

A Study of Situated Product Recommendations in Augmented Reality

Brandon Huynh, Adam Ibrahim, Yun Suk Chang, Tobias Höllerer, John O'Donovan
 Department of Computer Science
 University of California, Santa Barbara
 Santa Barbara, USA
 Email: {bhuynh, ai, ychang, holl, jod}@cs.ucsb.edu

Abstract—Augmented Reality interfaces increasingly utilize artificial intelligence systems to tailor content and experiences to the user. We explore the effects of one such system - a recommender system for online shopping - which allows customers to view personalized product recommendations in the physical spaces where they might be used. We describe results of a 2x3 condition exploratory study in which recommendation quality was varied across 3 user interface types. Our results highlight potential differences in user perception of the recommended objects in an AR environment. Specifically, users rate product recommendations significantly higher in AR and in a 3D browser interface, and show a significant increase in trust in the recommender system, compared to a web interface with 2D product images. Through semi-structured interviews, we gather participant feedback which suggests AR interfaces perform better due to their ability to view products within the physical context where they will be used.

Index Terms—Augmented Reality; Recommender Systems; User Interfaces; Recommendation Quality; Trust

I. INTRODUCTION

Recommender systems first emerged over two decades ago and have since become standard tools for dealing with information overload [1]–[3]. Major retail stores such as Amazon.com have a heavy focus on data-driven marketing, of which collaborative and content-based recommender systems are a core part. About 35% of sales on Amazon, and 75% of movies watched on Netflix are derived from recommendations [4]. The vast majority of recommendations for online retailers are delivered through email or in the traditional web browser interface. Interface technology, however, is developing rapidly: global revenues of Augmented Reality (AR) and Virtual Reality (VR) markets are expected to grow to over \$162 billion in 2020 [5]. Heavy investment in AR and VR by major companies such as Apple, Alphabet, Facebook, and Microsoft will mean that smaller, higher quality devices will become available at lower cost to consumers. Large retailers such as Amazon and IKEA are exploring and introducing new AR driven shopping experiences.

While there has been progress on in-store AR technology to improve shopping experiences, e.g. [6], less work has been done on the concept of in-home shoppers taking advantage of what we call ‘situated recommendations’, whereby *personalized* recommendations of products are placed virtually where the real product will be used. In particular, we are interested in how people *perceive* recommendations that are situated in

AR, and how this perception *differs* from that of traditional recommender system interfaces. We attempt to address the following specific research questions:

- 1) RQ1: Do users perceive product recommendations in AR differently than in a browser-based UI?
- 2) RQ2: Are there differences in recommender system trust when presented in AR versus a browser-based UI?
- 3) RQ3: What is the general sentiment towards an AR recommender system for in-home shopping?

To answer these questions, we conducted a 3 by 2 within-subjects lab study (N=31). The study examined the effects of 3 different interaction modalities: an Augmented Reality interface, a web browser interface with 3D view controls, and a web browser interface with 2D view controls. We also looked at how users respond to differences in recommendation algorithm quality (either high or low quality recommendations). We measured 2 key metrics, user ratings of each recommended object (also called perceived accuracy), and user trust in the recommender system. We collected subjective feedback on user perception of the modalities through a post study questionnaire and verbal interviews.

For the purposes of this study, we implemented a common online shopping user interface across all 3 modalities to allow for meaningful comparison. To avoid potential novelty effects, study participants undergo significant pre-study training sessions for each modality. Figure 1 shows an overview of our shopping interface. The right image shows a user wearing the HoloLens interacting with a virtual model of a recommended item and providing rating feedback to the system. The left images show the two web browser based interfaces that were tested in the study. In the web browser UIs, participants interact either by rotating the object with the mouse (3D), or clicking through static images (2D).

II. BACKGROUND

Our study combines facets from multiple research fields, including human computer interaction (HCI), recommender systems, and cognitive science. A discussion of the relevant literature in each area is presented here, to frame our contribution in the context of existing research.




Fig. 1. Left: Screen captures of the browser interfaces. For 2D browser, users can look through photos of the mug taken from different predefined angles. For 3D browser, users can freely rotate the mug and view it in any direction. Right: Mixed reality screen capture of a user providing feedback to a recommended item.

A. Augmented Reality in Retail Applications

Currently, there are many consumer applications for visualizing products in augmented reality. For example, IKEA uses a mobile AR app to place virtual models of their furniture in the physical world and Lego uses AR kiosks to visualize assembled Lego sets on the corresponding box [7]. Nike created a custom sneaker designer that uses projective AR to overlay designs onto a customer’s physical sneakers [8]. Recent work by Stoyonova et al. [9] reports on a cognitive study of purchase intent using AR in a shopping scenario, but in contrast to this work, does not have a focus on personalized recommendations, and is situated in a store as opposed to a home shopping scenario. Lu et al. [10] perform a study of AR for home shoppers, where selected products can be tried in AR before purchase. Olsson and colleagues [11] present a study of user experiences with AR in a shopping center context and report mainly positive feedback for mobile AR supported shopping. While there are many other examples of AR for improving shopping experiences [11], [12], to our knowledge there is no existing research that explores how users perceive *personalized recommendations* in this modality. We believe that our results can provide useful insight about this rapidly developing technology and its suitability as a channel for delivering personalized recommendations.

B. Augmented Reality and Recommendation

Many applications that integrate AR and recommendation use mobile platforms to perform location-based content recommendations. The Yelp monocle¹ for service recommendations is probably the most well-known example of this integration with AR. Balduini et al.’s Bottari system [13] provides personalized, location-based AR recommendations of social media content based on the Twitter network and evaluated the system in an urban area. While these approaches integrate AR and

¹https://www.yelp-support.com/article/What-is-Yelp-s-Monocle-feature?!=en_US

recommendation, they contrast with our approach in that they do not focus on evaluating perception of recommendations in AR compared to traditional UIs.

C. Interfaces and Decision Making

Prior research in recommender systems has a strong focus on algorithm performance. Recently however, more research attention is being paid to so called *user-aware* recommendation systems that attempt to improve the user’s experience with the recommender system by mechanisms that go beyond predictive accuracy, such as conversation [14], explanations [2], [15], [16], and various different flavors of user interfaces [17]–[19], interaction designs [20] and evaluations [21]–[23]. In this study, we are interested in a novel user interfaces aspect –that of the impact of placement of recommended content in physical contexts with augmented reality, on the metrics of accuracy and trust. We are also interested in how the interplay of AI performance (quality of the recommendation) with the choice of user interface influences these metrics. It is likely that user specific factors such as experience with visualizations, recommender systems, or multimodal display technology will impact the observed results. Nilashi et al. [24] performed a mixed-model evaluation of recommender system users on two real world e-commerce sites and analyzed the impact of many observed and latent factors on trust and purchase intention. Similar mixed model evaluations for recommenders were performed on a hybrid music prediction system by Knijnenburg et al. in [23] and in a system for analysis of commuter traffic data from microblogs by Schaffer et al. [22]. In this paper we also apply a mixed-model evaluation, designed to capture user-specific characteristics that impact our performance metrics.

III. SYSTEM ARCHITECTURE

To test our hypothesis, we implemented online shopping interfaces for each modality as well as a content recommendation system which generates a set of distinct high and low

quality recommendations based on user profile data. These recommendations are distributed evenly across the 3 modalities, where they are rendered using the Unity game engine. During the study, users interact with each modality and give ratings which are sent back to the server to be recorded.

A. Browser Interface

We implemented a simple e-commerce graphical interface in a web browser (see Figure 1). The interface shows the recommended object, the store logo, and generic text descriptions of the object. The browser interface is broken up into 2 presentation modalities: 2D browser and 3D browser. In the 2D browser modality, item recommendations are presented as a set of 2D pictures of the product taken from different angles. Users cycle through these pictures by clicking on the arrow buttons below the image. A pencil is shown in the images to provide a point of reference for scale. Users can rate the items by clicking on the radio buttons provided on the right side of the image. The 3D browser modality displays a 3D model of item recommendations that users can interact with. In this modality, users click and drag within the display window to rotate the object about its central X and Y axes. Users can provide ratings in the same way as the 2D browser. Note that both kinds of browser interactions take place on a traditional computer and monitor.

B. Augmented Reality Interface

For the AR interface, we use a Microsoft HoloLens device. Our application uses the devices' Spatial Mapping API to map the environment and situate virtual products and UI elements within the environment. We use HoloLens' World Anchor system to fix the recommended item and UI elements in the same position throughout the study. Users are able to walk around and look at the virtual items from different directions and provide feedback via the rating interface, presented through 2 panels as shown in Figure 1. The graphical interface is bare-bones, only displaying the store logo and generic text descriptions similar to the browser implementations. For interacting with the interface, we implemented a 3D cursor using a raycast formed by the user's head gaze direction. We define head gaze direction as the forward direction of the headset. Using the 3D cursor, users can aim and click on the rating panel. Although the HoloLens device supports hand gestures for clicking, we opted to use the HoloLens bluetooth clicker. This provides a fairer comparison to the browser interface, as there may be additional effects introduced by gesture-based interaction.

C. Content Based Recommendation

In order to generate personalized recommendations, we use an algorithm based on attribute Preference Elicitation (PE) and Multi-Attribute Utility Theory (MAUT). The item attributes considered by the recommender are *color*, *shape*, and *size*. We provide a validation for this choice of attributes in the Experimental Design section. For each attribute, we compute

TABLE I
ATTRIBUTES FOR ITEM CLASSIFICATION.

Attribute	High-level choice	Value
Shape	Non-cylindrical	-1
	Cylindrical	1
Size	Small	-1
	Large	1
Color	Disliked	-1
	Neutral	0
	Liked	1

the error between the recommender choice for that attribute, denoted as $recAttChoice$, and the user's preference, $userAttPref$.

If the attribute considered is a binary attribute (here, shape or size), let $recBinAtt \in \{-1, 1\}$ denote the recommender's choice for that attribute, where each possible value corresponds to a specific high level choice, as indicated in Table I. The reported user preference for that attribute, $userAttPref$, has values in $\llbracket 1, 5 \rrbracket$. In the pre-study questionnaire, values of 1 and 5 corresponded to a strong preference toward one of the possible values of the binary attribute, values of 2 and 4 to a slight preference, and 3 to no preference. The error can be computed from those two variables as:

$$Err_{recBinAtt} = L(recBinAtt, \text{sgn}(userAttPref - 3)) \cdot W_{recBinAtt} \quad (1)$$

where $L(\hat{y}, y)$ is the 0-1 binary loss function which equals 1 if $\hat{y} \neq y$ and 0 otherwise, and the weight is defined based on the importance given by the user to that particular attribute as $W_{recBinAtt} = (userAttPref - 3)^2$. Note that 3 is subtracted from the user's rating in order to turn the values ranging from 1 to 5 from the pre-study questionnaire into values in $\llbracket -2, 2 \rrbracket$. The sign function is then applied to map the user's rating to its corresponding value in Table I. This essentially decouples the user's preference into a binary choice and a weight.

Color was treated slightly differently: users were asked how much color weighed in their decision, and then asked to choose colors they liked and colors they disliked among 13 colors; colors not selected are considered neutral. The estimation of the error on a given color choice by the recommender $recColChoice$ therefore is:

$$Err_{recColChoice} = (\mathbb{1}_{DislikedCol}(recColChoice) + 0.5 \cdot \mathbb{1}_{NeutralCol}(recColChoice)) \cdot W_{col} \quad (2)$$

where $\mathbb{1}_A(x)$ is the indicator function on set A defined as 1 if $x \in A$ and 0 if $x \notin A$. The weight W_{col} has the same range of values as the weights used for the binary attributes, and the 0.5 factor for a neutral color ensures that picking a neutral color will yield an error superior to that of a liked color and inferior to that of a disliked color. The overall error is then obtain by summing the individual attribute errors:

$$Error = Err_{recColChoice} + Err_{recShapeChoice} + Err_{recSizeChoice} \quad (3)$$



Fig. 2. Example of recommendations for a user indicating preference for large, non-cylindrical mugs with navy, lime, cyan as liked colors and indigo, magenta, fuchsia as disliked colors.

It is worth noting that the errors can easily be turned into utility measurements by replacing L by $(1 - L)$ in Equation 1 and $\mathbb{1}_{DislikedCol}$ by $\mathbb{1}_{LikedCol}$ in Equation 2.

The personalized recommender system computes all the possible values of the total error based on the user weights, and stores for each value of the total error the incorrect attributes contributing to that value of the total error. There are $2^{\text{card}(\{W_{BinAtt}:W_{BinAtt} \neq 0\})}$ possible ways to get an error from potentially incorrect binary attributes that the user indicated having a preference for, and $3^{L(W_{col,0})}$ different possible values for the error from the color. Note that the error is degenerate; that is, different choices of incorrect attributes may yield the same value of the total error, which can be used to diversify the recommendations. The recommender system can then use a look-up table to show products in order of increasing error.

We define *high quality* recommendations as ones for which every attribute satisfies the user's preferences, and *low quality* recommendations as ones maximising the error, i.e. none of the attributes satisfy the user's preferences. Extreme values of the error were picked to avoid issues with parameter tuning in the recommender algorithm (e.g. power used in the calculation of the weights), which may depend on the granularity of the preference scales or the different possible interpretations of the scale labels by the users. In a longer sequence of interactions, or real-world deployment of the system, less strict parameters could be adopted to improve diversity and novelty of predicted items. An example of the two classes of recommendations are shown in Figure 2.

IV. EXPERIMENT DESIGN

Our main study had a 3 by 2 within subjects design with counterbalancing. The two independent variables were *UI modality* and *algorithm quality*, and the main dependant variables were item ratings (accuracy) and user trust in the recommender system. A preference profile was gathered from each participant in the experiment several days before the in-person study, via a Qualtrics² online survey. In this preference elicitation questionnaire, participants were asked for basic demographic information and experience with recommenders and AR/VR technology. They were also asked to select preferences for each of the classification dimensions for our domain items. These consisted of size (large or small), mug shape (cylindrical or non-cylindrical) and color preference. For color, participants were shown images of 13 coffee mugs of different colors and were asked to select their favorite 3 and

²<https://www.qualtrics.com/>

TABLE II
VALIDATION OF ITEM CLASSIFICATION FEATURES FROM 110 PARTICIPANTS IN AN ONLINE SURVEY.

	Mean	Std Dev
Color	57.84	25.49
Pattern	58.45	27.5
Size	71.27	23.15
Shape	60.75	26.17

least favorite 3. This information was stored on a server which computed sets of high and low quality recommendations for each user, based on the algorithm previously described.

To compare the effect of recommendation quality among the 3 different modalities, two different virtual retail stores were created: *4Buy* and *iMart*. *4Buy* always attempts to provide high quality recommendations and *iMart* always attempts to recommend items from the database that the user will dislike. Distinct logos for *4Buy* and *iMart* were visible in each modality (see Figure 1) to allow users to recognize which store they are in and form different perceptions of trust for each store.

A. Item-space Classification

As it is difficult to find free high-quality 3D models, we chose to modify the models on the fly to provide variance in recommendations. We began with a total of 18 different models, and applied transforms over size and color parameters to provide different virtual mugs for participants. The patterns on the mugs varied. To ensure that the pattern variable would not impact user preference more than the controlled features (size, shape and color), an MTurk study of 110 users was performed where each participant provided ratings between 1 to 100 for each of the four features. The mean and standard deviation for these ratings are found in Table II. We found no significant difference between pattern and the other features and so assume that manipulation of the other three features will be sufficient to provide good or bad recommendations based on the user profile. This is further confirmed in our results which show user ratings for good recommendations are significantly higher than those for bad ones. Different patterns in the mugs can contribute to *novelty* and *diversity* in recommendations, but overall quality can still be controlled in a meaningful way through manipulations on the other features.

B. Experimental Procedure

The experiment was conducted at an American university campus. Participants were assigned to particular orderings for each condition. Participants were given a brief introduction to the study by the experimenter. They were provided with a simple background story as follows: "You have just broken your coffee mug and are looking online to shop for a new one. You will shop at 2 different stores using a variety of their interfaces". For the AR condition, participants were given a training task where they had to observe several virtual items and use interaction in AR to provide feedback ratings. Once comfortable with the AR environment and rating procedure, they began the main rating phase. Here, they were shown a

sequence of 3 recommendations, either from iMart (low quality) or 4Buy (high quality). They were asked to walk around and inspect the items, and then provide a rating for how much they liked the recommended item on a scale of 1 to 5, and how much they trusted the system’s current ability to provide good recommendations on a scale of 1 to 10. There was no time limit imposed during the rating phase. Participants typically took less than 30 seconds to provide a rating, irregardless of the modality. Similar training steps were performed for the browser-based conditions. Participants complete all 3 conditions for a given store (and recommendation quality) first, before repeating the conditions in the same order for the other store. We alternated which store the participants start with. Once all conditions were complete, participants completed a post study questionnaire and were given a brief post study interview by the experimenter. In the post study questionnaire, participants were asked to rate their overall trust for each recommender, the helpfulness of the recommender, how much they liked interacting with each modality, and how much they liked each store overall. In the interview, participants were asked about their thoughts on the AR device, and whether they would choose to use it over the other modalities in a real world shopping scenario.

C. Novelty Effect

Since Augmented Reality is a new and emerging technology, and there is a “wow factor” with cutting edge devices such as the HoloLens, novelty effects will always be challenging to deal with. To mitigate novelty effects in the experiment, participants were allowed up to 10 minutes to familiarize themselves with each modality. In the AR condition, participants played with the built-in holograms application on the HoloLens device. Note that this familiarization period takes place before the training task begins. After the experiments, we compared performance between the participants who started with the AR condition versus those who started with the browser-based conditions. We ran paired t-tests on our key metrics but found no significant differences between the two groups, giving us confidence that our balancing and familiarization procedures were helpful in controlling novelty effects of the HoloLens device in the AR condition. This was further supported through post study interviews, during which participants reported that the familiarization period helped them to “get comfortable” using the AR headset.

V. RESULTS

To answer our research questions, we looked at user ratings for individual product recommendations and overall trust in the recommender system. We examined differences in ratings across each modality in order to assess relevant effects on user’s perception of recommendations. Additionally, we examine self-reported UX metrics from a post study questionnaire and verbal interview.

A. Participants

In total, 31 participants completed the in-person study. Data from 3 participants were removed due to being provided

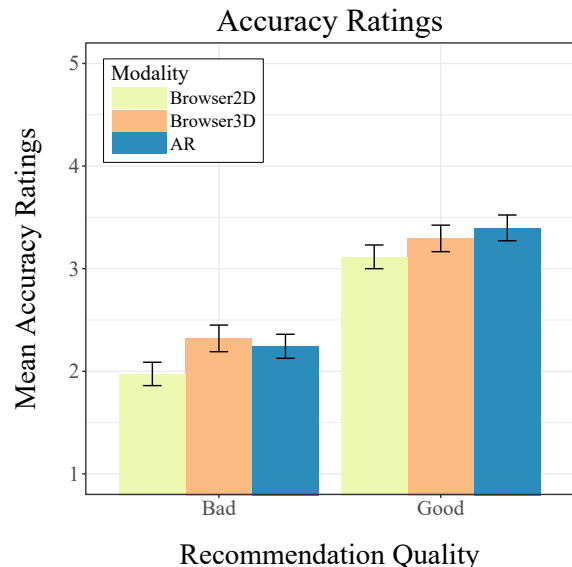


Fig. 3. Mean accuracy rating with standard error.

TABLE III
ACCURACY: PAIRWISE COMPARISON BETWEEN MODALITIES.

contrast	estimate	SE	df	t.ratio	p.value
Browser2D - Browser3D	-0.2641	0.1095	440.51	-2.413	0.0428
Browser2D - AR	-0.2791	0.1091	438.92	-2.557	0.0293
Browser3D - AR	-0.0150	0.1091	438.42	-0.138	0.9896

Results are averaged over the levels of: Recommendation Quality
P value adjustment: Tukey method for comparing a family of 3 estimates

incorrect instructions on the rating system. These participants misunderstood the task and rated other aspects such as the design of the logo. We also removed 2 additional participants due to system failure of the HoloLens during the experiment, leaving a total of 26 for analysis. Participants had a median age of 23, mean age of 27 with std. deviation of 9.58. 77% were male and 23% female. All had at least some college education. Participants were recruited through a user study pool at the university and were paid \$10 for the study, which lasted about 40 minutes.

B. Do users perceive product recommendations in AR differently than in a browser-based UI?

To begin our analysis, we looked at the average accuracy ratings within each condition. The resulting data is graphed in Figure 3. We tested for significance using paired t-tests.

For these ratings, our initial hypothesis was that increased reality and immersion provided by the AR modality would amplify users’ perception of recommendation accuracy. More realistic inspection methods might cause users to have a greater awareness of how well a product fits their preferences. Thus, we expected bad recommendations to be rated lower in AR compared to browser based methods, and likewise good recommendations would be rated higher in AR.

When looking at ratings in the bad recommender, we found a significant difference between the 2D modality ($\mu = 1.97$) and the 3D modality ($\mu = 2.32$) conditions; $p = 0.024$. There was almost significance between 2D modality and the AR modality ($\mu = 2.24$); $p = 0.064$. In the good recommender, we found significance between the 2D modality ($\mu = 3.12$) and the AR modality ($\mu = 3.4$); $p = 0.035$, but not between 2D and 3D modalities.

Additionally, we see a significant difference in ratings between the good and bad recommenders for all 3 modalities (all $p < 0.0005$). This gives us confidence that our recommendation algorithm is correctly providing high and low quality recommendations based on the user's preferences.

These results appear to reject our hypothesis. Irregardless of recommender quality, AR and 3D modalities seem to improve perception of recommendations. However, the signal does not appear consistently between bad and good recommenders. Thus, to look at the effects of each modality across both good and bad recommender conditions, we opted to perform further analysis using linear mixed effects models. Specifically, the modality type and recommendation quality are modeled as fixed effects, while participants and item design are modeled as random effects.

To validate this approach, we assessed the fit of our models using pseudo- R^2 values [25]. Marginal pseudo- R^2 was computed for fixed effects, and conditional pseudo- R^2 for random effects. For the accuracy model, the marginal pseudo- R^2 was 0.216 and the conditional pseudo- R^2 was 0.369. Additionally, mixed effects models assume that the residuals of the model are normally distributed. We plotted the residuals of each model as Q-Q plots to check this assumption and found that the residuals fall about a fairly straight line, suggesting normality. Finally, we created separate models where Modality and Recommendation Quality were modeled as having an interaction effect. We performed a likelihood ratio test against these to determine any significant interaction effects, but did not find any significant inter-dependence between them thus we did not include interaction effects in our models.

The full pairwise comparisons between each modality are shown in Table III. These tables describe the difference in ratings after averaging over the levels of recommendation quality and performing p-value adjustment using the Tukey method. Here, we can see a significant difference between Browser2D and the AR modalities ($p = 0.0293$), as well as between Browser2D and Browser3D ($p = 0.0428$). This provides further evidence that AR may improve user perception of product recommendations.

When comparing AR against the 3D interface, pairwise comparisons within our model did not show a significant difference in product rating. We believe that this result was due to a hidden variable created through differing levels of control in the interaction. In the 3D browser, users could rotate the items and view them from all angles. However, in the AR condition, the item was in a fixed position, and therefore could not be viewed from the bottom angle, since it was positioned on a table. During the verbal interview, 3 participants mentioned

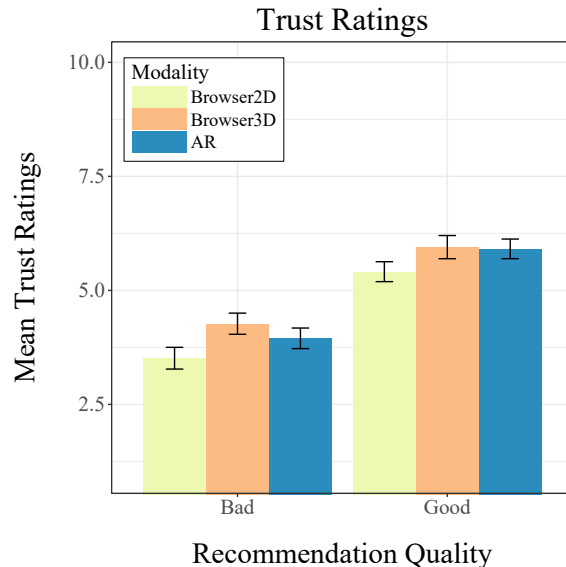


Fig. 4. Mean trust ratings with standard error.

TABLE IV
TRUST: PAIRWISE COMPARISON BETWEEN MODALITIES.

contrast	estimate	SE	df	t.ratio	p.value
Browser2D - Browser3D	-0.6474	0.1697	442	-3.815	0.0005
Browser2D - AR	-0.4679	0.1697	442	-2.757	0.0167
Browser3D - AR	0.1795	0.1697	442	1.058	0.5410

Results are averaged over the levels of: Recommendation Quality
P value adjustment: Tukey method for comparing a family of 3 estimates

they prefer the 3D view because it "allows you to see the mug in every possible orientation".

C. Are there differences in recommender system trust when presented in AR versus a browser-based UI?

This question focuses on the perception of algorithm quality within the different modalities. Similar to the product ratings, we hypothesized that the AR condition could help improve user awareness of a recommendation algorithm's performance, leading to lower ratings for the low quality recommender and higher ratings for the high quality recommender compared to the other modalities.

Our analysis on trust ratings mirrored the methods used for product ratings in the previous section. In Figure 4 you can see the graphed trust ratings. In the bad recommender, we found significant differences between 2D ($\mu = 3.51$) and 3D ($\mu = 4.27$); $p < 0.001$, as well as 2D and AR ($\mu = 3.95$); $p = 0.034$. In the good recommender, we also found significance between 2D ($\mu = 5.41$) and 3D ($\mu = 5.95$); $p = 0.005$, and also between 2D and AR ($\mu = 5.91$); $p = 0.008$. Again, we see a significant difference between the good and bad recommenders for all 3 modalities (all $p < 0.0005$). Figure 4 clearly show that users perceived a difference between good and bad algorithms in all conditions. For example, participants in the 2D browser condition rated trust in the iMart (low quality recommender

algorithm) at 3.51 and 4Buy at 5.41, which is a relative improvement of 54% over the iMart algorithm.

We again used linear mixed models to analyze trust ratings across recommendation quality. We performed the same steps to validate the model as in the previous section. For the trust model, marginal pseudo- R^2 was 0.184 and conditional pseudo- R^2 was 0.555. Table IV is the resulting pairwise comparisons. In particular, we highlight the differences between Browser2D and the AR modalities which are significant for trust ratings ($p = 0.0167$). Ultimately, the results we found did not support our hypothesis. Instead, our results suggest that Trust is improved for the AR and 3D modalities, despite the differences in recommender quality.

Additionally, we looked at trends over time within each modality to determine if users' trust ratings improved with increased interaction. We found that in the AR modality, trust exhibited a small positive slope upward from the first to third interaction with a recommendation, whereas the Browser2D and Browser3D both exhibited downward slopes. However, these slopes did not prove to be significant.

D. What is the general sentiment towards an AR recommender system for in-home shopping?

Our primary source of analysis for this research question are through a post-study questionnaire and semi-structured verbal interview conducted immediately after the experiment.

The results of the post questionnaire are shown in Figure 5. The leftmost plot shows the perceived trust in the system's recommendations broken down for each of the 6 conditions. Here the browser-based conditions are abbreviated to B2D and B3D, and algorithm quality is represented as HQ or LQ for high and low respectively. The first point to note is that the questionnaire responses for trust in the recommendations align well with the observed ratings during the experiment, with both AR-HQ and B3D-HQ showing a significant rating improvement of about 20% over the traditional UI B2D-HQ. There was no significant difference between the B3D and AR conditions. However, our post study interviews revealed that people either had a strong preference for the 3D-browser condition or the AR condition. Those who strongly preferred AR, tended to mention the value of being able to see the item in real-world context (situated recommendations), while those who preferred the 3D browser version tended to like the familiarity of the interface for shopping. Perceived helpfulness of the stores was also evaluated and showed a similar trend to trust, with AR and B3D having significant rating improvement over B2D for both recommender algorithms (LQ and HQ). However, the differences between recommender quality (LQ and HQ), was not as pronounced as it was on the trust metric. We believe this is an indication that users were considering other aspects than recommendation quality for their decisions on helpfulness, such as the quality of the UI design. Figure 5 also shows results for perceived quality of interaction with the system. As expected, both AR and B3D received very positive ratings. This is consistent with our interview feedback where participants preferred B3D almost as much as the AR condition due to better

inspection capabilities. Participants were also asked to rate each store overall. Here we see that participants did perceive the difference in algorithm quality across the two stores. The store with high quality recommendations (HQ) showed a 50% improvement over the LQ store. This was consistent with our observed ratings-based results.

In the verbal interviews, participants were highly positive, with all but 5 of the participants reporting that they would choose to use the AR system if it were available to them. Out of many different reasons cited, the most common was the 'try before you buy' reason –to visualize and interact with the item in the context where it is to be used. The participants who did not want to use it argued that the interaction was not sufficient and that the 3D browser version allowed for a better inspection of the item. Additionally, participants reported feelings of frustration and discomfort due to the limited field of view, weight of the device, and issues with color quality of the AR display.

VI. DISCUSSION

The results from our study contribute to an emerging body of work focused on understanding user perception of AR with predictive AI systems such as recommender systems. Throughout the study, AR and Browser3D modalities performed on par with each other, whereas both tended to improve ratings and other metrics compared to Browser2D. Participants generally fell into 2 camps, those preferring Browser3D and those preferring AR.

Many of the verbal interview responses seem to indicate that participants appreciate qualities from both mediums. In the case of AR, participants enjoy being able to visualize a product in a real world context and grasp the actual scale of the object. However, AR is marred by issues with a low quality display and headset discomfort. 3D on the other hand is quick and easy to use, and still allows users to view recommended products from a variety of viewing angles.

While some of these problems will be solved in future iterations of AR devices, it's important to understand what the role of interaction should be moving forward. It's clear that users are accustomed to browser based interaction methods. For many shopping experiences, they may prefer it over an AR experience. However, AR has potential to excel when delivering recommendations that have great impact on daily life, or where scale and contextual information is important, such as home appliances and interior design. These qualities should be emphasized and communicated when designing for the future of AR driven recommender systems.

VII. CONCLUSION

This paper presented a study that to our knowledge is the first empirical analysis of the effects of Augmented Reality interfaces on the perception of recommender systems. A 3 by 2 within subjects experiment assessed user perception of high and low quality personalized recommendations in three modalities: Augmented Reality with recommended items placed in a real world scene, web browser with 2D images, and




Fig. 5. Mean subjective ratings from the post study questionnaire with standard error. Participants were asked to rate how much they trust the recommendations, how helpful each store's interface was, the interaction quality of each modality, and overall preference for each store. Brackets show the level of significance between particular values (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$). Additionally, there was significance ($p < 0.01$) in recommendation trust between each HQ modality and their LQ counterparts.

web browser with 3D interaction. Quantitative metrics for product ratings and recommender trust were assessed, along with perception of the system through a post study questionnaire and verbal interview. Results of our main research questions show that overall product ratings for recommended objects, and trust in the recommender, are significantly higher in AR and interactive 3D than in a traditional browser UI. However, there is no significant difference in either metric between interactive 3D and AR modalities. Furthermore, people perceive differences between high and low quality algorithms in all 3 modalities, but there is no significant trend that suggests better awareness of quality differences in AR. Finally, a majority of participants preferred to use AR over browser based interfaces for product recommendations, finding it helpful for visualizing in the context where it will be used.

ACKNOWLEDGMENT

Parts of this research were supported by ONR grant #N00014-16-1-3002 and NSF grant IIS-1423676.

REFERENCES

- [1] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl, "GroupLens: An open architecture for collaborative filtering of netnews," in *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work*, ser. CSCW '94. New York, NY, USA: ACM, 1994, pp. 175–186. [Online]. Available: <http://doi.acm.org/10.1145/192844.192905>
- [2] J. L. Herlocker, J. A. Konstan, and J. Riedl, "Explaining collaborative filtering recommendations," in *ACM conference on Computer supported cooperative work*, 2000, pp. 241–250.
- [3] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, "Item-based collaborative filtering recommendation algorithms," in *Proceedings of the 10th International Conference on World Wide Web*, ser. WWW '01. New York, NY, USA: ACM, 2001, pp. 285–295. [Online]. Available: <http://doi.acm.org/10.1145/371920.372071>
- [4] V. Tamturk, *The ROI of Recommendation Engines*, 2017 (accessed April 1 2017), <http://bit.ly/2nW2aUz>.
- [5] *Is Apple Looking to Expand in the Augmented Reality World?*, Zacks Equity Research, 2017 (accessed April 1 2017), <http://bit.ly/2oTjwBv>.
- [6] S. Erickson, *Microsoft HoloLens and Lowe's, working to redefine your next home renovation*, 2017 (accessed April 2 2017), <http://bit.ly/2owh1ju>.
- [7] D. Williams, *3 Retail Giants Who Used Augmented Reality to Sell*, 2017 (accessed April 2 2017), <http://bit.ly/1Ufnaeq>.
- [8] K. N. Jr., *Nike Store in Paris lets customers test sneaker colors using augmented reality*, 2017 (accessed April 2 2017), <http://bit.ly/2jvcuSe>.
- [9] J. Stoyanova, R. Goncalves, A. Coelho, and P. Brito, "Real-time augmented reality shopping platform for studying consumer cognitive experiences," in *2013 2nd Experiment@ International Conference (exp.at'13)*. IEEE, sep 2013, pp. 194–195. [Online]. Available: <http://ieeexplore.ieee.org/document/6703069/>
- [10] Y. Lu and S. Smith, "Augmented Reality E-Commerce System: A Case Study," *Journal of Computing and Information Science in Engineering*, vol. 10, no. 2, p. 021005, 2010. [Online]. Available: <http://computingengineering.asmedigitalcollection.asme.org/article.aspx?articleid=1401892>
- [11] T. Olsson, E. Lagerstam, T. Kärkkäinen, and K. Väänänen-Vainio-Mattila, "Expected user experience of mobile augmented reality services: a user study in the context of shopping centres," *Personal and Ubiquitous Computing*, vol. 17, no. 2, pp. 287–304, feb 2013. [Online]. Available: <http://link.springer.com/10.1007/s00779-011-0494-x>
- [12] B. Wang, M. Ester, J. Bu, and D. Cai, "Who also likes it? generating the most persuasive social explanations in recommender systems," in *AAAI Conference on Artificial Intelligence*, 2014.
- [13] M. Balduini, I. Celino, D. Dell'Aglio, E. D. Valle, Y. Huang, T. Lee, S.-H. Kim, and V. Tresp, "BOTTARI: An augmented reality mobile application to deliver personalized and location-based recommendations by continuous analysis of social media streams," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 16, pp. 33–41, 2012. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S157082681200073X>
- [14] K. McCarthy, J. Reilly, L. McGinty, and B. Smyth, "Experiments in dynamic critiquing," in *Proceedings of the 10th International Conference on Intelligent User Interfaces*, ser. IUI '05. New York, NY, USA: ACM, 2005, pp. 175–182. [Online]. Available: <http://doi.acm.org/10.1145/1040830.1040871>
- [15] N. Tintarev and J. Masthoff, "Effective explanations of recommendations: User-centered design," in *Proceedings of the 2007 ACM Conference on Recommender Systems*, ser. RecSys '07. New York, NY, USA: ACM, 2007, pp. 153–156. [Online]. Available: <http://doi.acm.org/10.1145/1297231.1297259>
- [16] N. Tintarev and R. Kutlak, "Explanations - making plans scrutible with argumentation and natural language generation," in *Intelligent User Interfaces (demo track)*, 2014.
- [17] J. O'Donovan, B. Smyth, B. Gretarsson, S. Bostandjiev, and T. Höllerer, "Peerchooser: Visual interactive recommendation," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '08. New York, NY, USA: ACM, 2008, pp. 1085–1088. [Online]. Available: <http://doi.acm.org/10.1145/1357054.1357222>
- [18] K. Verbert, D. Parra, P. Brusilovsky, and E. Duval, "Visualizing recommendations to support exploration, transparency and controllability," in *International Conference on Intelligent User Interfaces*, ser. IUI '13, 2013, pp. 351–362.
- [19] D. Parra, P. Brusilovsky, and C. Trattner, "See what you want to see: Visual user-driven approach for hybrid recommendation," in *Proceedings of the 19th international conference on Intelligent*

- User Interfaces*. ACM, 2014, pp. 235–240. [Online]. Available: <http://web.ing.puc.cl/~dparra/pdfs/UI2014Setfusion.pdf>
- [20] S. Bostandjiev, J. O'Donovan, and T. Höllerer, "Linkedvis: exploring social and semantic career recommendations," in *18th International Conference on Intelligent User Interfaces, IUI 2013, Santa Monica, CA, USA, March 19-22, 2013*, J. Kim, J. Nichols, and P. A. Szekely, Eds. ACM, 2013, pp. 107–116. [Online]. Available: <http://doi.acm.org/10.1145/2449396>
- [21] J. L. Herlocker, J. A. Konstan, L. Terveen, and J. T. Riedl, "Evaluating collaborative filtering recommender systems," *ACM Trans. Inf. Syst.*, vol. 22, no. 1, pp. 5–53, 2004.
- [22] J. Schaffer, P. Giridhar, D. Jones, T. Höllerer, T. Abdelzaher, and J. O'Donovan, "Getting the message?: A study of explanation interfaces for microblog data analysis," in *Intelligent User Interfaces*, ser. IUI '15, 2015, pp. 345–356. [Online]. Available: <http://doi.acm.org/10.1145/2678025.2701406>
- [23] B. P. Knijnenburg, S. Bostandjiev, J. O'Donovan, and A. Kobsa, "Inspectability and control in social recommenders," in *Conference on Recommender Systems*, ser. RecSys '12, 2012, pp. 43–50. [Online]. Available: <http://doi.acm.org/10.1145/2365952.2365966>
- [24] M. Nilashi, D. Jannach, O. B. Ibrahim, M. D. Esfahani, and H. Ahmadi, "Recommendation quality, transparency, and website quality for trust-building in recommendation agents," *Electron. Commer. Rec. Appl.*, vol. 19, no. C, pp. 70–84, Sep. 2016. [Online]. Available: <https://doi.org/10.1016/j.elerap.2016.09.003>
- [25] S. Nakagawa and H. Schielzeth, "A general and simple method for obtaining r^2 from generalized linear mixed-effects models," *Methods in Ecology and Evolution*, vol. 4, no. 2, pp. 133–142, 2013.