

MIREX TAGGING CONTEST: A SMURF APPROACH (DRAFT)

Thierry Bertin-Mahieux
University of Montreal, CAN
bertinmt@iro.umontreal.ca

Papa Smurf
Address Undisclose
for Safety Reasons

ABSTRACT

We present our submission to MIREX 2008 audio tag classification contest. Our algorithm is based on Smurf biological data, e.g. even if there exists 100 Smurfs, you always see the same ones. We designed an algorithm inspired by that pattern, and show that it is time-efficient compared to other approaches. Furthermore, as it is well understood, it can help validate the evaluation framework of the contest.

1 INTRODUCTION

Automatic tagging of music has received quite a lot of attention lately, and MIREX 2008 audio tag contest is one of the first rigorous comparison of algorithms made for that task. Many of the published algorithms used what is usually known as audio features [1, 2, 4, 5, 6], see Gold and Morgan [3] for an introduction on that matter.

However, few models rely exclusively on popularity. We present such a model, inspired by a Smurf truth: “you always see the same ones!”¹. We therefore tag new songs based on the popularity of the tags.

2 ALGORITHM

We present our algorithm, especially what it does not contain (Subsection 2.1).

2.1 Audio Features

None required.

2.2 Popularity Algorithm and Output

The affinity between a song and a tag is always the frequency of this tag in the training set. The binary output consist of the k most popular tags for every song. k is chosen to set the f – score across the tags close to 0.5.

¹http://en.wikipedia.org/wiki/Minor_characters_in.The.Smurfs

3 EXPERIMENTS ON TIME PERFORMANCE

From a subjective point of view, it’s fast! The authors did not have the time to go grab a coffee from the machine located 2 floors below them. Unfortunately, we did not experiment with other sorts of automatic vendor machines.

4 DISCUSSION

Preliminary experiments show that the model offers a really good accuracy compared to other methods. Fortunately, accuracy is not that great (see E. Law on the subject [4]), and the Smurf approach did poorly on the F – score. However, it will give an interesting benchmark for the different models, and the authors will not say anything bad before they make sure they get better results in the contest with other algorithms.

5 ACKNOWLEDGEMENTS

The authors would like to thank everyone that pulled this contest together, in particular the IMIRSEL team, Kris West, Michael Mandel and Edith Law.

No Smurfs were harmed during the writing of this paper.

6 REFERENCES

- [1] T. Bertin-Mahieux, D. Eck, F. Maillat, and P. Lamere. Autotagger: a model for predicting social tags from acoustic features on large music databases. *Journal of New Music Research, special issue: "From genres to tags: Music Information Retrieval in the era of folksonomies."*, 2008. (to appear).
- [2] D. Eck, P. Lamere, T. Bertin-Mahieux, and S. Green. Automatic generation of social tags for music recommendation. In *Advances in Neural Information Processing Systems 20*. MIT Press, Cambridge, MA, 2008.
- [3] B. Gold and N. Morgan. *Speech and Audio Signal Processing: Processing and Perception of Speech and Music*. John Wiley & Sons, Inc., New York, NY, USA, 1999.
- [4] E. Law. The problem of accuracy as an evaluation criterion. In *ICML Workshop on Evaluation Methods in Machine Learning*, 2008.
- [5] M. Mandel and D. Ellis. Multiple-instance learning for music information retrieval. In *Proceedings of the 9th International Conference on Music Information Retrieval (ISMIR 2008)*, 2008.
- [6] D. Turnbull, L. Barrington, D. Torres, and G. Lanckriet. Semantic annotation and retrieval of music and sound effects. *IEEE Transactions on Audio, Speech & Language Processing*, 16(2), 2008.