Towards Ontology-based Explainable Classification of Rare Events

Franco Cardillo, Umberto Straccia

Towards Ontology-based Explainable Classification of Rare Events

Franco Alberto Cardillo[1], Umberto Straccia[2]
[1]ILC - CNR, Pisa
[2]ISTI-CNR, Pisa
[1] facardillo@ilc.cnr.it
[2] umberto.straccia@isti.cnr.it

Abstract

Rare events (e.g. major floods, violent conflicts) are events that have potentially widespread and/or disastrous impact on society. The overall goal is to build a framework capable to classify, predict and explain such rare events. To do so, we envisage the usage of a mixture of sub-symbolic Machine Learning (ML) and Ontology-based Statistical Relational Learning (OSRL) techniques to generate rare events classifiers and predictors, which additionally may be mapped into natural language to ease human interpretability of the decision process.

1 Introduction

Rare events are events that occur with low frequency and have potentially widespread and/or disastrous impact on society [King e Langche, 2001]. Rare events encompass natural phenomena (major earthquakes, tsunamis, hurricanes, major floods, etc.), anthropogenic hazards (warfare and related forms of violent conflict, acts of terrorism, industrial accidents, financial and commodity market crashes, etc.), as well as phenomena for which natural and anthropogenic factors interact in complex ways (spread of epidemic diseases, global warming-related changes in climate and weather, etc.). Rare events are discrete occurrences that are statistically 'improbable' in the sense that they are very infrequently observed. Despite being statistically improbable, such events are plausible insofar as historical instances of the event (or a similar event) have been documented [Morio e Balesdent, 2015].

One can make the distinction between two types of rareness [Van der Paal, 2014]:

1. *relative rareness*, also called unbalanced or imbalanced events. A set of event classes is said to be imbalanced when one class, *i.e.* the minority class or the class of interest, is much smaller than the other classes.
2. *absolute rareness*, that is essentially a small sample problem. Frequentist inference typically breaks down when sample sizes get too small.

Without going into technical details, let us remark that learning good classifier models of rare events is a very challenging task due to the relative and absolute rareness phenomena: for such cases usual ML techniques fail (see, e.g. [Van der Paal, 2014] for further insights).

The problem becomes even more challenging in case the *interpretability/explainability* dimension is added. That is, one does not want a classifier acting as a *black box*, whose decision process a human operator cannot explain, as typical in ML [Shalev-Shwartz e Ben-David, 2014], but a model whose final output can *be easily understood by humans*.

Notably, there is a growing demand on *explainable machine learning*, as pointed out in the more general context of *Explainable Artificial intelligence*. [1]

2 Towards explainable rare event classification

Our overall objective is to build a framework capable to classify, predict and explain such rare events. To do so, we envisage the usage of a mixture of sub-symbolic ML and (symbolic) OSRL [2] techniques [d'Amato *et al.*, 2010; Lehmann, 2009; Lisi e Straccia, 2013; Rettinger *et al.*, 2012; Rizzo *et al.*, 2014; Straccia e Mucci, 2015] to generate classifiers and predictors, which additionally may be mapped into natural language for human interpretability (see Figure 1).

The main ingredients of our approach are:

1. *Domain Ontologies* (DOs) that include the meaningful entities of the rare events under consideration.
   - Specifically, OWL and/or RDFS DOs are extended to deal with the inherent *fuzziness* and *uncertainty* of rare event terminology and data [Bobillo e Straccia, 2011; Lukasiewicz e Straccia, 2008; Straccia, 2013; Zimmermann *et al.*, 2012]. *Fuzzy Logic* is used to accommodate vague concepts such as, *heavy rain*, *major flood*, etc., that largely occur in rare events terminology. [3] On the other hand, as statistical analysis is a major player in rare event analysis, and in statistical relational learning, supporting *probabilistic logics* is mandatory

[1]https://en.wikipedia.org/wiki/Explainable_Artificial_Intelligence

[2]For an overview on Statistical Relational Learning, we refer the reader e.g. to [Raedt e Kersting, 2017].

[3]Strictly speaking, the term *rare* is itself fuzzy.

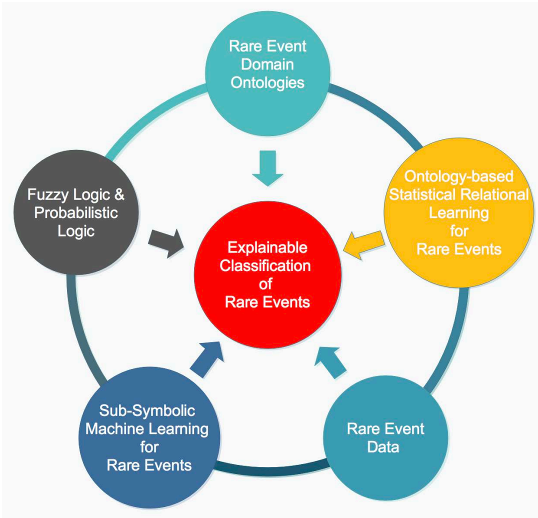

Figura 1: Explainable classification of rare events: ingredients.

to take full advantage of sub-symbolic and symbolic state of the art classifier induction methods and inference.

- Notably, the translation to natural language of a fuzzy and probabilistic OWL/RDFS statements is relatively straightforward.

2. Sub-symbolic classifiers [Van der Paal, 2014] to be used to populate atomic entities (classes and relations) of the DOs. Incorporating sub-logic classifiers as predicates, *i.e.* atomic entities, is an emerging approach as documented by [Manhaeve *et al.*, 2018].

3. Ontology-based classifiers (e.g. [Rettinger *et al.*, 2012; Straccia e Mucci, 2015]) to be adapted to rare event classification.

   - The integration of both symbolic and sub-symbolic representations and inference allows us to exploit the full expressiveness and strengths of both worlds.

4. Eventually, we consider fuzzy and probabilistic extensions of SPARQL that allows *data consumers* (users and applications) to fully access the conceptual, factual and learned entities of the domain.

3 Conclusione

We envisage the usage of a mixture of sub-symbolic machine learning and ontology-based statistical relational learning techniques, possibly combined with online machine learning methods, [4] to generate rare events classifiers and predictors, which additionally may be mapped into natural language for human interpretability of the decision process. To best of our knowledge, we are unaware of similar approaches so far.

[4]https://en.wikipedia.org/wiki/Online_machine_learning

Riferimenti bibliografici

[Agresti, 2012] A. Agresti. *Categorical Data Analysis*. Wiley, 2012.

[Allison, 2012] P. Allison. Logistic regression for rare events. Web, 2012.

[Bobillo e Straccia, 2011] F. Bobillo e U. Straccia. Fuzzy ontology representation using OWL 2. *International Journal of Approximate Reasoning*, 52:1073–1094, 2011.

[d'Amato *et al.*, 2010] C. d'Amato, N. Fanizzi, e F. Esposito. Inductive learning for the semantic web: What does it buy? *Semantic Web*, 1(1-2):53–59, 2010.

[King e Langche, 2001] G. King e Z. Langche. Logistic regression in rare events data. political analysis. *Political Analysis*, 9(2):137–163, 2001.

[Lehmann, 2009] J. Lehmann. DL-Learner: Learning concepts in description logics. *Journal of Machine Learning Research*, 10:2639–2642, 2009.

[Lisi e Straccia, 2013] F. A. Lisi e U. Straccia. A logic-based computational method for the automated induction of fuzzy ontology axioms. *Fundamenta Informaticae*, 124(4):503–519, 2013.

[Lukasiewicz e Straccia, 2008] T. Lukasiewicz e U. Straccia. Managing uncertainty and vagueness in description logics for the semantic web. *Journal of Web Semantics*, 6:291–308, 2008.

[Manhaeve *et al.*, 2018] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, e L. De Raedt. Deepproblog: Neural probabilistic logic programming. In *Advances in Neural Information Processing Systems*, 31: 3753–3763, 2018.

[Morio e Balesdent, 2015] J. Morio e M. Balesdent. *Estimation of Rare Event Probabilities in Complex Aerospace and Other Systems*. Elsevier, 2015.

[Raedt e Kersting, 2017] L. De Raedt e K. Kersting. Statistical relational learning. In *Encyclopedia of Machine Learning and Data Mining*, pages 1177–1187. Springer, 2017.

[Rettinger *et al.*, 2012] A. Rettinger, U. Lösch, V. Tresp, C. d'Amato, e N. Fanizzi. Mining the semantic web - statistical learning for next generation knowledge bases. *Data Min. Knowl. Discov.*, 24(3):613–662, 2012.

[Rizzo *et al.*, 2014] G. Rizzo, C. d'Amato, N. Fanizzi, e F. Esposito. Towards evidence-based terminological decision trees. In *Communications in Computer and Information Science* 442: 36–45. Springer, 2014.

[Shalev-Shwartz e Ben-David, 2014] S. Shalev-Shwartz e S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.

[Straccia e Mucci, 2015] U. Straccia e M. Mucci. pFOIL-DL: Learning (fuzzy) $\mathcal{EL}$ concept descriptions from crisp OWL data using a probabilistic ensemble estimation. In *Proceedings of SAC-15*, pages 345–352, 2015. ACM.

[Straccia, 2013] U. Straccia. *Foundations of Fuzzy Logic and Semantic Web Languages*. CRC Studies in Informatics Series. Chapman & Hall, 2013.

[Van der Paal, 2014] B. Van der Paal. A comparison of different methods for modelling rare events data. Master's thesis, Faculty of Sciences Department of Applied Mathematics, Computer Science and Statistics, University of Ghent, 2014.

[Zimmermann *et al.*, 2012] A. Zimmermann, N. Lopes, A. Polleres, e U. Straccia. A general framework for representing, reasoning and querying with annotated semantic web data. *Journal of Web Semantics*, 11:72–95, March 2012.