

# Visualization for Coreference Annotation

Andre Burkovski and Gunther Heidemann

Intelligent Systems Group

Institute for Visualization and Interactive Systems

University of Stuttgart

{firstname.lastname}@vis.uni-stuttgart.de

## Abstract

The annotation of documents with linguistic information requires time-consuming and therefore expensive manual annotation. Especially, a complex task, like coreference resolution, needs large data sets for the training of supervised machine learning methods. We present a tool which combines visualization techniques and unsupervised machine learning to support the annotation of documents with coreference information. Self-organizing Maps are used to cluster similar data and visualize the feature space. For link visualization, precise annotation, and error correction a matrix-based coreference visualization is used which exploits the transitive property of the coreference relation.

## 1 Introduction

The task of finding noun phrases which refer to the same discourse entity in a plain text is called coreference resolution. Many applications in information retrieval (Nicolov et al., 2008), machine translation, and text summarization (Mitkov et al., 2007) use coreference resolution to improve the results. Currently, the popular Jeopardy! winner machine “Watson” uses a coreference resolution module among other modules for question answering (Ferruci et al., 2010).

In many cases, coreference is resolved using supervised machine learning methods. These methods need large amounts of training samples. Yet, particularly for languages other than English, such data sets are rare and they mostly contain relatively few samples. Another problem is that supervised methods which are trained on a specific domain, such as news articles, may degrade on texts from other domains, like books or reports. E.g. Bakkenson and Soroka (2010) report on how the genre influences pronominal anaphora resolution.

The manual annotation of a document with coreference information is time consuming, because it depends not only on the background knowledge of the annotator, but also requires a high level of concentration to avoid annotation errors. Most annotation tools use text-based visualizations to show and highlight the noun phrases for coreference annotation. However, using only text-based visualizations tools has a major drawback. Annotators are prone to annotation errors because they often need to repeatedly read passages in a text for every new unlabeled noun phrase. Mitkov et al. (2000) introduce some strategies for annotators to reduce errors and to accelerate the annotation speed, but the annotation of large data sets still is a problem. Rule-based or unsupervised machine learning methods may be used to highlight probable coreference pairs and chains and to support the annotators in identifying coreferences. Such an approach further reduces the time spent on the search of suitable coreference candidates.

In our approach we address the annotation efficiency from the visualization and interaction point of view. We propose a combination of unsupervised machine learning method and visualizations for annotation purposes. We use the Self-organizing Map (SOM) to visualize groups of similar links of noun-phrases. These groups are further visualized with a coreference matrix, which takes advantage of the equivalence relation property of the coreference relationship. In such a way annotators are able to quickly identify conflicting annotations and to correct them.

## 2 Related Work

Generally, coreference annotation tools, like CorefDraw (Harabagiu et al., 2001), GATE (Cunningham et al., 2002), PALinkA (Orăsan, 2003), MMAX2 (Müller and Strube, 2006), or BART (Versley et al., 2008) provide only a text-based visualization of coreference information. The re-

cently introduced Reconcile tool (Stoyanov et al., 2009), which provide resources for developing a coreference resolution system, also uses plain text for presentation of coreference information.

All text-based visualizations present coreference information via a color scheme, link identification by indices, or visual edges between noun phrases. This does not show the the feature space or the similarity between links. Such visualizations are also limited by the size and the number of lines/colors a user can distinguish. This makes it difficult to analyze large chains, inter-document coreference or many links at once.

Advanced visualizations for coreference resolution exist. Witte and Tang (2007) present a graph-based visualization of coreferences. The framework manage coreferences as topic maps. The views used in the framework provide a good overview about the relationship of noun phrases in a link. The authors address the problem that the visualization of coreferences consists solely of highlighted plain text, possibly including edges for marking a coreference relation. We agree with them that textual visualization slows down the user and also makes cross document annotation difficult. Nonetheless, their representation only serves to visualize and navigate the space of already annotated coreferences. Another visualization created by Zeldes et al. (2009) add parse trees to plain text visualization to support annotation decisions. Still, visualization of coreference information is somewhat rare.

In contrast, we aim to visualize the actual coreferences as well as the coreference feature space and provide interaction methods for annotation in both visualizations. The basic idea is to use SOMs for clustering similar data and allow a fast and structured approach to annotation. Bekel et al. (2005) and Moehrmann et al. (2011) successfully applied this technique for image annotation. Instead of images we use pairs of noun phrases (links) as input. We extend the work of (Burkovski et al., 2011), where SOMs are used to create feature space visualizations of links and focus on methods for coreference annotation. We augment the SOM visualization with annotation information and additional matrix-based coreference visualization for a more precise annotation and annotation error detection supported by a visualization for links in a text-based manner. All visualizations are linked together via multiple coordinated views

(Roberts, 2007). Such an approach allows annotators to independently cycle through different representations of the data and systematically annotate different sets of links.

### 3 Coreference and Annotation

Coreference resolution is an active research area. Elango (2005) provide a detailed survey and Ng (2010) summarizes challenges and recent machine learning advances. In this work we use a pairwise model for coreferences: two noun phrases are considered to be coreferent if by replacing each other they do not change the meaning of a sentence. These two noun phrases form a link. The coreference relationship is reflexive, symmetric and transitive and therefore an equivalence relation. This property is used to automatically deduce and create additional annotations as well as to perform error detection of the manual annotation. A set of noun phrases which all refer to the same entity is called a chain. In a coreference chain every phrase is coreferent to every other phrase in that chain due to the transitive and symmetry property.

The main purpose of annotated data is to be used as a training and test set in supervised machine learning methods. One can argue, that it is enough to annotate only coreferent links to form coreference chains as one can create non-coreferent samples by creating links between disjunct chains. Such an approach would neglect the presence of non-referential noun phrases in a text, such as idioms, duration phrases, and others. Yet, links of these non-referential noun phrases are useful samples for supervised learning. Therefore, it is beneficial to annotate such links as well. Although a non-coreferent link does not guarantee that a noun phrase in the link is non-referential, it reduces the annotation errors. An erroneous annotation of a noun phrase as a non-referential would result in many false negative samples in the training set.

### 4 Visualization Methods

Visualization is an important tool to understand the underlying data (Shneiderman, 1996; Roberts, 2007; Keim et al., 2010) and in our case allow a systematic approach to annotation. Mitkov et al. (2000) developed systematic rules for annotators and the idea is to support this manual work with visualizations. First, we use a SOM to cluster similar links (Figure 1). We extract features

from links of a pairwise coreference model and let the SOM create groups of similar links. SOMs are easy to visualize and are rather intuitive. Second, the links in the selected clusters are visualized using the coreference matrix (Figure 2). The coreference matrix allows easy and systematic interaction for annotation. It shows annotations which contradict the transitive property of the coreference relationship and allows annotators to identify and resolve annotation errors. The visualizations are additionally supported by a traditional text-based visualization of links (not shown here).

#### 4.1 Self-organizing Maps

A SOM is an unsupervised machine learning method where artificial neurons are connected to each other by a low dimensional topology (Kohonen, 1990). This topology is applied to the high-dimensional data and the SOM learning algorithm tries to create a suitable projection of high-dimensional data. Thus, similar feature vectors in the feature space will be close in the projection space. Such coherence between the low dimensional map and high dimensional data allows creation of intuitive visualizations. The most popular visualization is the U-Matrix (Ultsch and Siemon, 1990) and its variants. The U-Matrix shows the low-dimensional grid with nodes and edges and color-codes them based on their distance in the high-dimensional feature space. Further, for a systematic annotation we use the component planes visualization introduced by Vesanto (1999). Component planes visualize the influence of one or more features to cluster formation.

To train the SOM, we use feature vectors created from pairs of noun phrases. We used a subset of features inspired by the popular feature set of Ng and Cardie (2002). Although using basic features, the SOM already clusters the data in a way, that annotators may annotate whole clusters with a low error rate. However, the SOM **does not** perform unsupervised coreference resolution. Instead, the SOM and its visualization provides methods to systematically select nodes and clusters with similar properties. For example, using component planes for string matching and Wordnet distance annotators are able to select areas of the SOM where links match in their head words and also are semantically close to each other. Interaction allows to annotate whole clusters as well as to show the actual links in selected clusters. For

the visualization of links we use the coreference matrix.

#### 4.2 Coreference Matrix

Transitive relations can be easily visualized in a matrix. Since coreference is symmetric, for annotation purposes we only need to show the upper triangular matrix. Each element of the matrix represents a link between the phrase in the row and the phrase in the column. Rows and columns are ordered by the text position of a phrase. Annotators can systematically cycle through the matrix entries and annotate the links accordingly. The contents of the links are visualized by displaying the surrounding text in another coordinated text-based visualization (not shown here). However, the key to an efficient annotation is an intelligent interaction technique. Annotators do not need to explore all entries in the matrix, but only entries in the upper triangular matrix close to the main diagonal, because of the transitive property.

For example, let  $i, j, k, l$  be phrases in the according column/row ordered by their position in the text. Let every link  $(i, j), (i, k), (i, l), (j, k), (j, l), (k, l)$  be coreferent. To annotate the links, the annotator only has to visit and label the cells  $(i, j), (j, k)$ , and  $(k, l)$ . Due to the transitive property, the annotation for the remaining links can be done automatically.

Another feature of the coreference matrix is to show annotation errors. The coreference matrix is able to highlight links that violate the transitive property. Such errors are created by a contradicting annotation by the annotator. A phrase  $a$  cannot be coreferent to a phrase  $b$  and non-coreferent to a phrase  $c$  when both phrases  $b$  and  $c$  are coreferent. In such cases, annotators are able to correct the links and modify the annotation accordingly.

The coreference matrix allow further strategy to annotate the links. Beginning with a row in the trace (diagonal) of the matrix annotators traverse the cells to the right until they find and annotate a coreferent noun phrase. Along the way, they annotate the cells between two noun phrases as non-coreferent. Now annotators follow either the chain, thereby jumping to the row of the next noun phrase in the current chain, or switch to the annotation of the next row, possibly creating a new coreference chain. The first interaction strategy allows a fast annotation of a single chain. It helps to create many coreferent samples (deduced from



the transitive property) and non-coreferent annotations to that chain in a single pass. The second interaction strategy allows the creation of multiple chains. From multiple chains non-coreferent samples are deduced by creating links with phrases from disjunct chains.

## 5 Annotation Strategies using Visualizations

The proposed visualizations allow additional annotation strategies to the guidelines presented by Mitkov et al. (2000).

SOMs can be trained for nominal to nominal noun phrase annotation only. In that case, clusters of similar noun phrases, as defined by the features, can lead to faster recognition of coreferent and non-coreferent links in matrix-based and text-based visualizations. After the annotation of nominal phrases, the SOM can be trained on nominal to pronominal links. By annotating a pronoun to a noun phrase in a chain, an annotation for all noun phrases in the chain is created automatically. Alternatively, the SOM can be trained with all kinds of noun phrases and additional features. Annotators are then able to use the component planes of the SOM to systematically investigate different combinations of noun phrases as shown by Burkovski et al. (2011).

Using the coreference matrix, annotators can follow two different strategies. First, annotators can use the SOM to select nodes with good indicators for coreferent or non-coreferent links. Such indicators depend of the features used. E.g. the head match feature is a good indicator for coreference. Annotators can easily identify regions with interesting features by using component planes of the SOM. Subsequently, by selecting the SOM nodes, the links are highlighted in the matrix and annotators are able to annotate them. In many cases most links are annotated automatically due to the transitive property. Second, starting with the first noun phrase, annotators can use the matrix to annotate some links in advance. The annotations will be reflected in the SOM, and allow annotators to see potentially coreferent and non-coreferent regions in the SOM. Also, using the SOM for general, and probably erroneous annotation the annotators will discover conflicts in the annotation. With the coreference matrix visualization it is easier to resolve conflicts in annotation than to annotate all links correctly in one single pass.

## 6 Conclusion

In this work we presented a visualization approach to coreference annotation. Instead of using traditional text-based visualization only, we propose Self-organizing Maps (SOM) and a matrix-based visualization of links in addition to text-based visualizations. SOM visualizes the feature space of the links where annotators may systematically choose regions with interesting links by utilizing component planes for feature space navigation. The matrix-based link visualization with interaction techniques allows a detailed and precise annotation of links. The coreference matrix exploits the transitive property of the coreference relation to detect and highlight annotation errors made in the process. Using visualizations as multiple coordinated views, annotators are able to systematically create coreference annotations. Our vision is that the next generation of annotation tools will employ more visualization techniques for a more efficient annotation of the data<sup>1</sup>.

Future work includes additional coreference resolution methods (other classifiers or rule-based systems) for the links in a SOM node which may provide confidence values. This can easily be visualized in the matrix or in the SOM by using a color gradient. For such links, annotators are able to quickly navigate to the matrix entries and inspect the proposed automatic annotation. Additional interaction techniques and graph-based visualizations will be investigated which may further improve and support a systematic annotation. In the future, we plan to show the efficiency of the tool by conducting a long term user study.

## References

- Michael Bakkenson and Barry I. Soroka. 2010. Importance of genre in pronominal anaphora resolution. In *International Conference on Artificial Intelligence*, pages 900–906.
- Holger Bekel, Gunther Heidemann, and Helge Ritter. 2005. Interactive image data labeling using self-organizing maps in an augmented reality scenario. *Neural Networks*, 18(5–6):566–574.
- Andre Burkovski, Wiltrud Kessler, Gunther Heidemann, Hamidreza Kobdani, and Hinrich Schütze. 2011. Self Organizing Maps in NLP: Exploration of Coreference Feature Space. In *Workshop on Self-Organizing Maps*, pages 228–237.

<sup>1</sup>The tool will be available at the project page: [www.vis.uni-stuttgart.de/suekre/](http://www.vis.uni-stuttgart.de/suekre/)

- Hammish Cunningham, Diana Maynard, Kalina Bontcheva, and Valentin Tablan. 2002. GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications. In *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics*, pages 168–175.
- Pradheep Elango. 2005. Coreference resolution: A survey. Technical report, University of Wisconsin Madison.
- David Ferruci, Eric Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya A. Kalyanpur, Adam Lally, J. William Murdock, John Prager, Nico Schlaefer, and Chris Welty. 2010. Building Watson: An Overview of the DeepQA Project. *AI Magazine, Association for the Advancement of Artificial Intelligence*, 31(3):59–79.
- Sanda M. Harabagiu, Razvan C. Bunescu, and Stefan Trausan-Matum. 2001. Corefdraw: A tool for annotation and visualization of coreference data. In *Proceedings of the 13th IEEE International Conference on Tools with Artificial Intelligence*, pages 273–279, Los Alamitos, CA, USA. IEEE Computer Society.
- Daniel A. Keim, Joern Kohlhammer, Geoffrey Ellis, and Florian Mansmann, editors. 2010. *Mastering The Information Age - Solving Problems with Visual Analytics*. Eurographics, to appear.
- Teuvo Kohonen. 1990. The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480.
- Ruslan Mitkov, Richard Evans, Constantin Orăsan, Catalina Barbu, Lisa Jones, and Violeta Sotirova. 2000. Coreference and anaphora: developing annotating tools, annotated resources and annotation strategies. In *Proceedings of the Discourse, Anaphora and Reference Resolution Conference*, pages 49–58.
- Ruslan Mitkov, Richard Evans, Constantin Orăsan, Le An Ha, and Viktor Pekar. 2007. Anaphora Resolution: To What Extent Does It Help NLP Applications? *Anaphora: Analysis, Algorithms and Applications*, 29:179 – 190.
- Julia Moehrmann, Stefan Bernstein, Thomas Schlegel, Günter Werner, and Gunther Heidemann. 2011. Improving the Usability of Hierarchical Representations for Interactively Labeling Large Image Data Sets. In *HCI International*, pages 618–627.
- Christoph Müller and Michael Strube. 2006. Multi-level annotation of linguistic data with MMAX2. In *Corpus Technology and Language Pedagogy: New Resources, New Tools, New Methods*, pages 197–214. Peter Lang.
- Vincent Ng and Claire Cardie. 2002. Improving machine learning approaches to coreference resolution. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 104–111.
- Vincent Ng. 2010. Supervised noun phrase coreference research: the first fifteen years. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1396–1411, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Nicolas Nicolov, Franco Salveti, and Steliana Ivanova. 2008. Sentiment Analysis: Does Coreference Matter? In *Proceedings of the Symposium on Affective Language in Human and Machine*, pages 37–40.
- Constantin Orăsan. 2003. PALinkA: A highly customizable tool for discourse annotation. In *4th SIG-dial Workshop on Discourse and Dialogue*, pages 39 – 43.
- Sameer S. Pradhan, Eduard Hovy, Mitch Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. 2007. Ontonotes: A unified relational semantic representation. In *International Conference on Semantic Computing*, pages 517–526.
- Jonathan C. Roberts. 2007. State of the art: Coordinated multiple views in exploratory visualization. In *Coordinated and Multiple Views in Exploratory Visualization, 2007. Fifth International Conference on Coordinated & Multiple Views.*, pages 61–71.
- Ben Shneiderman. 1996. The eyes have it: a task by data type taxonomy for information visualizations. In *IEEE Symposium on Visual Languages, 1996. Proceedings.*, pages 336–343.
- Veselin Stoyanov, Claire Cardie, Nathan Gilbert, Ellen Riloff, David Buttler, and David Bysom. 2009. Reconcile: A coreference resolution research platform. Technical report, Lawrence Livermore National Laboratory.
- Alfred Ultsch and H.P. Siemon. 1990. Kohonen’s self organizing feature maps for exploratory data analysis. In *Proceedings of International Neural Networks Conference*, pages 305–308.
- Yannick Versley, Simone Paolo Ponzetto, Massimo Poesio, Vladimir Eidelman, Alan Jern, Jason Smith, Xiaofeng Yang, and Alessandro Moschitti. 2008. Bart: a modular toolkit for coreference resolution. In *HLT '08: Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies*, pages 9–12, Morristown, NJ, USA. Association for Computational Linguistics.
- Juha Vesanto. 1999. SOM-Based data visualization methods. *Intelligent Data Analysis*, 3:111–126.
- René Witte and Ting Tang. 2007. Task-Dependent Visualization of Coreference Resolution Results. In *International Conference on Recent Advances in Natural Language Processing*.
- Amir Zeldes, Julia Ritz, Anke Lüdeling, and Christian Chiarcos. 2009. ANNIS: a search tool for multi-layer annotated corpora. In *Proceedings of Corpus Linguistics*.