Recommender Systems: Taking Advantage of Noisy Neighbours

Susan Craw

Smart Data Technologies Centre
Robert Gordon University, Aberdeen, UK
s.craw@rgu.ac.uk
http://www.rgu.ac.uk/dmstaff/craw-susan

1 Case-Based Recommenders

Millions of people use online services every day. With so many options, recommender systems are becoming increasingly important to help us make choices: which camera, computer, washing machine, hotel room etc best meets our needs; which clothing, furniture, fabrics, etc best fits our preferences; which book, movie, music best suits our taste; which destinations and activities match our vacation wishes; etc. Recommender Systems help us to make choices by presenting us with suggestions that may suit. Case-based recommenders [1] are an increasingly important application of case-based reasoning (CBR). There are two main types of recommender system: *content-based* systems suggest items whose features are most similar to the features requested or specified by the user; *collaborative filtering* systems find others with similar rating profiles to the user, and recommend items that they have liked. Both types of recommender use similarity matching: content-based suggests similar items; collaborative filtering finds similar users and recommends their items. A disadvantage of collaborative filtering is *cold start* where new items cannot appear in existing user profiles. A related issue is the *long tail* of recommendation where popular items are common in user profiles, but niche choices appear less frequently. Similar problems occur when new users do not have a profile to match, or users with niche tastes (*grey sheep* [2]) are excluded because no similar users exist. Although cold start and long tail are most commonly associated with collaborative filtering, they can also affect content-based recommendation when the features describing items are generated by users, as in social tagging or reviews. New/niche items will have no/few features describing them.

2 Noisy Recommendations

Case-based recommendation relies on having a case base that contains relevant items or relevant user profiles, and a representation for these items or profiles that captures features that are relevant for retrieval; e.g. features of a camera that customers identify as requirements, or different movie-rating profiles for family outings compared to adult viewing.

Noisy recommendations come from retrieving noisy neighbours. In CBR it is important to understand the local complexity around cases so that case base editing/maintenance removes noisy cases that provide incorrect solutions or poor recommendations. If the solutions in neighbours are similar to each other but different from the central case then the central case is likely to be noisy, is harmful to CBR, and should be removed. But if the solutions are evenly spread across the case and its neighbours, then it is probably near a decision boundary and should be retained [3, 4].

3 Music Recommenders

Recommender systems are essential to browse the tracks available through online music services. Users are looking for high quality recommendations, but also want to discover tracks and artists that they do not already know, newly released tracks, and the more niche music found in the 'long tail' of on-line music.

Texture (timbre) is one of the most powerful audio-based representations for music recommendation. We use the MFS Mel-Frequency Spectrum texture [5], a musical adaptation of the well-known Mel-Frequency-Cepstral-Coefficients (MFCC) texture. MFS is available through the Vamp audio analysis plugin system (www.vamp-plugins.org/download.html). However, tracks with similar audio features do not generally make good recommendations.

Many state-of-the-art recommender systems make use of social tagging that can provide useful semantic information such as genres, opinions, together with social and cultural information. However social tagging suffers from cold start for new tracks, and a popularity bias against niche tracks, so tagging can be sparse. The Million Song Dataset [6] includes a Last.fm tagged dataset[1] and its tagging is typical – 46% of its tracks have no tags at all!

One approach to overcome sparse tagging is to use hybrid recommenders that combine audio and tag representations/recommenders [7]. We have developed a new approach that exploits knowledge in noisy neighbours [8]! This Pseudo-Tag Hybrid learns tagging knowledge from the neighbours in the audio (noisy!) space, and uses these pseudo-tags to augment any existing tags. We have explored performance of the Pseudo-Tag Hybrid, with Horsburgh's dataset of 3k+ tracks by 750+ artists [9], through a user trial, and a separate off-line evaluation using user data [10]. Both evaluations demonstrate that the Pseudo-Tag Hybrid outperforms both tag and audio recommenders for quality recommendations, and moreover introduces the user to more unknown tracks [11].

4 Conclusions

In CBR, noisy neighbours in general provide wrong solutions or suggest incorrect decisions, and for recommenders, including music recommenders, this becomes the inclusion of poor recommendations. However, we have (naughtily) used the

[1] http://labrosa.ee.columbia.edu/millionsong/lastfm

term “noisy neighbours” to indicate neighbours in the audio space! Here, we have been able to exploit the knowledge in audio neighbourhoods, to augment the representation of music tracks, to achieve good recommendations, that also do not ignore new and niche tracks, for which limited non-audio data is available. So rather than removing tracks that are poor recommendations, we improve the representation of tracks with sparse information so that they may be recommended instead. In this way we have taken advantage of “noisy neighbours” in order to improve the representation of tracks, so that the recommender offers serendipitous, as well as good, recommendations.

References

1. Bridge, D., Goker, M.H., McGinty, L., Smyth, B.: Case-based recommender systems. Knowledge Engineering Review **20**(3):315–320 (2005)
2. Barragáns-Martínez, A.B., Costa-Montenegro, E., Burguillo, J.C., Rey-López, M., Mikic-Fonte, F.A., Peleteiro, A.: A hybrid content-based and item-based collaborative filtering approach to recommend TV programs enhanced with singular value decomposition. Information Sciences **180**(22):4290–4311 (2010)
3. Craw, S., Massie, S., Wiratunga, N.: Informed case base maintenance: A complexity profiling approach. In: Proceedings of the 22nd AAAI Conference on Artificial Intelligence, pp. 1618–1621. AAAI (2007)
4. Massie, S., Craw, S., Wiratunga, N.: When similar problems don’t have similar solutions. In: Proceedings of the 7th International Conference on Case-Based Reasoning. LNCS, vol. 4626, pp. 92–106, doi: 10.1007/978-3-540-74141-1_7. Springer, Heidelberg (2007)
5. Horsburgh, B., Craw, S., Massie, S.: Music-inspired texture representation. In: Proceedings of the 26th AAAI Conference on Artificial Intelligence, pp. 52–58, www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/5041. AAAI (2012)
6. Bertin-Mahieux, T., Ellis, D.P., Whitman, B., Lamere, P.: The million song dataset. In: Proceedings of the 12th International Society for Music Information Retrieval Conference, pp 591–596 (2011)
7. Horsburgh, B., Craw, S., Massie, S., Boswell, R.: Finding the hidden gems: Recommending untagged music. In: Proceedings of the 22nd International Joint Conference in Artificial Intelligence, pp. 2256–2261, doi: 10.5591/978-1-57735-516-8/IJCAI11-376. AAAI (2011)
8. Horsburgh, B., Craw, S., Massie, S.: Learning pseudo-tags to augment sparse tagging in hybrid music recommender systems. Artificial Intelligence **219** 25–39. doi: 10.1016/j.artint.2014.11.004 (2015)
9. Horsburgh, B.: Integrating content and semantic representations for music recommendation. PhD thesis, Robert Gordon University, http://hdl.handle.net/10059/859 (2013)
10. Craw, S., Horsburgh, B., Massie, S.: Music recommenders: User evaluation without real users? In: Proceedings of the 24th International Joint Conference in Artificial Intelligence, pp. 1749–1755 http://ijcai.org/papers15/Papers/IJCAI15-249.pdf. AAAI (2015)
11. Craw, S., Horsburgh, B., Massie, S.: Music recommendation: Audio neighbourhoods to discover music in the long tail. In: Proceedings of the 23rd International Conference on Case-Based Reasoning. LNAI, vol.9343, pp. 73–87, doi: 10.1007/978-3-319-24586-7_6. Springer (2015)