What are the points? What are the stances?

Decanting for question-driven retrieval and executive summarization

Jean-François Delannoy
University of Ottawa
Ottawa, Ontario, K1N 6N5
Canada
delannoy@site.uottawa.ca

Abstract

Decanter illustrates a heuristic approach to extraction for information retrieval and question answering. Generic information about argumentative text is found and stored, easing user-focused, question-driven access to the core information. The emphasis is placed on the argumentative dimension, to address in particular three types of questions: “What are the points?”, “Based on what?” “What are the comments?”. The areas of application of this approach include: question-answering, information retrieval, summarization, critical thinking and assistance to speed reading.

1 Introduction

Decanter is a prototype to detect and display high-level information from argumentative text. The game is one of situating and contextualizing.

1.1 Queries and Requests

Information requests can be classified by types of questions, bearing for example on: descriptive knowledge (“tell me about Pakistan”), narratives/updates (“what happened in Camp David?”), know-how (“how can I replace the ink cartridge on my XYZ printer?”), evaluation or advice (“Is Netscape 6 stable?”; “Should I install Netscape 6?”).

One can take them on face value or not. In explicitly argumentative, and in loaded topics (like politics) it is in the interest of the user to

have elements of context in the cognitive modeling he/she is doing of the text contents.

Paying due attention to argumentation contributes in two ways:

- by giving contexts to answers, helping qualify them for credibility
- By answering to questions about opinions and stances: what

1.2 Levels of Answering: on topic, on question, with justifications (and references), with a stance

Level zero is answering **on topic**. This has been the only concern of “classical IR” (and still, word-sense disambiguation is not quite there yet...).

Level one of question-answering is then to answer **to the point** semantically or pragmatically (depending of what kind of information need there is, relevance is of a different nature: in a nutshell, answering a practical question can require action-oriented information, but answers a la AskJeeves talking of travel agents when one just wants the distance from Paris to London are waylaid). As I stress heavily in my IR course (Delannoy 2001c) answers, and summaries alike, have to address **relations**, not just concepts. Answers should not just be “about” the keywords, but give the right kind of information: the height, the name, the colour, the description rather than the price, etc. (in many cases, a wholesale description may be judged satisfactory, but the user incurs a post-filtering overhead).

There is another dimension, though: **context**, in a broad sense. Context includes: who gives the answer; on what medium; what is the answer based on in terms of auctorial or demonstrative;

is it convincing for other reasons. A valuable answer is one given with good ... reason: the answer should be rational, i.e. plausible, checkable, supported by authority of source and/or good demonstration. There is here an idea of critical thinking.

Critical thinking (Toulmin 84, Little et al. 89, Mendenhall 90; Aristotle's *Organon*) is the study of formal argumentation, and of what can be accepted reasonably in not-so-formal argumentation (what dose of induction; what *auctoritas*). This is often met with derision by various brands of relativism in and outside academia, although people of this suasion too play the **argumentation game**: they offer elements of proof; they rarely fling totally non-logical-looking rhetoric. But bottom lines there are, and, in the apt words of the title of Little et al. 1989, good reasoning matters!

How to track it then? The next section is about the “decanting” done by our prototype; Section 3 is about retrieving what has been decanted, via questioning.

2 Decanting

The input consists of one or several texts, by one or several authors and possibly mentioning several “actors” (who are also often “utterers”, but).

The general workflow is the following.

1. segment
2. extract entities, in particular the actors
3. detect utterances
4. analyze them argumentatively in a simple way: links of claims-evidence, evidence-evidence (contribute, contrast)
5. infer underlying goals and values (e.g. prioritizing equity over efficiency; immediate goals vs stability)
6. detect polarity: for, against
7. link authors to utterances (who said what), and to points of view (what are the stances)

Topics are registered in a knowledge base (e.g. economics, war, elections) and issues (most efficient course of action, objective measurement of income or turnover, objectivity

of declarations by public figures, etc.). It is considered to implement a module of semi-automatic acquisition: the user, prompted with lists of potential entries, would select and edit them for incorporation into the knowledge base.

The output comes in multiple forms, as selected by the user:

- list of entities
- main structure of the claims
- quotes
- marked-up text (entities, reasoning)
- table of points and stances (and their holders) on the issues at hand.
- extractive summary based on claims rather than evidence) and, classically, position, importance cues and keyword density (Delannoy et al. 98).

3 Representation

3.1 Keys

Three keys give viewing angles on the information: actors, topics, and issues, correspond to basic factual questions a reader may have (Table 1).

key	attributes	question
An actor	utters quotes	What did X say?
	has stances on issues	What does X think on I, i.e. is s/he for or against?
	and stances on courses of action	What does X advocate/propose/support?
	has previsions	What does X think Will happen?
	prioritizes/foregr ounds some values over others	What are (allegedly) X's foremost values?
a topic	involves issues	What are the issues?
	involves courses of action	What are the possible courses of action?
an issue	involves actors situated pro and contra	What are the comments?

Table 1. Keys, attributes, and related questions

3.2 Background knowledge

Some knowledge is pre-encoded or reused from previous processing, and some is built during the analysis.

For repeated analysis of texts on the same topic, the knowledge built can of course be reused.

- list of topics and issues, and the corresponding heuristics used to determine which are expected to be relevant to a given text
- list text types, and associated heuristics
- values: e.g. equity, egalitarianism, vital minimum/income, safety, ethnic identity, personal freedom, access to information, democracy

3.3 Knowledge built with the processing:

- actors in the input text; other entities
- quotes in the text; their association with actors
- claims
- evidence

- association of actors with claims, evidence

4 Processing

The general working is the following.

4.1 Situate and segment the text

- guess text topic, from keywords situating know topics; this is done easily
- segment the text into clauses (the various clauses of the same utterance are then linked)

4.2 Extract elements

- extract entities, in particular the *actors*
- detect utterances

4.3 Assign relations:

- articulate utterance components (main relations: evidence-of, support, contrast)
- assign entity-to-utterance relations (who said what, textually)
- polarities (who is for/contra what; including the author of the document)
- infer underlying goals and values (e.g. identifying, if possible, whether an author prioritizes equity over efficiency; immediate goals vs stability)
- link authors to utterances (who said what), and to points of view (what are the stances).

The program uses a small **knowledge base** about the known topics (e.g. economics, war, elections) and issues (most efficient course of action, objective measurement of income or turnover, objectivity of declarations by public figures, etc.).

The processing uses heuristic rules and pattern-matching to recognize syntactic-semantic patterns, e.g.:

- entities regular expressions
- cues to topic
- syntactic patterns of direct and report speech, to assign quotes
- cues to polarity
- argumentation operators.

It is being considered to implement a module of semi-automatic acquisition: the user, prompted

with lists of potential entries, would select and edit them for incorporation into the knowledge base.

4.4 Querying/Questioning

Various questions can be asked and answered using the structures produced, and especially:

- What? -> What are the points made?
- Why so? -> What are the justifications?
- What are the points of view or comments? (including of the authors themselves)

5 Example 1: Results of Decanting (actual example)

From a simple input:

*Ehud Barak, the Israeli president, said "we want peace".
He added: "This is our main goal."
"We want peace too", OLP Leader Arafat answered.
Arafat added that Barak said that Israel may pull out of Gaza.
Because Barak and Arafat have different standpoints, the peace process is fragile, even though they both want a peaceful resolution.*

we derive the following structures.

ACTORS AND QUOTES

Context1

ref: textname="text1", utterer="John Doe", date="19990101"

```
{  
  Barak [assert]: "we want peace"  
  Barak [assert]: "this is our main goal"  
  Arafat [assert] "we want peace too".  
}
```

Context2 { utterer="Arafat"

```
  Barak [assert] "Israel may pull out of Gaza"  
}
```

```
The peace process will take time [cause_from]  
Barak and Arafat have different standpoints.  
The peace process will take time[detract] Barak  
and Arafat want a peaceful resolution.  
}
```

NB The utterer of the last assertions is the author of the input text. If we process multiple texts, we have to indicate it explicitly (author name)

STANCES

peace [pro] Barak

peace [pro] Arafat

PREDICTIONS

John Doe [predict] the peace process will take time

As of the submission of this article, the prototype detects the quotes but not the stances and contexts (which functionalities are under development).

6 Example 2: "What Are the Comments?" (manual study)

This example is to indicate the kind of comparative output targeted (but not implemented as yet), and the series of linguistic and modeling difficulties involved in producing it. It is based on an excerpt from a BBC bulletin board linked, at the time, from news.bbc.co.uk, called "BBC Talking point", at http://newsvote.bbc.co.uk/hi/english/talking_point.

The case in point was the desirable attitude towards the participation of Jörg Haider's Freedom Party (FPÖ) in Austria in a governmental coalition in February 2000.

Notes on the table

- ID: numbering, for convenience
- No d-author (author of the page or article) is mentioned, as all the texts in this example are from the same page.
- Author: author of the comment; identification if free (may be a pseudonym)
- Statement: original statement
- Marked up statement: statement after insertion of argumentation tags
- Summary, manual: freely rephrased (there is also a summary from the BBC editor, which we do not mention here)
- Arguments: main justifications, rephrased
- Orientation:: here, by convention, pro means "for" Haider's mandate and against sanctions; NOT necessarily in favour of Haider and his party.

id	Author	Statement	Summary, manual	Arguments	Orient.	Notes
1	Nico C. K., Austria	The EU is neither justified nor allowed to isolate Austria. Austria is, after all, a full member in good standing of the EU and its new government has not actually committed any acts contrary to EU principles. If the EU starts policing its members over the outcome of due democratic process, who will police the EU when it gets out of hand?	Sanctions are not justifiable, as Austria is a legitimate member and has done nothing wrong..	<ul style="list-style-type: none"> - sanctions are not justified nor legal - Austria is a member of the EU in good standing - no devious acts - Austria is master at home - counterfactual: if EU at large becomes devious, who will control it? 	p	Pb: the core point maybe less noteworthy or quote-worthy than a justification of it.
2	Jason H., USA	I do not believe the EU and America are over reacting. I think they see the Haider phenomenon as a "virus" that might infect other more important parts of Europe if it does not react strongly now to "quarantine" it.	Haider is like a virus. Yes, isolate Austria	virus epidemics metaphor	p	The author, implicitly, adopts the advice of quarantining the (metaphorical) virus or virus-bearer.
3	Ron, USA	The E.U. is over-reacting. There are two basic ideas Brussels does not get, freedom and liberty. If this "political censorship" is carried out, let us remember it came from the left and not the right. The EU is doomed if these sanctions are carried out. Brussels should let the people decide for once.	Freedom at national level has precedence.	<p>This amounts to political censorship.</p> <ul style="list-style-type: none"> - Decisions should not come from outside or above. 	p	Positively loaded terms: freedom, liberty. (In fact, rather redundant; and semantically pliable). Negatively loaded: political censorship; curiously, appears quoted. Paradox, from implicit knowledge that the left is normally more principled about liberty than the right. Implicit: The EU administration is often not heeding much other levels of decision.
7	Jaya N., India	I think the EU has reacted responsibly and followed through on its earlier statements. When one country acts in such a way as to promote leaders with outright prejudice, the rest of the Nations must do all in their power to subdue further action.	The EU is right, and has been acting consistently, because this is a clear case of prejudice.	Austria (or the FPÖ) is prejudiced	c	Loaded: "outright prejudice." Reasoning from general ("when one country") to particular. Rem: fails to distinguish between prejudice in the FPÖ's policy and supposed prejudice of the country as such or in majority.

7 Evaluation / Commentary

This is prototype work, but several original functionalities are already giving results:

- characterizing the topic, based on discriminating keywords – i.e. the system makes good guesses among a dozen topics including economics/finance, economic policy, conflict, social/labour relations, culture, electoral politics...
- from the topic, predicting typical issues on which stances articulate: for example for economic policy, one may expect stances about deregulation, globalization, interest rates, etc.
- extracting quotes in direct speech gives 60% good results; on indirect speech, this goes down to about 40%.
- stance assignment works at about 50% success (good positives).

Entity-extraction is not particularly original, like finding entities, classifying them, detecting naming equivalences for the entities.

8 Related work

8.1 Philosophy and Critical Thinking

Books on critical thinking (Little et al. 89, Mendenhall 90) use representations of argument structures (e.g. as diagrams) but give no hint as how to automate it, i.e. to go from text to model.

8.2 Linguistics and NLP

While research in linguistics has addressed several brands of “discourse analysis” as dialogue pragmatics and the search for underlying “ideology” or values, there is little in general linguistics about the study of argumentation proper.

Simone Teufel (1999) performs “argumentative zoning” on research papers, finding types of passages like: aim, background, own research, continuation. The result is a colour-coded display of the input, based on an XML markup. Bayes and ngrams are used to perform this classification task. (Interestingly, she finds good agreement between manual annotators, vs various research in summarization failing to

detect “golden standard” summaries.) This is argumentation in a rather specialized (scientific research in AI, i.e., largely, innovation in problem-solving) and shallow (no collation of the points themselves; one-level) sense. In contrast, Decanter is designed to deliver a representation of conclusions and justifications, from several utterers in parallel or in a nested fashion if applicable.

Some work on **summarization**, in particular by Daniel Marcu (Marcu 97) has looked at the “rhetoric” dimension of text, based on RST (Mann&Thompson 88). It produces a detailed and high-quality tree representing the articulation of the text, but it is qualitatively a hybrid: it does not separate argumentation from mere description or narration. The detailed user study and modeling done in (Endres-Nieggemeyer 97) gives little place to argumentation tracking in the summarization process.

(Barker et al. 94) process rules and examples legal text to produce a semantic output then fed to a machine learning system doing generalization and abstraction. Yet it does not consider contexts of utterance.

8.3 Information retrieval

Information has focused even less on argumentation. As indicated above, answering on-topic is useful, but often the user is in fact looking for information which answers a question, which is situated, and which may involve opinions. We know of no work in argumentation-based IR – all the overhead of high-level filtering of argument being left to the user.

8.4 Knowledge Representation and automated reasoning

Some authors in computational linguistics have approached **contexts**. Ballim & Wilks do knowledge representation with nested contexts with Fauconnier’s mental spaces. Moulin uses conceptual graphs to represent spatio-temporal contexts from text. (Recently, a student project in his department has addressed argumentation, it seems, but information is scarce). Recently, a contributor to the CG list, L. Misek-Falkoff, asked for tools to represent nested contexts in

tort/defamation; there were some answers pointing to tools, but not to tools capable of doing this.

Various studies of reasoning, on the legal domain like (Bench-Capon 97) or more general like (Zukerman et al. 99), represent sophisticated reasoning, without performing extraction from text.

(Delannoy 99) proposed an XML mark-up scheme for argumentation as such, the idea being to flag it inside the text besides producing a separate representation. Decanter is designed to do both.

9 Future work

Further work is intended to address a variety of robustness and scope issues, including reference resolution (neglected in IR) and the detection of lexicalized irony in the expression of stances.

9.1 More Manual Analysis

I am currently working on digesting several argumentative corpora on

- issues of drug legalization (in Delannoy 2001b)
- Colombia (Plan Colombia, conflict, violence)
- the Digital Divide (i.e. low access to the Internet by segments of Canadian or world population).

9.2 Reference resolution

This is another neglected topic in IR. Even medium-quality reference resolution would enhance performance in IR, including in our approach.

9.3 Indirect argumentation and irony

Indirect argumentation, especially irony

Irony is an ingredient of rhetoric and can be of use in tracking stance on topics, stances on other actors, and also style of course. In another study (Delannoy 2001b) I observe the alternating use of irony and indignation. Besides the direct interest as a study rhetoric, it shows the variance of one factor of enunciation the socio-psychological attitude, while the doxastic-epistemic attitude stays aligned (the stance).

From an IR point of view, one could try to differentiate ironic from non-ironic passages; also to normalize them into a “just the stance” form – a desalination device of sorts!

10 Conclusion

IR and NLP should pay due attention to question-focused information of course, but to other textual elements participating in the value of the returns, both 1) when it gives a useful characterization of the usability of the answer – as plausible, corroborated, demonstrated, novel, etc. 2) to begin to answer questions never addressed in IR and CL but definitely pervasive in user needs, either easily phrasable, in the style: “Is Netscape a good tool?”, “Is it advisable to buy Microsoft stock soon?”, or as a more underlying information goal: “So, what is Le Monde saying about the new developments of Plan Colombia and about the political reactions?”. This second type can be useful both to interested laymen and to professionals of information and politics.

Moreover, a matrix presentation as in example 2 can be quite useful and reusable. That is, to be even more useful, argumentation analysis should integrate information retrieval + analysis + aggregation.

In a Baconian vein: The information retriever and questioner has to use Invention (IR techniques) and Judgment (critical thinking) to tap into Memory (writing, library science) and Tradition (corpus of knowledge, opinions). Decanter opens the way to the necessary contribution of Judgment in Invention.

11 References

- Aristotle. *Organon*.
Ballim A & Wilks. 1991: *Artificial Believers*, Lawrence Erlbaum Associates
K. Barker, JF Delannoy, S. Matwin, S. Szpakowic. 1994: "From text to Horn clauses", *Proceedings of the Canadian Conference on Artificial Intelligence (AI/GI/CV '94)*, Banff, Alberta, Canada, March 1994, pp. 9-16
Bench-Capon, T.J.M. 1997. *Argument in Artificial Intelligence and Law Artificial*

Intelligence and Law, Vol 5 No 4., 1997, pp249-61.

Delannoy JF. 1999. "Argumentation Mark-Up: A Proposal", Workshop "Towards Standards and Tools for Discourse Tagging", *Conference of the Association for Computational Linguistics (ACL'99)*, U. Maryland, College Park, MD, June 22, 1999

Delannoy JF. 2001b. "Arguing about drugs", OSSA 2001 (conference of the Ontario Society for the Study of Argumentation), Windsor, Ontario, May 2001

Delannoy JF. 1200c. course material for CSI4107, Information Retrieval and the Internet, University of Ottawa:

<http://www.site.uottawa.ca/~delannoy/csi4107>

Endres-Nieggemeyer, Brigitte. 1997. *Summarizing Information*, Springer, 1997

Fauconnier, G. 1985. *Mental Spaces: Aspects of meaning construction in natural language*. MIT Press, Cambridge, MA.

Little J, Groarke L, Tindale C. 1989. *Good reasoning matters!*, McClelland & Stewart

Mann, W. & Thompson, S. 1988. "Rhetorical structure theory: Towards a functional theory of text organization", *Text* 8(3), 241-281

Marcu, Daniel. 1997. *The Rhetorical Parsing, Summarization, and Generation of Natural Language Texts*, Ph.D. Dissertation, U Toronto, 1997

Mendenhall V. 1990. *Une introduction à l'analyse du discours argumentatif*. Presses de l'Université d'Ottawa, 1990

Teufel, S. 1999. *Argumentative Zoning. Information Extraction from Argumentative Text*. Ph.D. Thesis, U. Edimburg, 1999

Toulmin S, Riek R, Janik A. 1984. *An Introduction to Reasoning*, MacMillan

Zukerman, I., McConachy, R., Korb, K. and Pickett, D. 1999. Exploratory Interaction with a Bayesian Argumentation System. *IJCAI99 Proceedings (16th International Joint Conference on Artificial Intelligence)*, pp. 1294-1299, Stockholm, Sweden, Morgan Kaufmann Publishers