

Think Before you Write: QA-Guided Reasoning for Character Descriptions in Books

Argyrios Papoudakis, Mirella Lapata, Frank Keller
School of Informatics, University of Edinburgh, UK



THE UNIVERSITY of EDINBURGH
informatics



Takeaways:


- For character description generation, enabling built-in reasoning degrades performance compared to no reasoning (empty trace).
- Decoupling reasoning from generation and training each component separately leads to more faithful, informative and better-grounded descriptions.

1. Character Description

- Understanding characters is important for narrative-focused applications (e.g., summarization, story generation).
- Character description task:** a short factual profile of a character (e.g., actions, relationships, attributes).

Character: Carrie Soto **Book:** Carrie Soto Is Back
Description: Carrie is the protagonist and narrator. Carrie values winning above all else. She wants to be the greatest, and she re-enters professional tennis after retirement just to keep her Grand Slam record. *This obsession began with her father, Javier. After her mother died, Javier threw himself into making Carrie the best tennis player he could, calling her Achilles ...* [truncated]


2. Built-in reasoning



→ Descriptions become more verbose, in some cases recall increases but precision degrades significantly leading to worse F1 score.

Performance drops across metrics and different settings (retrieval-based, hierarchical and incremental processing).

3. QA-guided Reasoning



Modular approach that decouples reasoning from generation. Two separate models, which are trained independently.

Structured reasoning trace: question-answering pairs with typed questions and explanations.

Reasoning model is trained with GRPO using silver-standard traces derived from gold descriptions. F1 score is used as a training reward.

Generation model is trained either with SFT or GRPO using gold-standard descriptions.

Training and Evaluation Datasets

Books paired with human-written character descriptions.


BookWorm (Papoudakis et al., 2024)
books (mostly novels and plays) from Project Gutenberg.

CroSS (Yuan et al., 2024)
Novels published in 2022 and 2023.


	BookWorm	CroSS
Books	324	126
Characters	9.74	6.56
Samples	5,869	824
Input Length	97,685	129,113
Output Length	88.79	295.27

4. Results

QA-guided reasoning outperforms no / built-in reasoning



GRPO training outperforms ZS and SFT



5. Limitations

- Different types of reasoning structure and free-form reasoning.
- Different RL algorithms, alternative training rewards and joint training.
- Human-based evaluation and character-focused evaluation metrics.

Questions? Reach out
a.papoudakis@sms.ed.ac.uk