

A Large-scale Exploration of Group Viewing Patterns

Allison J.B. Chaney*
Princeton University
achaney@cs.princeton.edu

Mike Gartrell*
University of Colorado Boulder
mike.gartrell@colorado.edu

Jake M. Hofman*
Microsoft Research
jmh@microsoft.com

John Guiver
joguiver@microsoft.com

Noam Koenigstein
noamko@microsoft.com

Pushmeet Kohli
pkohli@microsoft.com

Ulrich Paquet
ulripa@microsoft.com

ABSTRACT

We present a large-scale study of television viewing habits, focusing on how individuals adapt their preferences when consuming content with others. While there has been a great deal of research on modeling individual preferences, there has been considerably less work studying the preferences of groups, due mostly to the difficulty of collecting group data. In contrast to most past work that has relied either on small-scale surveys, prototypes, or a relatively limited amount of group preference data, we explore more than 4 million logged household views paired with individual-level demographic and co-viewing information. Our analysis reveals how engagement in group viewing varies by viewer and content type, and how viewing patterns shift across various group contexts. Furthermore, we leverage this large-scale dataset to directly estimate how individual preferences are combined in group settings, finding subtle deviations from traditional models of preference aggregation. We present a simple model which captures these effects and discuss the impact of these findings on the design of group recommendation systems.

ACM Classification Keywords

H.5.3. Group and Organization Interfaces: Collaborative Computing; H.3.3. Information Search and Retrieval: Information Filtering

Author Keywords

Group recommendation; group viewing patterns.

INTRODUCTION

We are in the midst of an industry-wide shift, wherein the primary means of home broadcast video entertainment is moving from traditional television sets to online and Web services (e.g., Netflix, Hulu, and Xbox) that contain a rapidly expanding catalogue of content. While there is a substantial body of work on understanding the preferences of individuals in these settings—largely for the purpose of aiding users in discovering relevant and novel content within these catalogues—there is a comparatively small amount of research on modeling

group viewing habits, mostly owing to the difficulty of collecting co-viewing data. Absent this data, group preferences are often modeled via simple aggregates of the underlying individual preferences. While such approaches are somewhat successful, they obscure more subtle group dynamics and interactions that affect group decision making—for instance, the preferences of a parent and child together may be difficult to determine from what each watches alone.

As reviewed below, previous studies often rely on small-scale, self-reported viewing data to draw qualitative conclusions about group viewing, and most existing large-scale log datasets contain group preference data for only several hundred groups [19, 18]. In contrast, we use a dataset that contains both individual and group viewing patterns from a representative panel of more than 50 million U.S. viewers—in over 50,000 groups—automatically recorded by Nielsen¹. Hence, our work presents one of the first attempts at understanding the relationship between viewing patterns of groups and their constituent individuals from direct, logged data at scale. Our findings indicate that group context substantially impacts viewer activity and that knowledge of the group’s composition can be informative in determining group interests.

Our study makes three key contributions: first, we provide a large-scale analysis of viewing patterns with an emphasis on differences between groups and individuals; we break down what users watch alone, how often they engage in group viewing, and how their preferences change in these contexts. Second, we analyze how individual preferences are combined in group settings. Finally, we propose an approach to group recommendations based on the demographic information of the group’s constituent individuals. By capturing interactions between the constituents’ preferences, our approach predicts group preferences more accurately than existing group recommendation methods. This calls for more sophisticated non-linear aggregation functions that can better estimate the interplay between individuals within a group.

RELATED WORK

Historic TV Viewing Studies. In the early eighties, Webster and Wakshlag [23] analyzed viewing patterns of groups and individuals. Groups that did not change their composition over time showed more program-type loyalty, similar to individual users. Group composition was not considered, however, and to our knowledge, this has remained unstudied.

Most historic studies of viewing behavior rely on self-reported diary surveys [8, 23] that have a few hundred respon-

*Authors contributed equally.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
TVX 2014, June 25–27, 2014, Newcastle Upon Tyne, UK.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2838-8/14/06 ...\$15.00.
<http://dx.doi.org/10.1145/2602299.2602309>

¹www.nielsen.com




Figure 1. (a) Cumulative distribution of user activity split by individual and group views. (b) Cumulative distribution of telecast popularity by number of viewers. (c) Number of views by group size.

dents, most likely resulting in subject selection biases. Later studies [14] show that television viewing behavior is affected by the demographic characteristics of viewers.

Recommendation Systems for Groups. The problem of group recommendation has been investigated in a number of works [1, 4, 5, 11, 16, 20, 22, 24]. Various techniques target different types of items (e.g., movies, TV programs, music) and groups (e.g., family, friends, dynamic social groups).

Most group recommendation techniques consider the preferences of individuals and propose different strategies to either combine the user profiles into a single group profile, and make recommendations for the pseudo user, or generate recommendation lists for individual group members and then merge these lists for the group. Jameson and Smyth’s three main strategies for merging individual recommendations are *average satisfaction*, *least misery*, and *maximum satisfaction* [11]. These form the bedrock of group recommendations [1, 5, 13] and we refer to them as “preference aggregation functions.” Average satisfaction assigns equal importance to each group member and is used in several group recommendation systems [4, 24, 25]; both average satisfaction and least misery are reasonable candidates for group decisions [13]. Different user weights, dissimilarity among group members, and social connections are also used in aggregation models [3, 1, 6]. If the group is guaranteed to remain static, the dynamic aspect of groups can be ignored [22]. All of this work involves relatively small-scale studies or prototypes; related research on group recommendations relies on synthetically generated data from the MovieLens dataset [2, 12, 15].

Smaller practical recommender systems include PolyLens, a group-based movie recommender that targets small, private, and persistent groups [16] and considers the nature of groups, rights of group members, and social value functions. PartyVote provides a democratic mechanism for selecting and playing music at social events [20].

Larger group preference datasets are beginning to emerge. The 2011 Challenge on Context-Aware Movie Recommendation (CAMRa 2011) used a dataset from the Moviepilot Web site consisting of about 170,000 users, 24,000 movies, and 4.4 million ratings [18]. This dataset also provides household membership identifiers, but this “group” component is substantially smaller: it accompanies a user’s rating for only 290 households. Many group recommendation approaches

have been proposed and evaluated using this dataset [7, 9, 10]. Similarly, a large-scale dataset from the BARB organization is used in [19], which consists of about 15,000 users, 6,400 households, and 30 million TV program views. However, only 136 of these households are used in [19], since the rest lack sufficient group activity.

In contrast to prior work, our work uses a dataset containing hundreds of thousands of implicit group preferences, along with substantial metadata for individuals, households, and programs. This dataset has been actively recorded, and contains detailed demographic information for a large representative sample of viewers. We present further details below.

DATASET

The Nielsen Company maintains a panel of U.S. households and collects TV viewing data through both electronic metering and paper diaries. In the month of June 2012, Nielsen recorded 4,331,851 program views by 75,329 users via their electronic People Meter system, which records both what content is being broadcast and who is consuming that content. We restrict this dataset to events where at least half of the program was viewed², resulting in a collection of 1,093,161 program views by 50,200 users. These views are comprised of 2,417 shows with 16,546 unique telecasts (e.g., individual series episodes, sports events, and movie broadcasts). Each program is associated with one of 34 genres and other metadata, including the distributor and optional sub-genre.

Users also have associated metadata, including age and gender, and are assigned to households, allowing a simple heuristic for identifying group viewing activity. We define a group view as one where at least two members of a household each watch at least half of the same telecast on the same day. There are 279,546 such group views in our dataset. When a user watches the majority of a telecast alone, we define this an individual view; 813,615 individual views are present. Due to the large number of views all viewing pattern results presented later in this paper are statistically significant.

The number of programs watched by users exhibit a heavy-tailed distribution, with many users viewing only a handful of telecasts while a few heavy users consume substantially more content. Figure 1a shows that roughly half of all users have viewed at least 5 telecasts individually; likewise, another

²This 50% threshold simplifies our analysis so that at most one telecast can be viewed by each user in a given time slot.

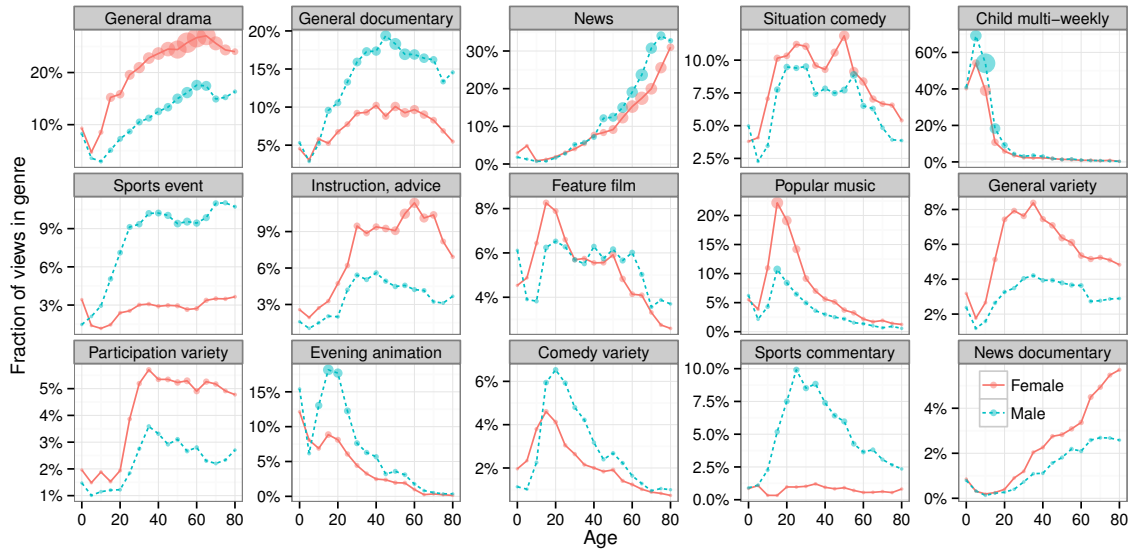


Figure 2. Distribution of views across genres by age and gender.

(probably overlapping) half of users have viewed at least 5 telecasts in a group. Similarly, most programs are watched relatively infrequently, with a few being very popular. For example, Figure 1b shows that less than 10% of telecasts have been viewed by at least 100 different users. We note that telecast popularity is slightly higher in group settings because each individual in a group view is counted separately here, so that a show watched by a pair of individuals is counted as two views for that broadcast. Finally, as shown in Figure 1c, upwards of 80% of co-viewing occurs in groups of size two, with larger groups occurring substantially less frequently. Most (78%) of couple views are by two adults, with 86% of such groups comprised of one male and one female.

INDIVIDUAL VIEWING PATTERNS

In this section, we analyze how individual viewing behavior varies with age and gender. For this purpose, we compute the genre-specific view counts in the context of demographic information. Figure 2 depicts how users of varying age and gender distribute their attention across genres at the aggregate level. Panels are ordered by decreasing overall genre popularity, and point size shows the relative fraction of overall views accounted for by each demographic group in each genre.

We observe strong age effects for the viewing of certain genres like general drama, child multi-weekly, evening animation, news, popular music, general variety and news documentary. For instance, we observe that older viewers spend a large fraction (about 20-30%) of their time watching news relative to teenagers, who consume little of this genre and devote substantially more of their attention to popular music shows. Likewise, general documentaries are more popular with adults and seniors than with children, while child multi-weekly programs are popular for children and much less popular with adults and seniors, as one would expect. General dramas are quite popular for every age and gender demographic we examined.

We also see gender differences in individual preferences, with females spending more of their time watching talk shows, drama, and music relative to males, who prefer animation, documentaries, and sports. Sports events tend to be more popular with males than with females, across all ages.




Figure 4. Fraction of views within a group by age and gender.

GROUP VIEWING PATTERNS

Having briefly investigated individual viewing activity, we turn to the main analysis of the paper and analyze group viewing patterns. We examine engagement in group viewing by group and program type, how groups of various types distribute their viewing time, and how individuals modify their viewing habits in group contexts.

Group Engagement

As noted above, roughly a quarter of all views in our dataset occur in groups of size two or larger, comprising a sizable fraction of total activity. To gain further insight into the composition of groups, Figure 4 shows the relative amount of group viewing by users of different ages and gender. The solid lines indicate the median fraction of group views for the specified demographic, with the top and bottom of the surrounding ribbon showing the upper and lower quartiles,




Figure 3. Fraction of views within a group by genre.

respectively. We see that younger users spend the majority (~75%) of their time viewing in groups compared to older viewers. Viewers in their 20s and 30s spend roughly equal amounts of time viewing alone and in groups, whereas older viewers generally spend slightly more time watching individually. We observe small gender effects for younger individuals and larger gender effects for older individuals, with younger females and older males displaying a higher rate of group views relative to their counterparts.

Next we investigate the type of content consumed by these groups. As shown in Figure 3, the relative fraction of group viewing varies significantly by genre. While more than a third of views on quiz shows, drama, and sports events are within groups, only 20% of music, news, and politics views occur in groups settings. We note that many of the genres that are likely to be viewed by groups comprise a relatively small fraction of total activity, as indicated by point size. For example, while upwards of a third of all award ceremony views are in groups, there are relatively few such views overall.

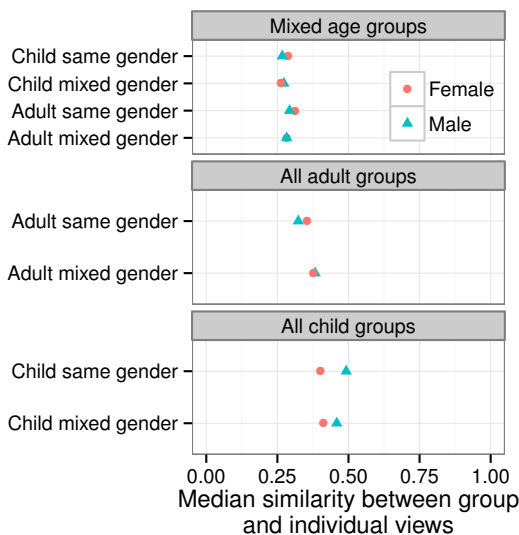


Figure 5. Similarity between group and individual viewing distributions.

Individual vs. Group Viewing

With this understanding of group engagement, we turn our attention to how individual viewing habits change in group settings. To do so, we compute viewing profiles for each

user in the dataset under various group contexts and compare their individual and group profiles. Specifically, we characterize each user as either an adult or child (over/under 18, respectively) and male or female; likewise, we categorize each group view by its gender (all male/mixed gender/all female) and age (all adult/mixed gender/all child) breakdowns. For each user, we compute the fraction of time they spend viewing each genre alone and in each of these nine possible group types. We then quantify the similarity between each user’s individual and group view profiles using Hellinger distance, a metric over probability distributions.³ Finally, we aggregate by user and group type and report the median similarity across users in each demographic when viewing in each group setting, as shown in Figure 5. From this plot we see that the similarity between individual and group viewing patterns varies substantially with the age composition of groups and less so with gender breakdown. For example, the bottom panel shows that activity by groups of all children looks most similar to views by individual children, compared to the mixed age groups in the top panel, which display the largest deviations from what members of those groups watch individually. Thus, the younger and more homogeneous the group, the higher the similarity between group and individual views.

For more details on how preferences shift in individual and group settings, Figures 6 and 7 show how attention is redistributed across genres with different age and gender audience compositions, respectively. For example, Figure 6 reveals that feature films are more popular among mixed age groups than they are either for individuals or groups of the same age. Likewise, we see that children devote substantially more of their time to child multi-weekly shows when viewing in groups (~50%) compared to viewing alone (~30%). Adults watch more dramas, documentaries, and sports events in groups with other adults, and are more likely to watch news, sports commentary, and advice shows alone. We also see that adults and children both compromise on certain genres: one group watching more than usual and the other watching less. This occurs for many genres, including dramas and documentaries, where adults watch less than usual and children watch more, as well as popular music and evening animation, where children watch less and adult watch more together than they do separately. We see little compromise for adults on sports events and participation shows, possibly due

³Hellinger distance is normalized to fall between 0 and 1; we measure similarity by the complement of Hellinger distance.

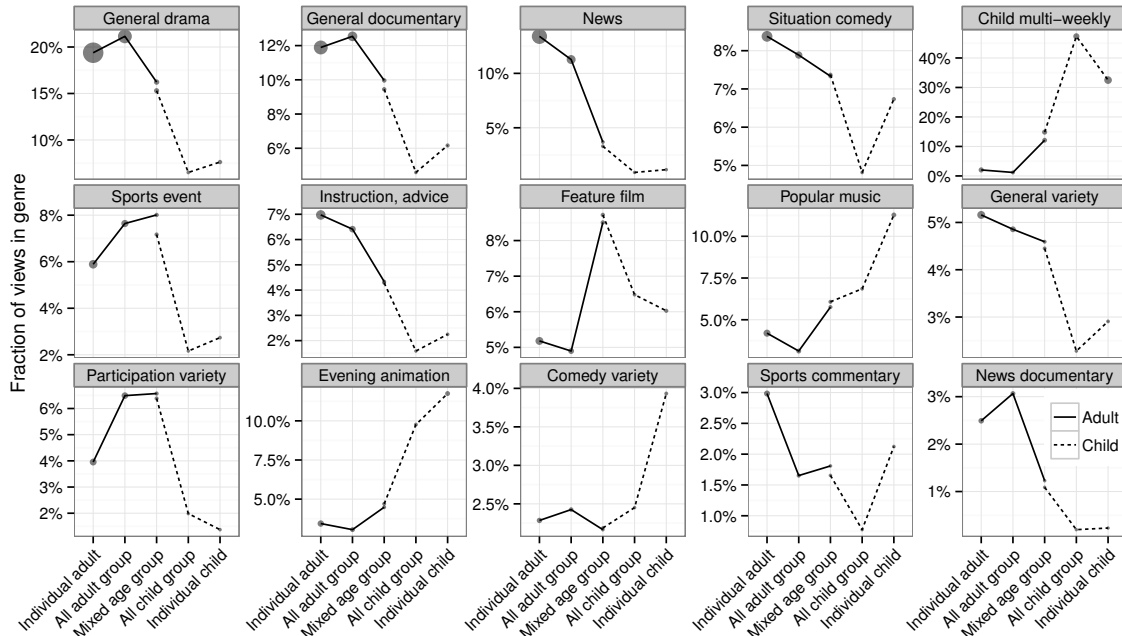


Figure 6. Distribution of views by genre for adults and children in different group contexts.

to time sensitivity; in both of these cases, adults watch just as much as they do in groups with other adults, and children watch far more than they otherwise would.

We also see substantial shifts in preferences as gender composition varies in Figure 7. For instance, feature films are more popular with same gender groups than they are with either individuals or mixed gender groups, whereas the opposite effect is seen for news, which is more popular amongst individual males and females than in groups. We also see that news is more popular in mixed gender groups than in same-gender groups. We speculate that this effect is attributed to passive viewing patterns of couples in the same household, rather than an active desire to watch news as a group. While these changes are fairly similar between men and women, we note that other genres show gender-specific effects. For example, groups of men spend nearly double the amount of their time watching sports compared to individual males, but no such difference is seen for females. Likewise, all female groups spend substantially more of their time viewing popular music shows than do individual females. Finally, as with age effects, mixed gender groups appear to compromise on many categories. For dramas, advice, and sitcoms, men watch more and women watch less together than when in homogeneous groups. We see the reverse effect for documentaries, evening animation, and sports shows, with women watching more and men watching less.

GROUP RECOMMENDATIONS

The previous section explores the differences between a group’s preferences and those of its individual constituents. While these effects are large at the aggregate level, both groups and individuals have substantial variability in their tastes, which can make modeling any particular group’s

preferences difficult. We investigate this problem in more detail—namely, assuming that we know what the members of a group like individually, how do we aggregate their preferences to predict what the group will view?

We approach this problem in two steps. First, we fit a matrix factorization model to approximate individual preferences, which demonstrates good empirical results in predicting individual views. Next, we evaluate popular baseline methods for aggregating each individual’s modeled preferences to predict group activity. We find that three of the traditional aggregation methods fail to capture subtle non-linearities and interactions between individual preferences, which we are able to estimate directly from our large-scale dataset. We propose a relatively simple model to account for these features that provides further insight into group decision making.

Modeling Individuals

To examine how to best combine preferences of individuals in a group, we first need a means of determining each individual user’s interest in each telecast in our dataset. We use the Matchbox recommendation system [21] without features, which fits a matrix factorization model to user’s individual viewing activity to approximate these preferences.

Fitting such a model requires information about both “positive examples”—the telecasts that a given individual viewed—and “negative examples”—telecasts that were available to individuals but not consumed. Unfortunately our dataset lacks negative examples, so we approximate this set as follows: for each telecast viewed by an individual, we consider all simultaneously broadcast telecasts on all channels in a user’s viewing history as potential negative examples. This results in roughly 16 negative examples for every positive example across the dataset. To keep a balanced number of posi-

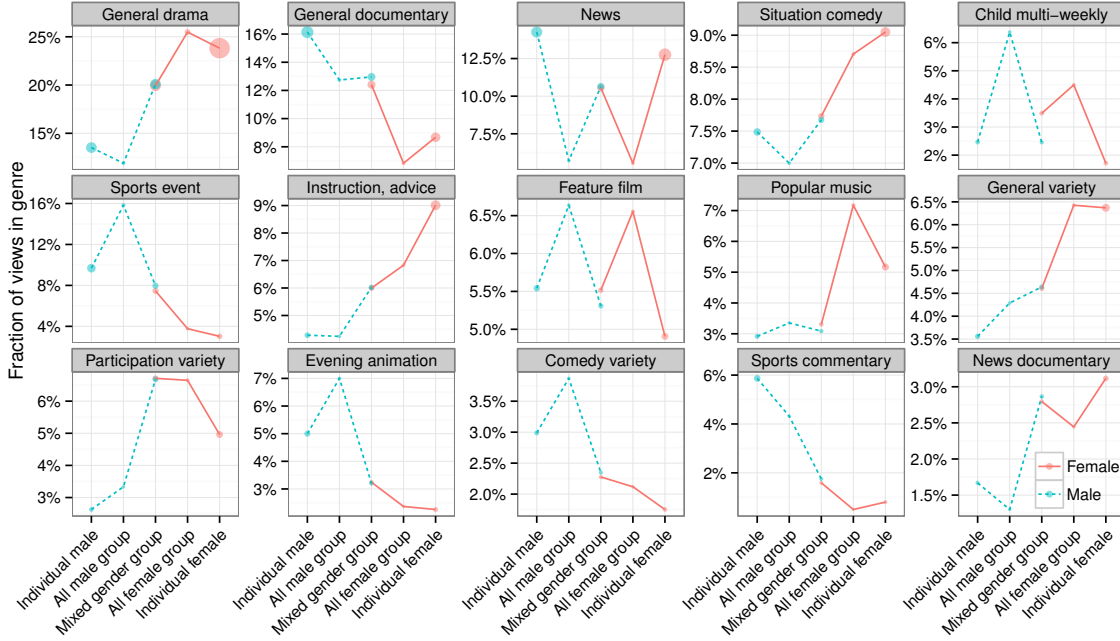


Figure 7. Distribution of views by genre for men and women in different group contexts.

tive and negative examples in our training set, we sample one negative example for each positive one, weighting telecasts by overall popularity [17].

We train Matchbox using this dataset with $K = 20$ latent trait dimensions on a randomly selected training set composed of 80% of the individual view data set, with the remaining 20% of individual views used for the test set. We set the user threshold prior and noise variances to 0, assuming a time-invariant threshold and a binary likelihood function. We place flexible priors on users and items by setting the user trait variance and item trait variance hyperparameters to $\frac{1}{\sqrt{K}}$, and the user bias variance and item bias variance hyperparameters to 1. The best-fit individual model found by Matchbox has an AUC of 88.3% on the held-out test set. Given this performance, we consider the model to be a reliable approximation to individual preferences and next investigate the group recommendation problem.

Preference Aggregation

As noted in our overview of related work, there are many approaches to aggregating individual preferences. Here we investigate three of the simplest, which are commonly used: least misery, average satisfaction, and maximum satisfaction. Denoting individual preference that user u has for item i by p_{ui} , these methods predict group preferences as follows:

$$\begin{aligned} \text{least misery} &: \min_{u \in G} p_{ui} \\ \text{average satisfaction} &: \frac{1}{|G|} \sum_{u \in G} p_{ui} \\ \text{max satisfaction} &: \max_{u \in G} p_{ui}. \end{aligned}$$

Least misery aims to minimize dissatisfaction of the least satisfied individual, maximum satisfaction to maximize enjoyment of the most satisfied, and average satisfaction takes an equal vote amongst all members.

After learning individual preferences with Matchbox, we evaluate each of these aggregation methods on all group views in our dataset. We find a strict ordering in terms of performance, with maximum satisfaction slightly outperforming average satisfaction, and both dominating least misery, across and within all group types. We find an overall AUC of 83.0% for maximum satisfaction, 82.6% for average satisfaction, and 79.7% for least misery. In further examining the quality of group predictions by group type, we see that mixed age and mixed gender group views are the most difficult to predict, with an AUC of 81.3%. Likewise, groups of all children are easiest to model, with performance on all male groups being considerably higher compared to all female groups (AUCs of 89.7% and 84.1%, respectively). Note that these results are obtained with maximum satisfaction and are largely consistent with the individual-to-group similarity comparison in Figure 5.

While some work on preference aggregation has been constrained to these relatively simple functions over individual preferences, our large-scale dataset of hundreds of thousands of group views enables us to conduct a direct examination of group preference landscapes. For simplicity, we limit this analysis to groups of only two members (which comprise 80% of all group views). For each group viewing event in our dataset, we bin the individual predicted probability for each member of the group to the nearest ten percent and aggregate views to examine the empirical probability of a group view within each bin. Panel 3 of Figure 8 shows the result for adult mixed gender couples, with the binned male’s and female’s preference on the x- and y-axis, respectively, and the probability of a group view on the z-axis. The predicted landscape for average satisfaction and maximum satisfaction are shown in the first two panels for comparison, from which it




Figure 8. Modeled and actual probability of group viewing as a function of individual viewing for 2-person, mixed-gender adult couples.

is clear that these traditional aggregation functions are overly simple, missing crucial interactions and non-linearities in the group preference landscape.

The empirical landscape appears to be a mixture of the average and maximum satisfaction functions, but differs from both of these functions along the diagonal, where users share identical individual preferences. For example, when both individuals equally dislike a program, there is a lower probability that the group will view the show than traditional approaches suggest. This difference is highlighted in Figure 9a, where the dotted line indicates the (identical) predictions made by average satisfaction, least misery, and maximum satisfaction, whereas the points show the empirical probabilities of group viewing. We see a similar deviation when matched preferences are large, showing a slightly higher likelihood of group viewing than naive methods predict. We also see that average satisfaction deals poorly with the extremes: for example, when one individual has a strong preference for a show while the other has a strong preference against it. One explanation for this behavior is a repeated bargaining scenario where groups alternate between satisfying a different individual in each instance.

In addition to differing from the three simple aggregation functions discussed above, the empirical landscape also deviates from predictions made by other popular aggregation methods [13]. For instance, the “average without misery” strategy corresponds to simply zeroing out the average satisfaction landscape below a certain predicted group preference, while the “multiplicative” method would result in a parabolic landscape.

To capture these subtleties, we fit a simple logistic regression with interactions to determine the probability of a group view (p_G) from the individual probabilities:

$$\log \frac{p_G}{1 - p_G} = \alpha_0 + \alpha_f p_f + \alpha_m p_m + \beta_f p_f^2 + \beta_m p_m^2 + \gamma_f p_f^3 + \gamma_m p_m^3 + \delta p_f p_m,$$

where p_f is the female’s probability of viewing the show individually and p_m is the male’s. The β and γ terms accommodate the non-linearities in the landscape, while the δ term accounts for multiplicative interactions. The resulting model fit for two-person, mixed-gender adult couples, shown in the fourth panel of Figure 8, provides an improved approximation to the empirical landscape, with an AUC of 83.1% compared to 82.9% and 82.7% for maximum satisfaction and average satisfaction, respectively, on a randomly selected 20%

held-out test set. Importantly, we note that while the differences in these aggregate metrics may seem insignificant, the model performs substantially better in crucial portions of the landscape—for example, traditional methods overpredict in regions where both group members dislike content (e.g., small individual values in Figure 9a), leading to potential dissatisfaction and possibly lost of trust in the recommender system. Aggregate metrics understate these improvements due to the non-uniform density of group views along the landscape.

Figure 9b shows further details of the model for mixed-gender adult couples, taken along slices of the landscape where either the male or female is indifferent, corresponding to a individual preference of 0.5. For instance, the blue curve in Figure 9b shows how the probability of a group view changes with the male’s individual preference when the female’s preference is held fixed at 0.5, and vice versa for the pink curve. This highlights two key observations: first, the modeled curves are far from (piecewise) linear, as traditional aggregation functions would suggest, and second, we see no obvious signs of gender dominance. We contrast this with Figure 9c, which shows the model fit for two-person mixed age groups. Here we see an asymmetry between adults and children, where the marginal increase in a child’s interest at low preference levels has higher impact than an adult’s.

We note that while we have discussed only mixed gender and age couples here, these same qualitative observations apply to other group types: a simple non-linear group model provides a better fit to the empirical group landscape compared to traditional aggregation functions, which translates to improved performance for group recommendations.

CONCLUSIONS

Throughout this study we have seen that groups are more complex than the sum of their parts. In particular, we saw that viewing habits shift substantially between individual and group contexts, and groups display markedly different preferences at the aggregate level depending on their demographic breakdowns. This led to a detailed investigation of preference aggregation functions for modeling group decision making. Owing to the unique nature of the large-scale observational dataset studied, we directly estimated how individual preferences are combined in group settings, and observed subtle deviations from traditional aggregation strategies.

While we were able to explain observed group behavior with a relatively simple model, these results raise nearly as many questions as they answer. For example, further investigation




Figure 9. (a) mixed gender adult couples with identical preferences, (b) mixed gender adult couples where one member is indifferent, (c) mixed age pairs where one member is indifferent.

is required to understand *why* these preference landscapes take the shape they do, with third-order non-linearities. Likewise, untangling the driving forces behind these observations requires more than simple observational data. On one hand, effects could be explained by direct influence of individuals on each other, while on the other hand these outcomes may be confounded with homophily, wherein individuals tend to preferentially participate in groups that share their tastes. We leave answers to these questions along with generalizations to arbitrary group settings as future work.

REFERENCES

- Amer-Yahia, S., Roy, S. B., Chawla, A., Das, G., and Yu, C. Group recommendation: Semantics and efficiency.
- Baltrunas, L., Makcinskis, T., and Ricci, F. Group recommendations with rank aggregation and collaborative filtering. In *RecSys* (2010).
- Berkovsky, S., Freyne, J., and Coombe, M. Aggregation trade offs in family based recommendations. In *Australasian Joint Conference on Advances in Artificial Intelligence* (2009).
- Boratto, L., Carta, S., Chessa, A., Agelli, M., and Clemente, M. L. Group recommendation with automatic identification of users communities. In *WI-IAT* (2009).
- de Campos, L., Fernández-Luna, J., Huete, J., and Rueda-Morales, M. Managing uncertainty in group recommending processes. *Journal of User Modeling and User-Adapted Interaction* 19 (2009).
- Gartrell, M., Xing, X., Lv, Q., Beach, A., Han, R., Mishra, S., and Seada, K. Enhancing group recommendation by incorporating social relationship interactions. In *GROUP* (2010).
- Gim, G., Jeong, H., Lee, H., and Yun, D. Group-aware prediction with exponential smoothing for collaborative filtering. In *CAMRa* (2011).
- Goodhardt, G. J., Ehrenberg, A. S. C., Collins, M. A., et al. *The television audience: patterns of viewing. An update*. No. Ed. 2. Gower Publishing Co. Ltd., 1987.
- Gorla, J., Lathia, N., Robertson, S., and Wang, J. Probabilistic group recommendation via information matching. In *WWW* (2013).
- Hu, X., Meng, X., and Wang, L. Svd-based group recommendation approaches: an experimental study of moviepilot. In *CAMRa* (2011).
- Jameson, A., and Smyth, B. Recommendation to groups. In *The adaptive web: methods and strategies of web personalization* (2007).
- Kim, H.-N., Rawashdeh, M., and El Saddik, A. Tailoring recommendations to groups of users: a graph walk-based approach. In *IUI* (2013).
- Masthoff, J. Group modeling: Selecting a sequence of television items to suit a group of viewers. *Journal of User Modeling and User-Adapted Interaction* 14, 1 (2004).
- McCarty, J. A., and Shrum, L. The role of personal values and demographics in predicting television viewing behavior: Implications for theory and application. *Journal of Advertising* 22, 4 (1993), 77–101.
- Ntoutsi, E., Stefanidis, K., Nørsvåg, K., and Kriegel, H.-P. Fast group recommendations by applying user clustering. In *Conceptual Modeling*. Springer, 2012, 126–140.
- OConnor, M., Cosley, D., Konstan, J. A., and Riedl, J. Polylens: A recommender system for groups of users. In *European Conference on Computer-Supported Cooperative Work* (2001).
- Paquet, U., and Koenigstein, N. One-class collaborative filtering with random graphs. In *WWW* (2013).
- Said, A., Berkovsky, S., and De Luca, E. W. Group recommendation in context. In *CAMRa* (2011).
- Senot, C., Kostadinov, D., Bouzid, M., Picault, J., Aghasaryan, A., and Bernier, C. Analysis of strategies for building group profiles. In *UMAP* (2010).
- Sprague, D., Wu, F., and Tory, M. Music selection using the partyvote democratic jukebox. In *AVI* (2008).
- Stern, D. H., Herbrich, R., and Graepel, T. Matchbox: large scale online bayesian recommendations. In *WWW* (2009).
- Vildjiounaite, E., Kyllnen, V., and Hannula, T. Unobtrusive dynamic modelling of tv programme preferences in a finnish household. *Journal of Multimedia Systems* 15 (2009).
- Webster, J. G., and Wakshlag, J. J. The impact of group viewing on patterns of television program choice. *Journal of Broadcasting & Electronic Media* 26, 1 (1982), 445–455.
- Yu, Z., Yu, Z., Zhou, X., and Nakamura, Y. Handling conditional preferences in recommender systems. In *IUI* (2009).
- Yu, Z., Zhou, X., Hao, Y., and Gu, J. Tv program recommendation for multiple viewers based on user profile merging. *Journal of User Modeling and User-Adapted Interaction* 16, 1 (2006).