A Short History of the RecSys Challenge

Alan Said

■ *The RecSys Challenge is a yearly recurring competition focusing on creating the best performing recommendation approach for a specific scenario. Over the years, the competition has drawn many participants from industry and academia and has become a key part of the ACM Conference on Recommender Systems series. This article presents a brief historical overview of the RecSys Challenge from its inception in 2010 until the seventh iteration in 2016.*

The first modern research paper on recommender systems (Resnick et al. 1994) was published in 1994 by Resnick, Iacovou, Suchak, Bergstrom, and Riedl. More than two decades have passed since, and the recommender systems research field has grown rapidly, and along with it a research community which, nowadays, is centered on the ACM Conference on Recommender Systems.[1] Ever since Resnick and colleagues published their paper, recommender systems have continuously adapted to current and upcoming trends and technologies. The recommendation approach proposed by Resnick and colleagues was revolutionary when it was introduced, utilizing collaborative filtering to predict scores that users would give to newsgroup articles in order to identify the most relevant articles, that is, those with the highest scores. Today, even though similar approaches are in use, they are usually just one part of complex recommendation approaches that can include large collections of algorithms and data sources.

In 2006, one initiative, the Netflix Prize,[2] created a focus on recommender systems and contributed to major advancements in the field during its three-year run. Similar initiatives have led to great improvements in related fields, for example, the Text Retrieval Conference[3] in information retrieval, and the KDD Cup[4] in data mining.

Following the success of these, the RecSys Challenge[5] is a yearly competition organized in conjunction with the ACM Conference on Recommender Systems.

  ISSN 0738-4602

A Brief History

The Netflix Prize was launched in 2006, the same year that the summer school on Recommender Systems was organized. By 2007, the Netflix Prize had attracted thousands of participating teams, and the first ACM Conference on Recommender Systems[6] was held. At the 2009 ACM RecSys conference, the Netflix Prize concluded. During the duration of the prize, significant advancements had been made in the recommender systems research field, for example, establishing matrix factorization methods such as SVD as state of the art in recommendation. At the 2010 ACM RecSys conference, the seed for what would become the RecSys Challenge was organized as the Challenge on Context-aware Movie Recommendation (CAMRa). CAMRa attracted a moderate number of participants, but contributed to establishing the RecSys Challenge series.

Challenge Structure, Yearly Overviews, and Future Trends

The RecSys Challenge has followed a similar structure since its inception: (1) a data set and problem are presented, (2) teams sign up and participate, (3) participants submit their solutions in time for a deadline, (4) participants submit papers outlining their approaches, (5) during a workshop at the ACM RecSys conference participants present their approaches and winners are announced. Only a selected number of teams present their work at the workshop due to time and space restrictions. One of the main features of the challenge is to make available a new real-world data set.

The first challenge to be held in conjunction with ACM RecSys, CAMRa (Adomavicius et al. 2010), focused on recommending movies in specific contexts, that is, during the Oscars and during Christmas, for a specific mood, and based on social ties between users. The challenge was organized as a collaboration of the Technische Universität Berlin (TU Berlin) and Moviepilot GmbH. Several data sets with various contextual dimensions were released for the challenge and 40 international teams participated.

In 2011, the second iteration of CAMRa was organized again by Moviepilot and the Technische Universität Berlin, focusing on contextual movie recommendations (Said et al. 2011) in the context of group recommendation. This time Moviepilot requested that the users of their service share information about who they shared a household with in order to accurately predict recommendations for groups of users. The number of participating teams grew to a total of 45.

The 2012 challenge was the first to focus on something else than movie recommendation (Manouselis et al. 2012) and to use the RecSys Challenge moniker. The challenge featured two tracks. The first was on creating Facebook campaigns, that is, the task of identifying which groups of users to recommend certain ad campaigns to. The data set was again provided by Moviepilot and was co-organized by TU Berlin. The second track focused on recommendation of scientific papers. The data set was provided by Mendeley and the track was co-organized by Agro-Know Technologies. The challenge attracted 30 participating teams.

In 2013, the second year under the new name, the challenge focused on rating prediction of venues on Yelp (Blomo, Ester, and Field 2013) and was co-organized by Simon Fraser University and Yelp who also provided the data. The challenge was organized as a Kaggle competition[7] and attracted 158 teams.

Following recent trends in recommender systems, the 2014 challenge did not focus on classical recommendation, but rather on prediction of user engagement, that is, what type of information will incentivize users to engage with it. The task was to identify which tweets containing ratings of movies would accumulate high numbers of either retweets or be favorited by other users on Twitter (Said et al. 2014). The challenge did not feature a data set provided by an industry partner; rather a data set procured by academic researchers from Ghent University and the Delft University Technolgy of was used. Participation grew to 225 teams.

The most popular challenge in terms of participants to date, 2015, attracted a record of 850 teams attempting to identify (1) which user visit sessions to an e-commerce site would result in a purchase, and (2) which items would be purchased. This iteration of the challenge introduced a new metric to measure the performance of the recommendation approaches proposed by the participants. The score, outlined by Ben-Shimon et al. (2015), combined both abovementioned goals. The challenge was a joint effort between YOUCHOOSE GmbH and Ben-Gurion University of the Negev.

After the introduction of a challenge-specific measure in 2015, the 2016 challenge followed suit and specified a measure[8] combining precision and recall. The measure is based on the key performance indicators used by XING GmbH (Abel et al. 2016), a career-oriented social networking site, which organized the challenge. The focus of this challenge was to predict which job postings users of the XING service interacted with and drew a total of participating 366 teams.

Over the years, the challenge has established itself as a benchmarking event for current recommender system research. It has attracted participants from academe and industry, allowing researchers and practitioners to learn from and cooperate with each other in a community-driven event. Each yearly instance takes on new research challenges based on ongoing trends in industry and academia. In the first years, the challenge focused on the then-popular movie

recommendation scenarios, but as recommender systems have grown mature, so have the problems covered by the challenge.

Current and future research in recommender systems acknowledges that certain recommender system settings require online evaluation, that is, an instantaneous feedback loop between the users of the system and the algorithm. There are already initiatives attempting to bring this setting to future instantiations of the RecSys Challenge.

Impact on the Community

The RecSys Challenge has become a key event at the ACM RecSys conference. It has attracted thousands of participants from industry and academia and has allowed researchers and practitioners to benchmark their work against each other in a friendly and open setting. The challenge has brought new recommendation settings into the public eye and given researchers the chance to work on never before seen, real-world problems and data sets.

Notes

1. recsys.acm.org.
2. www.netflixprize.com
3. trec.nist.gov.
4. www.kdd.org/kdd-cup.
5. www.recsyschallenge.com.
6. recsys.acm.org/recsys07/.
7. www.kaggle.com/c/yelp-recsys-2013.
8. github.com/recsyschallenge/2016/blob/master/EvaluationMeasure.md.

References

Abel, F.; Benczúr, A.; Kohlsdorf, D.; Larson, M.; and Pálovics, R. 2016. Recsys Challenge 2016: Job Recommendations. In *Proceedings of the 10th ACM Conference on Recommender Systems* (RecSys 2016), ed. W. Geyer, S. Sen, J. Freyne, and P. Castells. New York: Association for Computing Machinery. dx.doi.org/10.1145/2959100.2959207

Adomavicius, G.; Tuzhilin, A.; Berkovsky, S.; Luca, E. W. D.; and Said, A. 2010. Context-Awareness in Recommender Systems: Research Workshop and Movie Recommendation Challenge. In *Proceedings of the 2010 ACM Conference on Recommender Systems* (RecSys 2010), ed. X. Amatriain, M. Torrens, P. Resnick, and M. Zanker, 385–386. New York: Association for Computing Machinery.

Ben-Shimon, D.; Tsikinovsky, A.; Friedmann, M.; Shapira, B.; Rokach, L.; and Hoerle, J. 2015. Recsys Challenge 2015 and the YOOCHOOSE Dataset. In *Proceedings of the 9th ACM Conference on Recommender Systems,* (RecSys 2015), ed. H. Werthner, M. Zanker, J. Golbeck, and G. Semeraro, 357–358. New York: Association for Computing Machinery.

Blomo, J.; Ester, M.; and Field, M. 2013. Recsys Challenge 2013. In *Proceedings of the Seventh ACM Conference on Recommender Systems* (RecSys '13), ed. Q. Yang, I. King, Q. Li, P. Pu, and G. Karypis, 489–490. New York: Association for Computing Machinery. dx.doi.org/10.1145/2507157.2508008

Manouselis, N.; Said, A.; Tikk, D.; Hermanns, J.; Kille, B.; Drachsler, H.; Verbert, K.; and Jack, K. 2012. Recommender Systems Challenge 2012. In *Proceedings of the Sixth ACM Conference on Recommender Systems* (RecSys '12), ed. P. Cunningham, N. J. Hurley, I. Guy, and S. S. Anand, 353–354. New York: Association for Computing Machinery. dx.doi.org/10.1145/2365952.2366043

Resnick, P.; Iacovou, N.; Suchak, M.; Bergstrom, P.; and Riedl, J. 1994. Grouplens: An Open Architecture for Collaborative Filtering of Netnews. In *Proceedings of the Conference on Computer Supported Cooperative Work,* (CSCW '94), ed. J. B. Smith, F. D. Smith, and T. W. Malone, 175–186. New York: Association for Computing Machinery.

Said, A.; Berkovsky, S.; Luca, E. W. D.; and Hermanns, J. 2011. Challenge on Context-Aware Movie Recommendation: Camra2011. In *Proceedings of the 2011 ACM Conference on Recommender Systems* (RecSys 2011), ed. B. Mobasher, R. D. Burke, D. Jannach, and G. Adomavicius, 385–386. New York: Association for Computing Machinery. dx.doi.org/10.1145/2043932.20440

Said, A.; Dooms, S.; Loni, B.; and Tikk, D. 2014. Recommender Systems Challenge 2014. In *Proceedings of the Eighth ACM Conference on Recommender Systems* (RecSys '14) ed. A. Kobsa, M. X. Zhou, M. Ester, and Y. Koren , 387–388. New York: Association for Computing Machinery.

Alan Said is a senior lecturer at the University of Skövde, Sweden. He is the initiator of the RecSys Challenge and served as organizer in 2010, 2011, 2012, and 2014. His primary research focus is evaluation of personalized information systems.