

The Role of Global and Local Context in Named Entity Recognition



Arthur Amalvy^{*}, Vincent Labatut^{*}, Richard Dufour[†]

^{*} Laboratoire Informatique d'Avignon

[†] Laboratoire des Sciences du Numérique de Nantes

GitHub



contact

arthur.amalvy@univ-avignon.fr

Context in Named Entity Recognition (NER)

"Raoden stood, and as he did, his eyes fell on Elantris again." relevant context "Elantris is a large city."

Is Elantris a person (PER) or a location (ORG)? To answer, we need context.

- Context can either be local (close to the sentence) or global (document level)
- Transformers-based models limited range is problematic for global context
- Should models progress towards taking global context into account?

A Revised and Extended Literary Dataset

To study the role of global context, we correct and extend an existing English NER literary dataset from Dekker et al., 2019.

Dataset:

- The first chapter of 40 novels.
- Documents long enough to study global context
- Freely available (see QR code)

Contributions:

- We fix PER (persons) dataset errors by applying a consistent annotation scheme.
- We annotate LOC (locations) and ORG (organizations) NER classes.

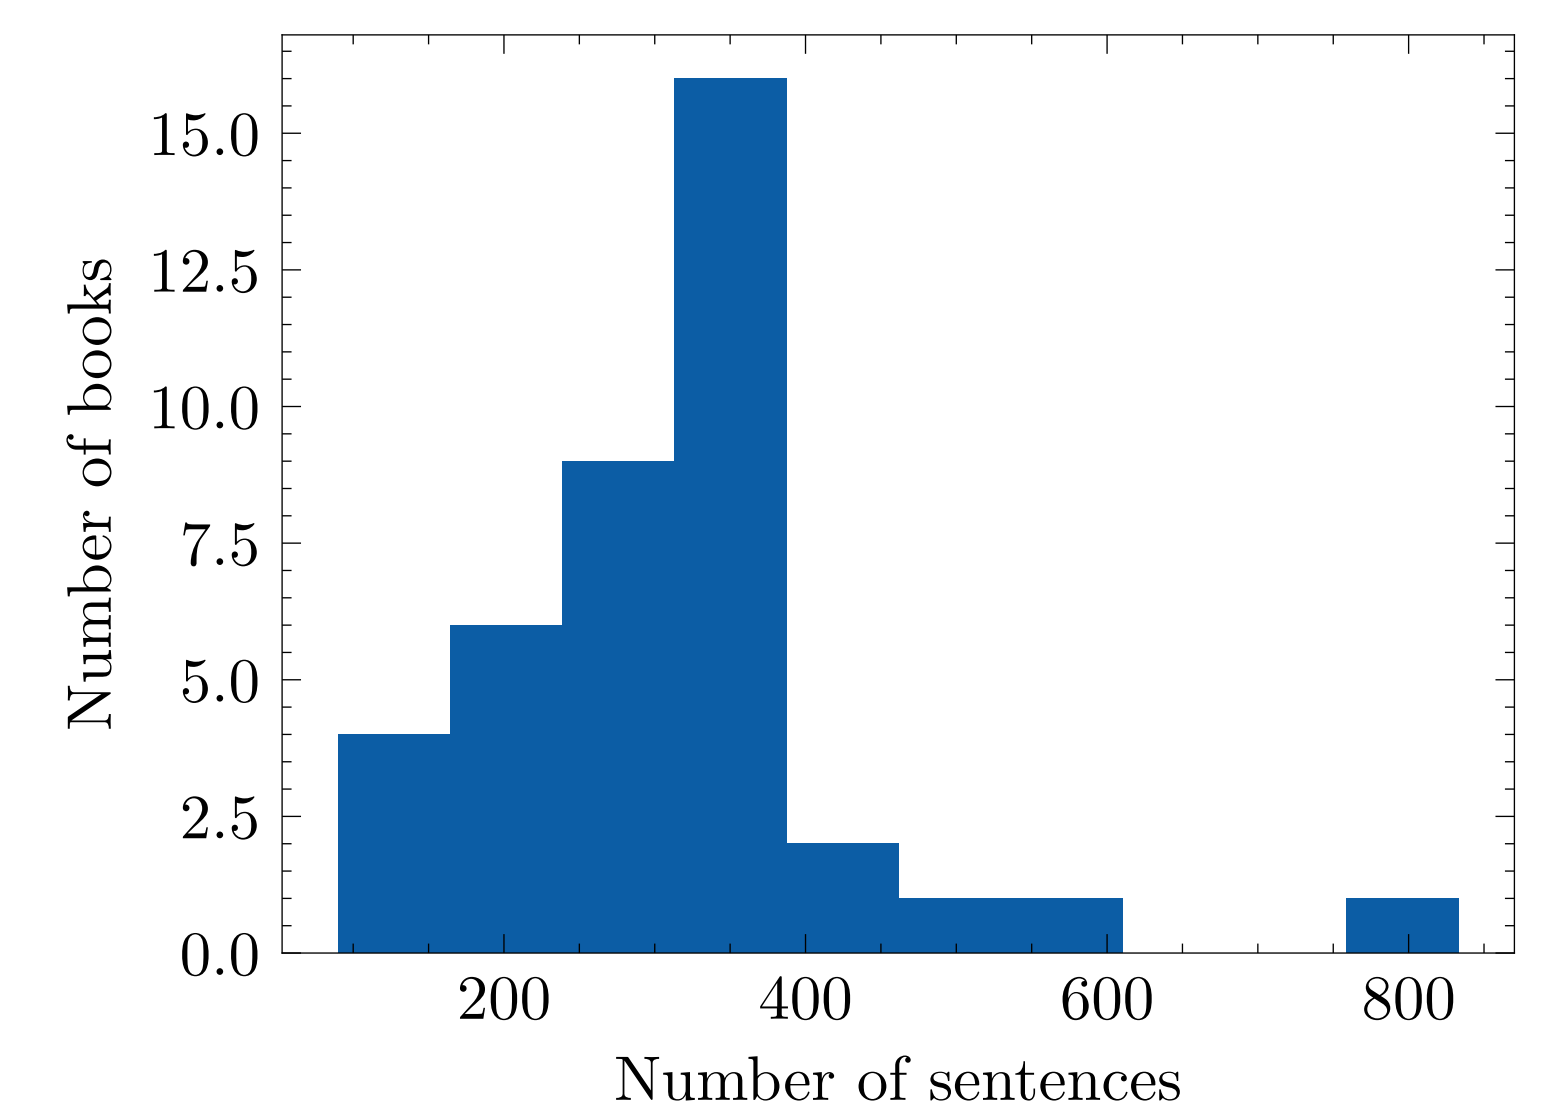


Figure 1: Distribution of the number of sentences per document in our dataset

Experiments

Comparing Local and Global Context

Is it useful to retrieve global context? We compare several simple retrieval heuristics.




Figure 2: Mean F1 score versus max number of retrieved sentences for all retrieval heuristics across 3 runs

Key takeaways:

- In general, context is beneficial
- Simple global heuristics offers competitive performance compared with local context
- Stronger methods may yield more performance gains

Evaluating Re-ranking Oracles

- We evaluate "retrieve and rerank" procedure with a perfect "oracle" re-ranker.
- This gives a performance upper-bound of such a retrieval method.




Figure 3: Mean F1 score versus max number of retrieved sentences for oracle versions of all retrieval heuristics across 3 runs

Key takeaways:

- Oracle re-rankers perform best when ranking over global context
- This result prompts for further research in training re-ranker models