Journal of Computational Literary Studies

OPEN ACCESS

conference version

BookNLP-fr, the French Versant of BookNLP

A Tailored Pipeline for 19th and 20th Century French Literature

Frédérique Mélanie-Becquet[1]
Jean Barré[1]
Olga Seminck[1]
Clément Plancq[2]
Marco Naguib[3]
Martial Pastor[4]
Thierry Poibeau[1]

1. Lattice UMR 8094, École Normale Supérieure - PSL - CNRS - Université Sorbonne Nouvelle, Montrouge, France.
2. MSH Val de Loire UAR 3501, CNRS - Université de Tours - Université d'Orléans, Tours, France.
3. LISN, Université Paris-Saclay and CNRS, Orsay, France.
4. Centre for Language Studies, Radboud University, Nijmegen, The Netherlands.

conference version

Citation
tba (2024). "BookNLP-fr, the French Versant of BookNLP. A Tailored Pipeline for 19th and 20th Century French Literature". In: *CCLS2024 Conference Preprints* 3 (1). 10.26083/tuprints-00027396

Date published 2024-05-28
Date accepted 2024-04-15
Date received 2024-01-25

Keywords
Natural Language Processing, Computational Literary Studies, French Literature, Coreference Resolution, Entity Recognition, Subgenre Classification

Note
This paper has been submitted to the conference track of JCLS. It has been peer reviewed and accepted for presentation and discussion at the 3rd Annual Conference of Computational Literary Studies at Vienna, Austria, in June 2024.

Abstract. This paper presents BookNLP-fr: the adaptation to French of BookNLP, an existing NLP pipeline tailored for literary texts in English. We provide an overview of the challenges involved in the adaptation of such a pipeline to a new language: from the challenges related to data annotation up to the development of specialized modules of entity recognition and coreference. Moving beyond the technical aspects, we explore practical applications of BookNLP-fr with a canonical task for computational literary studies: subgenre classification. We show that BookNLP-fr provides more relevant and – even more importantly – more interpretable features to perform automatic subgenre classification than the traditional *bag-of-words* approach. BookNLP-fr makes NLP techniques available to a larger public and constitutes a new toolkit to process large numbers of digitized books in French. This allows the field to gain a deeper literary understanding through the practice of distant reading.

1. Introduction

The domain known as Computational Humanities has recently emerged, with the availability of large corpora of literary texts in digitized format, and of transformer-based language models that are quick, robust and (generally) accurate (Devlin et al. 2019; Touvron et al. 2023, e.g.). This situation opened up new opportunities for exploration and analysis. For French, the collection *Literary fictions of Gallica* (Langlais 2021) includes 19,240 public domain documents from the digital platform of the French National Library, enabling researchers to navigate the wide diversity of literature with unprecedented ease.

The sheer volume of digitized texts presents a unique set of challenges. Traditional methods of literary analysis and interpretation are insufficient when confronted with such vast corpora. It is no longer feasible for individuals to manually analyze in close reading the entirety of these collections. This shift in scale necessitates the development of innovative tools and technologies, particularly Natural Language Processing (NLP). These tools are essential for extracting meaningful insights from digital corpora. They can illuminate patterns, trends, and connections that would be impractical or impossible for humans to discern within the vast amount of text data. This new technical paradigm opens up the possibility of conducting research through distant reading (Moretti 2000; Underwood 2019), enabling scholars to zoom in and out from the literary past, facilitating a more profound comprehension of trends and patterns that delineate the evolution of literature. The knowledge embedded in these digitized literary corpora is crucial not only for literary scholars but also for those interested in cultural analytics, defined as ”the analysis of massive cultural datasets and flows using computational and visualization techniques” by Manovich (2018), or more practical applications for example the automatic production of book summaries for catalogs (Zhang et al. 2019). The evolution of literature is intricately tied to the broader shifts in society, and digitized texts offer a unique opportunity to study these transformations.

To make the analysis of such large corpora possible, BookNLP (Bamman 2021) has been proposed as a specialized software solution adapted to literary texts. It includes the analysis of entities, coreference, events, and quotations within textual data. Originally conceived at the University of California, Berkeley in 2014 by David Bamman and his team, BookNLP has undergone continuous enhancements, aligning with the latest advancements in natural language processing. Notably, it has embraced emerging technologies such as integrated embeddings of large language models, more specifically BERT (Devlin et al. 2019) in early 2020.

The ongoing evolution of BookNLP extends beyond its initial scope, as efforts are underway to expand its applicability to five additional languages through the Multilingual BookNLP project (Bamman 2020). However, it's worth noting that French is not included in this extension. In response to this gap, it was decided in 2021, in coordination with Berkeley, to develop a dedicated French version of BookNLP. The goal is that researchers working with French literary data have access to basic tools required for the structured analysis of fiction. This paper thus presents the French BookNLP project, the related annotated corpus and the pieces of software defined within the project, as well as a specific study illustrating how BookNLP can be used for literary studies.

The structure of the paper is as follows: we start with a literature review in which we specify NLP tools and techniques that are of particular interest in a framework for distant reading (section 2). Special attention will be given to results of the English BookNLP project (subsection 2.2). In section 3, we provide a detailed description of how we elaborated the pipeline of BookNLP-fr: the training data, the annotation process and the software development. In Section 4, we give the evaluation scores of our pipeline on the subtasks of *entity recognition* and coreference resolution. Then, we will present a case study where we used BookNLP-fr for the classification of literary genre (section 5). We finish this article with a discussion about how the use of computational methods and the framework of distant reading using imperfect annotations affects the field of literary studies (subsection 6.1) and its perspectives in the era of *Large Language Models* (subsection 6.2) and finally summarize the paper in the conclusion (section 7).

2. Literature Review

2.1 Computational Methods Applied to Literary Text Analysis

Statistical methods have been used extensively in literary text analysis to identify patterns and trends in large amounts of textual data. Different pieces of software are available for this, for example: Quanteda (Benoit et al. 2018), stylo (Eder et al. 2016), TidyText (Silge and Robinson 2017) or Voyant tools (Rockwell and Sinclair 2016), to cite the most famous. They are available "off the shelf", which means that they can be used directly by scholars and researchers to analyze texts. These tools can handle raw text directly, or after basic NLP-processes such as lemmatization, part-of-speech-tagging, or other kinds of annotations. They offer various visualizations to interpret the texts, such as dendrograms to represent the 'distance' between various books of a corpus or charts that make it visible what type of vocabulary is typical to one author as opposed to another one.

There are clear benefits in using statistical methods to analyze literary texts, such as the ability to process and analyze large amounts of data quickly and efficiently, to identify patterns and trends that might not be apparent through traditional close reading methods, and to generate new research questions and hypotheses. But NLP is needed to better represent the content of the text, i.e. what the text says behind the words used. Natural language processing techniques can be used to annotate literary texts by providing syntactic and semantic annotations. NLP has become an increasingly important tool in the field of literary studies, providing new methods for analyzing and interpreting literary texts. NLP tools (e.g. NLTK (Bird et al. 2019) or Stanford tools (Manning et al. 2014)) have been used to perform a wide range of tasks, including part-of-speech tagging, syntactic analysis, named entity recognition, etc. In the following paragraphs, we will specify the linguistic analyses available by the BookNLP pipeline: entity recognition, coreference resolution, event recognition and quotation detection. The tools mentionned in the paragraph above do not propose these type of semantic analyses, and only use morphological and grammatical linguistic analyses. BookNLP thus occupies a special niche and provides more semantically-oriented annotations.

Entity Recognition. Entity recognition, along with coreference resolution, is of prominent importance, since it makes it possible to track characters, their actions and their relationships over time. Named entity recognition is a well-established task in NLP, referring to the recognition of persons, locations, companies and other institutions, etc. (Maynard et al. 2017) and systems exist for a wide array of languages (Emelyanov and Artemova 2019), with generally good performance, depending of course on the nature of the document to be analyzed and of the gap between training data and target data. Recognizing mentions referring to characters in a novel shares many features with named entity recognition, but is more varied (not all characters have a name, and a character can correspond to an animal, for example). Locations are also of the utmost importance to track the movements of characters (Ryan et al. 2016), but also to detect events. Note that performance may vary greatly depending on the nature of the novel and of the entities to be recognized, for example in the novel *Les Mystères de Paris* written between 1842 and 1843 by Eugène Sue, most characters have names that are similar to noun phrases, such as 'la Goualeuse' (meaning *the Street Singer*) or 'le Chourineur' (meaning *the Stabber*). Also science fiction, which is full of non-classical proper nouns, can be very challenging for the task (Dekker et al. 2019). A module able to predict, or at least, estimate performance from cues gathered in the text would be useful to process large collections of novels.

Coreference (especially linking together all the mentions in the text of a given character, although the task can involve all kinds of names, or even nouns) is challenging in nature. There is a long tradition of research in coreference resolution in NLP, and modules exist for different languages, with various levels of performance (Poesio et al. 2023). The quality of the different systems is still increasing (through end-to-end models (Lee et al. 2017) and then transformer-based language models (Joshi et al. 2019)), and coreference remains a very active field of research in NLP. The task is more challenging for French or Russian than for English, since the "it" pronoun limits ambiguity in English (whereas all nouns are masculine or feminine in French, not only human beings and are referred to with third person pronouns, as for instance in *"Marie veut qu'on lave la voiture, **elle** est sale."* (*"Marie wants that we wash the car, **it** is dirty*."), where *elle* refers to *the car*, but could theoretically also refer to *Marie*; there is no ambiguity from a human point of view in this sentence, but the analysis requires semantic information). When applied in literary studies, automatic coreference systems often break long coreference chains due to the fact that they use a fixed-sized sliding window. If a given character does not appear during a certain period of time (i.e. a certain number of pages), it makes it harder to retrieve its antecedent. Literature provides a good test bed for the coreference task, since novels are long, real, and complex texts on which performance can (and should) still improve a lot.

Event Recognition. Event recognition involves the automated identification and extraction of verbs and, more rarely, nouns referring to events. The task is difficult in that there is no clear definition of what an event is, and other features interact with the definition (among others: negation, adverbials and modals), and not all occurrences of verbs should be annotated (e.g. in *"I like to play tennis"*, play is an infinitive that refers to something I like, but it is generally considered that there is no event *per se* in the sentence). As for literary texts, there have been initiatives to annotate events (Sims et al. 2019), but most verbs and even some nouns can refer to events (Hogenboom et al. 2016; Sprugnoli and Tonelli 2016), which may lead to a too fine-grained annotation. There is thus a need to redefine the task and provide an intermediate level of annotation, between isolated events and the novel as a whole (Lotman 1977; Schmid 2010a,b), but higher level annotation (like the notion of scene) has also proven difficult to formalize, leading to very low accuracy in practical experiments (Zehe et al. 2021).

Quotation Recognition. Quotation recognition plays a crucial role in enhancing the understanding of textual content by identifying and isolating direct speech instances. This feature is instrumental in extracting and preserving the spoken words of characters, enabling a fine-grained analysis of dialogue patterns and character interactions (Durandard et al. 2023; Van Cranenburgh and Van Den Berg 2023). A crucial but complex part of the task consists in establishing what character is at the origin of a given utterance. A recent study has shown that performance on this task are still rather low and would need to improve to be realy usable in operational contexts (Vishnubhotla et al. 2023).

2.2 The BookNLP Project

BookNLP is a set of natural language processing modules designed specifically for the analysis of novels and other literary texts. Developed by D. Bamman (Bamman 2021; Bamman et al. 2014) and colleagues at the University of Berkeley, BookNLP employs a combination of machine learning and linguistic analysis techniques to extract information from text and perform tasks such as character recognition, coreference resolution, event recognition, and quotation extraction. Note that the Berkeley BookNLP suite currently is based upon BERT (Devlin et al. 2019, e.g.), but this could evolve as better language models continue to appear.

The annotated files that are available for training constitute the LitBank corpus (Bamman et al. 2020, 2019). This corpus is publicly available (see `https://paperswithcode.com/dataset/litbank`), which makes it possible to regularly retrain the system, as NLP continues to evolve rapidly (especially large language models)

Entity Recognition: One of the primary tasks of BookNLP is entity recognition, more specifically characters, locations and vehicles, showing the focus on the actions of characters. This information is used to study how mobile protagonist characters are and what kind of space male and female characters occupy (Soni et al. 2023). Character recognition is often coupled with other information (gender, attributes, relations between characters), that can be useful for sub-stream tasks.

Coreference Resolution: In the context of literature, coreference resolution often involves resolving pronouns and other referring expressions to specific characters or entities. BookNLP employs advanced linguistic analysis to identify and link references to the same entity, and the extra knowledge provided by large language models is especially useful for the task.

Event Recognition: Event recognition is another essential task performed by BookNLP. It should be crucial for analyzing the development of the storyline and identifying key plot points, but the huge number of verbs supporting actions make the annotation too prolific and not adapted to specific needs. The proper annotation of negation, adverb and modals is also an open problem. This is why event recognition has not been addressed as a priority in the context of the Multilingual BookNLP Project, that rather focus on entity recognition and coreference resolution.

Quotation Extraction: BookNLP is equipped with the capability to extract quotations from a text. This involves identifying and isolating the direct speech or quoted passages within the literary work. Accurate quotation extraction is vital for understanding character dialogue, the intentions of characters and develop further analyses. However, quotation recognition without speaker attribution is not so useful and, as we have seen before, speaker attribution remains an open question, as accuracy for the task remains low (Vishnubhotla et al. 2023).

The application of BookNLP for the analysis of novels and other literary works aims at providing a deeper understanding of narrative structures, character dynamics, and thematic elements in novels (Piper et al. 2021). The different modules are intended to assist researchers in literary analysis but also in digital humanities and cultural analytics.

3. French BookNLP

The French BookNLP project endeavors to construct a robust Natural Language Processing (NLP) pipeline specifically tailored for the comprehensive analysis of extensive French literary corpora of the 19th and 20th century. The ongoing MultiLingual BookNLP project (Bamman 2020), coordinated by Berkeley, seeks to update the initial pipeline (Bamman et al. 2014) and extend its capabilities to encompass four additional languages (Spanish, German, Russian and Japanese). In alignment with this initiative – even though we are not part of the Multilingual BookNLP project in itself, in the sense that we are independent from the research grant that the Berkeley's team obtained – we are actively engaged in the development of the necessary linguistic resources for the French language. Our collaborative efforts with the Berkeley project ensure a coordinated approach to this expansion, by sharing similar annotations and visualization tools, for example.

In line with the Multilingual BookNLP Project, we will mainly focus on entity recognition and coreference resolution. We have seen in the previous sections that annotating events entail a number of problems and may be too general, thus not be useful if it is not done with a specific goal in mind (which may entail some domain-specific annotations, with adapted categories, for example). We have also seen that quotation recognition with no proper speaker attribution algorithm is, for similar reasons, not really useful, but that speaker attribution remains an open problem (Zehe et al. 2021). In what follows, we will thus not address these two tasks (event and quotation recognition) for further investigation and concentrate on entity recognition and coreference resolution.

3.1 The Training Corpus and The Democrat Project

The "Democrat" project, led by Frédéric Landragin (2016; 2021) and funded by the French National Research Agency (ANR), aimed to develop an annotated corpus at the level of coreference chains in French. Before the Democrat project, no corpus of this kind existed. The project concluded in 2020.

One of the fundamental aspects of Democrat was the annotation of long texts, in contrast to the Ontonotes corpus (Weischedel et al. 2013) for example, which serves as a standard for English but is predominantly composed of short texts. Additionally, the Democrat project aimed to annotate a wide variety of text types, including chapters from novels, short stories, journalistic pieces, legal documents, encyclopedic entries, technical texts, and more. It also had a diachronic dimension, spanning from medieval French to contemporary French.

For the needs of the BookNLP-fr project, we focused on annotations related to novels and selected the texts spanning from the early 19th century to the early 20th century. Before this period, French is more prone to variation, and for the more recent period, texts are not freely shareable due to copyright issues. Lastly, to keep the annotation task manageable, each text in the Democrat corpus is actually composed of a 10,000-word excerpt (leaving us with 184,137 tokens). In addition to this selection from Democrat, we added two short stories from Balzac, good for 45,238 tokens. Information about these texts and those from Democrat can be found in Table 1.

Year	Author	Title	Source
1830	Honoré de Balzac	La maison du chat qui pelote	Full Text
1830	Honoré de Balzac	Sarrasine	Democrat 10 K
1836	Théophile Gautier	La morte amoureuse	Democrat 10 K
1837	Honoré de Balzac	La maison Nucingen	Full Text
1841	George Sand	Pauline	Democrat 10 K
1856	Victor Cousin	Madame de Hautefort	Democrat 10 K
1863	Théophile Gautier	Le capitaine Fracasse	Democrat 10 K
1873	Émile Zola	Le ventre de Paris	Democrat 10 K
1881	Gustave Flaubert	Bouvard et Pécuchet	Democrat 10 K
1882-1883	Guy de Maupassant	Mademoiselle Fifi, nouveaux contes (1)	Democrat 10 K
1882-1883	Guy de Maupassant	Mademoiselle Fifi, nouveaux contes (2)	Democrat 10 K
1882-1883	Guy de Maupassant	Mademoiselle Fifi, nouveaux contes (3)	Democrat 10 K
1901	Lucie Achard	Rosalie de Constant, sa famille et ses amis	Democrat 10 K
1903	Laure Conan	Élisabeth Seton	Democrat 10 K
1904-1912	Romain Rolland	Jean-Christophe (1)	Democrat 10 K
1904-1912	Romain Rolland	Jean-Christophe (2)	Democrat 10 K
1917	Adèle Bourgeois	Némoville	Democrat 10 K
1923	Raymond Radiguet	Le diable au corps	Democrat 10 K
1926	Marguerite Audoux	De la ville au moulin	Democrat 10 K
1937	Marguerite Audoux	Douce Lumière	Democrat 10 K

Table 1: The texts in the BookNLP-fr corpus.

3.2 Data Preparation and Annotation

Entities	#Occurrences
PER - Mentions	32,338
PER - Chain	3,006
FAC	2,325
TIME	1,836
LOC	1,040
GPE	928
VEH	475
ORG	205
TOTAL	**39,147**

Table 2: The number of occurrences per type of entity.

In the scope of the Democrat project, annotations have been applied to all types of coreference. However, for the BookNLP-fr project, our specific focus lies within a subset of these coreferences, corresponding to certain types of entities: persons, facilities, locations, geo-political entities, vehicles, organizations and denotations of time. Definitions from all these categories except for time are adapted from Bamman et al. (2019).

PER: According to Bamman et al. (2019): *"By person we describe a single person indicated by a proper name (**Tom Saywer**) or common entity (**the boy**); or set of people, such as **her daughters** and **the Ashburnhams**."*. Some examples from our corpus in (1), and (2):

(1) a. une de ces gentilhommières si communes en Gascogne, et que **les villageois** décorent du nom de château Le Capitaine Fracasse
b. one of those manors so common in Gascogne, and that **the villagers** decorated by the name of the castle of Captain Fracasse

(2) a. **Madame François**, adossée à une planchette contre **ses** légumes
b. **Madame François**, who leaning on a board next to **her** vegetables

Note that PER mentions are split into three parts to enable more fine-grained analyses, including proper nouns (PROP), common phrases (NOM), and pronouns (PRON). Pronouns account for the majority of mentions, specifically 59%, 32%, and 9%, respectively.

FAC: We follow Bamman's (2019) definition: *"For our purposes, a facility is defined as a "functional, primarily man-made structure" designed for human habitation (buildings, museums), storage (barns, parking garages), transportation infrastructure (streets, highways), and maintained outdoor spaces (gardens). We treat rooms and closets within a house as the smallest possible facility."*, see example (3):

(3) a. **Le chemin** qui menait de **la route** à **l'habitation** s'était réduit, par l'envahissement de la mousse et des végétations parasites
b. **The path** that led to **the road** to **the dwelling** was narrowed by the invasion of moss and parasitic vegetation

GPE: We followed Berkeley's guidelines for this category: *"Geo-political entities are single units that contain a population, government, physical location, and political boundaries."*, see example (4):

(4) a. Échappé de **Cayenne**, où les journées de décembre l'avaient jeté, rôdant depuis deux ans dans **la Guyane hollandaise**, avec l'envie folle du retour et la peur de la police impériale, il avait enfin devant lui **la chère grande ville**, tant regrettée, tant désirée.
b. Escaped from **Cayenne**, where the December days had thrown him, erring since two years in **Dutch Guyane**, with a crazy desire of returning and fear of the imperial police, he finally had before him the **dear big city**, so much regretted and desired.

LOC: As opposed to GPEs, locations are *"entities with physicality but without political organization* [...] *such as **the sea**, **the river**, **the country**, **the valley**, **the woods**, and **the forest**"* (Bamman et al. 2019). Two examples from our corpus:

(5) a. des moellons effrités aux pernicieuses influences de **la lune**
b. crumbling rubble masonry under the pernicious influences of **the moon**

(6) a. Poussez-moi ça dans **le ruisseau** !
b. Push this into **the stream** !

VEH: The definition for a vehicle is a *"physical device primarily designed to move an object from one location to another"* (Bamman et al. 2019). An example from our corpus:

(7) a. anciennement **des voitures** avaient passé par là

b. before, **carriages** had passed there

ORG: *"Organizations are defined by the criterion of formal association"* (Bamman et al. 2019), for example the church and the army. An example from our corpus:

(8) a. et la peur de **la police impériale**
b. and fear of **the imperial police**

TIME: This category is absent in the annotations of Bamman et al. (2019). We designed it to annotate temporal information, duration indications and moments of the day (*day*, *night*, *morning*).

(9) a. sous **le règne de Louis Xiii**,
b. under **the reign of Louis Xiii**,

(10) a. **Le soir**, il avait mangé un lapin.
b. **At night**, he had eaten a rabbit.

As part of the refinement process, the initial annotations required thorough revision and cleaning. We had multiple team discussions about many borderline cases, such as whether Gods and Greek heroes should be annotated as characters, the status of speaking animals and the exact distinction between GPE, FAC and LOC. We meticulously documented every choice made during the annotation process. This documentation is publicly available in an annotation guide[1], providing a valuable resource for understanding our decisions and methodologies in characterizing entities within the context of the BookNLP project, based on the initial ground provided by the Democrat project. Once the annotation guidelines were finished, the entire corpus was annotated by freshly trained annotators. Their first annotations (comprising 315 tags) produced during their training phase, featured an inter-annotator agreement score of Cohen's kappa = .38, meaning fair and almost moderate agreement (Cohen 1960) but showing that this is no trivial task. With better trained annotators, values between .76 and .75 were reached, which constitutes a reasonable basis for further training models. Most errors were due to forgotten mentions, and uncertainties about difficult cases (plurals, fuzzy expressions, non referential entities). Another look at the annotated files by another trained annotators makes a huge difference so as to get a better and more homogeneous coverage (esp. concerning forgotten entities during the initial annotation stage).

After annotation, to facilitate seamless integration with the BookNLP software, the annotations were transformed into a compatible format. We annotated the entity types in TXM (Heiden 2010) because the Democrat corpus is distributed in this format, and later migrated our annotations to brat (Stenetorp et al. 2012), the format used by Berkeley's team. The number of entities in each categorie can be found in Table Table 2.

1. See https://github.com/lattice-8094/fr-litbank.

3.3 Software Development

Large language models play now a prominent role in contemporary natural language processing. Our implementation of BookNLP-fr is built upon the software from the Multi-lingual BookNLP-project. For the two tasks that we perform (entity recognition and coreference resolution), two separated models are developed. Entity recognition is performed before coreference resolution.

Detecting the literary entities, a BiLSTM-CRF model (Bamman et al. 2020; Ju et al. 2018) is fed with contextual embeddings from the CamemBERT model (Martin et al. 2020), which is a BERT(Devlin et al. 2019) based architecture tailored for French.

For the coreference part, a BiLSTM is also fed with the embeddings from CamemBERT. Then, following (Bamman et al. 2020), who in their turn are following Lee (Lee et al. 2017), the BiLSTM architecture is attached to a feedforward network in which the probability of two mentions (detected entities) are coreferent with each other is evaluated. Mentions are linked to their highest scoring antecedent (a null-antecedent is always an option) and coreference chains are defined as the transitive closure of links.

For each model, we split the corpus into training (80%), development (10%) and test (10%) corpus, please see Section section 4 for the results.

While event annotation remains a focal point, challenges persist, primarily due to limitations in performance and the inherently ambiguous nature of defining events. The elusive nature of the concept makes it challenging to generate consistently relevant and usable results. As for quotation identification, we acknowledge the need to integrate speaker recognition for a more comprehensive understanding of textual nuances.

Given these considerations, we have more specifically directed our efforts toward optimizing modules for entity recognition and coreference resolution. This focus allows us to refine and train models that are specifically accurate in identifying and linking entities within a given text, contributing to the effectiveness of BookNLP-fr for downstream tasks (like subgenre classification, see section 5).

4. Results and Evaluation

In this section we give the results of our BookNLP-fr modules for entity recognition and coreference resolution on literary texts.

4.1 Named Entity Recognition Evaluation

Table 3 reports our results for entity recognition, measured traditionally through precision (the percentage of entities correctly recognized among those recognized) and recall (the percentage of entities correctly recognized among those to be recognized). Please note that ORG is absent from this evaluation, because due to an uneven distribution of this tag in different texts, it was only present 7 times in the test corpus, making estimation of precision and recall unreliable.

When assessing the model's performance, a higher precision relative to recall suggests that the model is more likely to make accurate predictions when identifying literary

	precision	recall	F_1
PER	85.0	92.1	88.4
LOC	59.4	54.3	56.8
FAC	73.4	66.0	69.5
TIME	75.3	36.4	49.1
VEH	68.9	63.6	66,1
GPE	68.2	52.9	59,6

Table 3: Entity recognition evaluation of BookNLP-fr on literary texts.

entities. Precision denotes the percentage of correctly predicted literary entities among all entities predicted by the model. High precision is advantageous, ensuring that the identified literary entities are more likely to be accurate, albeit at the potential cost of missing some relevant entities (lower recall). Prioritizing precision in this context aids in minimizing false positives, thereby enhancing the reliability of the identified literary entities. It is important to highlight that literary entities differ from typical Named Entities in Natural Language Processing (NLP), displaying a much larger range of possibilities. Consequently, the obtained results, though seemingly divergent from NLP standards, represent a pioneering achievement in the analysis of French fiction, as this is the first study of its kind.

Some scores may appear modest in comparison to the state-of-the-art, particularly regarding the recall for TIME expressions. This is due to the extensive diversity of time expressions in our corpus, which is far more varied than in the traditional news corpora typically used in NLP, coupled with the limited number of examples in the training corpus (see below, Table 4 for a comparison with a state-of-the art system). Nevertheless, we have opted to report these scores for the sake of comprehensiveness. In the near future, we will strive to expand the coverage of our system, aiming to achieve improved recall across various categories beyond PER.

As a baseline, we ran the CamemBERT-NER model[2], which is a NER model that was fine-tuned from camemBERT on wikiner-fr dataset.Table 4 shows baseline performance in comparison with BookNLP-fr. Results are showing that BookNLP-fr is as good as the fine-tuned model for proper name recognition, but it captures much more by including pronouns and common nouns, which the baseline does not handle at all. The F1 score for the detection of PROP/NOM/PRON mentions reaches 83.13, which is in line with the English BookNLP (88.3).

	BookNLP-fr			Camembert-NER		
pos_tag	precision	recall	F1 Score	precision	recall	F1 Score
PROP	82.5	79.2	**80.8**	91.85	72.05	80.75
NOM	74.9	74.7	**74.8**	96.32	14.17	24.70
PRON	86.3	89.5	**87.9**	100.00	0.10	0.20
ALL	82.39	83.88	**83.13**	92.58	7.92	14.59

Table 4: Comparison on litbank-fr for PER recognition performance between BookNLP-fr and Camembert-NER.

2. See `https://huggingface.co/Jean-Baptiste/camembert-ner`.

BookNLP-fr thus demonstrates its robustness for the classic task of proper name recognition, but the real value of our model lies in its ability to go beyond this to capture the full spectrum of what constitutes a character in novels. This aligns with Woloch (2003) concept of the character space as "the encounter between an individual human personality and a determined space and position within the narrative as a whole," allowing for the automatic detection and analysis of the distribution of character mentions throughout the narrative (Barré et al. 2023).

4.2 Coreference Resolution Evaluation

Table 5 presents the evaluation metrics for coreference resolution using BookNLP-fr on our test corpus. Three key metrics, namely *MUC*, B^3, and $CEAF_e$, are employed to assess its performance. As coreference chains are complex to modelize, different evaluation metrics are necessary to get a global image of systems performance. We refer to Luo and Pradhan (2016) for a comprehensible explanation of these metrics.

Our average F1 score, calculated as the mean of the three metrics, is presented as 76.4. The reported scores suggest a commendable performance, but the practical utility in the context of literary analysis should be further explored based on the specific goals of the research or application. Note that the English BookNLP yields 79.3 in performance for the same task.

Metrics	F_1	
MUC	88,0	
B^3	69,2	*Average* 76.4
$CEAF_e$	71.8	

Table 5: Coreference resolution evaluation of Fr-BookNLP on literary texts

The challenge of duplication arises when the model detects the same character multiple times within the analyzed text. In some instances, among the top five literary entities identified by the model, there may be cases where two or more main characters share the same name or attributes. While this duplication might raise concerns initially, for example, if one aims to study character networks (Perri et al. 2022) or the overall number of characters in novels, it may not pose a significant issue when the focus is on character characterization. For example, in studies about the representation of male and female characters, the output of BookNLP has been shown to be very useful (e.g. Gong et al. 2022; Hudspeth et al. n.d.; Naguib et al. 2022; Toro Isaza et al. 2023; Underwood et al. 2018; Vianne et al. 2023; Zundert et al. 2023).

Also in the following case study, the primary objective is not to pinpoint unique and distinct characters but rather to establish a proxy for characterization as a whole. Our goal is to capture the prevalence and significance of certain characters across various texts and literary works. Hence, the emphasis lies more on character representation and the overall impact of these characters on the literary landscape, rather than on identifying entirely separate and non-repeating characters.

5. Case Study: Genre Classification Using Booknlp-fr Features

5.1 Introduction

This case study aims to demonstrate that BookNLP-fr can be of significant assistance in the realm of computational literary studies (CLS). We illustrate this assertion through a canonical issue in CLS: the automatic detection of literary genres. Historically, the division of novels into specific sub-genres has been a classification practice employed by literary stakeholders such as librarians, editors, and critics. This practice is partly justified by a specific textual component that relates to the spatiotemporal framework, characters, themes, or narrative progression.

Genre is a central concept in poetics, defined successively from Aristotle to structuralists, through romantics and Russian formalists (Aristote 1990; Bachtin 2006; Genette 1986; Schlegel et al. 1996). From our computational standpoint, structuralists have offered intriguing definitions. For example, Schaeffer (1989) defines genericity as an "internalized norm that motivates the transition from a class of texts to an individual text conforming to certain traits of that class". There could be a set of textual procedures internal to works, and the mission of CLS would be to find the best ways to account for this fact. However, the norms or formal rules of sub-genres cannot be solely boiled down to formal or thematic rules. For instance, the sociological approach, as exemplified by Bourdieu (1979), tends to focus more on the "community of readers" with the study of power dynamics and accompanying aesthetic hierarchies. However, these norms do indeed exist, as they enable a work to align itself with the established and shared usage of a "horizon of expectations" (Jauß 1982) of the audience which might induce the authors to adhere to certain expected norms and styles.

Various studies have devised strategies to automatically identify subgenres. Selected studies have employed methods such as the bag of words (BoW) (Hettinger et al. 2016; Underwood 2019) or topic modeling (Schöch 2017; Zundert et al. 2022) to find subgenre similarities between texts. In addition to these basic features, researchers utilize machine learning techniques in a supervised setting, employing methods such as logistic regression or support vector machines when ground truth is available. However, the challenge often arises from the potential incompleteness or temporal bias of these ground truths. Unsupervised learning approaches and clustering methods have also enabled the exploration of hybrid texts that belong to multiple subgenres, as demonstrated by studies like (Calvo Tello 2021; Sobchuk and Šeļa 2023). In our case-study, we will rely on a corpus with predefined labels, while acknowledging the idea that sub-genres are not monolithic categories. Thus, the objective is not so much to demonstrate the validity of sub-genre labels, which are often incomplete or limiting in reality, but rather to show that the interpretability of errors in automatic classification can lead us to a more nuanced and comprehensive understanding of the subgenre phenomenon.

Despite recent advancements in NLP, the bag-of-words approach remains largely unchanged. This is because many tools, including document embeddings, are not easily interpretable and are optimized for short texts. In this context, we present in the next section a method that aims to find a balance between the use of state-of-the-art methods for literary text processing and their interpretability.

5.2 Method

5.2.1 Corpus and Subgenre Labels

Our case study is built upon one of the largest corpora for fiction in French: the "corpus Chapitres", a corpus of nearly 3000 French novels (Leblond 2022). The period concerned extends over two centuries of novel production, from the 19th to the 20th century, as can be seen in Figure 1.

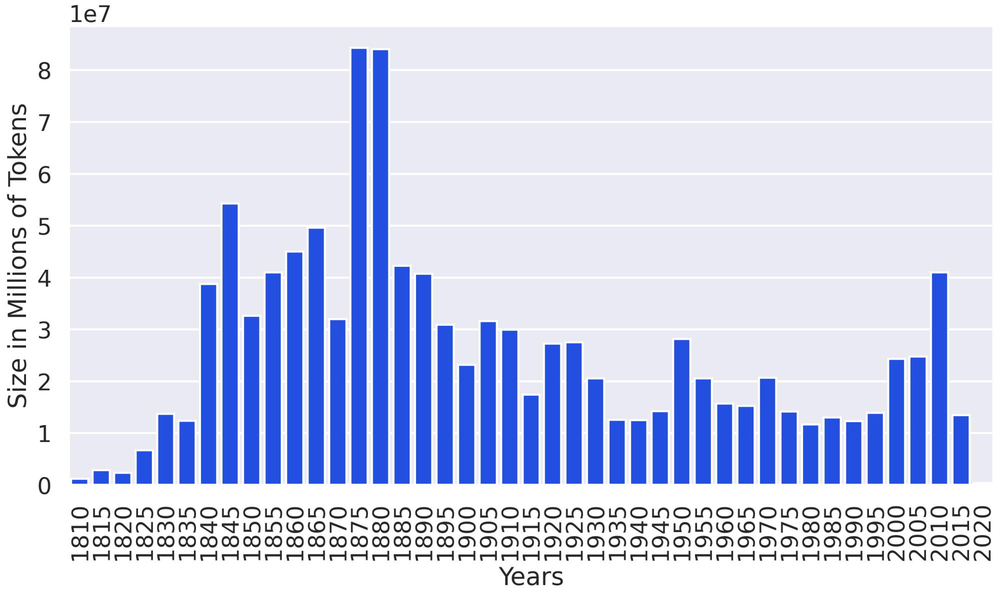

Figure 1: Distribution of the number of tokens over time.

Approximately two-thirds of Chapitres is annotated with sub-genre labels. This annotation is based on the classification of the French National Library. We choose to concentrate our analysis on the five most prevalent sub-genres within the corpus: adventure novels, romance, detective fiction, youth literature, and memoirs. The validity of these labels is not clearly established, as the practices of the BNF for assigning these labels have not been systematized nor standardized. Therefore, there is no "Ground Truth" per se, but our supervised approach described in subsubsection 5.2.3 aims precisely to understand the boundaries of subgenres.

5.2.2 Textual Features

The BoW method stands out as the default feature extraction technique, as it allows scholars to have an easy task to implement without requiring intensive computational resources (GPU, RAM). Underwood (2019) demonstrated that the BoW approach was highly effective in classifying subgenres such as Gothic, detective stories, and even science fiction.

Nevertheless, although this method proves valuable in specific contexts, it is not without two limitations. First, it does not consider the word order within the text. This limitation means that the sequential arrangement of words, which is crucial for capturing the nuances of literary elements like plot and narrative structure, is ignored. Second, there is a risk of overfitting to the idiolects of writers, particularly when emphasizing the most frequent words. Additionally, these tools may inadvertently capture chronolectal aspects, as it is established that the approximate writing date of a book can be predicted based on the prevalence of certain most frequent words (Seminck et al. 2022).

In this paper we rely on two distinct feature extraction approaches: the classic BoW as a control experiment, and the BookNLP-fr one, which we will implement as follows. The idea is based on a previous study (Kohlmeyer et al. 2021) where researchers demonstrated the limitations of traditional document embeddings (optimized for shorter texts) in capturing complex facets in novels (such as time, place, atmosphere, style, and plot). To address this problem, they propose to use multiple embeddings reflecting different facets, splitting the text semantically rather than sequentially. Inspired by these findings, we adapted their methodology to evaluate the impact of these features on subgenre classification when contrasted with the traditional BoW approach.

The method runs our BookNLP pipeline on our texts, allowing us to automatically retrieve, on the one hand, information related to space-time, notably with the set of LOC, FAC, GPE, TIME, and VEH. On the other hand, it provides information related to characterization, including all verbs for which characters are patients (PATIENT) or agents (AGENT), as well as the set of adjectives that will characterize them (ADJ).

Thus, two types of features are under consideration:

- For the BoW, we relied on the 600 most frequent lemmas, excluding the first 200, which comprise non-informative stop words not relevant to our subgenre case study. They could have been relevant if we wanted to acknowledge the authors who wrote in a specific subgenre, but it is not our goal here, and we will discuss how we handled this bias in Section 5.2.3.
- For the BookNLP-fr features, we compiled for each novel, lists of words extracted by BookNLP-fr. We then obtained vector representations using a Paragraph Vectors model (Le and Mikolov 2014) (Doc2Vec) trained on a subset of our novel dataset. Two vector embeddings of 300 dimensions were generated: one for characterization (AGENT, PATIENT, ADJ) and one for space-time (LOC, FAC, GPE, TIME, VEH).

Therefore we obtained two datasets for training, one with 600 dimensions representing the 600 most frequent lemmas, and the other with also 600 dimensions representing the two concatenated Doc2Vec vectors, one for the characterization and one for the space and time.

5.2.3 Modeling

We opted for an SVM as it has been demonstrated that these models obtain the best performance in classifying literary texts (Yu 2008), and more specifically literary subgenres (Hettinger et al. 2016). In this paper, we used the implementation of Pedregosa et al. (2011). The SVM doesn't perform multiclassification per se, but it classifies each subgenre against the others in binary classification and then aggregates the results.

Therefore, we don't have a single classification, but rather

$$\frac{n_{\text{classes}} \cdot (n_{\text{classes}} - 1)}{2}$$

With our 5 subgenres, this implementation results in 10 different classifications.

Considering our task of subgenre classification, we wanted to limit idiolectal bias, especially for the model trained on the BoW. To do so, we implemented Scikit-learn's Group strategy. All works by the same author (group) were placed in the same fold. Thus, each group will appear exactly once in the test set across all folds. Since SVM models are quite sensitive working with imbalanced classes, we re-balanced the classes before implementing the classification by randomly taking 130 novels for each subgenre. We implemented this selection a hundred times and for each resulting sample the model was run in a 5-fold cross-validation setting. The following results are aggregated from this process.

5.3 Results

5.3.1 BoW vs BookNLP-fr features

	Precision	Recall	F1-score	Support	Accuracy
Children	0.75	0.75	0.75	130	
Memoirs	0.79	0.82	0.80	130	
Detective	0.67	0.68	0.67	130	
Adventure	0.60	0.65	0.62	130	
Romance	0.84	0.72	0.80	130	
Full Dataset				650	**0.72**

Table 6: Classification Report for BoW

	Precision	Recall	F1-score	Support	Accuracy
Children	0.65	0.79	0.71	130	
Memoirs	0.78	0.89	0.84	130	
Detective	0.68	0.70	0.70	130	
Adventure	0.73	0.73	0.73	130	
Romance	0.90	0.65	0.75	130	
Full Dataset				650	**0.75**

Table 7: Classification Report for BookNLP-fr features.

Tables 6 and 7 display the classification report of the models' evaluation on the test set. Both models achieve good results: 72% for the BoW-based model and the BookNLP-based model achieves 75% accuracy. This means that our models are capable of correctly identifying the subgenre three out of four times, whereas a random baseline yields an accuracy score of 0.2. The main result here is that differences exist among our subgenres, whether from the perspective of text structure with MFW or from a semantic standpoint with BookNLP. The fact that the BookNLP-based model obtains an additional 3 points of accuracy might not be revolutionary, but the primary argument for this type of feature extraction lies more in the interpretation of features, as discussed in subsection 5.4.

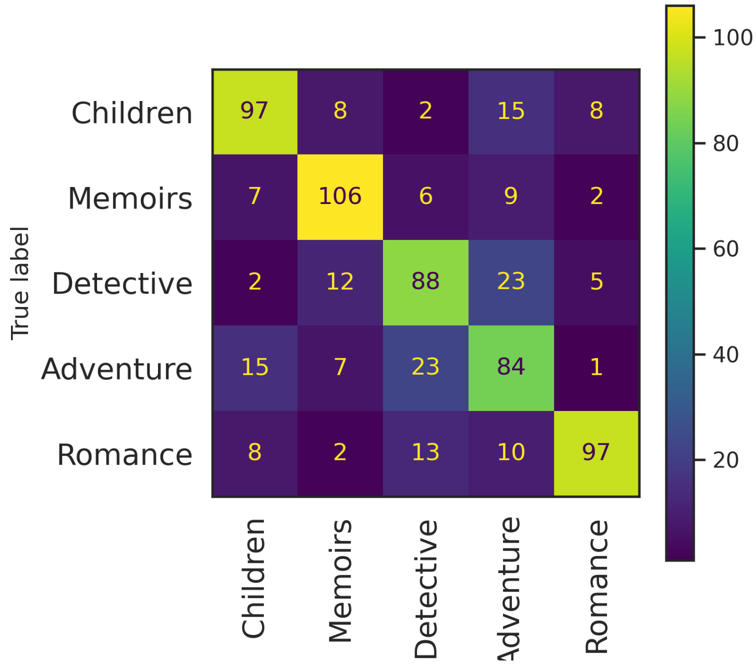

Figure 2: Confusion Matrix for BoW.

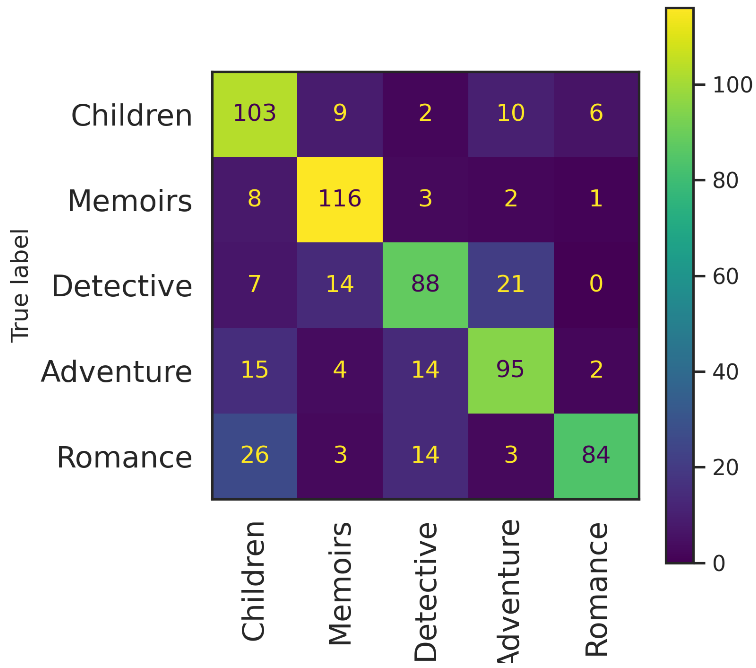

Figure 3: Confusion Matrix for BookNLP-fr features.

To enhance our comprehension of how the models behave and the nature of their errors, we visualize their confusion matrices in Figure 2 and Figure 3. The x-axis represents the predicted subgenre, while the y-axis represents the expected subgenre. A perfect classification would display a diagonal filled with 130 correct predictions for each subgenre.

We observe that both models have quite similar error patterns, and one distinct scenario stands out: Both models predict 'Adventure' instead of 'Detective' (23 errors for BoW, 21 for BookNLP). These common errors are quite understandable since these two subgenres share many similarities, including a penchant for suspense and violent action, which could confuse the models.

Another scenario seemed highly instructive for analysis: The errors made by the models when predicting the label 'Children', but the expected subgenre is 'Romance'. The BoW model performs quite well with 8 errors, but the BookNLP-based model makes 26 errors. The semantic model thus faces more challenges in distinguishing between these two subgenres, which makes sense, as both subgenres are characterized by themes centered around emotions and relationships between characters, common features to both subgenres.

5.3.2 BookNLP-fr Features Accuracy for Subgenre Classification

In this section, the objective is to evaluate, on the one hand, whether specific individual features from BookNLP can classify our subgenres, and on the other hand, we will attempt to interpret the differences in performance for each. Here, each pipeline is trained with a Doc2Vec vector of 300 dimensions for each type of feature.

BookNLP-fr features	Accuracy
LOC	0.45
FAC	0.59
VEH	0.42
GPE	0.47
TIME	0.50
PATIENT	0.52
AGENT	0.62
ADJ	0.50
Baseline	0.2

Table 8: BookNLP-fr features accuracy.

A first obvious observation is that all our models achieve results at least twice as good as the baseline. The information contained in each of these features is therefore highly relevant from the subgenre perspective. The 'VEH' class lags a bit behind (42% accuracy), which may suggest that vehicles are not decisively discriminating among our subgenres, but it is our least represented class in our texts, and therefore, there may not be enough data. Very good results are obtained for the 'FAC' (0.59) and 'AGENT' (0.62). This indicates that subgenres distinguish well in terms of mentioned buildings or verbs where the character is agentive, meaning that the type of action a character takes is specific to each subgenre.

Interestingly, the misclassifications (see the confusion matrices in the Appendix A for each individual feature), the same pattern emerges (misclassification of 'Adventure' instead of 'Detective' and 'Children' instead of 'Romance'), but the error rates vary depending on the features used. This can provide a lot of information about the differences and similarities between certain subgenres. The next section 5.4 offers an interpretation closely examining these anomalies.

5.4 Interpretability

This section explores the interpretation of the two SVM models, BoW-based and BookNLP-based. It focuses on the misclassifications of 'Adventure' instead of 'Detective'.

One of the advantages of the SVM pipeline is the ability to investigate the statistical inferences of the models when the kernel is in linear mode. The SVM searches for the plane in the latent space of words that best separates our two categories. Each dimension receives a coefficient, with a negative sign if the coefficient is used to predict a specific class and a positive sign for the other. For the BoW-based model, it's quite straightforward as a coefficient is assigned to each word, as can be seen in Figure 4.

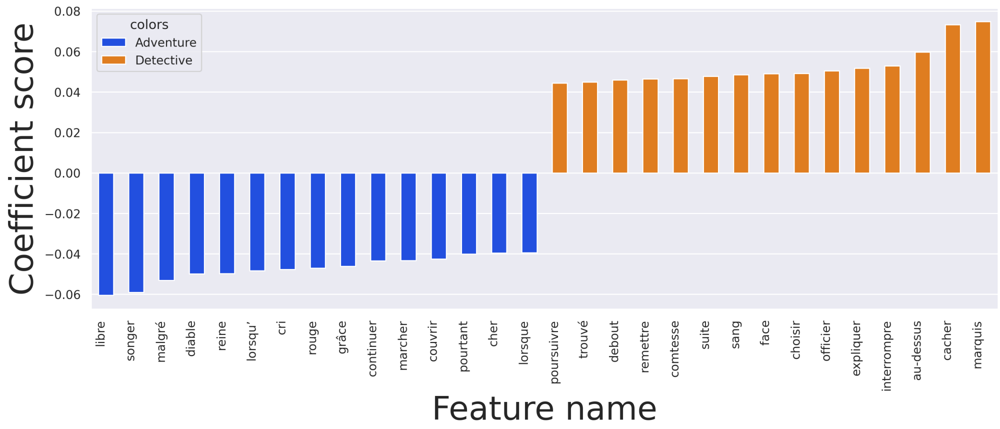

Figure 4: BoW discriminant features for Adventure vs Detective classification.

Looking at the coefficients assigned for the Adventure vs. Detective classification, we find some relevant elements, such as the presence of the word 'free' ('libre') as the most discriminant word for assigning the Adventure label. Apart from that, with perhaps 'cry' ('cri'), which could signify adventure, few clues remain. Verbs such as 'dream,' 'walk', 'continue', or conjunctions like 'when' ('lorsque'), 'despite' ('malgré'), and 'yet' ('pourtant') are not really characteristic of adventure novels. It is difficult to conclude, except that these less significant coefficients seem to indicate the model's difficulty in distinguishing between the two sub-genres.

For the BookNLP-based model, it's a bit more complex since the coefficients are assigned to each dimension of the Doc2Vec vectors. Therefore, we aggregated the coefficients by feature type to gain a more concrete overview of the results. Figure 5 illustrates the sum of all coefficients for each feature extracted by BookNLP-fr. We conducted a t-test to confirm that the difference between the means of the populations is statistically significant. Taking adjectives as an example (T-statistic: 28.7; P-value: 2.25×10^{-180}), we observe that the model relies more on these dimensions to assign the label 'detective' compared to 'adventure'.

This could be explained by the strong emphasis placed on character psychology in detective novels, especially those involving criminals and detectives. For instance, in *Maigret et le tueur* (1969), George Simenon's beloved detective (*Maigret*) is frequently characterized as 'wise', 'whimsical,' or even 'happy', while criminals are 'suspicious' or 'villainous'. This doesn't imply a lack of characterization in adventure novels but rather suggests that it is not a distinctive feature of the subgenre compared to detective novels.

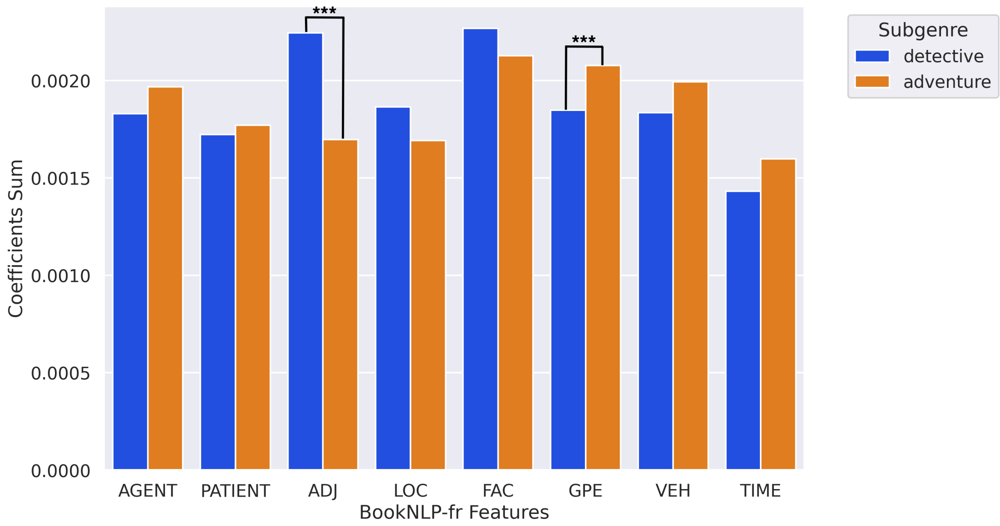

Figure 5: BookNLP-fr discriminant features for Adventures vs Detective classification. '***' meaning p<0.001.

Considering Geo-Political Entities (T-statistic: -21.0; P-value: 8.49×10^{-98}), the reasoning is inverse: the model relies slightly more on the dimensions of the GPE vector to assign the adventure label than the detective label. This makes sense when examining GPEs for example in *Les trappeurs de l'Arkansas* by Gustave Aimard (1857): 'Hermosillo', 'America', 'the New World', 'Guadalajara', 'Mexico', etc. The novel heavily emphasizes exotic locations and mentions places in the American or Mexican West for this purpose. GPEs in detective novels are more commonplace, as these novels often take place in France, typically in an urban setting.

Thus the model has learned that certain dimensions of characterization are more strongly associated with a particular subgenre (such as adjectives for detective novels), and that certain dimensions of the GPE or TIME vector are important for assigning the adventure label. Let's now generalize our approach to the entire classification process.

Examining the behavior of the coefficients when aggregated for the 10 classifications, we can observe the graph shown in Figure 6. This graph depicts the model coefficients after training based on the vectors of each facet, using a dataset of 2400 dimensions. We consider this graph as a dive into the model's inferences, where it will assign more weight to certain categories to assign a specific subgenre.

For example, it is observed that the value of 'FAC' is very high for the detective genre, indicating a particular specificity for this sub-genre. Details of locations, crime scenes, investigations in specific places, detective offices, interrogation rooms, etc., are distinguishing elements for this sub-genre. The same applies to 'GPE' for the adventure label, as seen previously, with an emphasis on exoticism that may play a role here, even though 'LOC' and 'FAC' do not show significant differentiation from this perspective. Conversely, for romance and the 'TIME' vector, where the coefficients for these vectors lag behind other sub-genres. Examples of time in romance novels may be used more to describe emotional moments or stages in relationships rather than to highlight complex temporal events. Consequently, the model might perceive that the 'TIME' vector is not as discriminative for this category.

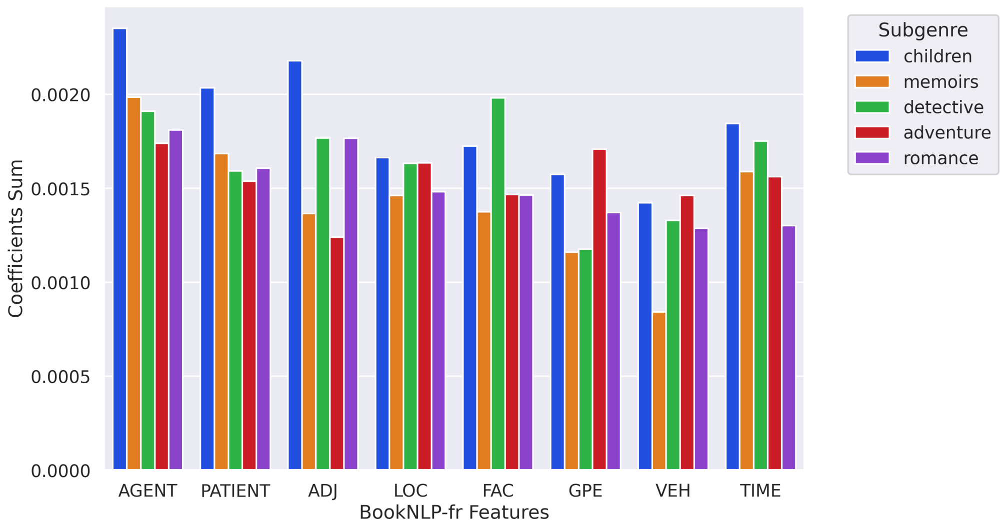

Figure 6: BookNLP-fr discriminant features for the classification.

We have thus demonstrated that the BoW-based classification approach is challenging to interpret, as certain highly discriminating words do not appear to bring about key distinctions between the subgenres. The BookNLP-fr-based method may offer an insightful understanding of the specificities that differentiate one subgenre from another. Both approaches do not completely substitute for each other since we are examining features of different nature (vocabulary vs semantic), but they can complement each other to enhance interpretability.

Diving into the model's indications, several types of features were observed to interpret the model's inferences. Many differences among the features were noticed, although we did not have the space to interpret all of them in this article. Much work remains to be done, and new experiments should be considered, for instance going beyond the SVM, including the use of deep neural networks and textual deconvolution saliency Vanni et al. (2018), which could facilitate the return to close reading based on the embeddings derived from BookNLP-fr data.

6. Discussion

6.1 Working with Imperfect Annotations

The utilization of computers for annotating literary texts has profoundly changed the landscape of literary studies, enabling the annotation of vast amounts of texts with unprecedented efficiency. This enables the community to address research questions that were out of reach before, such as a study at scale of characters with disabilities (Dubnicek et al. 2018) or the quantitative analysis of characters in fanfiction (Milli and Bamman 2016) and a quantitative, diachronic study of things appearing in fiction (Piper and Bagga 2022). However, this advancement is not without its challenges, particularly in the context of the inherent errors that may accompany automated annotation processes. This poses a twofold challenge for researchers engaged in the field of CLS.

Firstly, ensuring the reliability of studies based on imperfect annotations is a critical concern. Scholars must grapple with the task of guaranteeing that errors, though present, remain at a marginal level and do not compromise the validity of their research findings. This necessitates a careful balance between the benefits of computational efficiency and the maintenance of accuracy in annotations. Researchers are challenged to develop methodologies and quality control measures that safeguard against the potential pitfalls introduced by errors in the annotation process.

Secondly, the acceptance of computational approaches by literary scholars is not guaranteed, as the traditional paradigm within literary studies often revolves around meticulous, supposedly perfect annotations. The shift to working with non-perfect annotations, even if the errors are marginal, represents a departure from the established norm. This cultural shift within the academic community poses a psychological barrier, as literary scholars may be hesitant to fully embrace computational methods if they perceive a compromise in the level of precision to which they are accustomed.

Addressing these challenges requires not only the refinement of computational tools for annotation but also a broader cultural shift within the academic community. There is a need for transparent communication about the limitations of automated annotation processes, the establishment of best practices for mitigating errors, and the development of strategies to ensure that computational approaches align with the standards expected both in literary studies and in computer science.

6.2 Maintaining Annotations Tools in the Era of Large Language Models

The field of computational literary studies is currently grappling with a significant challenge due to the rapid evolution of natural language processing, particularly with the proliferation of large language models (LLMs). The continuous emergence of new LLMs has led to an accelerated pace of research in the domain. While this dynamism brings about positive outcomes, such as increased research activity, the introduction of novel tasks, and the generation of new results, it also presents several inherent dangers.

One primary challenge lies in the technical aspect of keeping annotation tools up to date amidst the constant production of new LLMs by the research community and the industry. There is a delicate balance to strike, ensuring that annotation systems remain up-to-date, without expending an excessive amount of resources on incessantly adapting to the latest trends in LLM development. The challenge here is not just about technological compatibility but also about efficiently managing the resources required for frequent updates and integrations, and to produce software that is usable by a large community (i.e. software should not be dependent on a unreasonably heavy computer infrastructure).

A more critical concern revolves around the need to guarantee the reproducibility of research outcomes. The rapid evolution of LLMs implies that a specific version in use today may become obsolete or unavailable tomorrow. This raises the risk that crucial details, such as the corpus utilized, configuration parameters, and hyperparameters of the model, may not be adequately documented in research reports. Ensuring reproducibility becomes a substantial challenge as the landscape of LLMs continues to evolve, necessitating a concerted effort to establish standardized practices for reporting model specifications and associated details.

In addressing these challenges, we believe it is crucial to focus not only on technical aspects but also on developing robust frameworks for documentation and reproducibility. Establishing clear guidelines for reporting model specifications, documenting corpus details, and archiving relevant information becomes paramount for the field.

7. Conclusion

In this paper, we introduced the BookNLP-fr pipeline, with a particular emphasis on entity recognition and coreference resolution. Demonstrating its practical utility, we illustrated how this software facilitates the analysis of extensive French literary corpora, relying on semantic features unique to the texts under examination. Through this study, we hope to show the potential of natural language processing in analyzing large literary corpora, to go beyond purely statistical approaches and to overcome bias by taking into account an unprecedented number of texts and not only the reduced set of texts of the literary canon. In concrete terms, we distinguish three research directions, all of which carry the above-described desire for large-scale generalization:

1. Studies on the characteristics of literary genre : BookNLP-en can be used to retrieve textual features of a semantic nature, in particular entities that provide information on the spatio-temporal setting of the story. The latter are very important for determining literary genres. For example, adventure novels have a very specific spatio-temporal setting (the emphasis is on the importance of geographical disorientation), while romance novels take place in a more urban, modern setting. The BookNLP-fr tools could thus be crucial for automatic classification.

2. Characterization: co-reference chains with mentions of a character allow us to recover how each character is portrayed. In this way, we can study the differences between certain types of characters on a large scale. For example, it's possible to report on how men and women have been characterized in literature over time (e.g. Naguib et al. 2022; Vianne et al. 2023) or what role secondary characters actually play in the narrative (Barré et al. 2023). To cite other examples: a tool like BookNLP makes it possible to study how characters with disabilities are presented (Dubnicek et al. 2018) or to carry out a quantitative analysis of characters in fan fiction (Milli and Bamman 2016).

3. Detection of specific scenes: BookNLP could be capable of detecting specific scenes in novels; these could be defined by one or more characters gravitating around a precise location and carrying out particular actions. This scene detection, understood as a minimal narrative unit, could enable us to better understand the workings of the plot by breaking down its layout over the course of the story.

Future work on the BookNLP-fr pipeline will include a renewed exploration of the concepts of events and scenes, aiming to establish an annotation framework that aligns with literary perspectives. Additionally, we plan to address the question of quotation analysis and attribution. Finally, a key focus will be on ensuring that results undergo scientific evaluation and that recent advancements in natural language processing can be continuously integrated, all while preserving the distinctive nature of literary works and literary studies. In that way, BookNLP-fr can play an significant role in the domains of automatic literary analysis and cultural analysis. Literary questions, one even more exciting and ambitious than the other, can finally be addressed automatically on a large scale.

8. Data Availability

Data can be found here: https://github.com/lattice-8094/fr-litbank.

9. Software Availability

Software can be found here: https://seafile.rlp.net/f/6c9d680114fe4583a89c/?dl=1.

10. Acknowledgements

This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA0001 (PRAIRIE 3IA Institute).

11. Author Contributions

Frédérique Mélanie-Becquet: Conceptualization, Data Curation, Supervision

Jean Barré: Formal Analysis, Writing – review & editing

Olga Seminck: Formal Analysis, Writing – review & editing

Clément Plancq: Conceptualization, Software

Marco Naguib: Conceptualization, Software

Martial Pastor: Conceptualization, Software

Thierry Poibeau: Conceptualization, Writing – original draft, review & editing, Supervision

References

Aristote (1990). *Poétique*. Le Livre de poche. Librairie générale française.

Bachtin, Michail Michajlovič (2006). *Esthétique et théorie du roman*. Collection Tel 120. Gallimard.

Bamman, David (2020). *Multilingual BookNLP: Building a Literary NLP Pipeline Across Languages*. https://apps.neh.gov/publicquery/main.aspx?f=1&gn=HAA-271654-20. Accessed: January 17, 2024.

— (2021). *BookNLP*. https://github.com/booknlp/booknlp.

Bamman, David, Olivia Lewke, and Anya Mansoor (2020). "An Annotated Dataset of Coreference in English Literature". In: *Proceedings of the Twelfth Language Resources and Evaluation Conference*. Ed. by Nicoletta Calzolari, Frédéric Béchet, Philippe Blache, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, Hélène Mazo, Asuncion Moreno, Jan Odijk, and Stelios Piperidis. Marseille, France: European Language Resources Association, 44–54. https://aclanthology.org/2020.lrec-1.6.

Bamman, David, Sejal Popat, and Sheng Shen (2019). "An annotated dataset of literary entities". In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota: Association for Computational Linguistics, 2138–2144. 10.18653/v1/N19-1220.

Bamman, David, Ted Underwood, and Noah A. Smith (2014). "A Bayesian Mixed Effects Model of Literary Character". In: *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*. ACL 2014. Association for Computational Linguistics, 370–379. 10.3115/v1/P14-1035.

Barré, Jean, Pedro Cabrera Ramírez, Frédérique Mélanie, and Ioanna Galleron (2023). "Pour une détection automatique de l'espace textuel des personnages romanesques". In: *Humanistica 2023*. Corpus. Association francophone des humanités numériques. Genève, Switzerland, 56–61. https://hal.science/hal-04105537.

Benoit, Kenneth, Kohei Watanabe, Haiyan Wang, Paul Nulty, Adam Obeng, Stefan Müller, and Akitaka Matsuo (2018). "quanteda: An R package for the quantitative analysis of textual data". In: *Journal of Open Source Software* 3.30, 774. 10.21105/joss.00774.

Bird, S., E. Klein, and E. Loper (2019). *Natural Language Processing with Python – Analyzing Text with the Natural Language Toolkit*. https://www.nltk.org/book/choo.html.

Bourdieu, Pierre (1979). *La distinction: critique sociale du jugement*. Le Sens commun 58. Éditions de Minuit.

Calvo Tello, José (Dec. 31, 2021). *The Novel in the Spanish Silver Age: A Digital Analysis of Genre Using Machine Learning*. Bielefeld University Press. 10.1515/9783839459256.

Cohen, Jacob (1960). "A coefficient of agreement for nominal scales". In: *Educational and psychological measurement* 20.1, 37–46.

Dekker, Niels, Tobias Kuhn, and Marieke van Erp (2019). "Evaluating named entity recognition tools for extracting social networks from novels". In: *PeerJ Computer Science* 5, e189. https://doi.org/10.7717/peerj-cs.189.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova (2019). "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding". In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Ed. by Jill Burstein, Christy Doran, and Thamar Solorio. Minneapolis, Minnesota: Association for Computational Linguistics, 4171–4186. 10.18653/v1/N19-1423.

Dubnicek, Ryan, Ted Underwood, and J Stephen Downie (2018). "Creating A Disability Corpus for Literary Analysis: Pilot Classification Experiments". In: *iConference 2018 Proceedings*.

Durandard, Noé, Viet Anh Tran, Gaspard Michel, and Elena Epure (2023). "Automatic Annotation of Direct Speech in Written French Narratives". In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long

Papers). Ed. by Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki. Association for Computational Linguistics, 7129–7147. 10.18653/v1/2023.acl-long.393.

Eder, Maciej, Jan Rybicki, and Mike Kestemont (2016). "Stylometry with R: a package for computational text analysis". In: *R Journal* 8.1, 107–121. https://journal.r-project.org/archive/2016/RJ-2016-007/index.html.

Emelyanov, A. and E. Artemova (2019). "Multilingual Named Entity Recognition Using Pretrained Embeddings, Attention Mechanism and NCRF". In: *Proc. of the 7th Workshop on Balto-Slavic Natural Language Processing*. Florence, Italy, 94–99.

Genette, Gérard (1986). "Introduction à l'architexte". In: *Théorie des genres*. Ed. by Gérard Genette and Tzvetan Todorov. Points 181. Éd. du Seuil.

Gong, Xiaoyun, Yuxi Lin, Ye Ding, and Lauren Klein (2022). "Gender and power in japanese light novels". In: *Proceedings http://ceur-ws. org ISSN* 1613, 0073.

Heiden, Serge (2010). "The TXM platform: Building open-source textual analysis software compatible with the TEI encoding scheme". In: *24th Pacific Asia conference on language, information and computation*. Vol. 2. 3. Institute for Digital Enhancement of Cognitive Development, Waseda University, 389–398.

Hettinger, Lena, Fotis Jannidis, Isabella Reger, and Andreas Hotho (2016). "Significance Testing for the Classification of Literary Subgenres". In: ADHO 2016 - Kraków. https://dh-abstracts.library.cmu.edu/works/2630 (visited on 01/18/2024).

Hogenboom, F., F. Frasincar, U. Kaymak, F. de Jong, and E. Caron (2016). "A survey of event extraction methods from text for decision support systems". In: *Decision Support Systems* 85.c, 12–22. Doi:%2010.1016/j.dss.2016.02.006.

Hudspeth, Marisa, Sam Kovaly, Minhwa Lee, Chau Pham, and Przemyslaw Grabowicz (n.d.). "Gender and Power in Latin Narratives". In: ().

Jauß, Hans Robert (1982). *Toward an aesthetic of reception*. Trans. by Timothy Bahti. Univ. of Minnesota Press.

Joshi, Mandar, Omer Levy, Luke Zettlemoyer, and Daniel Weld (2019). "BERT for Coreference Resolution: Baselines and Analysis". In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Ed. by Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan. Hong Kong, China: Association for Computational Linguistics, 5803–5808. 10.18653/v1/D19-1588.

Ju, Meizhi, Makoto Miwa, and Sophia Ananiadou (June 2018). "A Neural Layered Model for Nested Named Entity Recognition". In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Ed. by Marilyn Walker, Heng Ji, and Amanda Stent. New Orleans, Louisiana: Association for Computational Linguistics, 1446–1459. 10.18653/v1/N18-1131.

Kohlmeyer, Lasse, Tim Repke, and Ralf Krestel (2021). "Novel Views on Novels: Embedding Multiple Facets of Long Texts". In: *2021 Association for Computing Machinery.*

Landragin, Frédéric (2016). "Description, modélisation et détection automatique des chaînes de référence (DEMOCRAT)". In: *Bulletin de l'Association Française pour l'Intelligence Artificielle* 92, 11–15.

— (2021). "Le corpus Democrat et son exploitation. Présentation". In: *Langages* 4, 11–24.

Langlais, Pierre-Carl (May 2021). *Fictions littéraires de Gallica / Literary fictions of Gallica*. Version 1. Zenodo. 10.5281/zenodo.4751204.

Le, Quoc V. and Tomás Mikolov (2014). *Distributed Representations of Sentences and Documents*. arXiv: 1405.4053.

Leblond, Aude (2022). *Corpus Chapitres*. Version v1.0.0. 10.5281/zenodo.7446728.

Lee, Kenton, Luheng He, Mike Lewis, and Luke Zettlemoyer (2017). "End-to-end Neural Coreference Resolution". In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Ed. by Martha Palmer, Rebecca Hwa, and Sebastian Riedel. Copenhagen, Denmark: Association for Computational Linguistics, 188–197. 10.18653/v1/D17-1018. https://aclanthology.org/D17-1018.

Lotman, Yuri (1977). *The Structure of the Artistic Text*. Michigan Univ. Press (Michigan Slavic Contributions No. 7).

Luo, Xiaoqiang and Sameer Pradhan (2016). "Evaluation metrics". In: *Anaphora Resolution: Algorithms, Resources, and Applications*. Springer, 141–163.

Manning, C., M. Surdeanu, J. Bauer, J. Finkel, S. Bethard, and D. McClosky (2014). "The Stanford CoreNLP Natural Language Processing Toolkit". In: *Proc. of the 52nd Annual Meeting of the Association for Computational Linguistics (ACL): System Demonstrations*. Baltimore: ACL.

Manovich, Lev (2018). "The science of culture? Social computing, digital humanities and cultural analytics". In.

Martin, Louis, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric Villemonte de la Clergerie, Djamé Seddah, and Benoît Sagot (2020). "CamemBERT: a Tasty French Language Model". In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.

Maynard, D., K. Bontcheva, and I. Augenstein (2017). "Named Entity Recognition and Classification". In: *Natural Language Processing for the Semantic Web*. Synthesis Lectures on Data, Semantics, and Knowledge. Cham: Springer.

Milli, Smitha and David Bamman (2016). "Beyond Canonical Texts: A Computational Analysis of Fanfiction". In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Ed. by Jian Su, Kevin Duh, and Xavier Carreras. Austin, Texas: Association for Computational Linguistics, 2048–2053. 10.18653/v1/D16-1218. https://aclanthology.org/D16-1218.

Moretti, Franco (2000). "Conjectures on world literature". In: *New Left Review*.

Naguib, Marco, Marine Delaborde, Blandine Andrault, Anaïs Bekolo, and Olga Seminck (2022). "Romanciers et romancières du XIXème siècle : une étude automatique du genre sur le corpus GIRLS (Male and female novelists : an automatic study of gender of authors and their characters)". In: *Actes de la 29e Conférence sur le Traitement Automatique des Langues Naturelles. Atelier TAL et Humanités Numériques (TAL-HN)*. Ed. by Ludovic Moncla and Carmen Brando. Avignon, France: ATALA, 66–77. https://aclanthology.org/2022.jeptalnrecital-humanum.8.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay (2011). "Scikit-learn: Machine Learning in Python". In: *Journal of Machine Learning Research* 12, 2825–2830.

Perri, Vincenzo, Lisi Qarkaxhija, Albin Zehe, Andreas Hotho, and Ingo Scholtes (2022). *One Graph to Rule them All: Using NLP and Graph Neural Networks to analyse Tolkien's Legendarium*. arXiv: 2210.07871 [cs.CL].

Piper, Andrew and Sunyam Bagga (2022). "A Quantitative Study of Fictional Things". In: 268–279. https://ceur-ws.org/Vol-3290/long_paper1576.pdf.

Piper, Andrew, Richard Jean So, and David Bamman (2021). "Narrative Theory for Computational Narrative Understanding". In: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Ed. by Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics, 298–311. 10.18653/v1/2021.emnlp-main.26. https://aclanthology.org/2021.emnlp-main.26.

Poesio, Massimo, Juntao Yu, Silviu Paun, Abdulrahman Aloraini, Pengcheng Lu, Janosch Haber, and Derya Cokal (2023). "Computational models of anaphora". In: *Annual Review of Linguistics* 9, 561–587.

Rockwell, G. and S. Sinclair (2016). *Hermeneutica: Computer-Assisted Interpretation in the Humanities*. MIT Press.

Ryan, Marie-Laure, Kenneth Foote, and Maoz Azaryahu (2016). *Narrating space/spatializing narrative: Where narrative theory and geography meet*. The Ohio State University Press.

Schaeffer, Jean-Marie (1989). *Qu'est-ce qu'un genre littéraire?* Poétique. Seuil.

Schlegel, Friedrich, August Wilhelm Schlegel, August Ferdinand Bernhardi, and Wilhelm Dilthey (1996). *Critique et herméneutique dans le premier romantisme allemand : Textes de F. Schlegel, F. Schleiermacher, F. Ast, A.W. Schlegel, A.F. Bernhardi, W. Dilthey*. Trans. by Denis Thouard. Opuscules. Presses universitaires du Septentrion. https://books.openedition.org/septentrion/95397 (visited on 01/23/2024).

Schmid, W. (2010a). *Mental Events*. Hamburg University Press.

— (2010b). *Narratology. An introduction*. de Gruyter.

Schöch, Christof (2017). "Topic Modeling Genre: An Exploration of French Classical and Enlightenment Drama". In: *Digital Humanities Quarterly* 011.2.

Seminck, Olga, Philippe Gambette, Dominique Legallois, and Thierry Poibeau (2022). "The Evolution of the Idiolect over the Lifetime: A Quantitative and Qualitative Study of French 19th Century Literature". In: *Journal of Cultural Analytics* 7.3. 10.22148/001c.37588.

Silge, J. and D. Robinson (2017). *Text Mining with R: A Tidy Approach*. http://repo.darmajaya.ac.id/5417/ (visited on 09/19/2017).

Sims, M., J. Ho Park, and David Bamman (2019). "Literary Event Detection". In: *Proc. of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 3623–3634.

Sobchuk, Oleg and Artjoms Šeļa (2023). *Computational thematics: Comparing algorithms for clustering the genres of literary fiction*. arXiv: 2305.11251 [cs.CL].

Soni, Sandeep, Amanpreet Sihra, Elizabeth F Evans, Matthew Wilkens, and David Bamman (2023). "Grounding Characters and Places in Narrative Texts". In: *arXiv preprint arXiv:2305.17561*.

Sprugnoli, R. and S. Tonelli (2016). "Novel Event Detection and Classification for Historical Texts". In: *Computational Linguistics* 45.2, 229–265.

Stenetorp, Pontus, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun'ichi Tsujii (2012). "brat: a Web-based Tool for NLP-Assisted Text Annotation". In: *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*. Ed. by Frédérique Segond. Avignon, France: Association for Computational Linguistics, 102–107. https://aclanthology.org/E12-2021.

Toro Isaza, Paulina, Guangxuan Xu, Toye Oloko, Yufang Hou, Nanyun Peng, and Dakuo Wang (2023). “Are Fairy Tales Fair? Analyzing Gender Bias in Temporal Narrative Event Chains of Children’s Fairy Tales”. In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Ed. by Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki. Toronto, Canada: Association for Computational Linguistics, 6509–6531. 10.18653/v1/2023.acl-long.359.

Touvron, Hugo, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. (2023). “Llama 2: Open foundation and fine-tuned chat models”. In: *arXiv preprint arXiv:2307.09288*.

Underwood, Ted (2019). *Distant horizons: digital evidence and literary change*. The University of Chicago Press, 1–33.

Underwood, Ted, David Bamman, and Sabrina Lee (2018). “The Transformation of Gender in English-Language Fiction”. In: *Cultural Analytics* Feb 13 2018. 10.22148/16.019.

Van Cranenburgh, Andreas and Frank Van Den Berg (2023). “Direct Speech Quote Attribution for Dutch Literature”. In: *Proceedings of the 7th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*. Ed. by Stefania Degaetano-Ortlieb, Anna Kazantseva, Nils Reiter, and Stan Szpakowicz. Association for Computational Linguistics, 45–62. 10.18653/v1/2023.latechclfl-1.6.

Vanni, Laurent, Melanie Ducoffe, Carlos Aguilar, Frederic Precioso, and Damon Mayaffre (2018). “Textual Deconvolution Saliency (TDS) : a deep tool box for linguistic analysis”. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Ed. by Iryna Gurevych and Yusuke Miyao. Melbourne, Australia: Association for Computational Linguistics, 548–557. 10.18653/v1/P18-1051.

Vianne, Laurine, Yoann Dupont, and Jean Barré (2023). “Gender Bias in French Literature”. In: *Computational Humanities Research Conference*. CEUR Workshop Proceedings (CEUR-WS. org), 247–262. https://ceur-ws.org/Vol-3558/paper2449.pdf.

Vishnubhotla, Krishnapriya, Frank Rudzicz, Graeme Hirst, and Adam Hammond (2023). “Improving Automatic Quotation Attribution in Literary Novels”. In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Ed. by Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki. Toronto, Canada: Association for Computational Linguistics, 737–746. 10.18653/v1/2023.acl-short.64.

Weischedel, Ralph, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Nianwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, et al. (2013). “Ontonotes release 5.0 ldc2013t19”. In: *Linguistic Data Consortium, Philadelphia, PA* 23, 170.

Woloch, Alex (2003). *The One vs. the Many*. Princeton University Press.

Yu, B. (Sept. 5, 2008). “An evaluation of text classification methods for literary study”. In: *Literary and Linguistic Computing* 23.3, 327–343. 10.1093/llc/fqn015.

Zehe, Albin, Leonard Konle, Lea Katharina Dümpelmann, Evelyn Gius, Andreas Hotho, Fotis Jannidis, Lucas Kaufmann, Markus Krug, Frank Puppe, Nils Reiter, Annekea Schreiber, and Nathalie Wiedmer (2021). “Detecting Scenes in Fiction: A new Segmentation Task”. In: *Proceedings of the 16th Conference of the European Chapter of the*

Association for Computational Linguistics: Main Volume. Ed. by Paola Merlo, Jorg Tiedemann, and Reut Tsarfaty. Online: Association for Computational Linguistics, 3167–3177. 10.18653/v1/2021.eacl-main.276.

Zhang, Weiwei, Jackie Chi Kit Cheung, and Joel Oren (2019). "Generating character descriptions for automatic summarization of fiction". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 33. 01, 7476–7483.

Zundert, Joris van, Andreas van Cranenburgh, and Roel Smeets (2023). "Putting Dutchcoref to the Test: Character Detection and Gender Dynamics in Contemporary Dutch Novels". In: *CHR 2023: Computational Humanities Research Conference*. CEUR Workshop Proceedings (CEUR-WS. org). Paris, France, 757–771. https://ceur-ws.org/Vol-3558/paper9264.pdf.

Zundert, Joris van, Marijn Koolen, Julia Neugarten, Peter Boot, Willem van Hage, and Ole Mussmann (2022). "What Do We Talk About When We Talk About Topic?" In: *CHR 2022: Computational Humanities Research Conference*. Antwerp, Belgium. https://ceur-ws.org/Vol-3290/short_paper5533.pdf.

A. Appendix: Confusion matrices for BookNLP-fr-based models

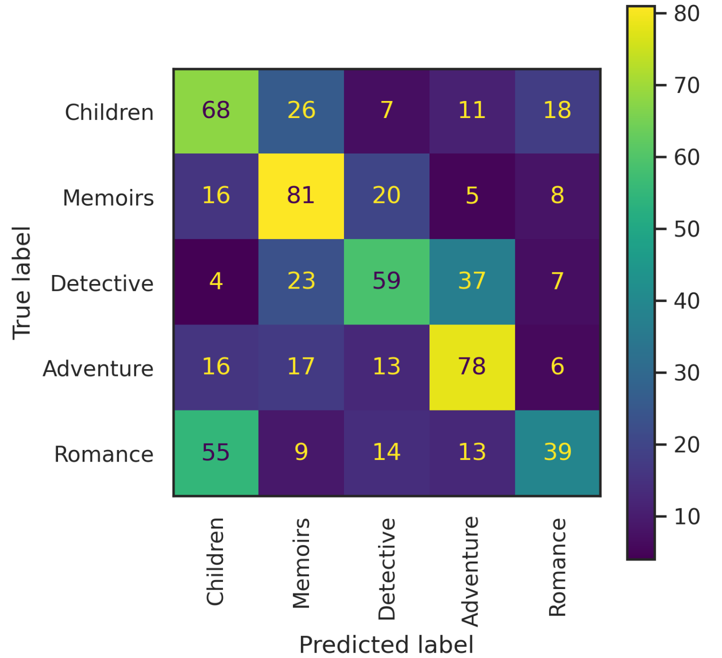

Figure 7: Confusion Matrix for ADJ features

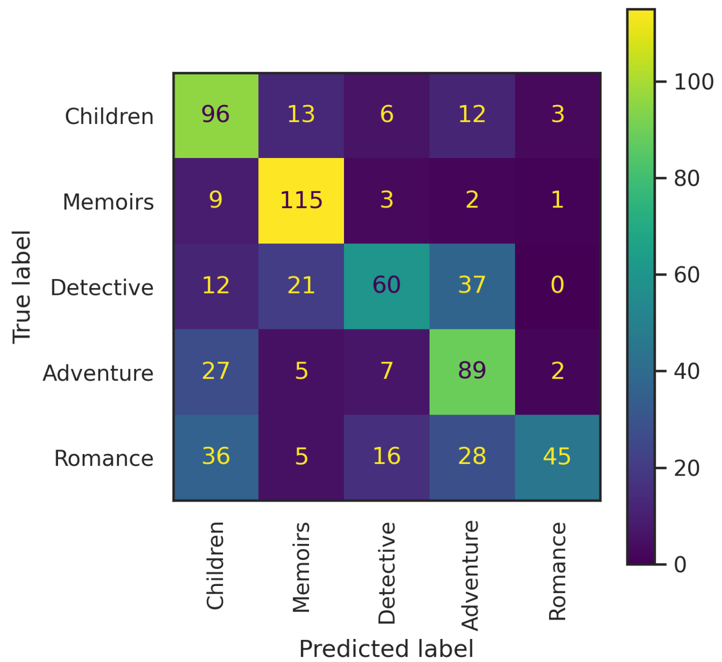

Figure 8: Confusion Matrix for AGENT features

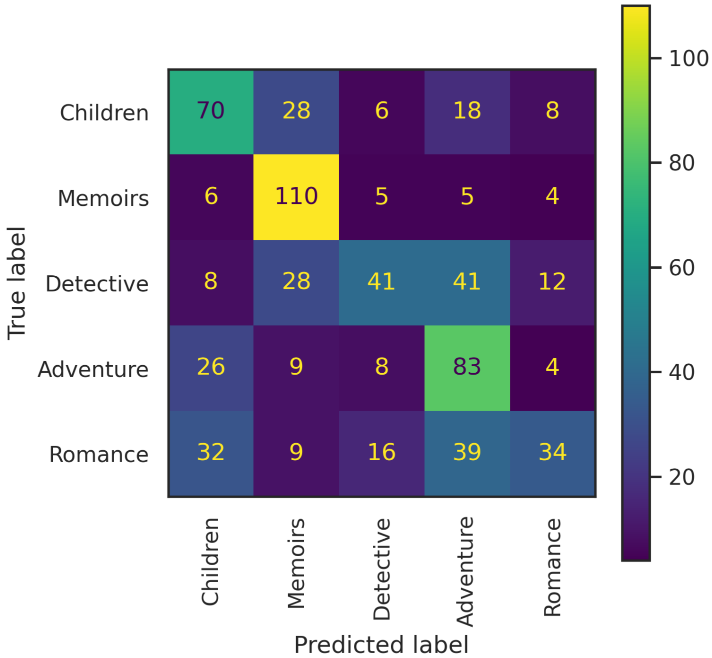

Figure 9: Confusion Matrix for PATIENT features

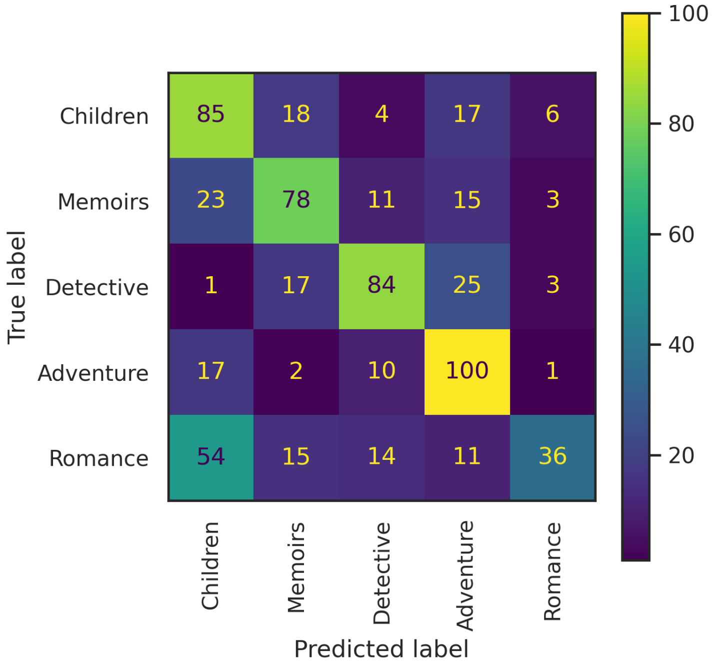

Figure 10: Confusion Matrix for FAC features

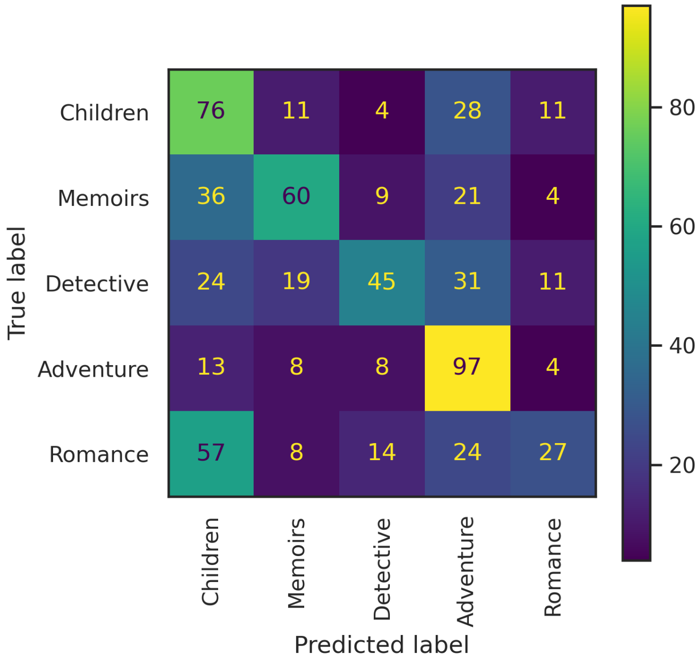

Figure 11: Confusion Matrix for GPE features

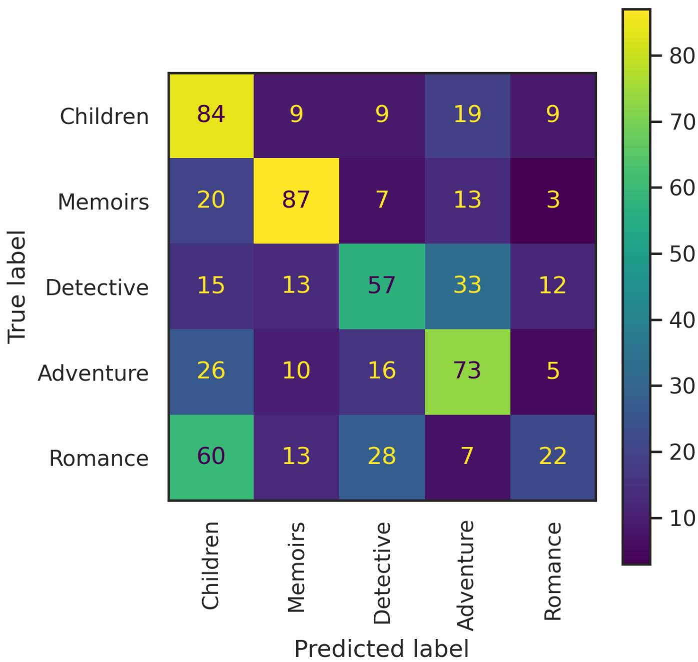

Figure 12: Confusion Matrix for TIME features

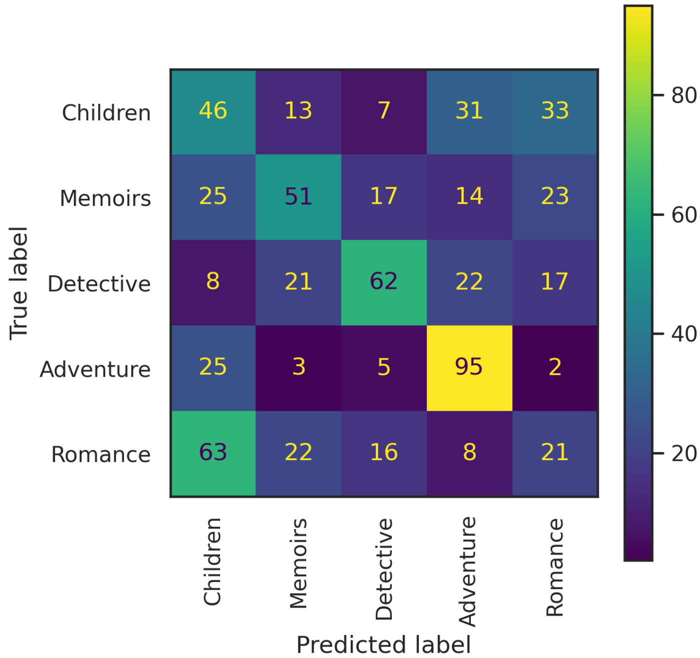

Figure 13: Confusion Matrix for VEH features

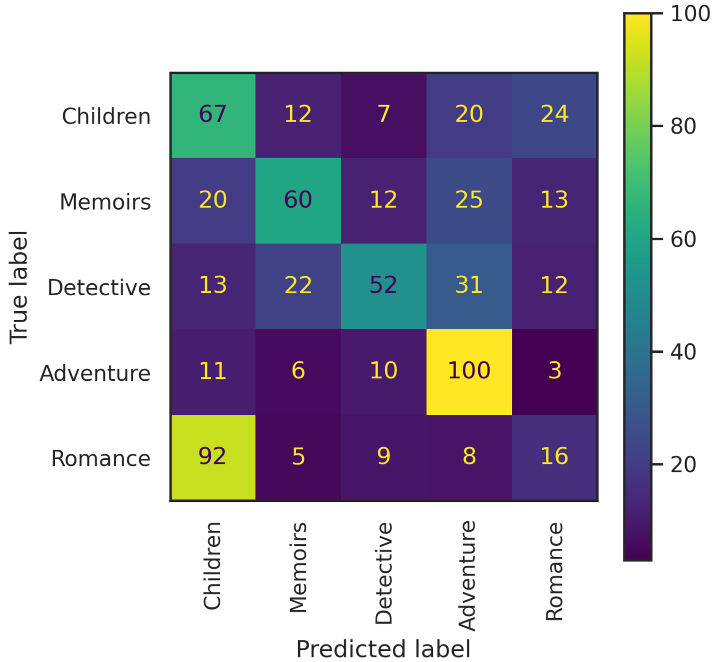

Figure 14: Confusion Matrix for LOC features