Rhetorical Figure Detection: the Case of Chiasmus

Marie Dubremetz
Uppsala University
Dept. of Linguistics and Philology
Uppsala, Sweden
`marie.dubremetz@lingfil.uu.se`

Joakim Nivre
Uppsala University
Dept. of Linguistics and Philology
Uppsala, Sweden
`joakim.nivre@lingfil.uu.se`

Abstract

We propose an approach to detecting the rhetorical figure called chiasmus, which involves the repetition of a pair of words in reverse order, as in "**all** for **one**, **one** for **all**". Although repetitions of words are common in natural language, true instances of chiasmus are rare, and the question is therefore whether a computer can effectively distinguish a chiasmus from a random criss-cross pattern. We argue that chiasmus should be treated as a graded phenomenon, which leads to the design of an engine that extracts all criss-cross patterns and ranks them on a scale from prototypical chiasmi to less and less likely instances. Using an evaluation inspired by information retrieval, we demonstrate that our system achieves an average precision of 61%. As a by-product of the evaluation we also construct the first annotated corpus of chiasmi.

1 Introduction

Natural language processing (NLP) automates different tasks: translation, information retrieval, genre classification. Today, these technologies definitely provide valuable assistance for humans even if they are not perfect. But the automatic tools become inappropriate to use when we need to generate, translate or evaluate texts with stylistic quality, such as great political discourse, novels, or pleadings. Indeed, one is reluctant to trust computer assistance when it comes to judging the rhetoric of a text. As expressed by Harris and DiMarco (2009):

> Too much attention has been placed on semantics at the expense of rhetoric (in-

cluding stylistics, pragmatics, and sentiment). While computational approaches to language have occasionally deployed the word 'rhetoric', even in quite central ways (such as Mann and Thompsons Rhetorical Structure Theory (1988)), the deep resources of the millenia-long research tradition of rhetoric have only been tapped to a vanishingly small degree.

Even though the situation has improved slightly during the last years (see Section 2), the gap underlined by Harris and DiMarco in the treatment of traditional rhetoric is still important. This study is a contribution aimed at filling this gap. We will focus on the task of automatically identifying a rhetorical device already studied in the first century before Christ by Quintilian (Greene, 2012, art. antimetabole), but rarely studied in computational linguistics: the chiasmus.

Chiasmi are a family of figures that consist in repeating linguistic elements in reverse order. It is named by the classics after the Greek letter χ because of the cross this letter symbolises (see Figure 1). If the name 'chiasmus' seems specific to rhetorical studies, the figure in itself is known to everybody through proverbs like (1) or quotations like (2).

(1) **Live** not to **eat**, but **eat** to **live**.

(2) Ask not what your **country** can do for **you**; ask what **you** can do for your **country**.

One can talk about chiasmus of letters, sounds, concepts or even syntactic structures as in (3).

Twist **facts** to suit **theories**,
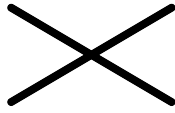
not **theories** to suit **facts**.

Figure 1: Schema of a chiasmus.

(3) Each **throat** was **parched**, and **glazed** each **eye**.

Strictly speaking, we will only be concerned with one type of chiasmus: the chiasmus of identical words (also called antimetabole) or of identical lemmas, as exemplified in (4) and (5), respectively.

(4) A comedian does **funny things**,
a good comedian does **things funny**.

(5) A **wit** with **dunces** and a **dunce** with **wits**

From now on, for the sake of simplicity, we will restrict the term 'chiasmus' to exclusively chiasmus of words that share identity of lemma.

We see different reasons why NLP should pay attention to chiasmi. First, it may be useful for a literary analysis: tools for supporting studies of literature exist but mostly belong to textometry. Thus, they mainly identify word frequency and some syntactic patterns, not figures of speech. Second, the chiasmus has interesting linguistic properties: it expresses semantic inversion thanks to syntax inversion, as in (2) above. Horvei (1985, p.49) points out that chiasmus can be used to emphasize antagonism as in (6) as well as reciprocity as in (7).

(6) **Portugal** has no rivers that flow into **Spain**, but **Spain** does have rivers that flow into **Portugal**.

(7) Hippocrates said that **food** should be our **medicine** and **medicine** our **food**.

To see whether we can detect such interplay of semantics and syntax is an interesting test case for computational linguistics.

Chiasmus is extremely rare. To see this, one can read a book like *River War* by Winston Churchill.

Indeed, despite the fact that this book comes from a politician recognized for his high rhetorical abilities and despite the length of the text (more than one hundred thousand words) we could find only one chiasmus:

(8) **Ambition** stirs **imagination** nearly as much as **imagination** excites **ambition**.

Such rareness is a challenge for our discipline. NLP is accustomed to treating common linguistic phenomena (multiword expressions, anaphora, named entities), for which statistical models work well. We will see that chiasmus is a needle in the haystack problem. For the case of chiasmus, we have a double-fold challenge: we must not only perform well at classifying the majority of wrong instances but above all perform well in finding the infrequent phenomenon.

This paper will promote a new approach to chiasmus detection that takes chiasmus as a graded phenomenon. According to our evaluation, inspired from information retrieval methodology, our system gets up to 61% average precision. At the end of this paper the reader will have a list of precise features that can be used in order to rank chiasmus. As a side effect we provide a partially annotated tuning set with 1200 extracts manually annotated as true or false chiasmus instances.

2 Related Work

The identification of chiasmus is not very common in computational linguistics, although it has sometimes been included in the task of detecting figure of repetition (Gawryjolek, 2009; Harris and DiMarco, 2009; Hromada, 2011). Gawryjolek (2009) proposes to extract every pair of words repeated in reverse order in the text, but this method quickly becomes impractical with big corpora. When we try it on a book (*River War* by Churchill, 130,000 words) it outputs 66,000 inversions: as we shall see later, we have strong reason to believe that only one of them is a real chiasmus.

At the opposite end of the spectrum, Hromada (2011) makes a highly precise detection by detecting three pairs of words repeated in reverse order without any variation in the intervening material as illustrated in (9), but thus he limits the variety of chi-

asmus pattern that can be found and limits the recall (Dubremetz, 2013).

(9) You don't need [**intelligence**] [to have] [**luck**], but you do need [**luck**] [to have] [**intelligence**].

In our example, *River War*, the only chiasmus of the book (Sentence (8)) does not follow this pattern of three identical words. Thus, with Hromada's algorithm we obtain no false positive but no true positive either: we have got an empty output. Finally, Dubremetz (2013) proposes an algorithm in between, based on a stopword filter and the occurrence of punctuation, but this algorithm still has low precision. With this method we found one true instance in *River War* but only after the manual annotation of 300 instances. Thus the question is: can we build a system for chiasmus detection that has a better trade-off between recall and precision?

3 A New Approach to Chiasmus Detection

To make progress, we need to move beyond the binary definition of chiasmus as simply a pair of inverted words, which is oversimplistic. The repetition of words is an extremely common phenomenon. Defining a figure of speech by just the position of word repetitions is not enough (Gawryjolek, 2009; Dubremetz, 2013). To become a real rhetorical device, the repetition of words must be "a use of language that creates a literary effect".[1] This element of the definition requires us to distinguish between the false positives, or accidental inversions of words, and the (true) chiasmi, that is, when the inversion of words explicitly provokes a figure of speech. Sentence (10) is an example of false positive (here with 'the' and 'application'). It contrasts with Sentence (11) which is a true positive.

(10) My government respects the **application** of **the** European directive and **the application** of the 35-hour law.

(11) **One** for **all**, **all** for **one**.

However, the distinction between chiasmus and accidental inversion is not trivial to draw. Some cases

are obvious for every reader, some are not. For instance, it is easy to say that there is chiasmus when the figure constitutes all the sentence and shows a perfect symmetry in the syntax:

(12) Chuck **Norris** does not fear **death**, **death** fears Chuck **Norris**.

But how about cases where the chiasmus is just one clause in a sentence or is not as symmetric as our canonical examples:

(13) We are all **European citizens** and we are all **citizens** of a **European** Union which is underpinned.

The existence of borderline cases indicates the need for a detector that does not only eliminate the most flagrant false positives, but above all establishes a ranking from prototypical chiasmi to less and less likely instances. In this way, the tool becomes a real help for the human because it selects interesting cases of repetitions and leaves it to the human to evaluate unclear cases.

A serious bottleneck in the creation of a system for ranking chiasmus candidates is the lack of annotated data. Chiasmus is not a frequent rhetorical figure. Such rareness is the first reason why there is no huge corpus of annotated chiasmi. It would require annotators to read millions of words in order to arrive at a decent sample. Another difficulty comes from the requirement for qualified annotators when it comes to literature-specific devices. Hammond et al. (2013), for example, do not rely on crowd sourcing when it comes to annotating changing voices. They use first year literature students. Annotating chiasmi is likely to require the same kind of annotators which are not the most easy to hire on crowdsourcing platforms. If we want to create annotations we have to face the following problem: most of the inversions made in natural language are not chiasmi (see the example of *River War*, Section 2). Thus, the annotation of every inversion in a corpus would be a massive, repetitive and expensive task for human annotators.

This also explains why there is no large scale evaluation of the performance of the chiasmus detectors through literature. To overcome this problem we can seek inspiration from another field of computational linguistics: information retrieval targeted at

---

[1]Definition of 'rhetorical device' given by Princeton wordnet: https://wordnet.princeton.edu/

the world wide web, because the web cannot be fully annotated and a very small percentage of the web pages is relevant to a given request. As described already fifteen years ago by Chowdhury (1999, p.213), in such a situation calculating the absolute recall is impossible. However, we can get a rough estimate of the recall by comparing different search engines. For instance Clarke and Willett (1997, p.186), working with Altavista, Lycos and Excite, made a pool of relevant documents for a particular query by merging the outputs of the three engines. We will base our evaluation system on the same principle: through our experiments our different "chiasmus retrieval engines" will return different hits. We annotate manually the top hundreds of those hits and obtain a pool of relevant (and irrelevant) inversions. In this way both precision and recall can be estimated without a massive work of annotation effort. In addition, we will produce a corpus of annotated (true and false) instances that can later be used as training data.

The definition of chiasmus (any pair of inversion that provokes literary effect) is rather floating (Rabatel, 2008, p.21). No clear discriminative test has been stated by linguists. This forces us to rely our annotation on human intuition. However, we keep transparency of our annotation and evaluation by publishing all the chiasmi we consider positive as well as samples of false and borderline cases (see Appendix).

4   A Ranking Model for Chiasmus

A mathematical model should be defined for our ranking system. The way we compute the chiasmus score is the following. We define a standard linear model by the function:

$$f(r) = \sum_{i=1}^{n} x_i \cdot w_i$$

Where r is a string containing a pair of inverted words, x_i is a set of feature values, and w_i is the weight associated with each features. Given two inversions r_1 and r_2, $f(r_1) > f(r_2)$ means that the inversion r_1 is more likely to be a chiasmus than r_2.

4.1   Features

We experiment with four different categories of features that can be easily encoded. Rabatel (2008),

Horvei (1985), García-Page (1991), Nordahl (1971), and Diderot and D'Alembert (1782) deliver examples of canonical chiasmi as well as useful observations. Our features are inspired by them.

1. Basic features: stopwords and punctuation

2. Size-related features

3. N-gram features

4. Lexical clues

We group in a first category what has been tested in previous research (Dubremetz, 2013). They are indeed expected and not hard to motivate. For instance, following the Zipf law, we can expect that most of the false positives are caused by the grammatical words ('a', 'the', etc.) which justifies the use of a stopword list. As well, even if nothing forbids an author to make chiasmi that cross sentences, we hypothesize that punctuation, parentheses, quotation marks are definitely a clue that characterises some false positives, for instance in the false positive (14).

(14) They could **use** the format *:* 'Such-and-such **assurance** , reliable ' , so that the citizens will know that the **assurance** undertaking **uses** its funds well .

We see in (14) that the inversion 'use/assurance' is interrupted by quotation marks, commas and colon.

The second feature type is the one related to size or the number of words. We can expect that a too long distance between main terms or a huge size difference between clauses is an easy-to-compute false positive characteristic as in this false positive:

(15) It is strange that other **committees** <u>can apparently acquire</u> **secretariats** and well-equipped **secretariats** <u>with many staff, but that this is not possible for the Women' s</u> **Committee**.

In 15, indeed, the too long distance between 'secretariats' and 'Committee' in the second clause breaks the axial symmetry prescribed by Morier (1961, p.113)

The third category of features follows the intuition of Hromada (2011) when he looks for three pairs of

inverted words instead of two: we will not just check if there are two pairs of similar words but as well if some other words are repeated. As we see in (16), similar contexts are a characteristic pattern of chiasmi (Fontanier, 1827, p.381). We will evaluate the similarity by looking at the overlapping N-grams.

(16) *In* peace, **sons** *bury their* **fathers**; *in* war, **fathers** *bury their* **sons**.

The last group consists of the lexical clues. These are language specific and, in contrast to stopwords in basic features, this is a category of positive features. We can expect from this category to encode features like conjunction (Rabatel, 2008, p.28) because they underline the axial symmetry (Morier, 1961) as in (17).

(17) **Evidence** of **absence** *or* **absence** of **evidence**?

We capture negations as well. Indeed, Fontanier (1827, p.160) stresses that prototypical chiasmi are used to underline opposition. As well Horvei (1985) concludes that chiasmi often appear in general atmosphere of antagonism. We can observe such negations in (1), (6), (9), and (12).

4.2 Setting the Weights

As observed in Section 3, there is no existing set of annotated data. This excludes traditional supervised machine learning to set the weights. So, we proceeded empirically, by observation of our successive results on our tuning set. We make available on the web the results of our trainings.

5 Experiments

5.1 Corpus

We perform experiments on English. We choose a corpus of politics often used in NLP: Europarl.[2] From this corpus, we take two extracts of two million words each. One is used as a tuning corpus to test our hypotheses on weigths and features. The other one is the test corpus which is only used for the final evaluation. In Section 5.3 we only present results based on this test corpus.

5.2 Implementation

Our program[3] takes as an input a text file and gives as output the chiasmus candidates with a score that allows us to rank them: higher scores at the top, lower scores at the bottom. To do so, it tokenises and lemmatises the text with TreeTagger (Schmid, 1994). Once this basic processing is done it extracts every pair of repeated and inverted lemmas within a window of 30 tokens and passes them to the scoring function.

In our experiments we implemented twenty features. They are grouped into the four categories described in Section 4.1. We present all the features and weights in Table 1 using the notation from Figure 2.

5.3 Evaluation

In order to evaluate our features we perform four experiments. We start with only the basic features. Then, for each new experiment, we add the next group of features in the same order as in Table 1 (size, similarity and lexical clues). Thus the fourth and last experiment (called '+Lexical clues') accumulates all the features. Each time we run an experiment on the test set, we manually annotate as True or False the top 200 hits given by the machine. Thanks to this manual annotation, we present the evaluation in a precision-recall graph (Figure 3).
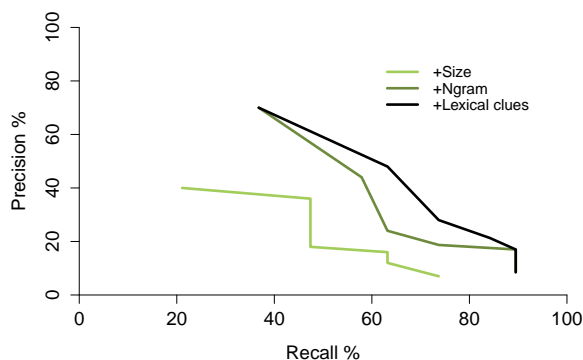

Figure 3: Precision-Recall graphs for the top two hundred candidates in each experiment. (The 'basic' experiment is not included because of too low precision.)

In prehistoric times **women** resembled **men** , and **men** resembled **women**.

$$\underbrace{\text{In prehistoric times}}_{C_{\text{Left}}} \underbrace{\textbf{women}}_{W_a} \underbrace{\text{resembled}}_{C_{ab}} \underbrace{\textbf{men}}_{W_b} \underbrace{\text{,}}_{C_{bb}} \underbrace{\text{and}}_{W'_b} \underbrace{\textbf{men}}_{} \underbrace{\text{resembled}}_{C_{ba}} \underbrace{\textbf{women}}_{W'_a}.$$

Figure 2: Schematic representation of chiasmus, C stands for context, W for word.

Feature	Description	Weight
Basic		
#punct	Number of hard punctuation marks and parentheses in C_{ab} and C_{ba}	-10
#softPunct	Number of commas in C_{ab} and C_{ba}	-10
#centralPunct	Number of hard punctuation marks and parentheses in C_{bb}	-5
isInStopListA	W_a is a stopword	-10
isInStopListB	W_b is a stopword	-10
#mainRep	Number of additional repetitions of W_a or W_b	-5
Size		
#diffSize	Difference in number of tokens between C_{ab} and C_{ba}	-1
#toksInBC	Position of W'_a minus position of W_b	-1
Similarity		
exactMatch	True if C_{ab} and C_{ba} are identical	5
#sameTok	Number of identical lemmatized tokens in C_{ab} and in C_{ba}	1
simScore	#sameTok but normalised	10
#sameBigram	Number of bigrams that are identical in C_{ab} and C_{ba}	2
#sameTrigram	Number of trigrams that are identical in C_{ab} and C_{ba}	4
#sameCont	Number of tokens that are identical in C_{Left} and C_{bb}	1
Lexical clues		
hasConj	True if C_{bb} contains one of the conjunctions 'and', 'as', 'because', 'for', 'yet', 'nor', 'so', 'or', 'but'	2
hasNeg	True if the chiasmus candidate contains one of the negative words 'no', 'not', 'never', 'nothing'	2
hasTo	True if the expression "from . . . to" appears in the chiasmus candidate or 'to' or 'into' are repeated in C_{ab} and C_{ba}	2

Table 1: The four groups of features used to rank chiasmus candidates

Precision at candidate	+Size	+Ngram	+Lex. clues
10	40	70	70
50	18	24	28
100	12	17	17
200	7	9	9
Ave. P.	34	52	61

Table 2: Average precision, and precision at a given top rank, for each experiment.

In Figure 3, recall is based on 19 true positives. They were found through the annotation of the 200 top hits of our 4 different experiments. On this graph

the curves end when the position 200 is reach. For instance, the curve of '+Size' experiment stops at 7% precision because at candidates number 200 only 14 chiasmi were found. The 'basic' experiment is not present because of too low precision. Indeed, 16 out of the 19 true positives were ranked by the 'basic' features as first but within a tie of 1180 other criss-cross patterns (less than 2% precision).

When we add size features, the algorithm outputs the majority of the chiasmi within 200 hits (14 out of 19, or a recall of 74%), but the average precision is below 35%(Table 2). The recall can get significantly better. We notice, indeed, a significant progression of the number of chiasmi if we use N-gram

similarity features (17 out of 19 chiasmi). Finally the lexical clues do not permit us to find more chiasmi (the maximum recall is still 90% for both the third and the fourth experiment) but the precision improves slightly (plus 9% for '+Lexical clues' experiment Table 2).

We never reach 100% recall. This means that 2 of the 19 chiasmi we found were not identifiable by the best algorithm ('+Lexical clues'). They are ranked more highly by the non-optimal algorithms. It can be that our features are too shallow, but it can be as well that the current weights are not optimal. Since our tuning is manual, we have not tried every combination of weights possible.

Chiasmus is a graded phenomenon, our manual annotation ranks three levels of chiasmi: true, borderline, and false cases. Borderline cases are by definition controversial, thus we do not count them in our table of results.[4] Duplicates are not counted either.[5]

Comparing our model to previous research is not straightforward. Our output is ranked, unlike Gawryjolek (2009) and Hromada (2011). We know already that Gawryjolek (2009) extracts every crisscross pattern with no exception and thus obtains 100% recall but for a precision close to 0% (see Section 2). We run the experiment of Hromada (2011) on our test set.[6] It outputs 6 true positives for a total of only 9 candidates. In order to give a fair comparison with Hromada (2011), the 3 systems will be compared only for the nine first candidates (Table 3).

	+Lex. clues	Gawryjolek (2009)	Hromada (2011)
Precision	78	0	67
Recall	37	0	32
F_1-Score	59	0	43

Table 3: Precision, recall, and F-measure at candidate number 9. Comparison with previous works.

---

[4]We invite our reader to read them in Appendix B and at http://stp.lingfil.uu.se/~marie/chiasme.htm

[5]For example, if the machine extracts both: "**All for** one, one **for all**", "**All** for **one**, **one** for **all**" we take into account only the second case even if both extracts belong to a chiasmus.

[6]Program furnished by D. Hromada in the email of the 10th of February. We provide this regex at http://stp.lingfil.uu.se/~marie/chiasme.htm.

Finally, the efficiency: our algorithm takes less than three minutes and a half for one million words (214 seconds). It is three times more than Hromada (2011) (78 seconds per million words) but still reasonable.

6 Conclusion

The aim of this research is to detect a rare linguistic phenomenon called chiasmus. For that task, we have no annotated corpus and thus no possibility of supervised machine learning, and no possibility to evaluate the absolute recall of our system. Our first contribution is to redefine the task itself. Based on linguistic observations, we propose to rank the chiasmi instead of classifying them as true or false. Our second contribution was to offer an evaluation scheme and carry it out. Our third and main contribution was to propose a basic linear model with four categories of features that can solve the problem. At the moment, because of the small amount of positive examples in our data set, only tentative conclusions can be drawn. The results might not be statistically significant. Nevertheless, on this data set, this model gives the best F-score ever obtained. We kept track of all annotations and thus our fourth contribution is a set of 1200 criss-cross patterns manually annotated as true, false and borderline cases, which can be used for training or evaluation or both.

Future work could aim at gathering more data. This would allow the use of machine learning techniques in order to set the weights. There are still linguistic choices to make as well. Syntactic patterns seem to emerge in our examples. Identifying those patterns would allow the addition of structural features in our algorithm such as the symmetrical swap of syntactic roles.

A Chiasmi Annotated as True

1. Hippocrates said that **food** should be our **medicine** and **medicine** our **food**.

2. It is much better to bring **work** to **people** than to take **people** to **work**.

3. The annual reports and the promised follow-up report in January 2004 may well prove that this report, this agreement, is not the **beginning** of the **end** but the **end** of the **beginning**.

4. The basic problem is and remains: social State or liberal State? More **State** and less **market** or more **market** and less **State** ?

5. She wants to feed **chickens** to **pigs** and **pigs** to **chickens**.

6. There is no doubt that **doping** is changing **sport** and that **sport** will be changed due to **doping**, which means it will become absolutely ludicrous if we do nothing to prevent it.

7. Portugal and Spain are in quite unequal positions, given that we are a downstream country, in other words, **Portugal** has no rivers that flow into **Spain**, but **Spain** does have rivers that flow into **Portugal**.

8. Mr President, some **Eurosceptics** are in favour of the Treaty of Nice because **federalists** protest against it. **federalists** are in favour of it because **Eurosceptics** are opposed to it.

9. Companies must be enabled to find a solution to sending their products from the **company** to the **railway** and once the destination is reached, from the **railway** to the **company**.

10. That is the first difficulty in this exercise, since each of the camps expects first of all that we side with it against its enemy, following the old adage that my enemy 's **enemy** is my **friend**, but my enemy 's **friend** is my **enemy**.

11. It is high time that **animals** were no longer adapted to suit their **environment**, but the **environment** to suit the **animals**.

12. That is, it must not be a matter of **Europe** following the **Member** States or of **Member** States following **Europe**.

13. I would like to say once again that the European **Research** Area is not simply the European **Framework** Programme. The European **Framework** Programme is an aspect of the European **Research** Area.

14. We also need to clarify another point, i.e. that we are calling for an international **solution** because this is an international **problem** and in-

ternational **problems** require international **solutions**.

15. Finally, I would like to discuss this commitment from all in the sense that, as Bertrand Russell said, in order to be happy we need three things: the courage to accept the **things** that we cannot **change**, enough determination to **change** the **things** that we can change and the wisdom to know the difference between the two.

16. Perhaps we should adapt the Community **policies** to these **regions** instead of expecting the **regions** to adapt to an elitist European **policy**.

17. What is to prevent **national** parliamentarians appearing more regularly in the **European** Parliament and **European** parliamentarians more regularly in the **national** parliaments ?

18. In my opinion, it would be appropriate if the **European** political parties took part in **national** elections, rather than having **national** political parties take part in **European** elections.

19. The directive entrenches a situation in which **protected** companies can take over **unprotected** ones, yet **unprotected** companies cannot take over **protected** ones.

B Chiasmi Annotated as Borderline Cases

Random sample out of a total of 10 borderline cases.

1. also ensure that beef and **veal** are safer than ever for **consumers** and that **consumer** confidence is restored in beef and **veal**.

2. If all this is not **helped** by a **fund** , the **fund** is no **help** at all.

3. Both men and **women** can be **victims** , but the main **victims** are **women** [...].

4. The **Commission** should **monitor** the agency and we should **monitor** the **Commission**.

5. The more harmless **questions** have been **answered** , but the evasive or inadequate **answers** to other **questions** are unacceptable.

C Chiasmi Annotated as False

Random sample out of 390 negative instances.

1. the charging of environmental and **resource** costs associated with water **use** are aimed at those Member States that still make excessive **use** of and pollute their water **resources** , and therefore not

2. at 3 p.m.) Council **priorities** for the meeting of the United Nations Human Rights Commission in Geneva The next item is the **Council** and Commission statements on the **Council priorities** for the meeting of

3. President , the two basic **issues** are whether we intend to harmonise social **policy** and whether the power of the Commission will be extended to cover social **policy issues** . I will start

4. of us within the Union **agree** on every **issue** , but on this **issue** we are **agreed** . When we are

5. appear that the reference to **regulation** 2082/ 90 may be a departure from the **directive** and that the **directive** and the **regulation** could be considered to

References

Gobinda Chowdhury. 1999. The internet and information retrieval research: a brief review. *Journal of Documentation*, 55(2):209–225.

Sarah J. Clarke and Peter Willett. 1997. Estimating the recall performance of Web search engines. *Proceedings of Aslib*, 49(7):184–189.

Denis Diderot and Jean le Rond D'Alembert. 1782. *Encyclopédie méthodique: ou par ordre de matières, volume 66*. Panckoucke.

Marie Dubremetz. 2013. Vers une identification automatique du chiasme de mots. In *Actes de la 15e Rencontres des Étudiants Chercheurs en Informatique pour le Traitement Automatique des Langues (RECITAL'2013)*, pages 150–163, Les Sables d'Olonne, France.

Pierre Fontanier. 1827. *Les Figures du discours*. Flammarion, 1977 edition.

Mario García-Page. 1991. El "retruécano léxico" y sus límites. *Archivum: Revista de la Facultad de Filología de Oviedo*, 41-42:173–203.

Jakub J. Gawryjolek. 2009. *Automated Annotation and Visualization of Rhetorical Figures*. Master thesis, Universty of Waterloo.

Roland Greene. 2012. *The Princeton Encyclopedia of Poetry and Poetics: Fourth Edition*. Princeton University Press.

Adam Hammond, Julian Brooke, and Graeme Hirst. 2013. A Tale of Two Cultures: Bringing Literary Analysis and Computational Linguistics Together. In *Proceedings of the Workshop on Computational Linguistics for Literature*, pages 1–8, Atlanta, Georgia, June. Association for Computational Linguistics.

Randy Harris and Chrysanne DiMarco. 2009. Constructing a Rhetorical Figuration Ontology. In *Persuasive Technology and Digital Behaviour Intervention Symposium*, pages 47–52, Edinburgh, Scotland.

Harald Horvei. 1985. *The Changing Fortunes of a Rhetorical Term: The History of the Chiasmus*. The Author.

Daniel Devatman Hromada. 2011. Initial Experiments with Multilingual Extraction of Rhetoric Figures by means of PERL-compatible Regular Expressions. In *Proceedings of the Second Student Research Workshop associated with RANLP 2011*, pages 85–90, Hissar, Bulgaria.

William C. Mann and Sandra A. Thompson. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text*, 8(3):243–281.

Henri Morier. 1961. *Dictionnaire de poétique et de rhétorique*. Presses Universitaires de France.

Helge Nordahl. 1971. Variantes chiasmiques. Essai de description formelle. *Revue Romane*, 6:219–232.

Alain Rabatel. 2008. Points de vue en confrontation dans les antimétaboles PLUS et MOINS. *Langue française*, 160(4):21–36.

Helmut Schmid. 1994. Probabilistic Part-of-Speech Tagging Using Decision Trees. In *Proceedings of International Conference on New Methods in Language Processing*, pages 44–49, Manchester, Great Britain.