Using Large Language Models to Study Fiction

Diving into Computational Literary Studies

Jean Barré & Pedro Cabrera Ramirez

4 mars 2024

PSL Intensive Week DHAI

Computational Literary Studies

- Machine Learning & Text Mining to model concepts in large literary corpus.
- A key concept : Distant Reading - Franco Moretti.
- Our Project : using LLMs as annotation power / as plot predictors
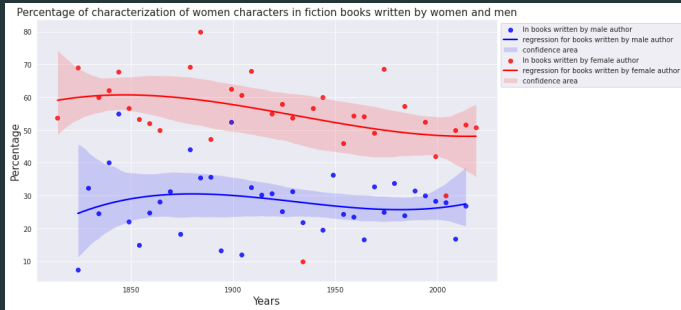- **Stakes** : Evaluation - Interpretation

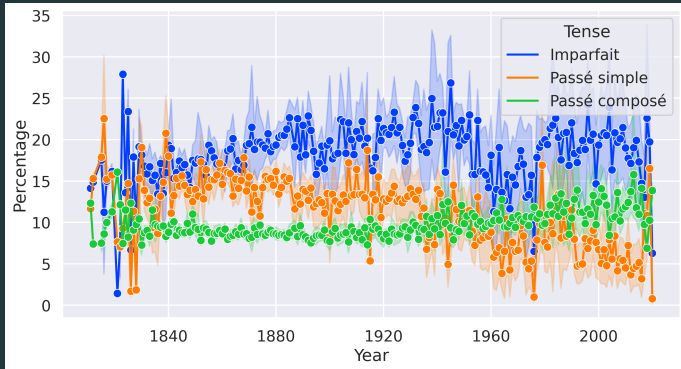**Figure 1 :** Evolution of women's screentime over time, broken out by author gender

Figure 2 : Evolution of verb tense over time

Group Project 2 : Measuring predictability and surprise in adventure and police fiction

LLMs as a tool to study plot construction

- Plot - the Achilles' heel of the discipline
- LLMs & plot - (Underwood, 2024)
  `https://tedunderwood.com/2024/01/05/`
  `can-language-models-predict-the-next-twist-in-a-stor`
- Modeling Predictability & Surprise

Main Research Questions

Are LLMs relevant to CLS ?

- Can LLMs effectively predict the unfolding events in a novel ?
- How can we assess their performance ?
- What changes in results occur when specializing (fine-tuning) LLMs on subgenres ?

Background : Language Models are Next Word Predictors

LLMs

- Pretrained language models take as input text and output the most likely next token(s) of that text.
- LLMs as an annotation tool - subgenre tags (recall problem)
- Finetuning - Specialization of LLMs on specific tasks

Subgenres are powerful cultural artefacts

- Plot is highly dependent on subgenres
- Adventure & Detective novels - Similarity / Differences?
- Predictability benchmark of subgenre == Plot codification?

Work Packages

- Data Annotation - Stereotypical Scenes + Plot completion
- Evaluation - text similarity : cosine + perplexity
- Scripts - Data preparation, Completion, Embeddings, Evaluation, Machine learning
- Analyses - Interpretation & Presentation

Work Packages - dated

Day	Indicative Work Packages
Monday	Project Presentation, Data Annotation - start.
	Notebooks - llms_annotation + llms_fine_tuning
Tuesday	Data Annotation.
	Notebooks - get_embeddings - text_similarity
Wednesday	Data Annotation - end.
	Notebooks - machine learning.
Thursday	Results -> Analyses + Start building Presentation
Friday	Morning - Writing Results, Finalizing Presentation
	Afternoon - Presentation

Questions ?

Github of the Project : `https://github.com/`
`crazyjeannot/DHAI_2024_TD_LLM`