King's Research Portal

DOI:

[10.1007/978-3-662-44845-8_33](https://doi.org/10.1007/978-3-662-44845-8_33)

Document Version

Early version, also known as pre-print

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Kuiper, J., Marshall, I., Wallace, B. C., & Swertz, MA. (2014). Spá: A Web-Based Viewer for Text Mining in Evidence Based Medicine. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2014, Nancy, France, September 15-19, 2014. Proceedings, Part III* (pp. 452-455). (Lecture Notes in Computer Science; Vol. 8726). Springer. https://doi.org/10.1007/978-3-662-44845-8_33

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Spá: a web-based viewer for text mining in Evidence Based Medicine

Kuiper, J¹., Marshall, I.J.², Wallace, B.C.³, and Swertz, M.A.¹

¹ University of Groningen P.O. Box 30001, 9700 RB Groningen
{joel.kuiper,m.a.swertz}@rug.nl

² King's College London, London SE1 3QD, UK
iain.marshall@kcl.ac.uk

³ University of Texas at Austin, Austin, TX 78712, USA
byron.wallace@utexas.edu

Abstract. Summarizing the evidence about medical interventions is an immense undertaking, in part because unstructured Portable Document Format (PDF) documents remain the main vehicle for disseminating scientific findings. Clinicians and researchers must therefore manually extract and synthesise information from these PDFs. We introduce Spá,^{4,5} a web-based viewer that enables automated annotation and summarisation of PDFs via machine learning. To illustrate its functionality, we use Spá to semi-automate the assessment of bias in clinical trials. Spá has a modular architecture, therefore the tool may be widely useful in other domains with a PDF-based literature, including law, physics, and biology.

1 Introduction

Imposing structure on full-text documents is an important and practical task in natural language processing and machine learning. *Systematic reviews* are an instructive example. Such reviews aim to answer clinical questions by providing an exhaustive synthesis of all the current evidence in published literature. They are fundamental tools in Evidence-based Medicine (EBM) [9,10]. Data must be manually extracted from the literature to produce the systematic reviews. These extraction tasks are extremely laborious, but could potentially be assisted by machine learning approaches.

As an example we consider risk of bias assessment. Here reviewers assess, e.g., whether study participants and personnel were properly blinded [5]. Assessing risk of bias is a time-consuming task. A single trial typically takes a domain expert ten minutes [4], and a single review typically includes several dozen trials. Making matters worse, due to low rates of reviewer agreement it is regarded as best practice to have each study assessed twice by independent reviewers who later come to a consensus [3].

---

⁴From the Old Norse word spá or spæ referring to prophesying (prophecy)

⁵Source code available under GPLv3 at <https://github.com/joelkuiper/spa> [6]; demo available at <http://spa.clinici.co/>

Machine learning methods could provide the machinery to automate such extractions; as they can effectively impose the desired structure onto PDFs. But if such technologies are to be practically useful, we need tools that visualize these model predictions and annotations. Here we describe Spá, which aspires to realize this aim.

Spá is an open-source, web-based tool that can incorporate machine learning to automatically annotate PDF articles. As a practical demonstration of this technology, we have built a machine learning system that automatically annotates PDFs to aid EBM. This tool is unique in that it leverages state-of-the-art machine learning (ML) models applied to full-text articles to assist practitioners of EBM.

Fig. 1. Screenshot of a PDF with highlighted risk of bias. Here the risk of bias is assessed to be low, for example, and one of the supporting sentences for this assessment describes the randomization procedure (highlighted in green).

While our application of interest is EBM, we emphasize that the visualization tool can be used for any domain in which one wants to annotate PDFs, e.g. genome-wide association studies or jurisprudence. Thus the contribution of this work is two-fold, as we present: (1) a practical tool that incorporates machine learning to help researchers rapidly assess the risk of biases in published biomedical articles, and, (2) a general open-source system for visualizing the predictions of trained models from full-text articles on the web.

2 Case study: Risk of Bias in Evidence-based Medicine

2.1 Machine Learning Approaches

To automatically assess the study risk of bias, we have leveraged the Cochrane Database of Systematic Reviews (CDSR) in lieu of manually annotated data, which would be expensive to collect. The CDSR contains descriptions and data about clinical trials reported in existing systematic reviews. We match the full-texts of studies to entries in the CDSR, which contains risk of bias assessment; providing document level labels. The CDSR also contains quotations that reviewers indicated as supporting their assessments. We match these strings to substrings in the PDFs to provide sentence-level supervision. This can be viewed as a *distantly supervised* [7,8] approach.

From a ML vantage, we have two tasks for a given article: (1) predict the overall risk of bias for each of the domains, and (2) extract the sentences that support these assessments. For both tasks we leverage standard bag-of-words text encoding and linear-kernel Support Vector Machines. Because the risk of bias predictions are correlated across domains, we take a *multi-task* [2] approach to classification and jointly learn a model for the domains. We accomplish this by way of a feature space construction that includes both shared and domain-specific terms, similar to the domain adaptation approach in [1]. Specifically, we first make sentence level prediction, and then insert features representing the tokens in the predicted sentences for exploitation by the document level classifier. Figure 1 shows the system in use.

3 Spá Architecture Overview

Spá relies on Mozilla pdf.js⁶ for visualization of the document and text extraction. The results of the text extraction are processed server-side by a variety of processing topologies. Results are communicated back to the browser and displayed using React components.⁷

For each of the annotations the relevant nodes in the document are highlighted. A custom scrollbar, inspired by substance.io, that acts as a ‘mini-map’ is projected to show where annotations reside within the document. The user can interactively activate and inspect specific results.

4 Future work

We have presented a web-based tool for visualization of annotations and marginalia for PDF documents. Furthermore, we have demonstrated the use of this system within the context of Evidence-based Medicine by automatically extracting potential risks of bias.

---

⁶<http://mozilla.github.io/pdf.js>

⁷<http://facebook.github.io/react>

We believe the tool to be potentially useful for a much wider range of machine learning applications. Currently we are developing a pluggable system for processing topologies, allowing developers to quickly plug in new systems for automated PDF annotation. Furthermore, we are working to allow users to persist annotations and marginalia, possibly embedded within the document itself, for sharing and off-line use. The vision is to have an extensible system for machine assisted data extraction that will greatly increase both the quality and the reproducibility (i.e. data provenance) of current Evidence-based Medicine.

Acknowledgments Part of this research was funded by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 261433. The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-662-44845-8_33

References

1. Daumé III, H.: Frustratingly easy domain adaptation. In: Association for Computational Linguistics (ACL). vol. 1785 (2007)
2. Evgeniou, T., Pontil, M.: Regularized multi-task learning. In: Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 109–117. KDD '04, ACM, New York, NY, USA (2004)
3. Hartling, L., Ospina, M., Liang, Y.: Risk of bias versus quality assessment of randomised controlled trials: cross sectional study. *BMJ* 339, b4012 (2009), <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2764034/>
4. Hartling, L., Bond, K., Vandermeer, B., Seida, J., Dryden, D.M., Rowe, B.H.: Applying the risk of bias tool in a systematic review of combination long-acting beta-agonists and inhaled corticosteroids for persistent asthma. *PloS one* 6(2), e17242 (Jan 2011), <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3044729&tool=pmcentrez&rendertype=abstract>
5. Higgins, J., Altman, D., Gotzsche, P., Juni, P., Moher, D., Oxman, A., Savovic, J., Schulz, K., Weeks, L., Sterne, J.: The Cochrane Collaboration’s tool for assessing risk of bias in randomised trials. *BMJ* 343, d5928 (Oct 2011), <http://www.bmj.com/cgi/doi/10.1136/bmj.d5928>
6. Kuiper, J., Wallace, B.C., Marshall, I.J.: Spa. <http://figshare.com/articles/Spa/997707> (2014)
7. Mintz, M., Bills, S., Snow, R., Jurafsky, D.: Distant supervision for relation extraction without labeled data. In: Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL. pp. 1003–1011. Association for Computational Linguistics (2009)
8. Nguyen, T., Moschitti, A.: End-to-end relation extraction using distant supervision from external semantic repositories. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics. pp. 277–282. Association for Computational Linguistics (2011)
9. Sackett, D.L., Rosenberg, W.M., Gray, J., Haynes, R.B., Richardson, W.S.: Evidence based medicine: what it is and what it isn’t. *BMJ: British Medical Journal* 312(7023), 71–72 (1996)

10. Valkenhoef, G., Tervonen, T., Brock, B., Hillege, H.: Deficiencies in the transfer and availability of clinical trials evidence: a review of existing systems and standards. BMC Medical Informatics and Decision Making 12(1), 95 (2012), <http://www.biomedcentral.com/1472-6947/12/95>