Computational Models (of Narrative) for Literary Studies

di Antonio Lieto

In the last decades a growing body of literature in Artificial Intelligence (AI) and Cognitive Science (CS) has approached the problem of narrative understanding by means of computational systems. Narrative, in fact, is an ubiquitous element in our everyday activity and the ability to generate and understand stories, and their structures, is a crucial cue of our intelligence. However, despite the fact that – from an historical standpoint – narrative (and narrative structures) have been an important topic of investigation in both these areas, a more comprehensive approach coupling them with narratology, digital humanities and literary studies was still lacking.

With the aim of covering this empty space, in the last years, a multidisciplinary effort has been made in order to create an international meeting open to computer scientist, psychologists, digital humanists, linguists, narratologists etc. This event has been named CMN (for Computational Models of Narrative) and was launched in the 2009 by the MIT scholars Mark A. Finlayson and Patrick H. Winston[1].

From a technological and cognitive perspective, the original goal of the CMN community (see Finlayson et al. 2015) is to explain intelligence through the understanding of how narrative elements and structures are stored, manipulated and processed by natural and artificial minds. In the last years, however, with the explicit goal of extending the classical approach to narrative studies, the CMN community has converged towards a renewed research framework aiming at additionally investigate the cross-relationships with sister disciplines through the development of a mutual loop of common interests. In particular: while cognitive science and artificial intelligence can provide computational models of narrative, i.e. in terms of reader's modeling (cf. cognitive narratology); on the other hand, narratological and literary studies can provide relevant insights for modelling artificial systems, i.e. they could furnish access to the "keys" adopted by expert scholars for interpreting the literary world and reasoning about it (e.g. «reading between the lines», for example, is a crucial capability that artificial systems are not yet able to perform and that, on the other hand, scholars in literary studies, given their background, are able to do without particular difficulties).

For sake of clarity, an illustrative selection of questions of interest for the CMN community is reported below:

- How can computational narratives be studied from a humanities point of view?
- Are generative models of narrative texts, movies or video games possible, desirable, and useful?
- What comprises the set of possible narrative arcs? Is there such a set? How many possible story lines are there?
- Is narrative structure universal, or are there systematic differences in narratives from different cultures?

- How are narratives affected by the media used to convey them?
- How are narratives indexed and retrieved? Is there a universal scheme for encoding episodes?
- What impact do the purpose, function, and genre of a narrative have on its form and content?
- Are there systematic differences in the formal properties of narratives from different cultures?
- Is there a recipe (a` la Joseph Campbell or Vladimir Propp) for generating narratives?
- What are the appropriate representations of narrative? What representations underlie the extraction of narrative schemas from the blooming, buzzing confusion of the world?
- How should we evaluate computational models of narrative?

By leaving aside some aspects that are of specific interest for the AI and Cognitive Science communities, in the following I will argue that some of the computational methods and tools adopted by the CMN community can be beneficial for the scholars working in the area of Digital Humanities (in general) and, in particular, for those interested in Literary Criticisms studies. An important remark: arguing in favour of the use of some techniques and modelling tools for literary criticism does not imply, on my side, a general endorsement to a strong view of the Distant Reading approach (Moretti, 2013). On the contrary, in my opinion, computational systems and techniques coming from AI, and adopted by the CMN community, could be used to support (and not to substitute) literary scholars and experts in their exploratory or analytical phases of research. The interpretation of the obtained results, however, is something that cannot be currently demanded to autonomous artificial systems.

Additionally, another important aspect to take into account regards the problems and the issues that arise when one tries to apply computational method developed for quite different domains to the analysis of humanities and literary objects and phenomena. I will not provide here any kind of epistemological discussion about such problems (that, however, exist). In my understanding the only way to deal with these complex issues is through a strong cross-collaboration between computational scholars, humanists and digital humanists.

With this goal in mind, I will provide some examples of the methods and techniques coming from AI that

have been applied to the analysis of literary "texts". I will use the term "text" or "textual" here in a wide semiotic perspective (Fabbri and Marrone, 2000, pp. 7-11). Therefore also movies, paintings, pictures and not just books o written productions, can be considered as such.

Computational Methods for Literary Texts

One of the main areas of research of the CMN community regards the application of Natural Language Processing techniques for the semantic comprehension of written texts. For this specific purpose, a plethora of methods and techniques have been developed coming from the Computational Linguistics arena and from Machine Learning. Such methods include hierarchical decomposition via grammars (P. Thorndyke, 1977), distributional semantics (Lenci, 2008) and many others. In the following I will provide a very short example concerning the use of the so called «Text Mining» techniques.

Text Mining

Text Mining (see Feldman & Sanger, 2007) is a particular discipline of Computational Linguistics aiming at processing and automatically extracting relevant information from large amounts of unstructured textual data. Text mining methods and tools have been and are currently used for tasks such as: genre recognition, plagiarism detection, computation of similarity between documents, etc. There are several existing algorithms developed for dealing with this aspects and able to formally process textual documents. For the purposes of this article I will not focus on them since they are extensively treated in the specialistic literature. However, I will just sketch the general idea behind these approaches in order to suggest in which way they could help literary scholars in their everyday activities.

In Text Mining applications, usually, textual documents (e.g books, poetry, novels etc.) are transformed in vector representations (see figure 1). Such representations allow the creation of clusters of documents (i.e. the automatic creation of groups of documents sharing more or less a similar content according to the analysis provided at the linguistic level).
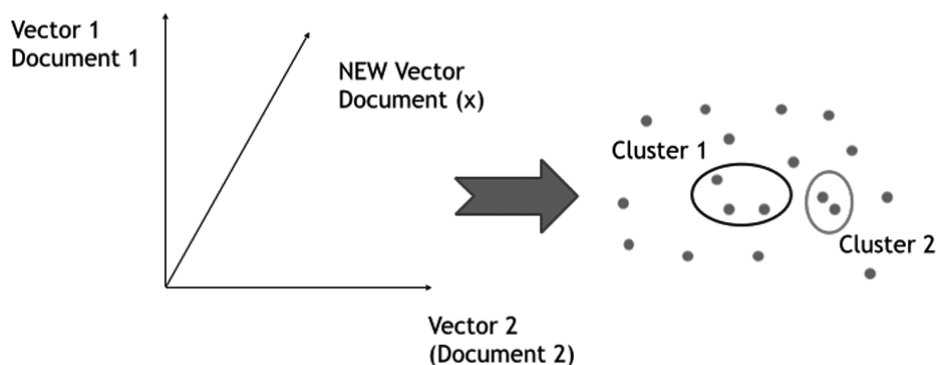
Figure 1. A pictorial account of a Text Mining Pipeline.

The figure 1 below shows the general idea of this approach. In input we have texts. Such texts/documents are then usually transformed in vector representations where each vector reports the co-occourence of each word within the document (e.g. the Vector 1 represents the co-occourrence of the words it contains, the same for Vector 2 and so on). This kind of representations (usually called BOW: bag-of-words representations) allow to compare, in a given collection, each document with each other by using some vector-similarity metrics (e.g. the cosine similarity or other known metrics). Additionally, they allow to perform some «inferences» about which document is «semantically» closer to which one by resorting to a mix of clustering techniques. For example: in Figure 1 the result of the application of clustering techniques to bag of word representations of texts allows to group some documents. In this case, the Cluster 1, in figure 1, connects 3 texts/documents[2] sharing the same topic (e.g. let us assume that they are horror novels). On the other hand, Cluster 2 is composed by 2 documents (e.g. let us assume that they are Noir novels since the two clusters are quite close each other).

This kind of very simple analysis, operating only at the surface level of words, is used for many tasks. As mentioned, it could be useful for literary scholars for: genre recognition, stylistic comparisons, authorship attribution and plagiarism detection (these two tasks are two sides of the same coin) etc.. In this way it could be possible, in fact, to process a huge amount of textual information and compare the obtained output with a more detailed and hand-crafted analysis conducted by literary scholars.

Social Network Analysis

Another crucial method adopted by the CMN community for the narrative understating of literary texts with computational tools is the Social Network Analysis (SNA). In narratology, SNA has been used mostly as a new instrument for the study of plot evolution. By the extraction of the interactional networks of characters from narrative works and the subsequent synthesis of the obtained data in network graphs it is possible to open a whole new perspectives to better comprehend the dynamics and the structure of a narrative plot. Even excluding the numerous quantitative analysis options available, the mere rearrangement of the narration from the written context to a visual and understandable display represents a powerful explanatory enhancement. SNA analytical approach has already been employed with a fairly large selection of different literary text, ranging from Shakespeare's tragedies (Moretti, 2011) to Lewis Carroll's Alice in Wonderland (Agarwal et al. 2012) including the whole Marvel comics universe (Alberich et al. 2002).

Recently the SNA has been employed for studying a text as structured and as complex as I Promessi Sposi (Bolioli et al. 2014). This work, in particular, focused on the extraction and visualisation of «conversational edges», where an edge is formed between two characters/nodes every time the studied text features an explicit dialogue between aforementioned characters (the text of the dialogues from which the visualisation network is created was obtained by the authors combining both manual and semi-automatically extracted annotations).
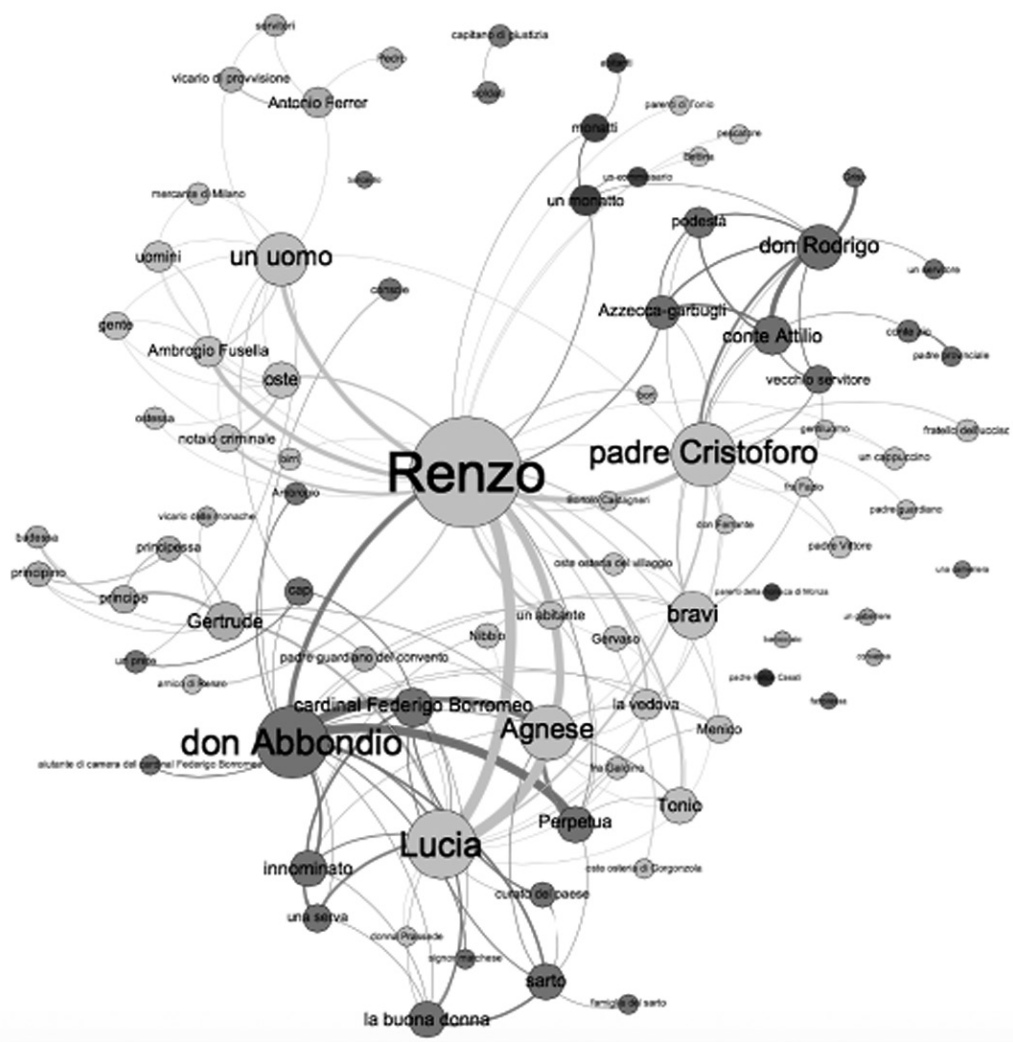
Computational Models (of Narrative) for Literary Studies

Figure 2. Conversational network of "I Promessi Sposi" obtained through SNA.

The graph in Figure 2 represents the complete conversational network of *I Promessi Sposi*, elaborated with the SNA visualization software Gephi[3]. The dimension of each node is due to the number of conversational interactions («degree» in SNA terminology) in which that node's character is involved: in fact the largest node, in the very middle of the graph, embodies Renzo Tramaglino, one the main characters of the novel. The thickness of each 'edge' is directly correlated to the number of interactions between the relative couple of nodes: if a dyad of characters shares many dialogues, the edge that tie their node will be more thick. Finally, different colors distinguish each community of characters.

The example provided above shows how the interpretation of the obtained results by expert scholars in literary studies is crucial for the understanding of the output processed by such kind of systems. In this case, in fact, a 'wrong' conclusion would be that Renzo is the most important character of the whole plot. This, however, would be trivially false. In fact: the system, in this case, only visualizes the network of dialogues and discussions between characters. In other words: being more 'silent' does not imply to be less important from a narrative perspective (in fact, usually, the contrary is true).

Ontologies

Another method developed from the AI tradition and of potential interest for literary scholars is represented by the possibility of using 'formal ontologies' as conceptual models upon which to interpret and organize the informational content extracted from narrative texts (i.e. texts having a plot based on characters, character's roles, stories, complex events etc.). In the last century many attempts have been made to classify contents according to their narrative motifs, patterns or grammars, witnesses the importance of the narrative theory for indexing and organizing literary and non literary works. Some notable efforts are given by Thompson's motif index of folk literature (Thompson, 1955), by Propp's well known work on Russian fairy tales, by Polti's attempt to systematize the patterns of drama. In cultural studies, an attempt at content classification has been accomplished by Warburg, whose BilderAtlas (Warburg, 2008), though based on visual representation, encompasses a number of narrative themes many of which issued from classical mythology.

In AI, and more specifically in its subfield known as Knowledge Representation (KR), a large variety of approaches have been proposed as well to represent formal models able to encode narrative structures (e.g. rules, frames, scripts, semantic nets, conceptual graphs, etc). In this specific field, the term «ontology» is referred to «an engineering artifact, constituted by a specific vocabulary used to describe a certain reality, plus a set of explicit assumptions regarding the intended meaning of the vocabulary words» (Guarino, 1998).

The main building blocks of ontological models are, therefore, concepts (or classes), roles (or properties), and individuals describing a given domain. For example, by using this alphabet, given an individual called 'Dante Alighieri' belonging to the class 'HUMANS' it is possible to describe, through the predicate 'was_in_ love', that he felt in love with another individual called 'Beatrice' and so on.

In other words: ontologies provide an explicit reference domain model (where a domain can be also a fictional one: e.g. it is possible to build an ontological model of the Harry Potter's world). Such model is used to interpret and organise the information coming from 'Textual' Data and to perform simple forms of automatic reasoning (as specified before here the term 'text' is used in its semiotic interpretation). The main

automatic reasoning tasks performed in the ontologies are: categorization (the process regarding the class membership assignment to specific individuals: e.g. it is possible to infer that the individual 'Beatrice' belongs to the class 'HUMANS' even if this information is not explicitly provided) and classification (the process through which new subclass relations are inferred: e.g. in our simple example MEN and WOMEN can be automatically recognised as subclasses of HUMANS).

The figure 3 below shows the general idea of this approach. We can consider that different types of "texts" (e.g. *La Divina Commedia* or the painting *Minotauromachia* of Pablo Picasso) need to be annotated (manually, semi-automatically or automatically) through a predefined metadata schema (e.g. Dublin Core or others) and embedded within a given ontological model in order to populate it with real data. Such embedding allows the populated ontological model to perform some simple forms of automatic reasoning.

An example of a running ontology-based system used for narrative content organization and fruition is Labyrinth[4]. Labyrinth relies on an ontology of Archetypes (i.e. general 'narrative structures' adopted to provide different 'views' on the same textual object) describing a set of related stories, characters, locations and objects which share some symbolic meaning. (see Damiano and Lieto 2013 for details). Typical examples of archetypes are: the 'labyrinth', the 'journey', the 'hero', etc.. Such archetypes allows to ground different multimedia textual elements contained in a standard Database (annotated according to the Dublin Core standards) to a common narrative ontological model. This grounding allows to organise, and to explore, the cultural multimedia archive (containing books, videos, pictures etc.) according to the shared narrative elements emerging by means of the forms of automatic reasoning performed by the ontology (e.g. in the example in figure 4 both the paintings *The Death of Marat* and *Minotauromachia* have been associated to narrative action of 'killing' since they both displays killing stories). The resulting framework lends itself to the creation of personalized navigation paths in cultural repositories, represented in digital form, for the sake of exploration and study.

Interestingly, the output produced by such kind of ontological systems can also be presented through different kinds of visualization interfaces (see Damiano, Lieto and Lombardo, 2014; Damiano, Lombardo, Li-
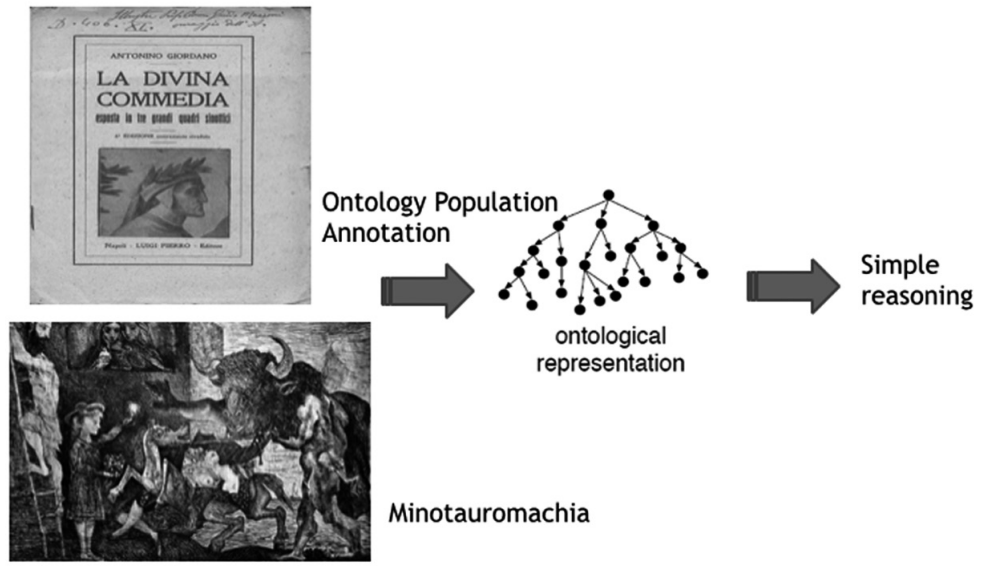
Figure 3. Pipeline for the adoption of ontology based systems in a literary context.

Figure 4. An example of the Labyrinth Interface.

eto, 2015) thus providing a potential service similar to that one provided by the SNA methods.

Additionally, the same framework could be plausibly applied to the analysis of written texts such as: books, poetry collections, novels etc. thus providing an useful support for literary scholars and students. In this case the annotation process would interest the paragraphs and the chapters of each text and would allow an enhanced semantic search within them.

Summary and Future Work

In this paper I have showed some examples of how computational AI methods and techniques adopted by the interdisciplinary group composing the Computational Models of Narrative community can be usefully applied to topics of interest for literary critics scholars. In particular, I have tried to provide some evidences regarding the usefulness of such methods to support the research activity of experts in literary studies since, in my opinion, time has come for a real mutual collaboration between disciplines!

As a future fork, a common ground of particular interest that I envisage is represented by the possibility of combining the computational efforts coming from the Digital Humanities community with that ones adopted within the AI frameworks. In this respect a great potential is represented by the possible integration of the TEI (Text Encoding Initiative) with Ontological Schemas[5]. In particular: the realisation of a communication grid between TEI and Ontological Schemas would open the possibility to automatically mapping the textual information marked in TEI within ontological models, thus allowing, *de facto*, a massive population of such models. Such population would open the literary texts to many kinds of potential automatic analysis that could help humanists, and literary scholars, in shading new lights on some elements remained hidden and that, if unveiled, could lead to the development of novel and original research lines.

References

Agarwal A., Corvalan A., Jensen J., Rambow O., «Social network analysis of Alice in Wonderland» in *Proceedings of the NAACL-HLT 2012 Workshop on Computational Linguistics for Literature*, Association for Computational Linguistics, 2012, pp. 88-96.

Alberich R., Miro-Julia J., Rosselló F., *Universe looks almost like a real social network,* arxiv Preprint 2002. http://arxiv.org/abs/cond-mat/0202174)

Bolioli A., Casu M., Lana M., Roda R. (2013), «Exploring the Betrothed Lovers» in *Proceedings of Computational Models of Narrative 2013, CMN'13* (2013), pp. 30-5.

Damiano R., Lieto A., «Ontological representation of narratives: a case study on stories and actions», in *Proceedings of Computational Models of Narrative 2013, CMN'13* satellite workshop of the 35th Meeting of the Cognitive Science Society CogSci 2013, Hamburg 4-6 August, 2013. Hamburg.

Damiano R., Lieto A., Lombardo V. (2014, July). «Ontology-based visualisation of cultural heritage», in *Complex, Intelligent and Software Intensive Systems (CISIS), 2014 Eighth International Conference* IEEE, pp. 558-63.

Damiano R., Lombardo V., Lieto A., *Visual metaphors for semantic cultural heritage,* in «EAI Endorsed Trans. Future Intellig. Educat. Env.» 1(4)/e4 (2015).

Fabbri P., Marrone G., *Semiotica in nuce*, vol. I*, fondamenti e l'epistemologia strutturale*, Roma, Meltemi 2000.

Feldman R., Sanger J., *The text mining handbook: advanced approaches in analyzing unstructured data*, Cambridge, Cambridge University Press 2007.

Finlayson M.A., Miller B., Lieto A., Ronfard R., *Proceedings of Computational Models of Narrative, CMN 2015*, Atlanta, Georgia, USA in *OASIcs-OpenAccess Series in Informatics*, Schloss Dagstuhl-Leibniz-Zentrum für Informatik 45 (2015).

Guarino N., «Formal Onthology in Information Systems», in *Proceedings of the First International Conference* (FOIS'98), June 6-8, Trento, Italy, volume 46 (1998), IOS press.

Lenci A., «Distributional semantics in linguistic and cognitive research. From context to meaning» in *Distributional models of the lexicon in linguistics and cognitive science*, special issue of «The Italian Journal of Linguistics» 20/1 (2008), pp. 1-31.

Moretti F., *Network theory, plot analysis*, in «New Left Review» 68 (2011), pp. 80-102.

Moretti F., *Distant reading*, London, Verso Books 2013.

Ore C.E., Eide Ø., *TEI and cultural heritage ontologies: Exchange of information?*, in «Literary and Linguistic Computing», 24/2 (2009), pp. 161-72.

Thompson S., *Myths and folktales*, in «The Journal of American Folklore», 68/270 (2005), pp. 482-8.

Thorndyke Perry W., *Cognitive structures in comprehension and memory of narrative discourse*, in «Cognitive Psychology», 9/1 (1977), pp. 7-110.

Warburg A., *Der Bilderatlas Mnemosyne*, vol. I, Berlin, Akademie 2008.

Note

[1] The Seventh Workshop on Computational Models of Narratives (CMN 2016) will be co-located with the International Conference on Digital Humanities in Krakow (DH 2016). For more information see: http://narrative.csail.mit.edu/cmn16/

[2] After the application of clustering techniques, the 'points' in the cluster synthesize the information contained in the vectors.

[3] The software is downloadable free of charge at: http://gephi.github.io

4   A demo page of the system is available at http://app.laby-rinth-project.it:8080/LabyrinthTest/#. The Labyrinth project (2012-2014) has been supported by Regione Piemonte, Poli di Innovazione, Polo per la Creatività Digitale e la Multime-dialità, POR_FESR 2007-2013 (P.I. Rossana Damiano) and the resulting system has been developed by the University of

Turin (Department of Computer Science).

5   Despite some initial efforts have been done in this direction (see C.E. Ore, Ø. Eide, *TEI and cultural heritage ontologies: Exchange of information?*, in «Literary and Linguistic Compu-ting», 24/2 (2009), pp. 161-72) they are not satisfactory and had no impact on the current practices.

The Mechanic Reader