A Field Study of Related Video Recommendations: Newest, Most Similar, or Most Relevant?

Yifan Zhong
University of Minnesota
zhong371@umn.edu

Tahir Lazaro Sousa Menezes
University of Minnesota
sousa009@umn.edu

Vikas Kumar
University of Minnesota
kumar093@umn.edu

Qian Zhao
University of Minnesota
zhaox331@umn.edu

F. Maxwell Harper
University of Minnesota
max@umn.edu

ABSTRACT

Many video sites recommend videos related to the one a user is watching. These recommendations have been shown to influence what users end up exploring and are an important part of a recommender system. Plenty of methods have been proposed to recommend related videos, but there has been relatively little work that compares competing strategies. We describe a field study of related video recommendations, where we deploy algorithms to recommend related movie trailers. Our results show that recency- and similarity-based algorithms yield the highest click-through rates, and that the recency-based algorithm leads to the most trailer-level engagement. Our findings suggest the potential to design non-personalized yet effective related item recommendation strategies.

CCS CONCEPTS

• Information systems → Recommender systems; • Human-centered computing → Field studies;

KEYWORDS

recommender systems; related item recommendations; item similarity; movie trailers; field study.

ACM Reference Format:

Yifan Zhong, Tahir Lazaro Sousa Menezes, Vikas Kumar, Qian Zhao, and F. Maxwell Harper. 2018. A Field Study of Related Video Recommendations: Newest, Most Similar, or Most Relevant?. In *Twelfth ACM Conference on Recommender Systems (RecSys '18)*, October 2–7, 2018, Vancouver, BC, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3240323.3240395>

1 INTRODUCTION

Online video consumption has grown to tremendous volumes. YouTube, for instance, reported in early 2017 that its users consume a billion hours of video on its service each day [7].

Online video services take several forms, including sharing sites like YouTube, social networks like Facebook, and media companies like Comedy Central. Each of these services offers a version of an

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 ACM 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.
RecSys '18, October 2–7, 2018, Vancouver, BC, Canada
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5901-6/18/10...\$15.00
<https://doi.org/10.1145/3240323.3240395>

Figure 1: Trailer viewing interface. (a) The current trailer, *Moana*, is highlighted on the bottom left; recommended trailers extend to the right. Users can (b) add the current movie to their wishlist, (c) like or dislike the trailer, and (d) play the next trailer.

“up next” recommender that displays related videos. The algorithms that recommend these related videos impact what people end up watching [17]. Therefore, understanding what strategies work well is important for improving users’ experiences.

Related video recommendation involves a trade-off among relevance to the seed item, relevance to the user, and non-personalized criteria. The most similar videos might not be the most interesting videos to the user. Some methods focus on identifying similar items by analyzing their visual, audio, or textual features [12]; some aim to provide recommendations tailored to users’ personal preferences [2, 5]; other methods incorporate non-personalized elements to recommend the most popular items or new releases. Yet, there is comparatively little work that directly compares algorithms in this context to better understand this trade-off.

Therefore, we conduct a field experiment where we deploy several algorithms to recommend movie trailers. The algorithms recommend additional trailers to watch based on a seed trailer that a user is watching. Two of the algorithms rank candidates by non-personalized criteria (recency and similarity), while the third ranks candidates using a personalized strategy (predicted rating). We aim to answer the research question:

In the context of related video recommendation, how do algorithms prioritizing recency, similarity, and predicted rating affect users in terms of their propensity to click and their subsequent engagement?

In this paper, we present results from an 18-month field experiment, with a focus on how users respond to the different recommendation strategies. Our results show that the non-personalized algorithms yield higher click-through rates, and that the algorithm prioritizing recency is the strongest in promoting trailer-level user engagement. Our findings suggest the potential for designing non-personalized yet effective related item recommendation strategies.

2 RELATED WORK

Related item recommendation is an important feature in many recommendation systems. For example, Zhou et al. [17] find that related videos on YouTube are a major source of video views. Related item recommendations are also used extensively in music recommenders [10], movie recommenders [16], e-commerce [8], and news sites [1] to enhance users' experiences.

Some methods that recommend related items focus on identifying items similar to a seed item. For example, Mei et al. [12] propose a related video recommendation strategy that identifies similar videos through their visual, audio and textual features. Other methods recommend items tailored to users' personal preferences such as their search histories or viewing habits [2, 5]. Bendersky et al. [3] demonstrate that an algorithm that incorporates users' personal preferences leads to longer video watch time than a content-based information retrieval approach.

3 METHODS

To learn about related video recommendations, we deployed a field experiment on MovieLens (movielens.org), a non-commercial movie recommendation website. MovieLens has an interface for viewing movie trailers, short videos that can help people find movies matching their interests [13] and influence decision-making [9].

Prior to this study, the trailer viewing interface was a minimal full-screen modal overlay. We enhanced the overlay to include several new features: "related trailer" recommendations, "like/dislike" buttons, and the ability to wishlist the movie from the trailer interface (see Figure 1). Like most video watching sites, we auto-advance through the list of recommendations as the user finishes each trailer. Users can open the trailer viewing interface by clicking a play button that we added to movie poster graphics and movie detail pages throughout the system.

3.1 Recommendation Algorithms

To generate related trailer recommendations, we developed and deployed three experimental algorithms and one random baseline algorithm. The three related item algorithms each generate trailer recommendations based on a particular trailer that a user is watching. These algorithms first identify a set of the 250 most similar movies to the current movie. To determine similarity, we use a content-based metric called the tag genome [15]. The tag genome computes similarity scores between pairs of movies based on the similarity between latent feature vectors generated by a supervised machine learning process. This process of first filtering to similar items, then applying a ranking function, is a typical approach to related item recommendations [4, 8].

The three related-item algorithms then rank the 250 most similar movies as follows:

- *TagSimilarity*. This non-personalized algorithm ranks movies in order of their similarity to the seed movie.
- *FilmReleaseDate*. This non-personalized algorithm ranks movies by release date, newest first.
- *PredictedRating*. This personalized algorithm ranks movies by predicted rating for the current user. We compute predicted ratings using item-item collaborative filtering [14] based on 27 million ratings.

The *Baseline* algorithm generates a set of random, non-personalized recommendations, drawn from all movies in the database. These recommendations are not necessarily related to the seed movie.

3.2 Field Experiment and Metrics

We designed an online within-subjects experiment to evaluate the experimental algorithms. Users in the experiment were randomly assigned one of the four algorithms each time they logged in to the system. The same user could be assigned to different algorithms in different sessions. We chose this design because we gain more observations for each algorithm, relying on mixed-effects statistical analysis to take into account the dependency of multiple observations from the same user.

We evaluate our algorithms by measuring their effects on user within-session engagement. We examine users' interactions with two groups of trailers: 1) all trailers that users viewed, and 2) trailers that were recommended and clicked on, which we call recommendation-click (RC) trailers.

We use two kinds of user engagement metrics:

- Click-through rate (CTR): the ratio of clicks on recommendations over total trailer views.
- Trailer-level user engagement: we measure how frequently users take the following actions when viewing a trailer (see Figure 1 for the user interface): (1) *TrailerLiked*, *TrailerDisliked*, and *WatchedMoreThanHalf* indicate the quality of a recommendation – whether the trailer is enjoyable to watch. (2) *WishlistedMovie* measures users' interest in the movie.

In addition, we use three metrics to measure the type of recommendations made by each algorithm:

- *popularityLastYear*: The number of times the movie was rated in the past year. Higher numbers indicate more popular movies.
- *avgRating*: Average rating of the movie by our users on a 0.5-5 star scale with half-star increments.
- *ageMonth*: Difference in months between the time of measurement and the release date of the movie. Smaller numbers indicate newer movies.

We choose these metrics because they capture three different aspects of a movie and are easy to interpret and compare.

4 RESULTS

We deployed the new interface and the four algorithms on May 5th, 2016 and collected data until January 17th, 2018. During this time, 39,400 users logged in a total of 482,963 times (median 2 logins/user). These users viewed trailers 166,959 times and clicked on recommendations 9,142 times. In subsequent analysis, we restrict analysis to those users who viewed at least one trailer.

Table 1: An overview of users' interactions with all trailers and with recommendation-click (RC) trailers in our experiment.

	Trailer Views	TrailerLiked	TrailerDisliked	WatchedMoreThanHalf	WishlistedMovie
Count (All)	166,959	2,260 (1.35%)	998 (0.59%)	101,059 (60.53%)	3,164 (1.89%)
Count (RC)	9,142	223 (2.44%)	53 (0.58%)	4,078 (44.60%)	572 (6.26%)

4.1 Descriptive Statistics: Recommended Trailers

See Table 1 for a summary of trailer viewing and engagement data. An interesting observation is that users are more likely to wishlist the movie in the trailer interface after they click on a trailer recommendation (6.26% vs. 1.89%).

The first item in the recommendation list is the most likely to be clicked. About 18% of trailers that users clicked were in the first position, while items at later positions all have a lower chance. This result is consistent with several studies on YouTube related video recommendations[11, 17]

4.2 Recommendation Algorithms

The four recommendation algorithms produce different lists of items. We summarize the contents of these recommenders in Figure 2. Unsurprisingly, FilmReleaseDate recommends much newer items than the other algorithms. We also observe that PredictedRating tends to recommend items with higher average ratings and higher popularities than the other algorithms.

Based on overall usage, FilmReleaseDate has the highest click-through rate (CTR), closely followed by TagSimilarity (see Figure 3). Because users can be assigned to different algorithms in different sessions, more active users could dominate our data and skew our observations. Therefore, we perform a statistical test to better understand the variations among the similarity-based algorithms. Specifically, we build a mixed-effect logistic model to predict whether each recommendation in the list will be clicked. This model excludes within-users effects as random effects. The independent variables are the similarity-based algorithms, properties (avgRating, popularityLastYear, age) of the seed video, and position in the list (left-to-right). We exclude the properties of recommendations because they are correlated with the algorithms. Since we are primarily concerned with comparing the three experimental algorithms, we do not include observations from the baseline algorithm in this model. Table 2 displays a summary of the model.

The statistical model has a good fit with AUC = 0.8879. With PredictedRating as the reference group, TagSimilarity and FilmReleaseDate are both more likely to generate recommendations that users will click when users' random effects are excluded. In our pairwise comparisons, Baseline has poorer performance than PredictedRating, and FilmReleaseDate is not statistically different from TagSimilarity. The model also confirms that position has a significant impact on whether users click on a recommendation. The properties of the seed video have only a slight impact on the result.

We further compare the algorithms based on trailer-level engagement metrics (TrailerLiked, TrailerDisliked, WatchedMoreThanHalf, and WishlistedMovie). Based on overall usage (see Figure 4), both

Table 2: Summary of the mixed effect logistic model. PredictedRating is the reference group. Note that the last three predictors are properties of seed videos and their coefficients are standardized. * $p < 0.001$, ** $p < 0.01$, * $p < 0.05$**

Predictor	Coefficient	S.E.	P Value
Intercept	-8.1700	0.1503	< 0.0001***
TagSimilarity	0.1892	0.0487	0.0078**
FilmReleaseDate	0.2202	0.0495	0.0124***
position	-0.1093	0.0045	< 0.0001***
age_seedmovie	0.0001	0.0207	0.9935
avgRating_seedmovie	-0.0479	0.0210	0.0227*
popularity_seedmovie	-0.0194	0.0219	0.3771

FilmReleaseDate and PredictedRating did a better job in recommending enjoyable trailers and encouraging users to wishlist movies than TagSimilarity. We build four mixed effects logistic regression models to incorporate user-level effects. The four trailer-level metrics are the dependent variables, the algorithms and properties of seed trailers are fixed effects, userID is the random effect, and we vary which algorithm we use as the reference condition. The results are shown in Table 3. We find that FilmReleaseDate performs better than PredictedRating ($p < 0.05$) and TagSimilarity ($p < 0.05$) in terms of WatchedMoreThanHalf and WishlistedMovie. TagSimilarity has negative coefficients with respect to PredictedRating, although this result is not statistically significant. We don't find significant results for TrailerLiked and TrailerDisliked because the data is limited.

Table 3: Coefficients of the mixed effect logistic models that compare algorithms in pairs. Standard errors are included in parentheses. We use algorithm abbreviations: PR (PredictedRating), TS (TagSimilarity), and FRD (FilmReleaseDate). ** $p < 0.01$, * $p < 0.05$.

Predictor	WatchedMoreThanHalf	WishlistedMovie
PR(Base) vs TS	-0.0666 (0.0902)	-0.1205 (0.2760)
PR(Base) vs FRD	0.2496 (0.0906)**	0.4675 (0.2216)*
TS(Base) vs FRD	0.2671 (0.0915)**	0.5186 (0.2207)*

5 DISCUSSION AND CONCLUSION

In this research, we conduct a field study to learn which of three approaches to ranking similar movie trailers yields the most click-throughs and trailer-level actions.

We find that ranking related items with non-personalized attributes (similarity and recency) leads to more clicks than a personalized strategy based on predicted rating. This is a surprising

Figure 2: Distributions of the three properties of recommendations across the four algorithms: average rating, age (in months), and popularity last year. The y axis represents kernel density – the probability of X falling into a certain range is the area under the curve within this range.

Figure 3: How often the recommendations from each algorithm were clicked. The x axis represents the ratio of clicks on recommendations over total trailer views.

Figure 4: Several measures of utility of the different recommendations, based on the frequency of different actions a user may take after clicking on a recommended trailer. The action frequency measures the percentage of time that users take the given action in this context.

result, given the historical reliance on predicted ratings to order recommendations in collaborative filtering systems [14]. One implication is that (at least in this domain) users prioritize item relevance and recency over personal relevance. It also seems that the higher

average ratings or popularities from PredictedRating (see Figure 2) do not necessarily make its recommendations more appealing to users. Davidson et al. [5] contributed a similar finding in YouTube, showing that related videos based on similarity have higher click-through rates than top rated ones. The success of ranking by recency may be domain-specific: this may be the result of industry efforts to optimize the appealing qualities of modern movie trailers [6], or because users of our experimental system are simply most interested in using movie trailers to learn about new releases. This result speaks to the importance of incorporating domain-specific insights in the development of recommendation algorithms. The success of ranking by similarity is also surprising. In prior research, similar items are usually identified so that they can be passed to a more sophisticated, user-personalized ranking method [5, 8].

Looking one step past click-through rates, we see a somewhat different picture. Users are more likely to watch more than half of the trailer or wishlist the movie when the recommendation comes from the FilmReleaseDate algorithm. PredictedRating and TagSimilarity perform similarly with respect to these metrics in the statistical test (though PredictedRating has an edge in activity counts). Therefore, ranking by similarity may be successful only superficially: it appears to be better at boosting the click-through rate than at finding interesting content.

These results suggest the potential for building non-personalized related item recommenders. There are many ways to build on this work. The algorithms we test are simple; it is future work to test more sophisticated similarity or ranking algorithms. Our results may be domain specific; it is future work to test the effectiveness of prioritizing recency in other domains. We think the results presented here contribute an interesting data point to the investigation of related item recommendation algorithms, and we look forward to seeing more empirical results in this research space.

6 ACKNOWLEDGMENTS

We thank Joseph Konstan and Loren Terveen for their advice, members of GroupLens research for their support and feedback, and MovieLens members for their participation. This material is based on work supported by the National Science Foundation under grants IIS-0964695, IIS-1017697, IIS-1111201, IIS-1210863, IIS-1218826, and IIS-1319382 and by grants from Amazon and Google.

REFERENCES

- [1] Deepak Agarwal, Bee-Chung Chen, Pradheep Elango, and Raghu Ramakrishnan. 2013. Content Recommendation on Web Portals. *Commun. ACM* 56, 6 (June 2013), 92–101. <https://doi.org/10.1145/2461256.2461277>
- [2] Shumeet Baluja, Rohan Seth, D. Sivakumar, Yushi Jing, Jay Yagnik, Shankar Kumar, Deepak Ravichandran, and Mohamed Aly. 2008. Video Suggestion and Discovery for Youtube: Taking Random Walks Through the View Graph. In *Proceedings of the 17th International Conference on World Wide Web (WWW '08)*. ACM, New York, NY, USA, 895–904. <https://doi.org/10.1145/1367497.1367618>
- [3] Michael Bendersky, Lluís García-Pueyo, Jeremiah Harmsen, Vanja Josifovski, and Dima Lepikhin. 2014. Up Next: Retrieval Methods for Large Scale Related Video Suggestion. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '14)*. ACM, New York, NY, USA, 1769–1778. <https://doi.org/10.1145/2623330.2623344>
- [4] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16)*. ACM, New York, NY, USA, 191–198. <https://doi.org/10.1145/2959100.2959190>
- [5] James Davidson, Benjamin Liebald, Junning Liu, Palash Nandy, Taylor Van Vleet, Ullas Gargi, Sujoy Gupta, Yu He, Mike Lambert, Blake Livingston, and Dasarathi Sampath. 2010. The YouTube Video Recommendation System. In *Proceedings of the Fourth ACM Conference on Recommender Systems (RecSys '10)*. ACM, New York, NY, USA, 293–296. <https://doi.org/10.1145/1864708.1864770>
- [6] Josh Dzeiza. 2015. The Star Wars history of trailers. <https://www.theverge.com/2015/12/10/9882404/star-wars-trailers-movie-marketing-youtube-disney>
- [7] Cristos Goodrow. 2017. You know what's cool? A billion hours. <https://youtube.googleblog.com/2017/02/you-know-whats-cool-billion-hours.html>
- [8] J. Katukuri, T. Kōnik, R. Mukherjee, and S. Kolay. 2014. Recommending similar items in large-scale online marketplaces. In *2014 IEEE International Conference on Big Data (Big Data)*. 868–876. <https://doi.org/10.1109/BigData.2014.7004317>
- [9] Lisa Kernan. 2004. *Coming Attractions: Reading American Movie Trailers*. University of Texas Press. Google-Books-ID: 7gu0ZU7K834C.
- [10] Peter Knees and Markus Schedl. 2013. A Survey of Music Similarity and Recommendation from Music Context Data. *ACM Trans. Multimedia Comput. Commun. Appl.* 10, 1 (Dec. 2013), 2:1–2:21. <https://doi.org/10.1145/2542205.2542206>
- [11] Dilip Kumar Krishnappa, Michael Zink, Carsten Griwodz, and Pål Halvorsen. 2015. Cache-Centric Video Recommendation: An Approach to Improve the Efficiency of YouTube Caches. *ACM Trans. Multimedia Comput. Commun. Appl.* 11, 4 (June 2015), 48:1–48:20. <https://doi.org/10.1145/2716310>
- [12] Tao Mei, Bo Yang, Xian-Sheng Hua, and Shipeng Li. 2011. Contextual Video Recommendation by Multimodal Relevance and User Feedback. *ACM Trans. Inf. Syst.* 29, 2 (April 2011), 10:1–10:24. <https://doi.org/10.1145/1961209.1961213>
- [13] Theodora Nanou, George Lekakos, and Konstantinos Fouskas. 2010. The effects of recommendations' presentation on persuasion and satisfaction in a movie recommender system. *Multimedia Systems* 16, 4-5 (Aug. 2010), 219–230. <https://doi.org/10.1007/s00530-010-0190-0>
- [14] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based Collaborative Filtering Recommendation Algorithms. In *Proceedings of the 10th International Conference on World Wide Web (WWW '01)*. ACM, New York, NY, USA, 285–295. <https://doi.org/10.1145/371920.372071>
- [15] Jesse Vig, Shilad Sen, and John Riedl. 2012. The Tag Genome: Encoding Community Knowledge to Support Novel Interaction. *ACM Trans. Interact. Intell. Syst.* 2, 3 (Sept. 2012), 13:1–13:44. <https://doi.org/10.1145/2362394.2362395>
- [16] Yuan Yao and F. Maxwell Harper. 2018. Judging Similarity: A User-Centric Study of Related Item Recommendations. In *Proceedings of the Twelfth ACM Conference on Recommender Systems (RecSys '18)*. ACM, New York, NY, USA. <https://doi.org/10.1145/3240323.3240351>
- [17] Renjie Zhou, Samamon Khemmarat, and Lixin Gao. 2010. The Impact of YouTube Recommendation System on Video Views. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement (IMC '10)*. ACM, New York, NY, USA, 404–410. <https://doi.org/10.1145/1879141.1879193>