

My Climate Copilot: An Application of NLP in Climate Adaptation for Agriculture

Vincent Nguyen¹, Sarvnaz Karimi¹, Willow Hallgren¹, Ashley Harkin², Mahesh Prakesh¹

CSIRO¹, Bureau of Meteorology²

Vincent.Nguyen@csiro.au

https://ngu.vin



Motivation

What did we do?

We designed a co-pilot for farm advisors that helps review relevant management practices suggested in the literature.

Why is it important?

These management practices should help farmers adapt their farms to climate change, reducing their financial risk while improving national food security.

How do we define relevant management advice?

- Commodity and location-specific
- Considers significant climate factors
- Grounded in recent and trustworthy literature

Data sources


Where can we get information?

Trustworthy literature and data are needed to ensure that we can provide scientifically robust management advice to farm advisors.

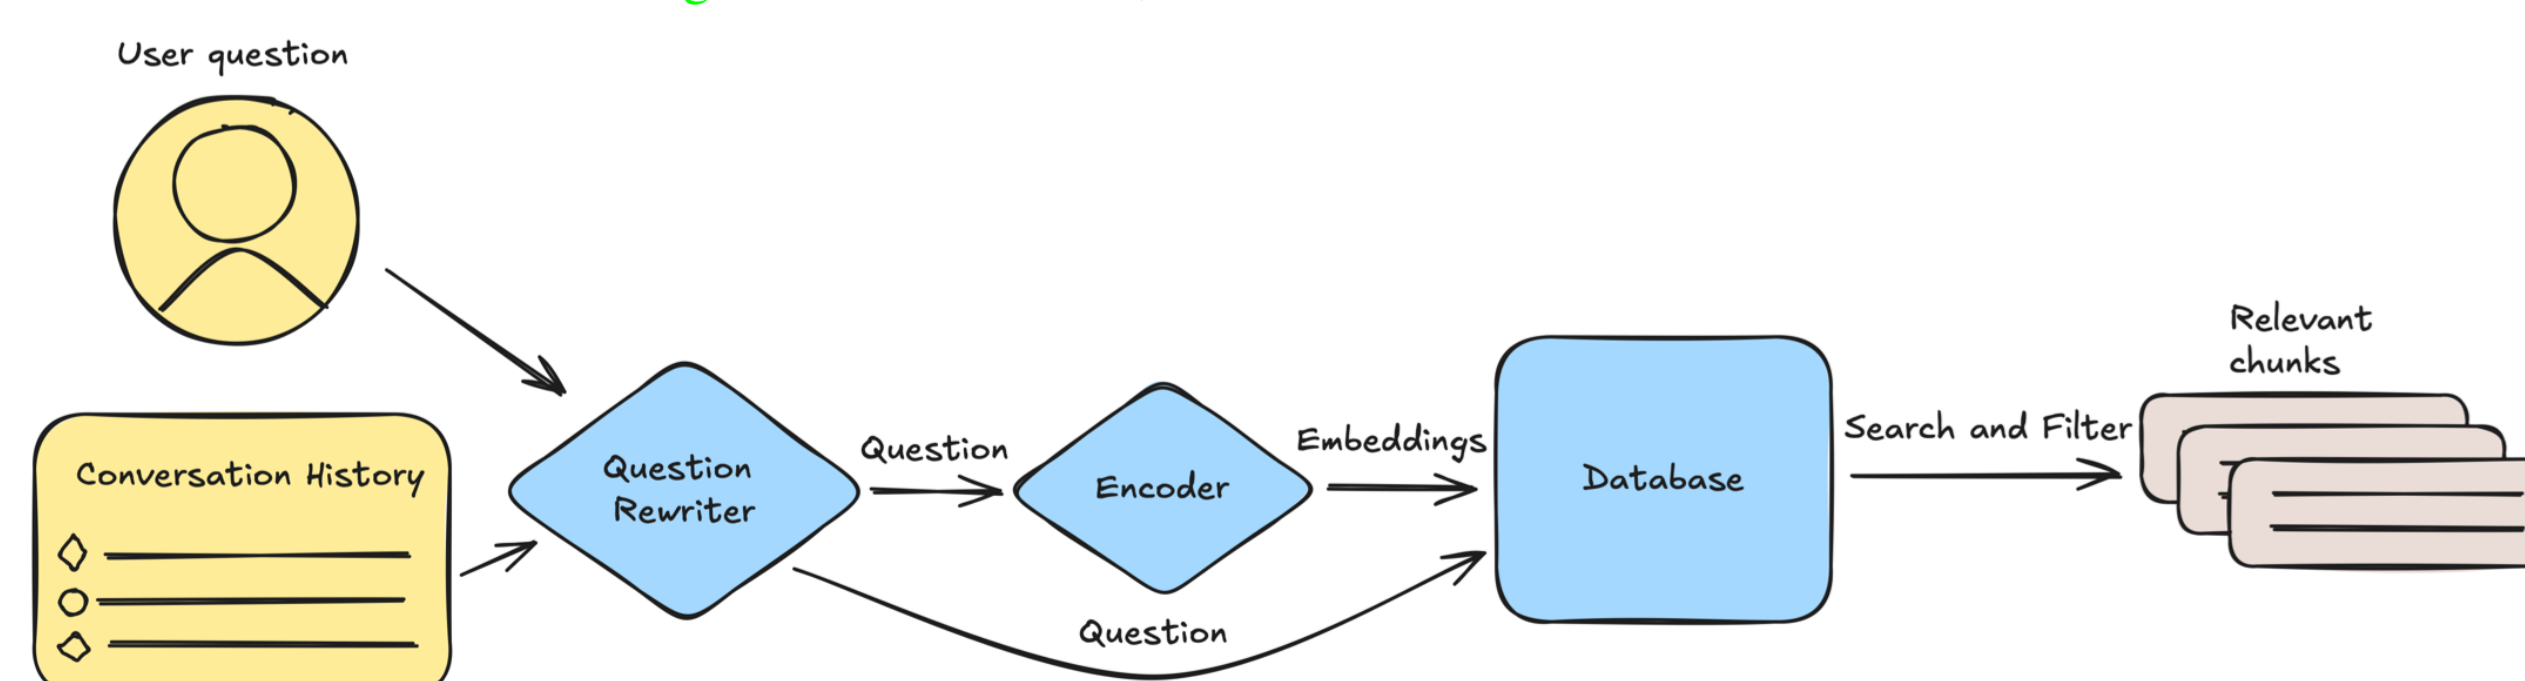
1. Agricultural literature from academia
2. Agricultural literature from the top 100 agriculture journals
3. Australian industry grey literature

Corpus	# Documents	# Chunks (C=400)	Size (GB)
Agricultural literature	1.36M	30.6M	124
Top agriculture journals	126K	221K	8.3
Grey literature	28	1513	0.008

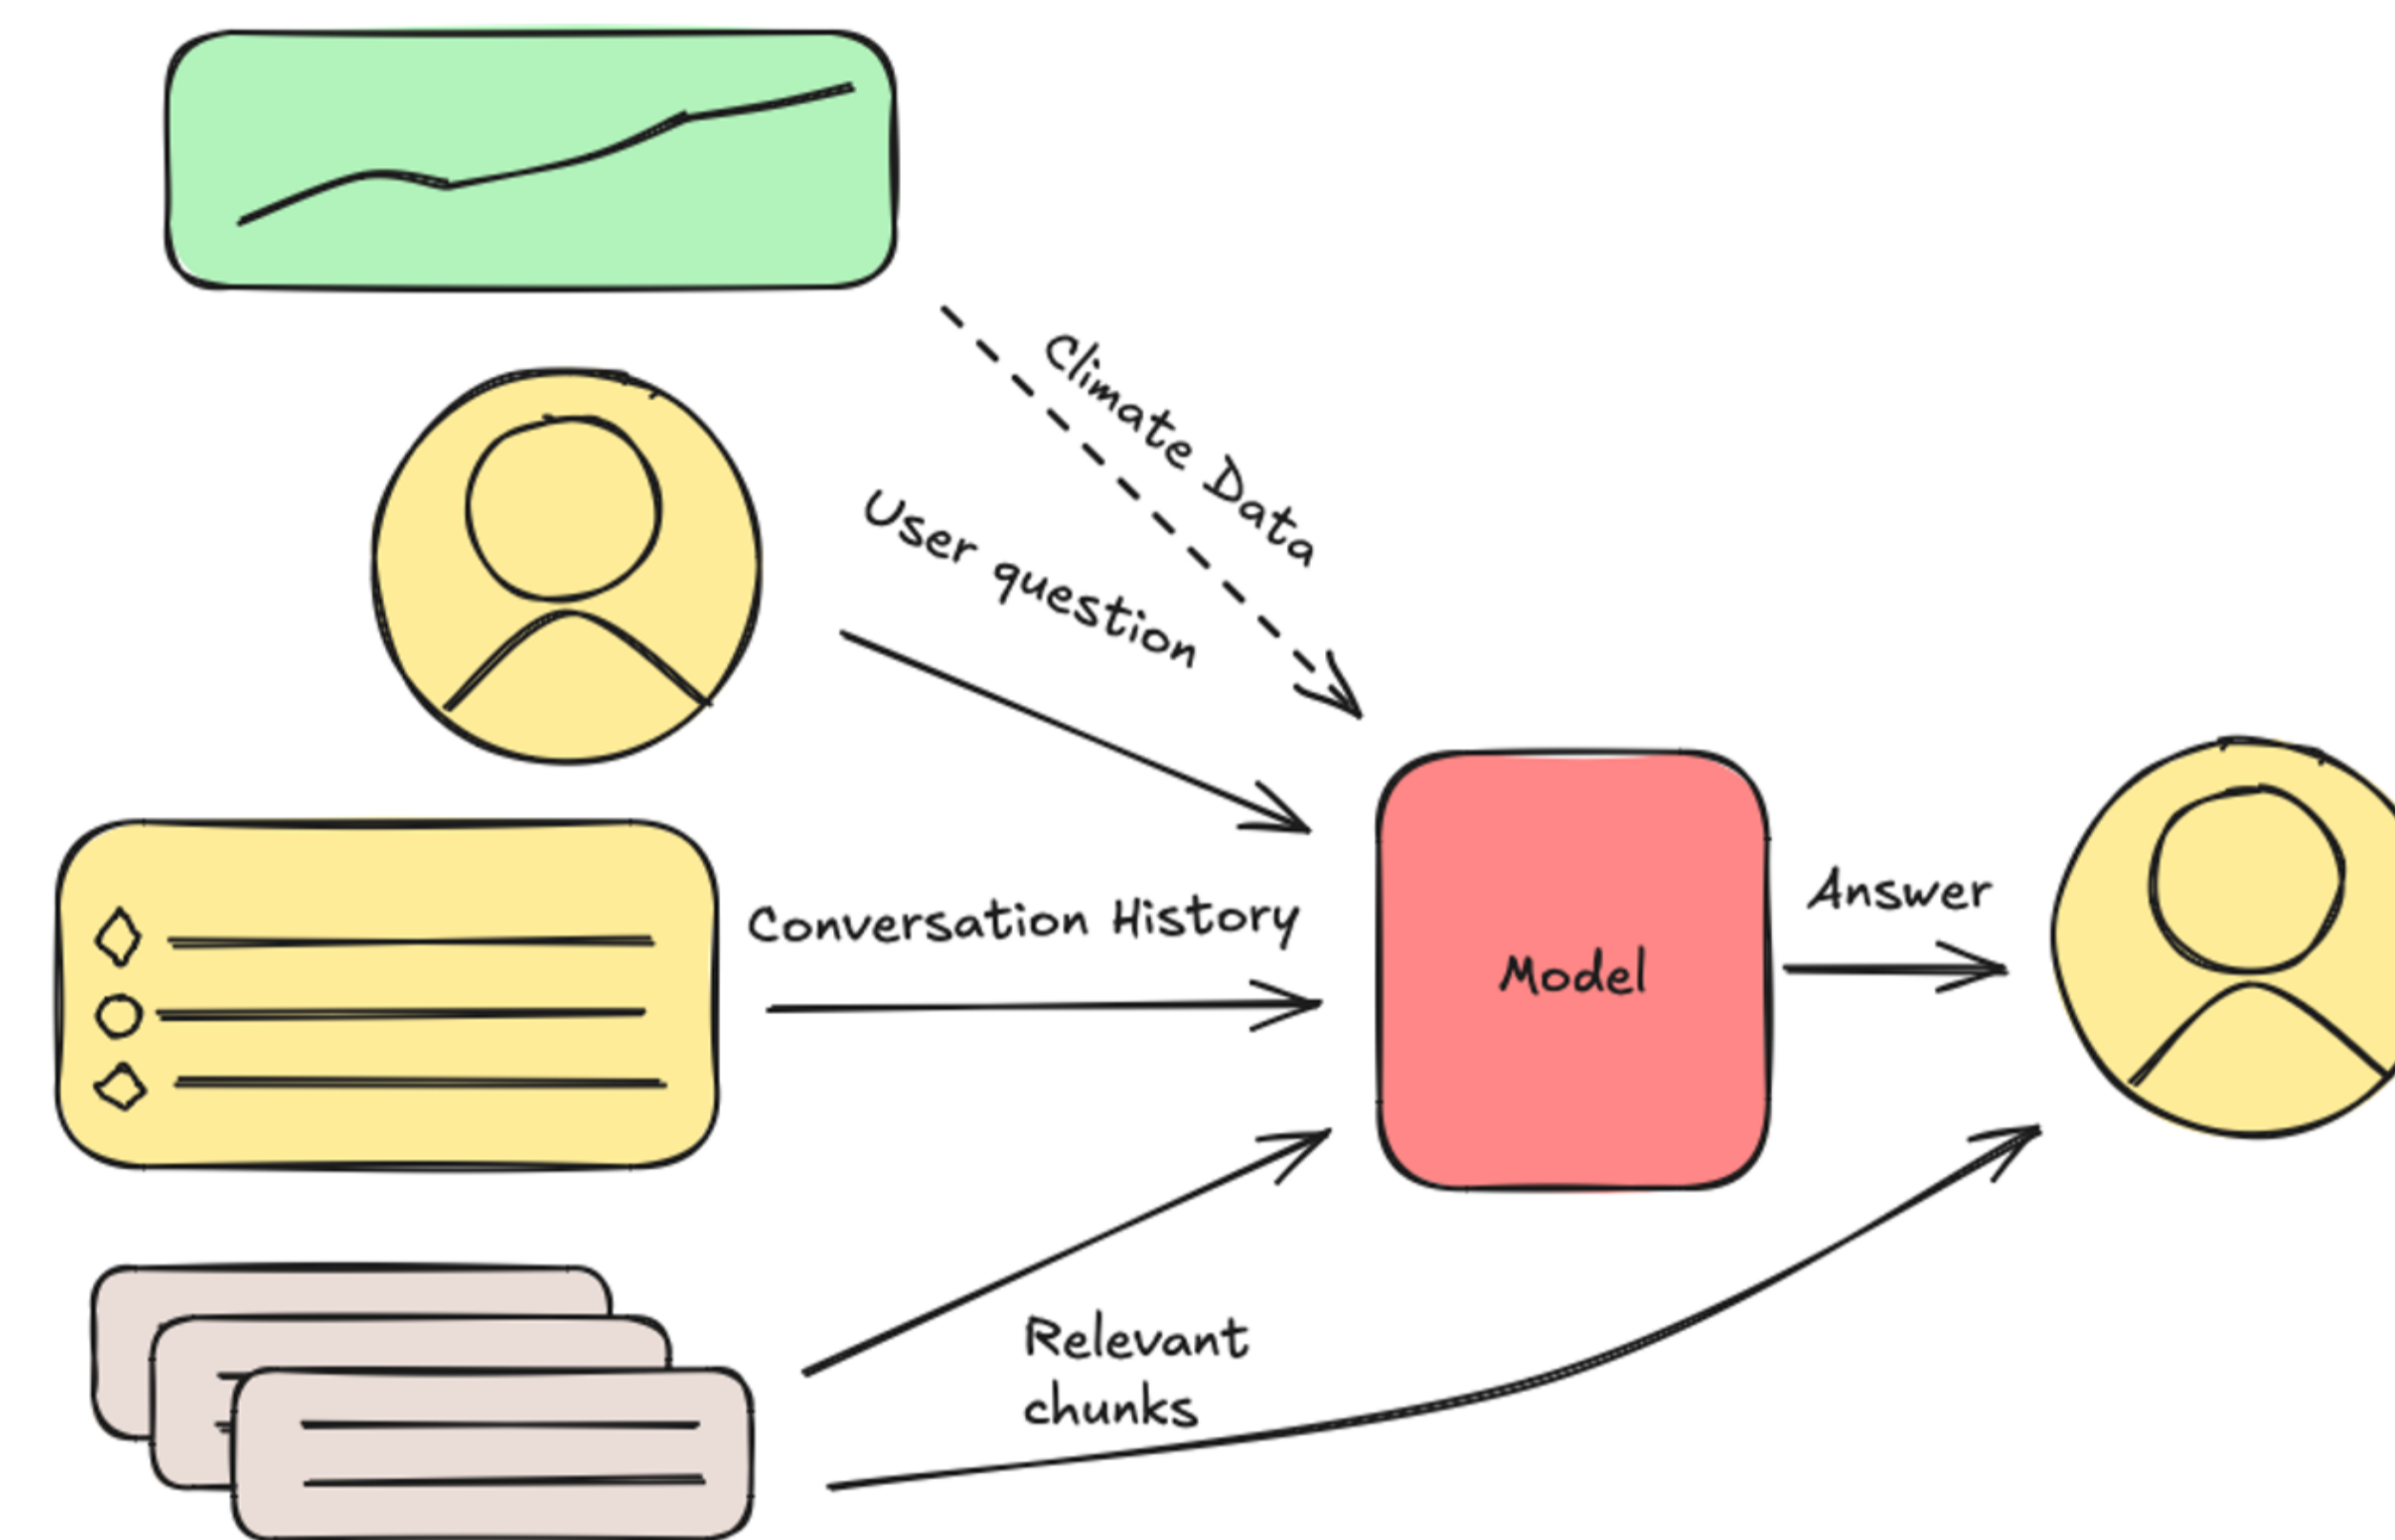
Systems Architecture



When indexing of the data sources, we store both text and vectors in Elasticsearch.



To ensure context is considered, the conversation history in addition to the question is used to draft the queries for literature retrieval.



We use climate data relevant to the user's commodity and location, alongside trustworthy literature to help answer the user's question.

Evaluation

- 15 climate adaptation questions were provided by domain experts
- 12 model combinations were used to generate responses.
- 7 metrics were used to evaluate each response from the models.

Evaluation Criteria


Style

- **Context:** Does the LLM provide enough background information to understand its response?
- **Readability/Structure:** Is the response of the LLM easy to read?
- **Specificity:** Is the information in the response relevant? For instance, to location, time and commodity in question?
- **Comprehensiveness:** Does the LLM respond with a complete answer?
- **Language:** Does the LLM use industry terminology fluently?

Scientific robustness

- **Scientific accuracy:** Is the information correct, given the source material?
- **Provenance/Citation:** Does the LLM provide relevant citations to its answers?

Results



Evaluation of proprietary and open-source models for climate adaptation questions.

Key findings

- **Competitive against proprietary offerings:** Performs well compared to GPT-4 and Claude.
- **Performance:** Falls behind flagship models but outperforms public offerings like ChatGPT and Gemini.
- **Stylistic Preferences:** While answers are accurate, presentation could be improved.
- **Ambiguous Location Names:**
 - Will my rainfall continue to increase in variability in Northern NSW?
 - How will climate change impact cherry production in Young?
- **Hallucinations and Safeguards:** Issues can occur in proprietary offerings.
 - What varieties of apples are more tolerant to sunburn?
 - Gemini 1.0 pro: "Apples do not get sunburn."
 - GPT-4: The response was filtered due to the prompt triggering Azure OpenAI's content management policy. Please modify your prompt and retry.

Future Work

- Evaluating on a larger sample of questions, and with a more diverse group of annotators.
- Evaluating the capability of climate data and scientific literature in other domains such as health and energy.

Acknowledgments

Special thanks to the CSIRO Climate services for agriculture team and the language technology team for their assistance in the work.