

On the Evaluation of Outlier Detection: Measures, Datasets, and an Empirical Study Continued

Guilherme O. Campos¹, Arthur Zimek², Jörg Sander³,
Ricardo J. G. B. Campello¹, Barbora Micenková⁴, Erich Schubert⁵,
Ira Assent⁴, and Michael E. Houle⁶

¹ University of São Paulo, {gocampos, campello}@icmc.usp.br

² University of Southern Denmark, zimek@imada.sdu.dk

³ University of Alberta, jsander@ualberta.ca

⁴ Aarhus University, {barbora, ira}@cs.au.dk

⁵ Ludwig-Maximilians-Universität München, schube@dbis.lmu.de

⁶ National Institute of Informatics meh@nii.ac.jp

The evaluation of unsupervised outlier detection algorithms is a constant challenge in data mining research. Little is known regarding the strengths and weaknesses of different standard outlier detection models, and the impact of parameter choices for these algorithms. The scarcity of appropriate benchmark datasets with ground truth annotation is a significant impediment to the evaluation of outlier methods. Even when labeled datasets are available, their suitability for the outlier detection task is typically unknown. Furthermore, the biases of commonly-used evaluation measures are not fully understood. It is thus difficult to ascertain the extent to which newly-proposed outlier detection methods improve over established methods. We performed an extensive experimental study [1] on the performance of a representative set of standard k nearest neighborhood-based methods for unsupervised outlier detection, across a wide variety of datasets prepared for this purpose. Based on the overall performance of the outlier detection methods, we provide a characterization of the datasets themselves, and discuss their suitability as outlier detection benchmark sets. We also examine the most commonly-used measures for comparing the performance of different methods, and suggest adaptations that are more suitable for the evaluation of outlier detection results.

We present the results from our previous publication [1] as well as additional observations and measures available at the outlier benchmark data repository: <http://www.dbs.ifi.lmu.de/research/outlier-evaluation/>

References

- [1] G. O. Campos, A. Zimek, J. Sander, R. J. G. B. Campello, B. Micenková, E. Schubert, I. Assent, and M. E. Houle. “On the Evaluation of Unsupervised Outlier Detection: Measures, Datasets, and an Empirical Study”. In: *Data Mining and Knowledge Discovery* 30 (4 2016), pp. 891–927. DOI: 10.1007/s10618-015-0444-8.