InfoLossQA: Characterizing and Recovering Information Loss in Text Simplification

Jan Trienes, Sebastian Joseph, Jörg Schlötterer, Christin Seifert, Kyle Lo, Wei Xu, Byron C. Wallace, Junyi Jessy Li

ACL 2024, Bangkok

Information loss is inherent to text simplification

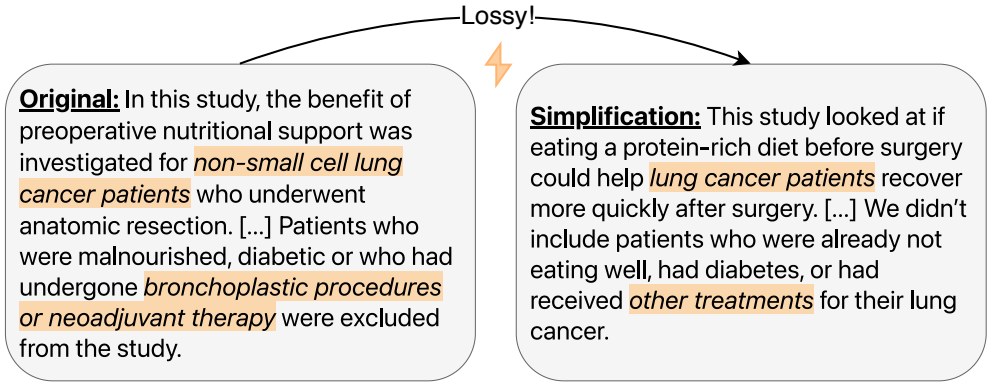

Why is this a problem? May lead to…

- Reduced reader comprehension (Agrawal and Carpuat, TACL 2024)
- Factuality errors (Devaraj et al., ACL 2022)

Information loss is inherent to text simplification

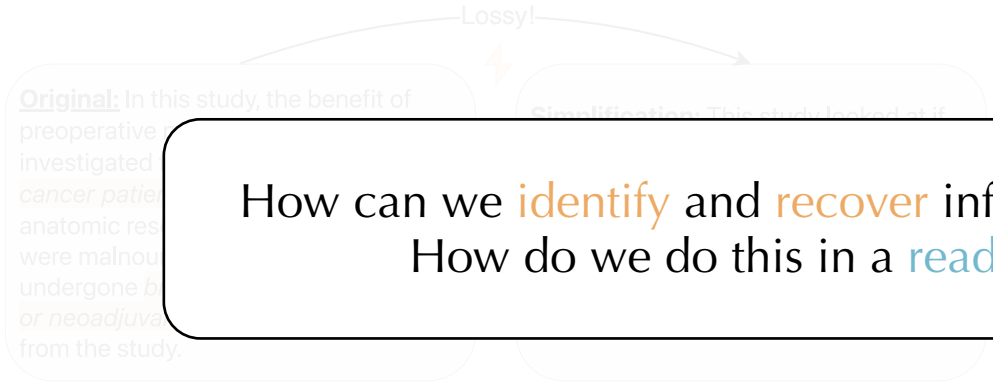

How can we identify and recover info loss in simplification?
How do we do this in a reader-centric way?

Why is this a problem? May lead to…

- Reduced reader comprehension (Agrawal and Carpuat, TACL 2024)
- Misinterpretation or factuality errors (Devaraj et al., ACL 2022)

We propose InfoLossQA

Task: generate questions that recover information loss

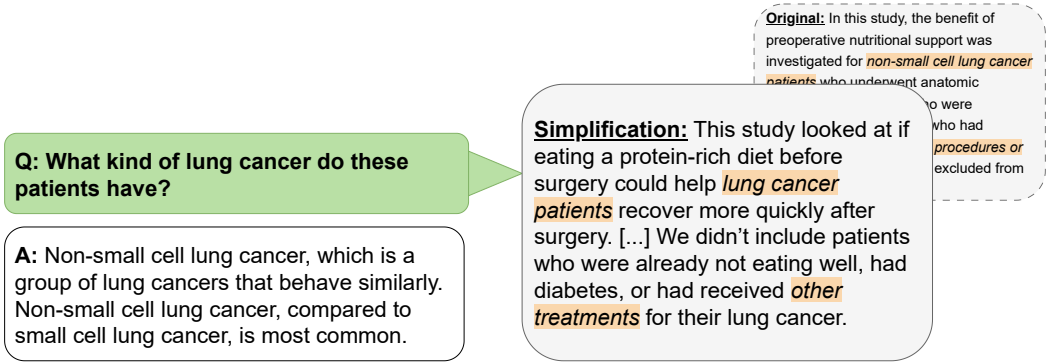

Q asks for missing information;
A provides it in lay language based on original

QA builds on Questions Under Discussion (QUD) theory

Qs are rooted in simple text, w/o assuming reader access to original

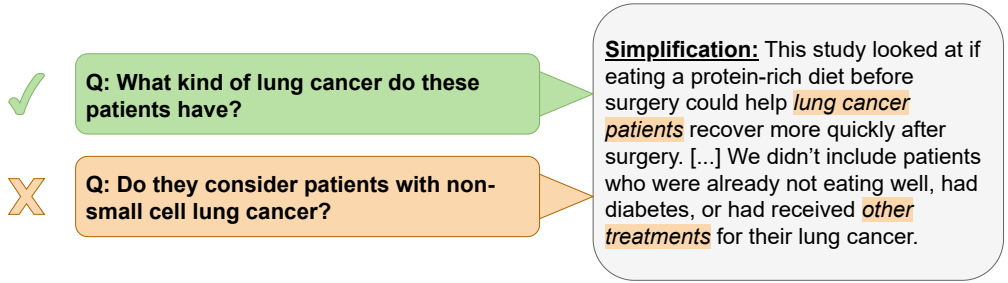

QUD: Benz and Jasinskaja (2017)

Data: Linguists annotate info loss in LLM simplifications

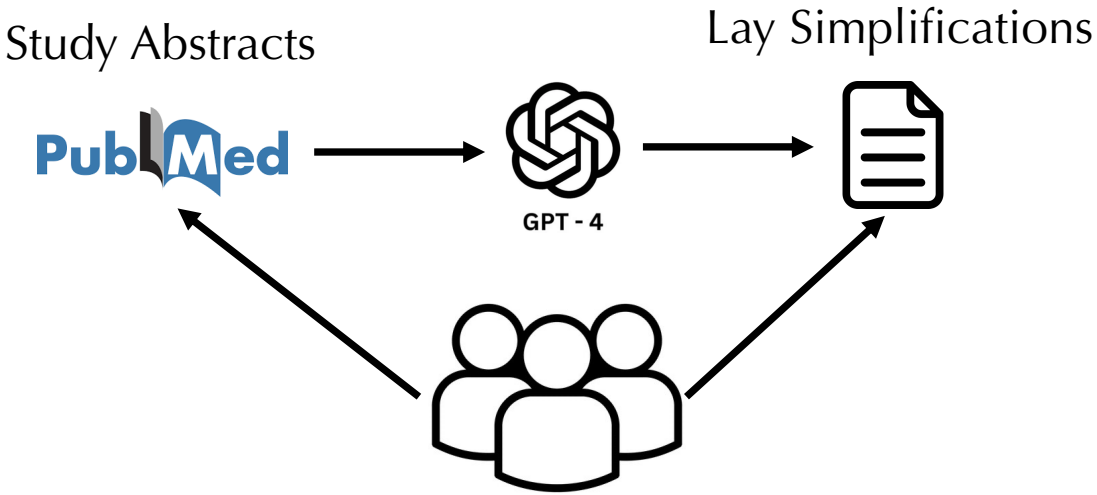

Key Stats

- 104 parallel simplifications
- |Original| = 312 tokens
- |Simplified| = 271 tokens
- 4 – 6 QAs/annotator/doc
- 1,000 QAs total

Automatically identifying information loss

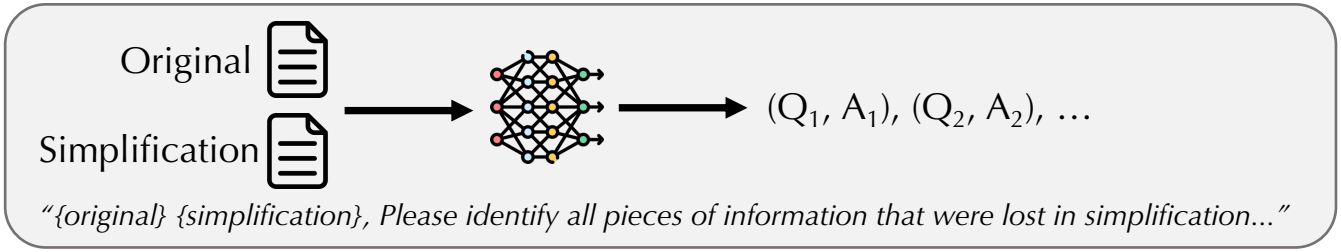

Method 1.
End-to-end prompting

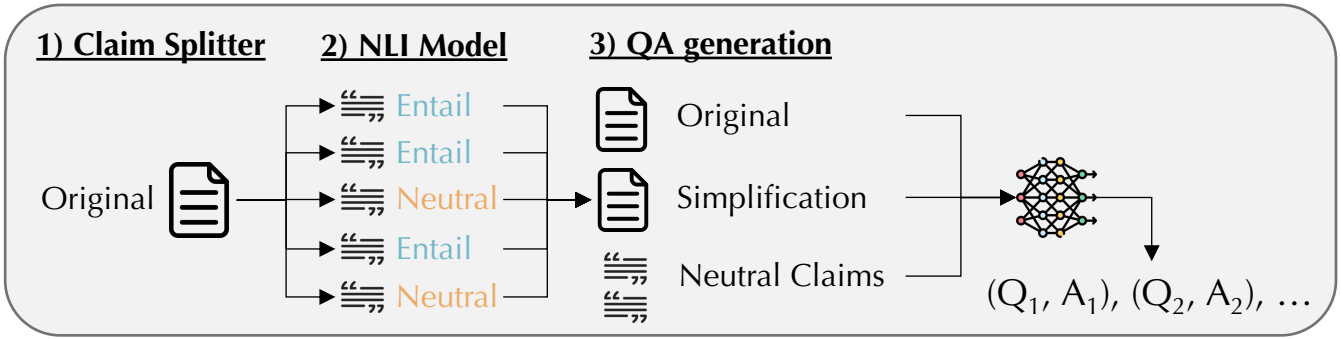

Method 2.
NLI pipeline

QAs provide a high-level summary of what's lost

Question Type	Frequency	Examples
Procedural	34.3%	How did they measure the patients opioid medication needs? How did the study control for bias?
Concept	25.7%	What kind of hip surgery were patients undergoing? What type of mental illnesses are being studied?
Extent	17%	How many patients were in each group? How much lurasidone was given to patients?
Comparison	8.3%	How much did abnormal blood vessels reduce compared to placebo?
…	…	…

Question type taxonomy by Cao and Wang (ACL, 2021)

Models are not reliable at identifying information loss

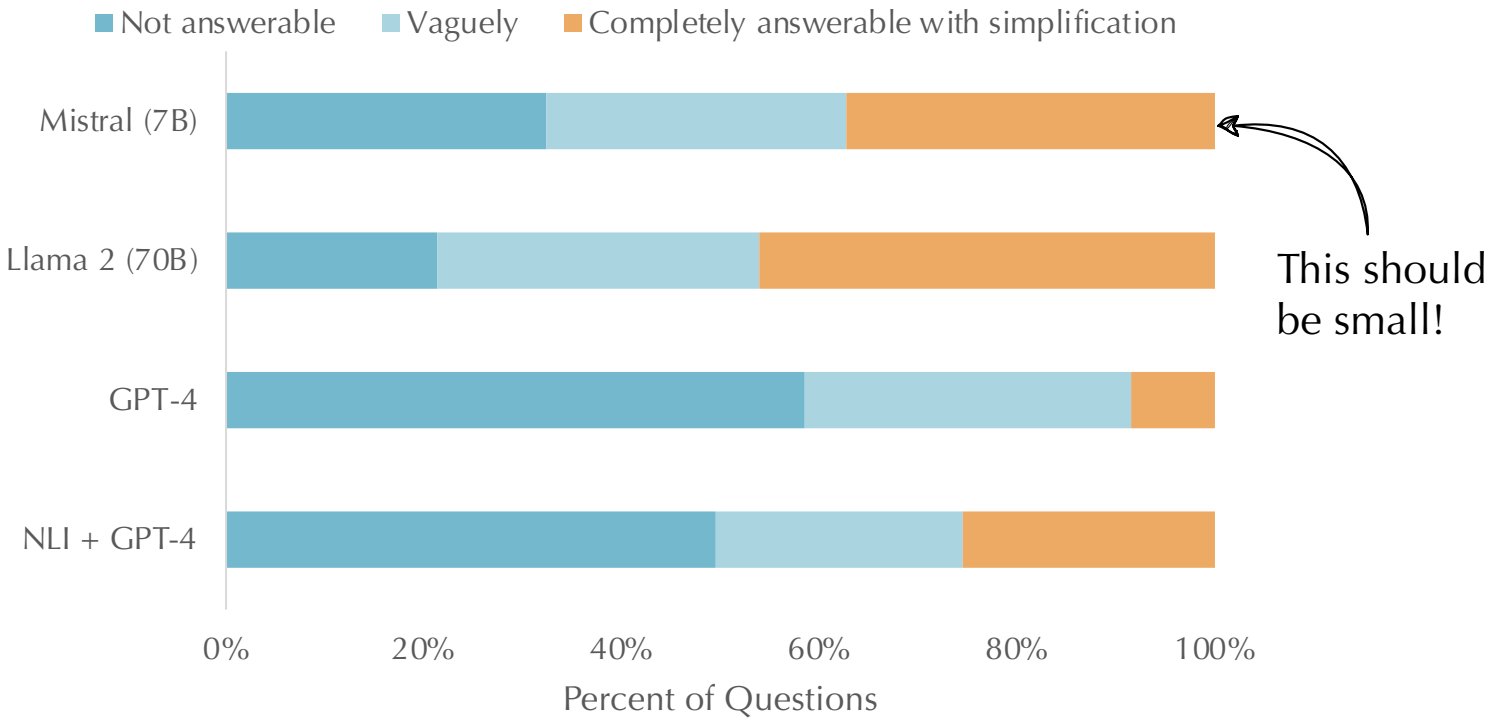

Models often fail to rationalize where info was lost

Original: [...] the benefit of preoperative nutritional support was investigated for *non-small cell lung cancer patients* who underwent anatomic resection.

Simplification: This study looked at if eating a protein-rich diet before surgery could help *lung cancer patients* recover more quickly after surgery.

Evidence

… mostly correct

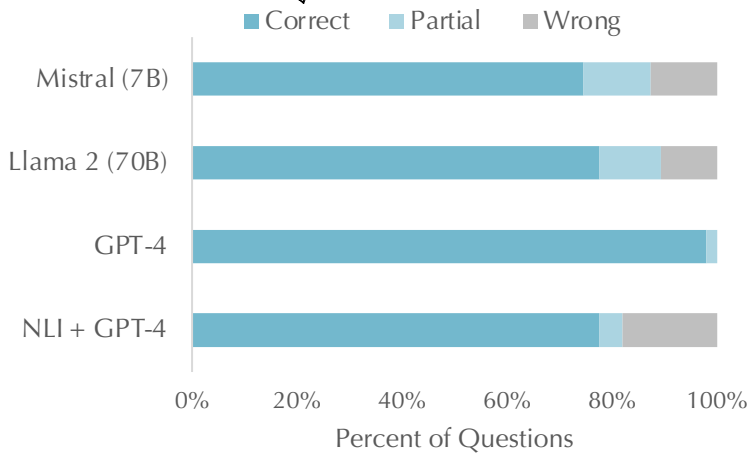

Rationale

… is often missing

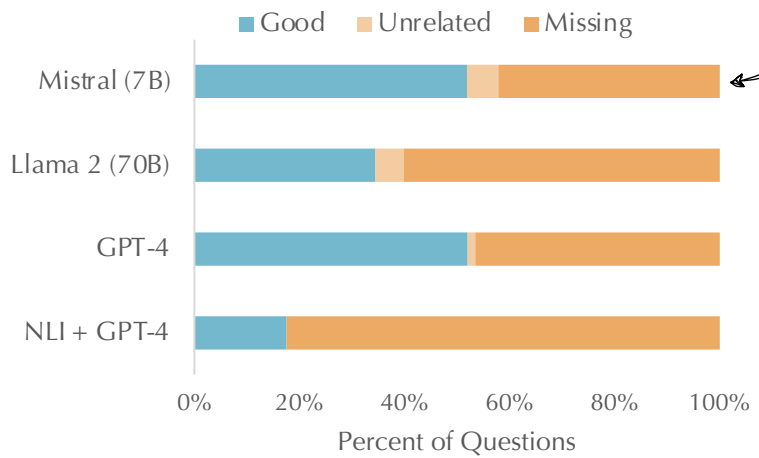

Models have low recall of human-generated QAs

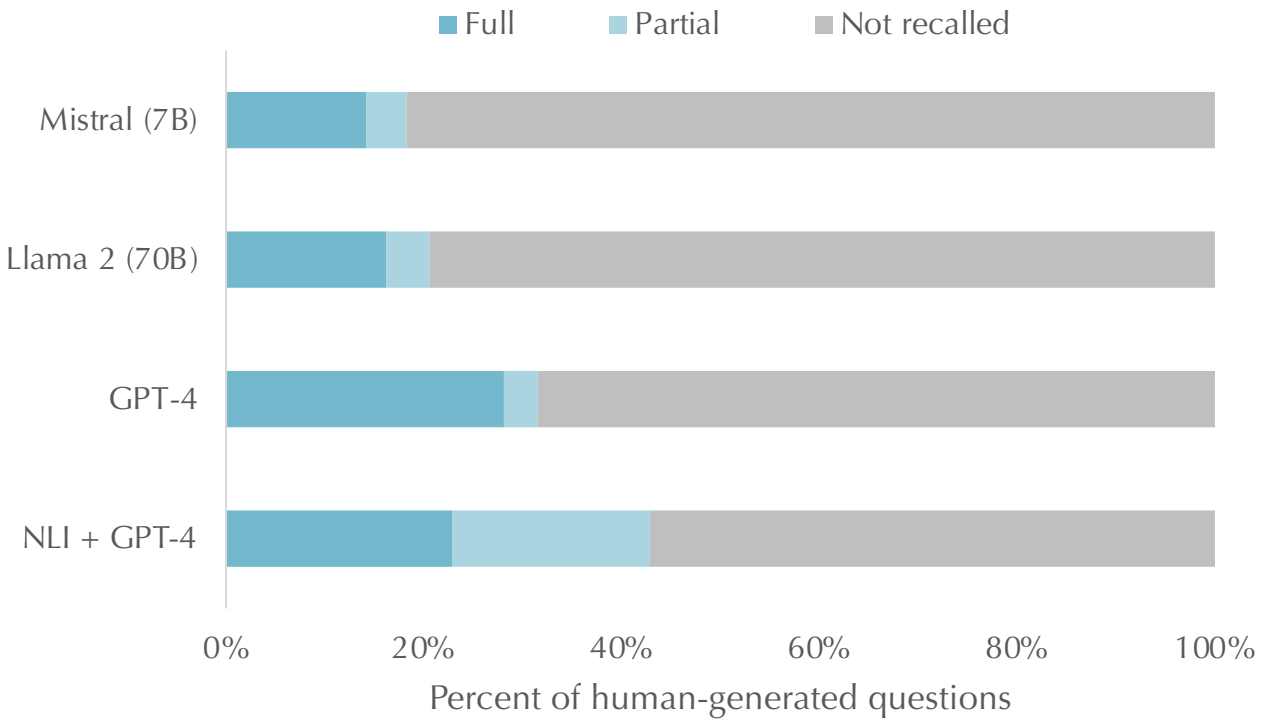

Takeaways

- Simplification is inherently lossy which may affect comprehension
- Solution: use QAs to **characterize** and **recover** info loss in a reader-centric way
- Current methods struggle to reliably identify info loss

Thanks! Code + Data available

github.com/jantrienes/InfoLossQA

jan.trienes@uni-marburg.de