of online dating websites, for many the 'approach' has changed from walking across a crowded bar to sliding one's finger across a cell phone screen to indicate one's interest in dating. The initial question about how we make these split-second judgments about the desirability of a partner has changed as well, as our judgments are often solely based on photographs with minimal additional information. While studies have shown that our judgments 'on paper' (from merely looking at the online dating profile of an individual) more often than not fail to predict how much we like them in person (Finkel et al., 2012), millions of online dating users are doing the initial selection of who to contact precisely in this rapid fashion. Since 2000, the population of online dating users has grown significantly. According to unofficial reports, Tinder, a popular dating app, has over 50 million users. Tinder reports that since their launch in 2012 their app has matched more than 12 million people and that theirs users make billions of decisions every day about whether they are interested in potentially meeting someone by swiping right or left. Online dating is changing the dating world, and it is becoming more and more important to understand the specific mechanisms of mate selection that are in place in the unique social environment of online dating websites.

B. Goal: Predicting choice in online dating from target features and individual differences

The choice to contact a partner in online dating involves two sets of factors: those relating to the individual making the decisions (or individual variations in preferences) and those characterizing the 'target' of the decision (or target features). One interest is in developing better algorithms for online dating websites to recommend potential matches to their users. If we understand how to accurately predict a given user's ideal partner choices, we could help them find someone they like much faster. In this study, I aim to use both target features and individual differences to predict users' choices to contact a given partner in online dating or not.

C. Why is asking users what the are looking for not enough?

How does one measure a user's partner preferences? One way is to simply ask a user what qualities they would want their ideal partner to have (ideal partner preferences), or to state how important different traits of a potential partner are for their decision (I will refer to this as stated preferences). Wood & Brumbaugh (2009) for example asked participants to indicate how important a set of characteristics were in how attractive they found people (they labeled this as 'direct preference measure'). Then they asked participants how attractive they found 98 people based on their photos (they labeled this as 'revealed preferences'). They found that across all groups of participants (which included participants of different genders, sexual orientations and ethnicities) revealed preferences were not strongly correlated with direct preference measures.

D. Neural activation as a measure of individual preferences

Neuroscience may offer another individual differences measure, namely, how the brain responds to the photographs of different potential partners. For many years scientists have been interested in neural processes that help us make decisions. Over the past few decades, a large body of studies using methods such as functional magnetic resonance imaging (FMRI) in humans, as well as electrophysiological recording, electrical and optogenetic stimulation, and pharmacological inhibition in animal models, has implicated the dopaminergic system of the brain as essential for estimating the expected value of a choice in diverse decision tasks. Specifically, research suggests that ventral striatal activity (and nucleus accumbens or NAcc activity in particular) correlates with the value of anticipated rewards (see Fig.1) (similar to the economic notion of expected value; Breiter et al. 2001, Knutson et al. 2001), and can predict choice (e.g., Knutson & Greer, 2008).

Consistent with the notion that the dopaminergic system plays a critical role in initiating motivated behavior, multiple studies have shown activity in dopamine projection areas such as the striatum and ventral tegmental area related to love (Aron et al. 2005, Cacioppo et al. 2008, Ortigue et al. 2010, Fisher et al. 2005) and sexual desire (Buhler et al. 2008, Cacioppo et al. 2012, Hamann et al. 2004, Hanes et al. 2004). Both Cacioppo et al. (2012) in a large meta analysis of neuroimaging studies on love and desire, and Aron et al. (2005) suggested that in addition to striatal regions, the insula was commonly tracked feelings of love and desire.

While there is a large body of research on attraction and early stage romantic love, how people initially select a romantic partner remains an open question. Cooper et. al (2012) attempted to answer this question in the context of speed dating. Prior to attending a series of speed dating events, participants were asked to rate the attractiveness of and their dating interest towards a group of potential partners (some of which they would meet during the speed dating events) while in a MRI scanner. After 3 speed dating events, participants submitted their final choices. The investigators found that activity in the anterior cingulate cortex, medial prefrontal cortex (MPFC), and cerebellum scaled with consensus judgments of attractiveness (based on average desirability of a person), whereas activity in the dorsolateral prefrontal cortex (DMPFC) tracked individual preferences and predicted choice to date targets.

II. DATA COLLECTION AND PROCESSING

A. Stimuli

We recruited 36 Caucasian[1] females between the ages of 21 and 25 via social media and mailing lists to submit each two identical photos of themselves ('selfies'), such that their face was fully visible. On one of the photos they were instructed to smile and one the other one they were instructed to keep a neutral facial expression. The 72 photos (2 per female for the 36 females) were cropped to ensure that the size of the face, size of the photo and the resolution of the photo were the same across photos. The photos were cropped to a square, so that only the face and part of the hair was in the photo. We also ensured that the photos were comparable in brightness.

B. Target ratings

On the Qualtrics survey site, we recruited a panel of 50 Caucasian males between the ages of 18 and 28 who were romantically interested in females and not in a committed relationship at the time. Participants were asked to rate how 'physically attractive', and how 'happy' each target appeared. The ratings were on a 4-point scale: *To a very small extent, To a small extent, To a moderate extent, To a large extent*. Each participant only saw one photo per female (therefore, each photo was rated by 25 participants). The order of the photos was randomized for each participant.

C. Facial features

I then used the Microsoft Cognitive Services Face API to obtain the coordinates of several important points on the face of each target. I then estimated 6 features of the face as follows:

- nose width: distance between the coordinated of the left and right side of the root of the nose
- nose length: distance between the root of the nose and the tip of the nose
- eye length: distance between the inner and outer end of the eye (calculated for the left eye)
- eye width: distance between the top lid and bottom lid of the eye (calculated for the left eye)
- eye distance: distance between the inner side of the two eyes
- upper lip width: distance between the top of the upper lip and the bottom of the upper lip

[1]We chose to only include Caucasian females as targets, because of strong effects of race on preferences documented in previous research (Hitsh et al. 2010, Robnett &Feliciano. 2011). Including the effects of race would have required a much larger database of photos and could have introduced additional variability that might obfuscate the effects of smiling, which were the initial focus of the study. We hope to replicate the study in the future with other combinations of gender, race, and sexual orientation.

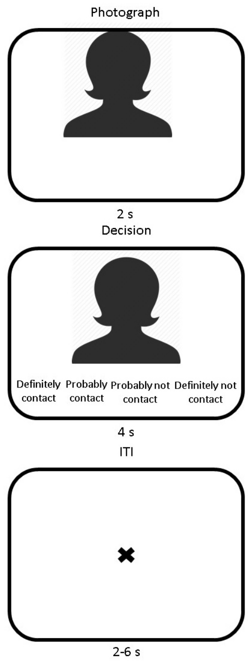

Fig. 1. Structure of the online dating task for the neuroimaging study.

- lower lip width: distance between the top of the lower lip and the bottom of the lower lip

In addition to these features I also manually scored each photo on three additional sets of features: hair color, hair texture, and hairstyle. The hair color feature has three variations: blond, brown and red. Each target received a 1 for the heir color she had in the photo and 0 for the other two colors. Similarly, the hair texture feature had two variations: straight and curly, and the hairstyle feature had two features: hair up (in a ponytail or a bun) or hair down.

D. Task

The online dating task was designed to re-create the experience of choosing to contact a partner on an online dating mobile application, such as Tinder while subjects underwent scanning. During each trial of the task, subjects first saw a photograph of a female target (2 s; the 'cue period'). Subjects then indicated whether they wanted to contact that target or not (4 s) (see Fig. 1; the choice period). In data analyses, negative responses ('Definitely not contact', 'Probably not contact') were combined into one category ('Not contact'), and positive responses ('Definitely contact' and 'Probably contact') were combined into a second category ('Contact'). Finally, subjects fixated on a cross for a variable intertrial interval (26 s).

E. Participants

We recruited participants for a study on stock markets that was run prior to the online dating task. Twenty-three males, romantically interested in females and between the ages of 21 and 31 (average age 26) also agreed to participate in the online dating task.

F. Neuroimaging data and regions of interest

Magnetic resonance imaging (MRI) data was acquired with a 3.0-T General Electric MRI scanner. I acquired the standardized (using z-scores) MR signal from 6 regions of interest: nucleus accumbens, which has been associated with predicting choice (Knutson & Greer, 2008), early stage romantic love (Aron et al. 2005, Cacioppo et al. 2008, Ortigue et al. 2010, Fisher et al. 2005) and sexual desire (Buhler et al. 2008, Cacioppo et al. 2012, Hamann et al. 2004, Hanes et al. 2004), anterior insula, which has also been observed to track feelings of love and desire (Cacioppo et al., 2012, Aron et al., 2005), medial prefrontal cortex (MPFC), anterior cingulate cortex, and dorsolateral prefrontal cortex (DLPFC), which were associated in partner choice in a speed dating study (Cooper et. al 2012), and the fusiform gyrus, which is involved in perceiving faces (McCarthy et. al 1997). The coordinates of the ROIs were determined based on previous meta analyses (Knutson & Greer, 2008).

G. Selecting training and testing examples

I removed all trials in which the participant failed to respond, as well as trials in which the MR signal in any of the regions of interest was above 4 standard deviations from the average. In order to accurately estimate the testing error for each model, I performed 20 iterations, at each randomly choosing 3*36 trials as testing examples, and the remaining 20*36 trials as testing examples. I then report the average error in classifying the testing examples as 'Contact' or 'Not Contact' based on the input.

III. PREDICTING CHOICE FROM NEURAL ACTIVITY AND TARGET RATINGS

A. Support Vector Machine: three models

I trained a kernelized Support Vector Machine (SVM) on three sets of inputs one including only neural data (neural model), one including only average happiness and attractiveness ratings of the 72 target photos (target model) and one including both neural data and target ratings (combined model). The SVM algorithm was adapted from John Duchi, Machine Learning (CS 229) Supplemental Lecture Notes. I used a Gaussian kernel function:

$$K(x, z) = exp(-\frac{1}{2\tau}||x - z||_2^2)$$

The prediction for each example was given by the equation:

$$sign(\sum_{i=1}^{m} exp(-\frac{1}{2\tau}||x^{(i)} - x||_2^2)\alpha_i)$$

Where the $x^{(i)}$s are given by the inputs of the training examples and the x was given by the input of the example we are attempting to predict. In the case of the neural model the input was the signal change from the beginning of the trial to 6 s after the onset of the trial. In the case of the target model, the input was the average attractiveness and happiness rating that the photo received from the Raters group. In the combined model the input was the combination of the inputs of the neural and the target models. The term

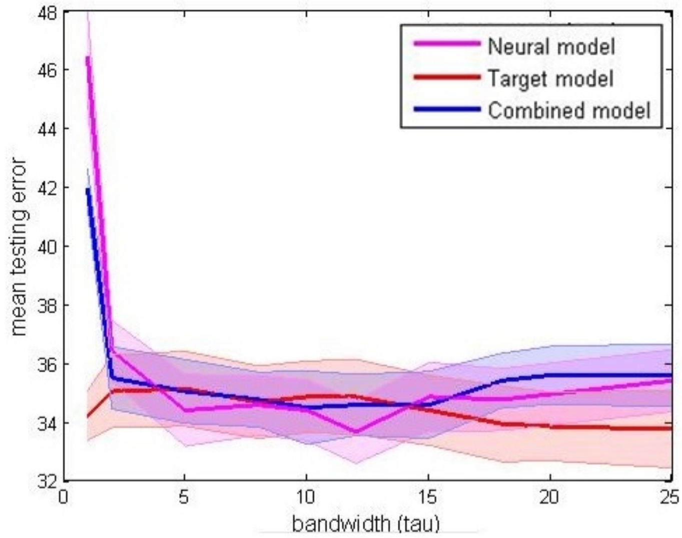

Fig. 2. Comparing the effects of bandwidth on the neural, target and combined models.

$x^{(i)} - x_2^2$ measures how similar the input of the example we are trying to predict to the i-th training example. The term α was determined using stochastic gradient descent on the training examples. The algorithm outputs -1 for trials when it predicts that the subject chose not to contact the female and +1 on trials when it predicts the subject chose to contact the target.

B. Finding the optimal bandwidth

I ran the three SVM algorithms (neural, target and combined) with different bandwidth values ($\tau = \{1, 2, 3, 5, 8, 10, 12, 15, 18, 20, 25\}$). Interestingly the target model was not significantly affected by the choice of bandwidth, whereas for the neural model the choice of bandwidth made a dramatic difference (see Fig. 2). This is likely because the neural data is noisier and needs more smoothing. This is in part because biological signals tend to be noisy, and in part because the behavioral data comes from average ratings from 52 raters. The optimal bandwidth for the neural model was $\tau = 12$. At the optimal bandwidth for the neural model there was no significant difference between the models, however, the neural model was slightly better (average testing error neural model: 33.6697%, average testing error target model: 34.8624%, average testing error combined model: 34.5872%).
The lack of a notable difference between the combined model and the other two models could suggest that activity in the 6 regions of interest is simply a neural representation of the two features of the target model.

C. Removing features from the neural model

In order to see if there is a single region in the brain (out of the 6 ones discussed above) that provides especially important information for predicting choice, I tested the neural model at bandwidth 12 while removing each of the

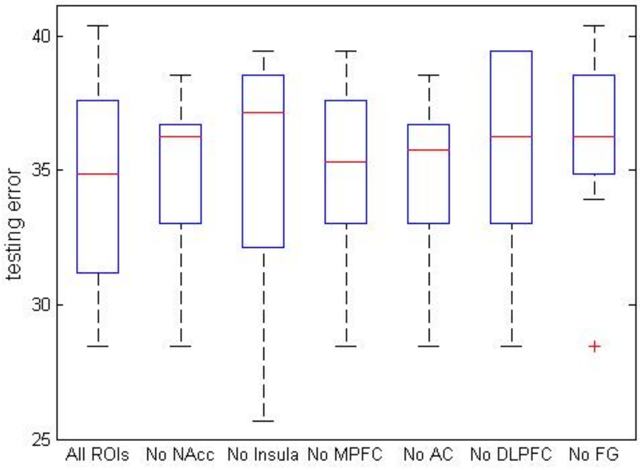

Fig. 3. Eliminating features from the neural model.

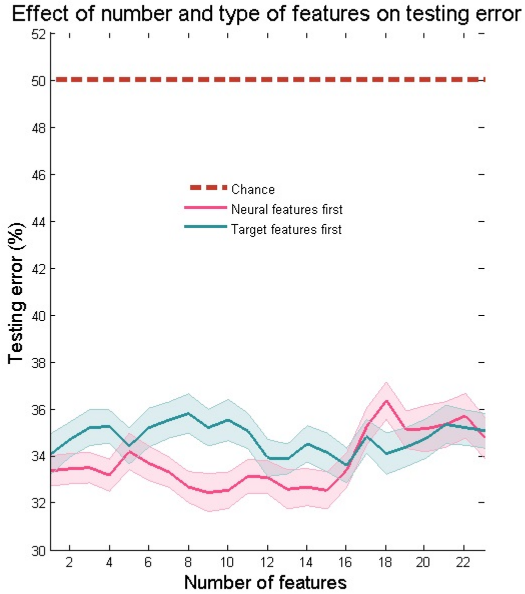

Fig. 4. Effect of adding photo features to the model.

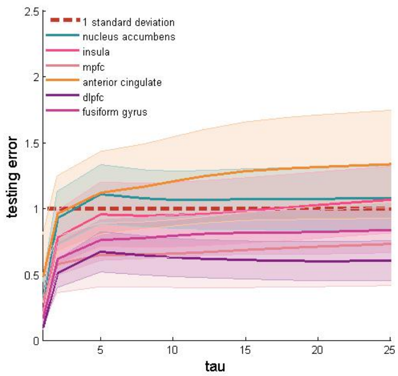

Fig. 5. Weighted linear regression based on photo features.

regions of interest. Removing any one of the regions of interest did not seem to have a significant effect on the accuracy of the model (see Fig. 3). This is possibly due to the fact that activity in some of the regions is correlated, and thus removing a single one of them would be compensated by the presence of the remaining regions.

D. *Effect of number of features on test error*

In order to see whether adding more features to the model would increase the accuracy of prediction, I sequentially added one feature at a time in two different orders. The first order started with the target ratings, followed by the facial features (see section C in methods) and finally adding the neural features. The second order started with the neural features, was then followed by the target ratings and the facial features last. (see Fig. 4). Adding more facial features appeared to increase the error (features 3-17 for the teal line and features 9-23 for the pink line). This pattern is interesting and somewhat intuitive. While it might seem that knowing the exact length of the nose or the size of the eyes of a target would help us predict the user's choice better, including these characteristics might lead the model to over-fit to the target and discount the valuable information about the individual preferences of each user. In contract, the neural response and the attractiveness ratings would not help us to identify which target the user is making decisions about, however they give us more information about the preference for that target (of the individual and of the group). Information, which is arguably more closely linked to choice. Including the neural features with attractiveness and happiness ratings lead to a smaller testing error (32.4771%) compared to adding the neural features last.

E. *Predicting choice: conclusions*

While being able to predict a binary outcome with about 68% accuracy might seem far from perfect, it is important to note that we are predicting human decisions about other humans. Quite a few studies of economic decision-making, where it is possible (and in fact quite easy) to come up with an estimation for the value of each choice, report similar success rates. Therefore, for a social psychology experiment, where the value of each choice is noisy, and affected by a large set of unknown individual differences and external factors, 68% accuracy is quite high, and perhaps as far as we can go given the data set.

IV. PREDICTING NEURAL SIGNAL IN RESPONSE TO TARGET

For the second part of my project I attempted to predict the signal change in response to seeing a given target photo using a locally weighted linear regression. I defined my weights in two ways. In the first model the weights were only based on the photo features (average ratings on happiness and attractiveness as well as facial features), whereas in the

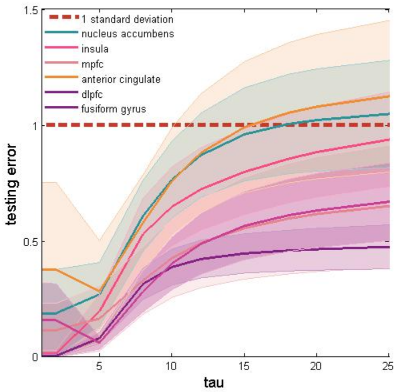

Fig. 6. Weighted linear regression based on photo features and average pattern of neural response.

second model the weights were also based on the average pattern of neural response for the 6 regions of interest for the participant the training example came from. To determine the similarity between the features I used the formula from the lecture notes, where the weight of the i-th training example was given by:

$$w^{(i)} = exp(\frac{(x^{(i)} - x)^2}{2\tau^2})$$

As in the case of SVM, the model consisting only of photo features did best under small bandwidth values, and had similar performance for larger bandwidth values (see Fig. 5). Again, this is likely due to the fact that photo features are subject to less noise than the biological signals from the brain. The model that included neural data for the prediction did best for a bandwidth value of 5 for most regions (with the exception of the DLPFC and the Nucleus Accumbens, for which the prediction was most accurate at bandwidth 1).Overall, the predictions for the model accounting for the similarity in the average neural response performed slightly better (see Fig. 6).

V. CONCLUSION

Using neural data in predicting choice in online dating allowed for slightly more accurate predictions than merely using photo features as previous prediction algorithms had done. By doing so I was able to account for participant's individual preferences. In addition, introducing similarity in the average pattern of neural response when predicting signal in the 6 regions of interest increased our accuracy of prediction slightly.

REFERENCES

[1] Aron, A., Fisher, H., Mashek, D. J., Strong, G., Li, H., & Brown, L. L. (2005). Reward, motivation, and emotion systems associated with early-stage intense romantic love. Journal of neurophysiology,94(1), 32737.

[2] Bhler, M., Vollstdt-Klein, S., Klemen, J., & Smolka, M. (2008). Does erotic stimulus presentation design affect brain activation patterns? Event-related vs. blocked fMRI designs. Behavioral and Brain Functions, 4(1), 30.

[3] Cloutier, J., Heatherton, T. F., Whalen, P. J., & Kelley, W. M. (2008). Are attractive people rewarding? Sex differences in the neural substrates of facial attractiveness. Journal of Cognitive Neuroscience, 20(6), 941-951.

[4] Cooper, J. C., Dunne, S., Furey, T., & O'Doherty, J. P. (2012). Dorsomedial prefrontal cortex mediates rapid evaluations predicting the outcome of romantic interactions. The Journal of Neuroscience, 32(45), 15647-15656.

[5] Cacioppo, S., Bianchi-Demicheli, F., Frum, C., Pfaus, J., & Lewis, J. (2012). The Common Neural Bases Between Sexual Desire and Love: A Multilevel Kernel Density fMRI Analysis. The Journal of Sexual Medicine, 9(4), 1048.

[6] Eastwick, P. W., & Finkel, E. J. (2008). Sex differences in mate preferences revisited: do people know what they initially desire in a romantic partner?.Journal of personality and social psychology, 94(2), 245.

[7] Finkel, E. J., Eastwick, P. W., Karney, B. R., Reis, H. T., & Sprecher, S. (2012). Online Dating A Critical Analysis From the Perspective of Psychological Science. Psychological Science in the Public Interest, 13(1), 3-66.

[8] Fiore, A. T., & Donath, J. S. (2005, April). Homophily in online dating: when do you like someone like yourself?. In CHI'05 Extended Abstracts on Human Factors in Computing Systems (pp. 1371-1374). ACM.

[9] Fisher, H., Aron, A., & Brown, L. (2005). Romantic love: An fMRI study of a neural mechanism for mate choice. The Journal of Comparative Neurology, 493(1), 58.

[10] Hamann, S., Herman, R. A., Nolan, C. L., & Wallen, K. (2004). Men and women differ in amygdala response to visual sexual stimuli. Nature neuroscience, 7(4), 411416. Nature Publishing Group.

[11] Hitsch, G. J., Hortasu, A., & Ariely, D. (2010). What makes you click?Mate preferences in online dating. Quantitative marketing and Economics,8(4), 393-427.

[12] Knutson, B., Adams, C. M., Fong, G. W., & Hommer, D. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. J Neurosci, 21(16), RC159.

[13] Knutson, B., Fong, G. W., Bennett, S. M., Adams, C. M., & Hommer, D. (2003). A region of mesial prefrontal cortex tracks monetarily rewarding outcomes: characterization with rapid event-related fMRI. Neuroimage, 18(2), 263-272.

[14] Knutson, B., & Greer, S. M. (2008). Anticipatory affect: neural correlates and consequences for choice. Philosophical Transactions of the Royal Society of London B: Biological Sciences, 363(1511), 3771-3786.

[15] McCarthy, G., Puce, A., Gore, J. C., & Allison, T. (1997). Face-specific processing in the human fusiform gyrus. Journal of cognitive neuroscience,9(5), 605-610.

[16] O'Doherty, J., Winston, J., Critchley, H., Perrett, D., Burt, D. M., & Dolan, R. J. (2003). Beauty in a smile: the role of medial orbitofrontal cortex in facial attractiveness. Neuropsychologia, 41(2), 147-155.

[17] Ortigue, S., BianchiDemicheli, F., Patel, N., Frum, C., & Lewis, J. W. (2010). Neuroimaging of love: fMRI metaanalysis evidence toward new perspectives in sexual medicine. The journal of sexual medicine, 7(11), 3541-3552.

[18] Robnett, B., & Feliciano, C. (2011). Patterns of racial-ethnic exclusion by internet daters. Social Forces, 89(3), 807-828.

[19] Wood, D., & Brumbaugh, C. C. (2009). Using revealed mate preferences to evaluate market force and differential preference explanations for mate selection. Journal of personality and social psychology, 96(6), 1226.