

Melita: Active Document Enrichment using Adaptive Information Extraction from Text

Fabio Ciravegna¹, Alexiei Dingli¹ and Daniela Petrelli²

¹Department of Computer Science ²Department of Information Studies,
University of Sheffield, Regent Court, 211 Portobello Street,
S1 4DP Sheffield, UK

F.Ciravegna@dcs.shef.ac.uk , A.Dingli@dcs.shef.ac.uk, D.petrelli@shef.ac.uk

The traditional process of document annotation for knowledge identification and extraction in the Semantic Web (SW) is complex and time consuming, as it requires human manual annotation. There is currently a strong interest in Text Mining technologies (and in particular in Human Language-based Technologies), for reducing the burden of text annotation e.g. for Knowledge Management [Maybury2001]. In this poster we present Melita, an annotation interface that uses Adaptive Information Extraction from texts (IE) for reducing the burden of text annotation. In Melita, adaptation starts with the definition of a scenario, including a tag set for annotation (possibly organized as an ontology) and a corpus of texts to be annotated. Annotations are inserted by first selecting a tag from the ontology and then identifying the text area to annotate with the mouse. Differently from similar annotation tools [Day1997, Cunningham2001], Melita actively supports corpus annotation using Amilcare, an adaptive Information Extraction (IE) tool based on the (LP)² algorithm [Ciravegna2001]. While users annotate texts, Amilcare runs in the background learning how to reproduce the inserted annotation. Induced rules are silently applied to new texts and their results are compared with the user annotation. When its rules reach a (user-defined) level of accuracy, Melita presents new texts with a preliminary annotation derived by the rule application. In this case users have just to correct mistakes and add missing annotations. User corrections are inputted back to the learner for re-training. This technique focuses the slow and expensive user activity on uncovered cases, avoiding requiring annotating cases where a satisfying effectiveness is already reached. Moreover validating extracted information is a much simpler task than tagging bare texts (and also less error prone), speeding up the process considerably. If the IE based annotation becomes very reliable, the user can decide to let Melita produce further annotation automatically. Differently from other SW annotation tools that use IE as a backbone for user support (e.g. Open University's MnM or Karlsruhe University's Ontomat), Melita provides non-intrusive and just in time support for annotation. While other tools use a batch approach (training is performed after a number of texts has been

has been annotated), in Melita training is continuously performed while user annotates texts. It is non-intrusive because users can fully customize the level of support the interface provides (pervasive, very active, active, lazy or very lazy), allowing users to tune the intrusivity of the suggestions to their needs and tastes.

In some experiments we have simulated the user annotation of two manually tagged corpora (the CMU seminar announcement and the Austin Jobs corpus). Melita showed to be able to drastically reduce the quantity of user tagging. After annotating less than 30 texts it was able to reproduce correctly about 90% of the annotation including domain specific time expressions and specific location names. More complex information required the annotation of up to 100 examples.

This work is carried on as part of the AKT project (<http://www.aktors.org>), an Interdisciplinary Research Collaboration (IRC) sponsored by the UK Engineering and Physical Sciences Research Council (grant GR/N15764/01). AKT involves the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University. A description of Melita and a link to all its related publications is available at www.dcs.shef.ac.uk/~alexiei/Melita.htm

References

- [Ciravegna 2001] F. Ciravegna: "Adaptive Information Extraction from Text by Rule Induction and Generalisation" in Proceedings of 17th International Joint Conference on Artificial Intelligence (IJCAI 2001), Seattle, August 2001.
- [Cunningham 2001] H. Cunningham, D. Maynard, V. Tablan, C. Ursu, K. Bontcheva: "Developing Language Processing Components with GATE", www.gate.ac.uk
- [Day 1997] D. Day, J. Aberdeen, L. Hirschman, R. Kozierok, P. Robinson and M. Vilain *Mixed-initiative development of language processing systems*. In Proc. of the Fifth Conference on Applied Natural Language Processing, Washington, 1997.
- [Maybury 2001] M. Maybury (ed.): Proc. of the 2001 *EACL/ACL Workshop on Human Language Technology and Knowledge Management*, at the 39th ACL meeting, Toulouse, July 2001.