

Netflix Prize

From Wikipedia, the free encyclopedia

The **Netflix Prize** was an open competition for the best collaborative filtering algorithm to predict user ratings for films, based on previous ratings.

The competition was held by Netflix, an online DVD-rental service, and was open to anyone not connected with Netflix (current and former employees, agents, close relatives of Netflix employees, etc.) or a resident of Cuba, Iran, Syria, North Korea, Myanmar, or Sudan. On 21 september of 2009, the grand prize of US\$1,000,000 was given to the BellKor's Pragmatic Chaos team which bested Netflix's own algorithm for predicting ratings by 10%.^[1]

Contents

- 1 Problem and data sets
- 2 Prizes
- 3 Progress over the years
 - 3.1 2007 Progress Prize
 - 3.2 2008 Progress Prize
 - 3.3 2009
- 4 Sequel
- 5 Privacy concerns
- 6 References
- 7 External links

Problem and data sets

Netflix provided a *training* data set of 100,480,507 million ratings that 480,189 users gave to 17,770 movies. Each training rating is a quadruplet of the form `<user, movie, date of grade, grade>`. The user and movie fields are integer IDs, while grades are from 1 to 5 (integral) stars.^[2]

The *qualifying* data set contains over 2,817,131 triplets of the form `<user, movie, date of grade>`, with grades known only to the jury. A participating team's algorithm must predict grades on the entire qualifying set, but they are only informed of the score for half of the data, the *quiz* set of 1,408,342 ratings. The other half is the *test* set of 1,408,789, and performance on this is used by the jury to determine potential prize winners. Only the judges know which ratings are in the quiz set, and which are in the test set—this arrangement is intended to make it difficult to hill climb on the test set. Submitted predictions are scored against the true grades in terms of root mean squared error (RMSE), and the goal is to reduce this error as much as possible. Note that while the actual grades are integers in the range 1 to 5, submitted predictions need not be. Netflix also identified a *probe* subset of 1,408,395 ratings within the *training* data set. The *probe*, *quiz*, and *test* data sets were chosen to have similar statistical properties.

For each movie, title and year of release are provided in a separate dataset. No information at all is provided about users. In order to protect the privacy of customers, "some of the rating data for some customers in the training and qualifying sets have been deliberately perturbed in one or more of the following ways: deleting ratings; inserting alternative ratings and dates; and modifying rating dates".^[1]

The training set is such that the average user rated over 200 movies, and the average movie was rated by over 5000 users. But there is wide variance in the data—some movies in the training set have as little as 3 ratings,^[3] while one user rated over 17,000 movies.^[4]

There was some controversy as to the choice of RMSE as the defining metric. Would a reduction of the RMSE by 10% really benefit the users? It has been claimed that even as small an improvement as 1% RMSE result is a significant difference in the ranking of the "top-10" most recommended movies for a user.^[5]

Prizes

Prizes were based on improvement over the Cinematch algorithm or the previous year's score if a team has made improvement beyond a certain threshold. A trivial algorithm that predicts for each movie in the quiz set its average grade from the training data produces an RMSE of 1.0540. Netflix's own algorithm, called *Cinematch*, uses "straightforward statistical linear models with a lot of data conditioning".^[6]

Using only the training data, Cinematch scores an RMSE of 0.9514 on the quiz data, roughly a 10% improvement over the trivial algorithm. Cinematch has a similar performance on the test set, 0.9525. In order to win the grand prize of \$1,000,000, a participating team had to improve this by another 10%, to achieve 0.8572 on the test set.^[1] Such an improvement on the quiz set corresponds to an RMSE of 0.8563.

As long as no team won the grand prize, a *progress* prize of \$50,000 was awarded every year for the best result

thus far. However, in order to win this prize, an algorithm had to improve the RMSE on the quiz set by at least 1% over the previous progress prize winner (or over Cinematch, the first year). If no submission succeeded, the progress prize was not to be awarded for that year.

To win a progress or grand prize a participant had to provide source code and a description of the algorithm to the jury within one week after being contacted by them. Following verification the winner also had to provide a non-exclusive license to Netflix. Netflix would publish only the description, not the source code, of the system. A team could choose to not claim a prize, in order to keep their algorithm and source code secret. The jury also kept their predictions secret from other participants. A team could send as many attempts to predict grades as they wish. Originally submissions were limited to once a week, but the interval was quickly modified to once a day. A team's best submission so far counted as their current submission.

Once one of the teams succeeded to improve the RMSE by 10% or more, the jury would issue a *last call*, giving all teams 30 days to send their submissions. Only then, the team with best submission was asked for the algorithm description, source code, and non-exclusive license, and, after successful verification; declared a grand prize winner.

The contest would last until the grand prize winner was declared. Had no one received the grand prize, it would have lasted for at least five years (until October 2, 2011). After that date, the contest could have been terminated at any time at Netflix's sole discretion.

Progress over the years

The competition began on October 2, 2006. By October 8, a team called WXYZConsulting had already beaten Cinematch's results.^[7]

By October 15, there were three teams who had beaten Cinematch, one of them by 1.06%, enough to qualify for the annual progress prize.^[8] By June 2007 over 20,000 teams had registered for the competition from over 150 countries. 2,000 teams had submitted over 13,000 prediction sets.^[2]

Over the first year of the competition, a handful of front-runners traded first place. The more prominent ones were:^[9]

- WXYZConsulting, a team by Yi Zhang and Wei Xu. (A front runner during Nov-Dec 2006.)
- ML@UToronto A, a team from the University of Toronto led by Prof. Geoffrey Hinton. (A front runner during parts of Oct-Dec 2006.)
- Gravity, a team of four scientists from the Budapest University of Technology (A front runner during Jan-May 2007.)
- BellKor, a group of scientists from AT&T Labs. (A front runner since May 2007.)

On August 12, 2007, many contestants gathered at the KDD Cup and Workshop 2007, held at San Jose, California.^[10] During the workshop all four of the top teams on the leaderboard at that time presented their techniques.

On September 2, 2007, the competition entered the "last call" period for the 2007 Progress Prize. Teams had thirty days to tender submissions for consideration. At the beginning of this period the leading team was BellKor, with an RMSE of 0.8728 (8.26% improvement). followed by Dinosaur Planet (RMSE = 0.8769; 7.83% improvement), and Gravity (RMSE = 0.8785; 7.66% improvement). In the last hour of the last call period, an entry by "KorBell" took first place. This turned out to be an alternate name for Team BellKor.

Over the second year of the competition, only three teams reached the leading position:

- BellKor, a group of scientists from AT&T Labs. (front runner during May 2007 - Sept 2008.)
- BigChaos, a team of Austrian scientists from commendo research & consulting (single team front runner since Oct 2008)
- BellKor in BigChaos, a joint team of the two leading single teams (A front runner since Sept. 2008)

2007 Progress Prize

On November 13, 2007, team KorBell (aka BellKor) was declared the winner of the \$50,000 Progress Prize with an RMSE of 0.8712 (8.43% improvement).^[11] The team consisted of three researchers from AT&T Labs, Yehuda Koren, Robert Bell, and Chris Volinsky.^[12] As required, they published a description of their algorithm.^[13]

2008 Progress Prize

The 2008 Progress Prize was awarded to the team BellKor in BigChaos. The winning submission achieved a 9.44% improvement over Cinematch (an RMSE of 0.8616).^[14] The joint-team consisted of two researchers from commendo research & consulting GmbH, Andreas Töschler and Michael Jahrer (originally team BigChaos) and three researchers from AT&T Labs, Yehuda Koren, Robert Bell, and Chris Volinsky (originally team BellKor).^[15] As required, they published a description of their algorithm.^{[16][17]}

This was the final Progress Prize because obtaining the required 1% improvement over the 2008 Progress Prize would be sufficient to qualify for the Grand Prize.

2009

On June 26, 2009 the team "BellKor's Pragmatic Chaos", a merger of teams "Bellkor in BigChaos" and "Pragmatic Theory", achieved a 10.05% improvement over Cinematch (an RMSE of 0.8558). The Netflix Prize competition did then enter the "last call" period for the Grand Prize. In accord with the Rules, teams had thirty (30) days, until July 26, 2009 18:42:37 UTC, to make submissions that will be considered for this Prize.^[18]

On July 25, 2009 the team "The Ensemble", a merger of the teams "Grand Prize Team" and "Opera Solutions and Vandelay United", achieved a 10.09% improvement over Cinematch (an RMSE of 0.8554).^[19]

On July 26, 2009, Netflix stopped gathering submissions for the Netflix Prize contest.^[20]

The final standing of the Leaderboard at that time showed that two teams met the minimum requirements for the Grand Prize. "The Ensemble" with a 10.10% improvement over Cinematch on the Qualifying set (an RMSE of 0.8553), and "BellKor's Pragmatic Chaos" with a 10.09% improvement over Cinematch on the Qualifying set (an RMSE of 0.8554).^[21] The Grand Prize winner was to be the one with the better performance on the Test set.

On September 18, 2009, Netflix announced team "BellKor's Pragmatic Chaos" as the prize winner, and the prize was awarded to the team in a ceremony on September 21, 2009.^[22] "The Ensemble" team had in fact succeeded to match BellKor's result, but since BellKor submitted their results 20 minutes earlier, the rules award the prize to them.^[23]

The joint-team "BellKor's Pragmatic Chaos" consisted of two Austrian researchers from Commendo Research & Consulting GmbH, Andreas Töschler and Michael Jahrer (originally team BigChaos), two researchers from AT&T Labs, Robert Bell, and Chris Volinsky, Yehuda Koren from Yahoo! (originally team BellKor) and two researchers from Pragmatic Theory, Martin Piotte and Martin Chabbert.^[24] As required, they published a description of their algorithm.^[25]

The team reported to have achieved the "dubious honors" (*sic* Netflix) of the worst RMSEs on the *quiz* and *test* data sets from among the 44,014 submissions made by 5,169 teams was "Lanterne Rouge", led by J.M. Linacre, who also a member of "The Ensemble" team.

Sequel

In August 2009, Netflix announced it would run a second competition, with shorter time spans of 6 months and 18 months to prize awarding.^[26] According to the New York Times, the new challenge will use demographic data rather than previous ratings as data to base predictions on.^[23]

Privacy concerns

Although the data sets were changed to preserve customer privacy, the Prize has been criticized by privacy advocates. In 2007 two researchers from the University of Texas were able to identify individual users by matching the data sets with film ratings on the Internet Movie Database.^[27]

In December 2009, an anonymous Netflix user sued Netflix in *Doe v. Netflix*, alleging that Netflix had violated U.S. fair trade laws and the Video Privacy Protection Act by releasing the datasets.^[28]

References

- [^] ^{*a b c*} "The Netflix Prize Rules" (<http://www.netflixprize.com/rules>) . <http://www.netflixprize.com/rules>. Retrieved 2007-08-21.
- [^] ^{*a b*} James Bennett; Stan Lanning (August 12, 2007). "The Netflix Prize" (http://www.netflixprize.com/assets/NetflixPrizeKDD_to_appear.pdf) . Proceedings of KDD Cup and Workshop 2007. http://www.netflixprize.com/assets/NetflixPrizeKDD_to_appear.pdf. Retrieved 2007-08-25.
- [^] Sigmoid Curve (2006-10-08). "Miss Congeniality" (<http://www.netflixprize.com/community/viewtopic.php?id=164>) . *Netflix Prize Forum*. <http://www.netflixprize.com/community/viewtopic.php?id=164>. Retrieved 2007-08-25.
- [^] "prodigious" (2006-10-06). "A single customer that rated 17,000 movies" (<http://www.netflixprize.com/community/viewtopic.php?id=141>) . *Netflix Prize Forum*. <http://www.netflixprize.com/community/viewtopic.php?id=141>. Retrieved 2007-08-25.
- [^] YehudaKoren (2007-12-18). "How useful is a lower RMSE?" (<http://www.netflixprize.com/community/viewtopic.php?id=828>) . *Netflix Prize Forum*. <http://www.netflixprize.com/community/viewtopic.php?id=828>.
- [^] "Netflix Prize Frequently Asked Questions" (<http://www.netflixprize.com/faq>) . <http://www.netflixprize.com/faq>. Retrieved 2007-08-21.
- [^] "Netflix Prize Rankings" (http://www.hackingnetflix.com/2006/10/netflix_prize_r.html) . *Hacking Netflix*. October 9, 2006. http://www.hackingnetflix.com/2006/10/netflix_prize_r.html. Retrieved 2007-08-21.
- [^] "Netflix Prize (I tried to resist, but...)" (<http://jsnell.iki.fi/blog/archive/2006-10-15-netflix-prize.html>) . *Juho Snellman's Weblog*. October 15, 2006. <http://jsnell.iki.fi/blog/archive/2006-10-15-netflix-prize.html>. Retrieved 2007-08-21.
- [^] "Top contenders for Progress Prize 2007 chart" (<http://www.research.att.com/~volinsky/netflix/leaders.gif>) . <http://www.research.att.com/~volinsky/netflix/leaders.gif>.
- [^] "The KDD Cup and Workshop 2007" (<http://www.cs.uic.edu/~liub/Netflix-KDD-Cup-2007.html>) . <http://www.cs.uic.edu/~liub/Netflix-KDD-Cup-2007.html>.
- [^] Prizemaster (2007-11-13). "Netflix Progress Prize 2007 awarded to team KorBell"

- (<http://www.netflixprize.com/community/viewtopic.php?id=799>) . *Netflix Prize Forum*.
<http://www.netflixprize.com/community/viewtopic.php?id=799>.
12. ^ "\$50,000 Progress Prize is Awarded on First Anniversary of \$1 Million Netflix Prize" (<http://www.netflix.com/MediaCenter?id=6422>) . <http://www.netflix.com/MediaCenter?id=6422>.
 13. ^ R. Bell, Y. Koren, C. Volinsky (2007). ""The BellKor solution to the Netflix Prize"" (http://www.netflixprize.com/assets/ProgressPrize2007_KorBell.pdf) .
http://www.netflixprize.com/assets/ProgressPrize2007_KorBell.pdf.
 14. ^ Prizemaster (2008-12-10). "Netflix Progress Prize 2008 to Team BellKor in BigChaos" (<http://www.netflixprize.com/community/viewtopic.php?id=1193>) . *Netflix Prize Forum*.
<http://www.netflixprize.com/community/viewtopic.php?id=1193>.
 15. ^ "Netflix Awards \$50,000 Progress Prize in Year Two of Multi-Year, Multi-National Netflix Prize Competition" (<http://netflix.mediaroom.com/index.php?s=43&item=295>) . <http://netflix.mediaroom.com/index.php?s=43&item=295>.
 16. ^ A. Töschner, M. Jahrer (2008). ""The BigChaos solution to the Netflix Prize 2008"" (http://www.netflixprize.com/assets/ProgressPrize2008_BigChaos.pdf) .
http://www.netflixprize.com/assets/ProgressPrize2008_BigChaos.pdf.
 17. ^ R. Bell, Y. Koren, C. Volinsky (2008). ""The BellKor solution to the Netflix Prize 2008"" (http://www.netflixprize.com/assets/ProgressPrize2008_BellKor.pdf) .
http://www.netflixprize.com/assets/ProgressPrize2008_BellKor.pdf.
 18. ^ "BellKor's Pragmatic Chaos" (<http://www.research.att.com/~volinsky/netflix/bpc.html>) . 2009-06-26.
<http://www.research.att.com/~volinsky/netflix/bpc.html>.
 19. ^ "The Ensemble" (<http://www.the-ensemble.com/>) . 2009-07-25. <http://www.the-ensemble.com/>.
 20. ^ "Contest Closed" (<http://www.netflixprize.com/closed>) . 2009-07-26. <http://www.netflixprize.com/closed>.
 21. ^ "The Netflix Prize Comes To A Buzzer-Beater, Nailbiting Finish" (<http://www.techcrunch.com/2009/07/26/the-netflix-prize-comes-to-a-buzzer-beater-nailbiting-finish/>) . 2009-07-26.
<http://www.techcrunch.com/2009/07/26/the-netflix-prize-comes-to-a-buzzer-beater-nailbiting-finish/>.
 22. ^ "Grand Prize awarded to team BellKor's Pragmatic Chaos" (<http://www.netflixprize.com/community/viewtopic.php?id=1537>) . Netflix Prize Forum. 2009-09-21.
<http://www.netflixprize.com/community/viewtopic.php?id=1537>.
 23. ^ ^a ^b Steve Lohr (2009-09-21). "A \$1 Million Research Bargain for Netflix, and Maybe a Model for Others" (<http://www.nytimes.com/2009/09/22/technology/internet/22netflix.html>) . New York Times.
<http://www.nytimes.com/2009/09/22/technology/internet/22netflix.html>.
 24. ^ "Netflix Awards \$1 Million Netflix Prize and Announces Second \$1 Million Challenge" (<http://netflix.mediaroom.com/index.php?s=43&item=327>) . <http://netflix.mediaroom.com/index.php?s=43&item=327>.
 25. ^ Andreas Töschner and Michael Jahrer (2009-09-21). "The BigChaos Solution to the Netflix Grand Prize." (<http://www.commendo.at/en-publications-48.html>) . commendo. <http://www.commendo.at/en-publications-48.html>.
 26. ^ "Netflix Prize 2 (Yes, a sequel!)" (<http://www.netflixprize.com/community/viewtopic.php?id=1520>) . Netflix Prize Forum. 2009-08-06. <http://www.netflixprize.com/community/viewtopic.php?id=1520>.
 27. ^ http://www.wired.com/science/discoveries/news/2007/03/72963_Rise_of_the_Netflix_Hackers
 28. ^ http://www.wired.com/threatlevel/2009/12/netflix-privacy-lawsuit/?utm_source=feedburner&utm_medium=feed&utm_campaign=Feed%3A+ wired%2Findex+%28Wired%3A+Index+3+%28Top+Stories+2%29%29&utm_content=Google+Feedfetcher
Netflix Spilled Your Brokeback Mountain Secret, Lawsuit Claims

External links

- Official website (<http://www.netflixprize.com/>)
- Kate Greene (2006-10-06). "The \$1 million Netflix challenge" (http://www.technologyreview.com/read_article.aspx?id=17587&ch=biztech) . *Technology Review*.
http://www.technologyreview.com/read_article.aspx?id=17587&ch=biztech.
- Robert M. Bell, Jim Bennett, Yehuda Koren, and Chris Volinsky (May 2009). "The Million Dollar Programming Prize" (<http://www.spectrum.ieee.org/may09/8788>) . *IEEE Spectrum*.
<http://www.spectrum.ieee.org/may09/8788>.
- Robust De-anonymization of Large Sparse Datasets by Arvind Narayanan and Vitaly Shmatikov (http://www.cs.utexas.edu/~shmat/shmat_oak08netflix.pdf)

Retrieved from "http://en.wikipedia.org/wiki/Netflix_Prize"

Categories: Computer science competitions | Crowdsourcing

- This page was last modified on 24 January 2010 at 20:35.
 - Text is available under the Creative Commons Attribution-ShareAlike License; additional terms may apply. See Terms of Use for details.
- Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.