Don’t Add, don’t Miss: Effective Content Preserving Generation from Pre-Selected Text Spans

Aviv Slobodkin, Avi Caciularu, Eran Hirsch, Ido Dagan

- Recently-introduced a new task:
  - **Controlled Text Reduction (CTR)**
  - Input: document + highlights
  - Output: custom summary (covering all and only highlights)

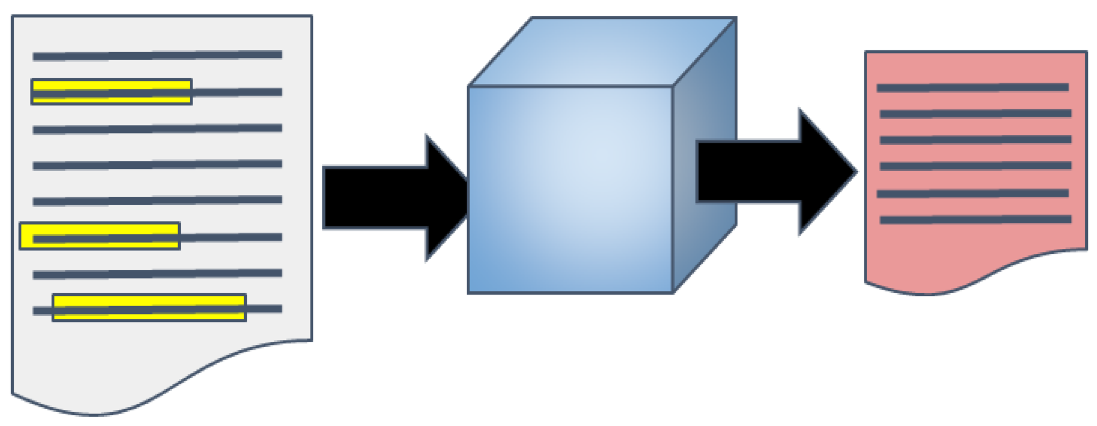

Motivation:

- **Modularity** of research and architectures:
  - Targeted research → better modelling
  - Modular architectures (reusable models)
  - More control over content generation
- **Human-in-the-loop**.

Problem:

No high-quality CTR model

- Existing baseline – only around 50% adherence to highlights

Improve baseline CTR model via:

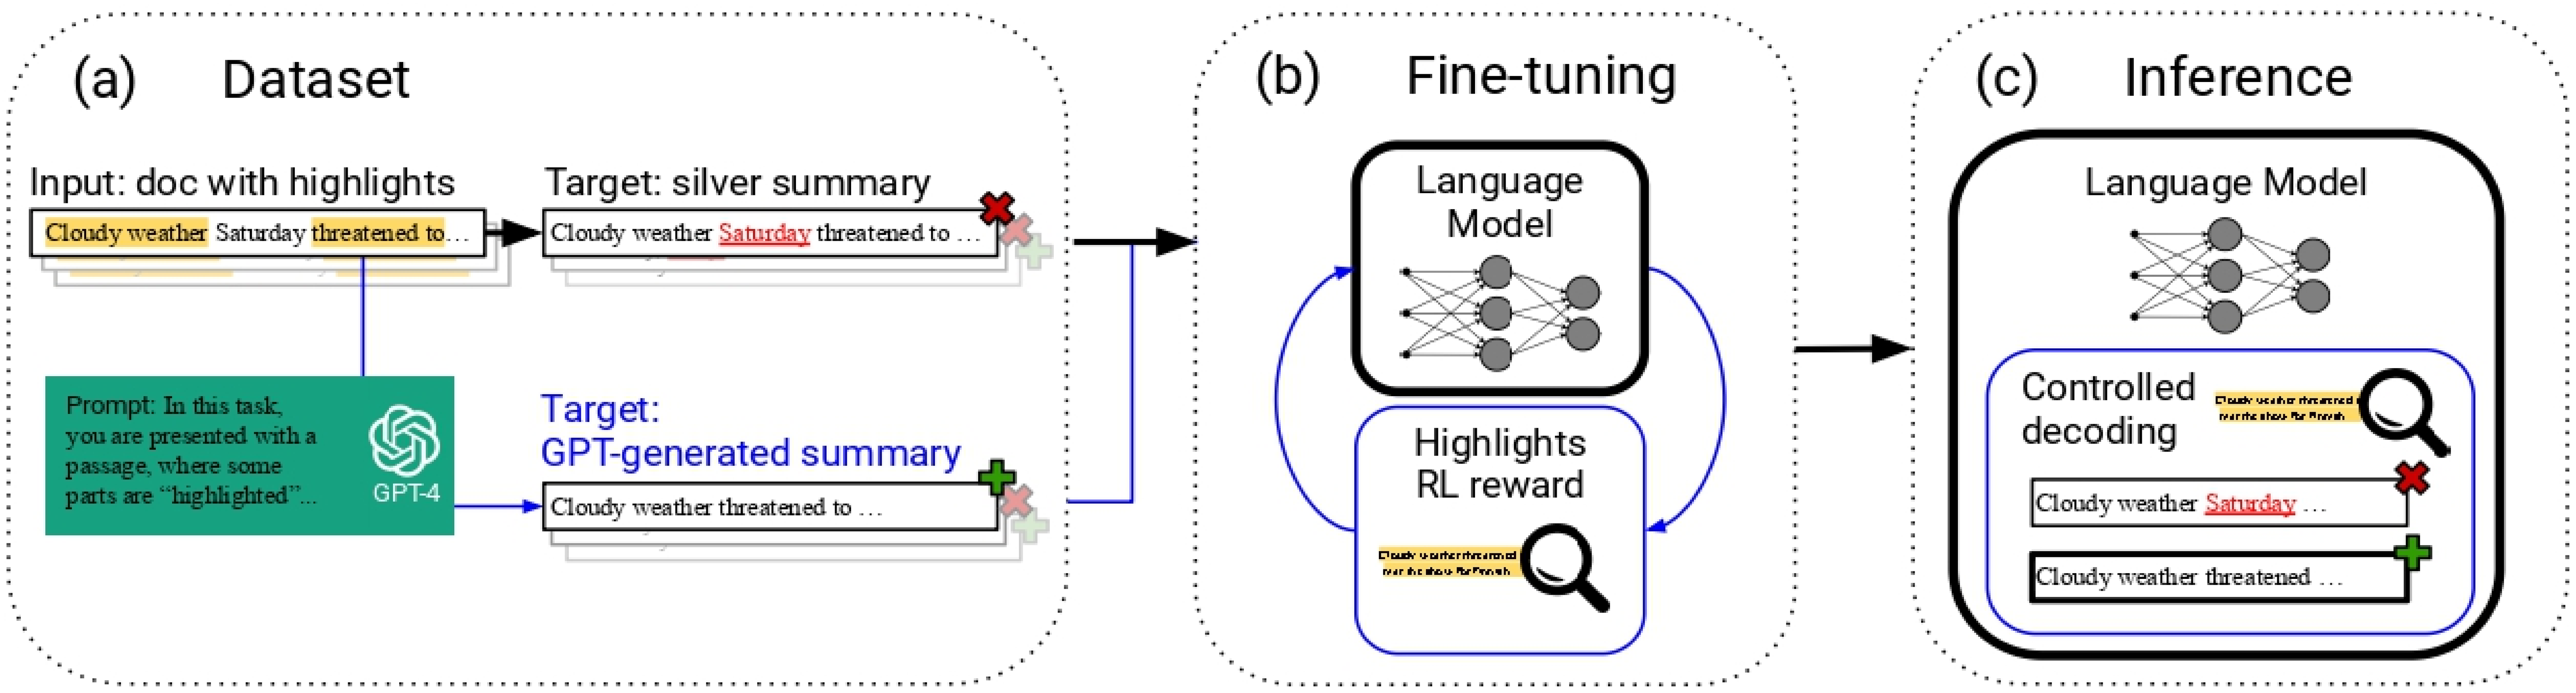

Improving Dataset Quality via Distillation

- Generate better outputs for highlights using GPT-4
  - Via in-context-learning (2 examples) and a CoT-like prompting
- Finetune on new dataset

RL Finetuning

- Deploy the Quark algorithm - a reinforced (un)learning algorithm
- Reward - alternating between ROUGE-L precision and recall, compared to the highlights concatenation

Controlled Decoding

- Adapt Beam-search
- At every step - lookahead mechanism
  - Complete potential summary for each candidate
- Beam selection: combination of LLM score and highlights adherence
  - Calculated by ROUGE-L between lookahead completion and highlights concatenation

Results

Model	Faithfulness (P)	Coverage (R)	F-1
Flan-T5	71.1	74.0	72.5
Flan-T5 (cleaned)	79.1	90.8	84.6
+RL	81.3	93.4	86.9
+Con. Decoding	85.6	91.3	88.3
+RL +Con. Decoding	83.4	92.3	87.7

- Key points:
  - Distillation – substantial improvements
  - Best coverage – distillation + RL
  - Best faithfulness – distillation + controlled decoding
  - RL + controlled decoding – trade-off effect