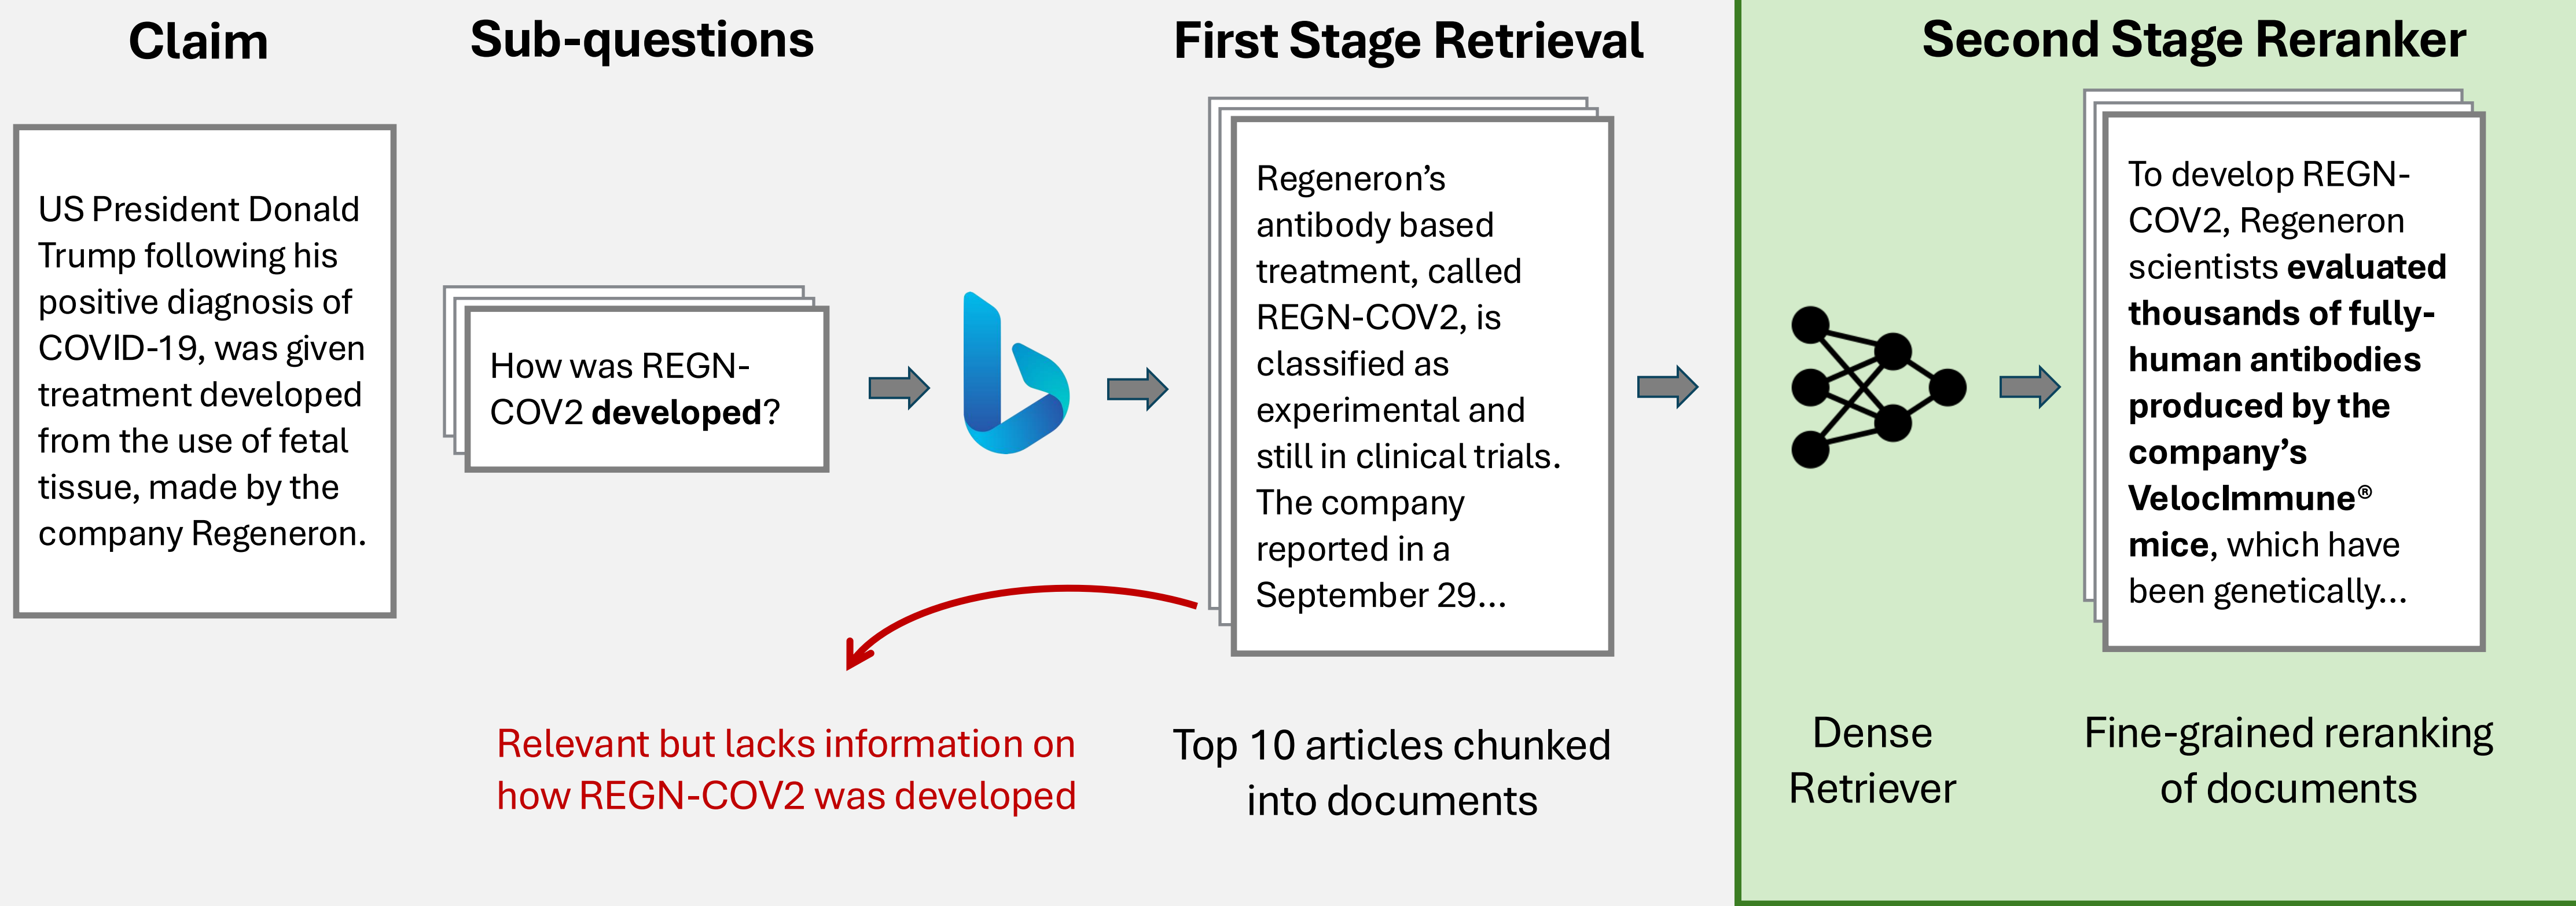




Task: Retrieve documents to fact-check claims

Relevant documents should address questions in-context with the claim



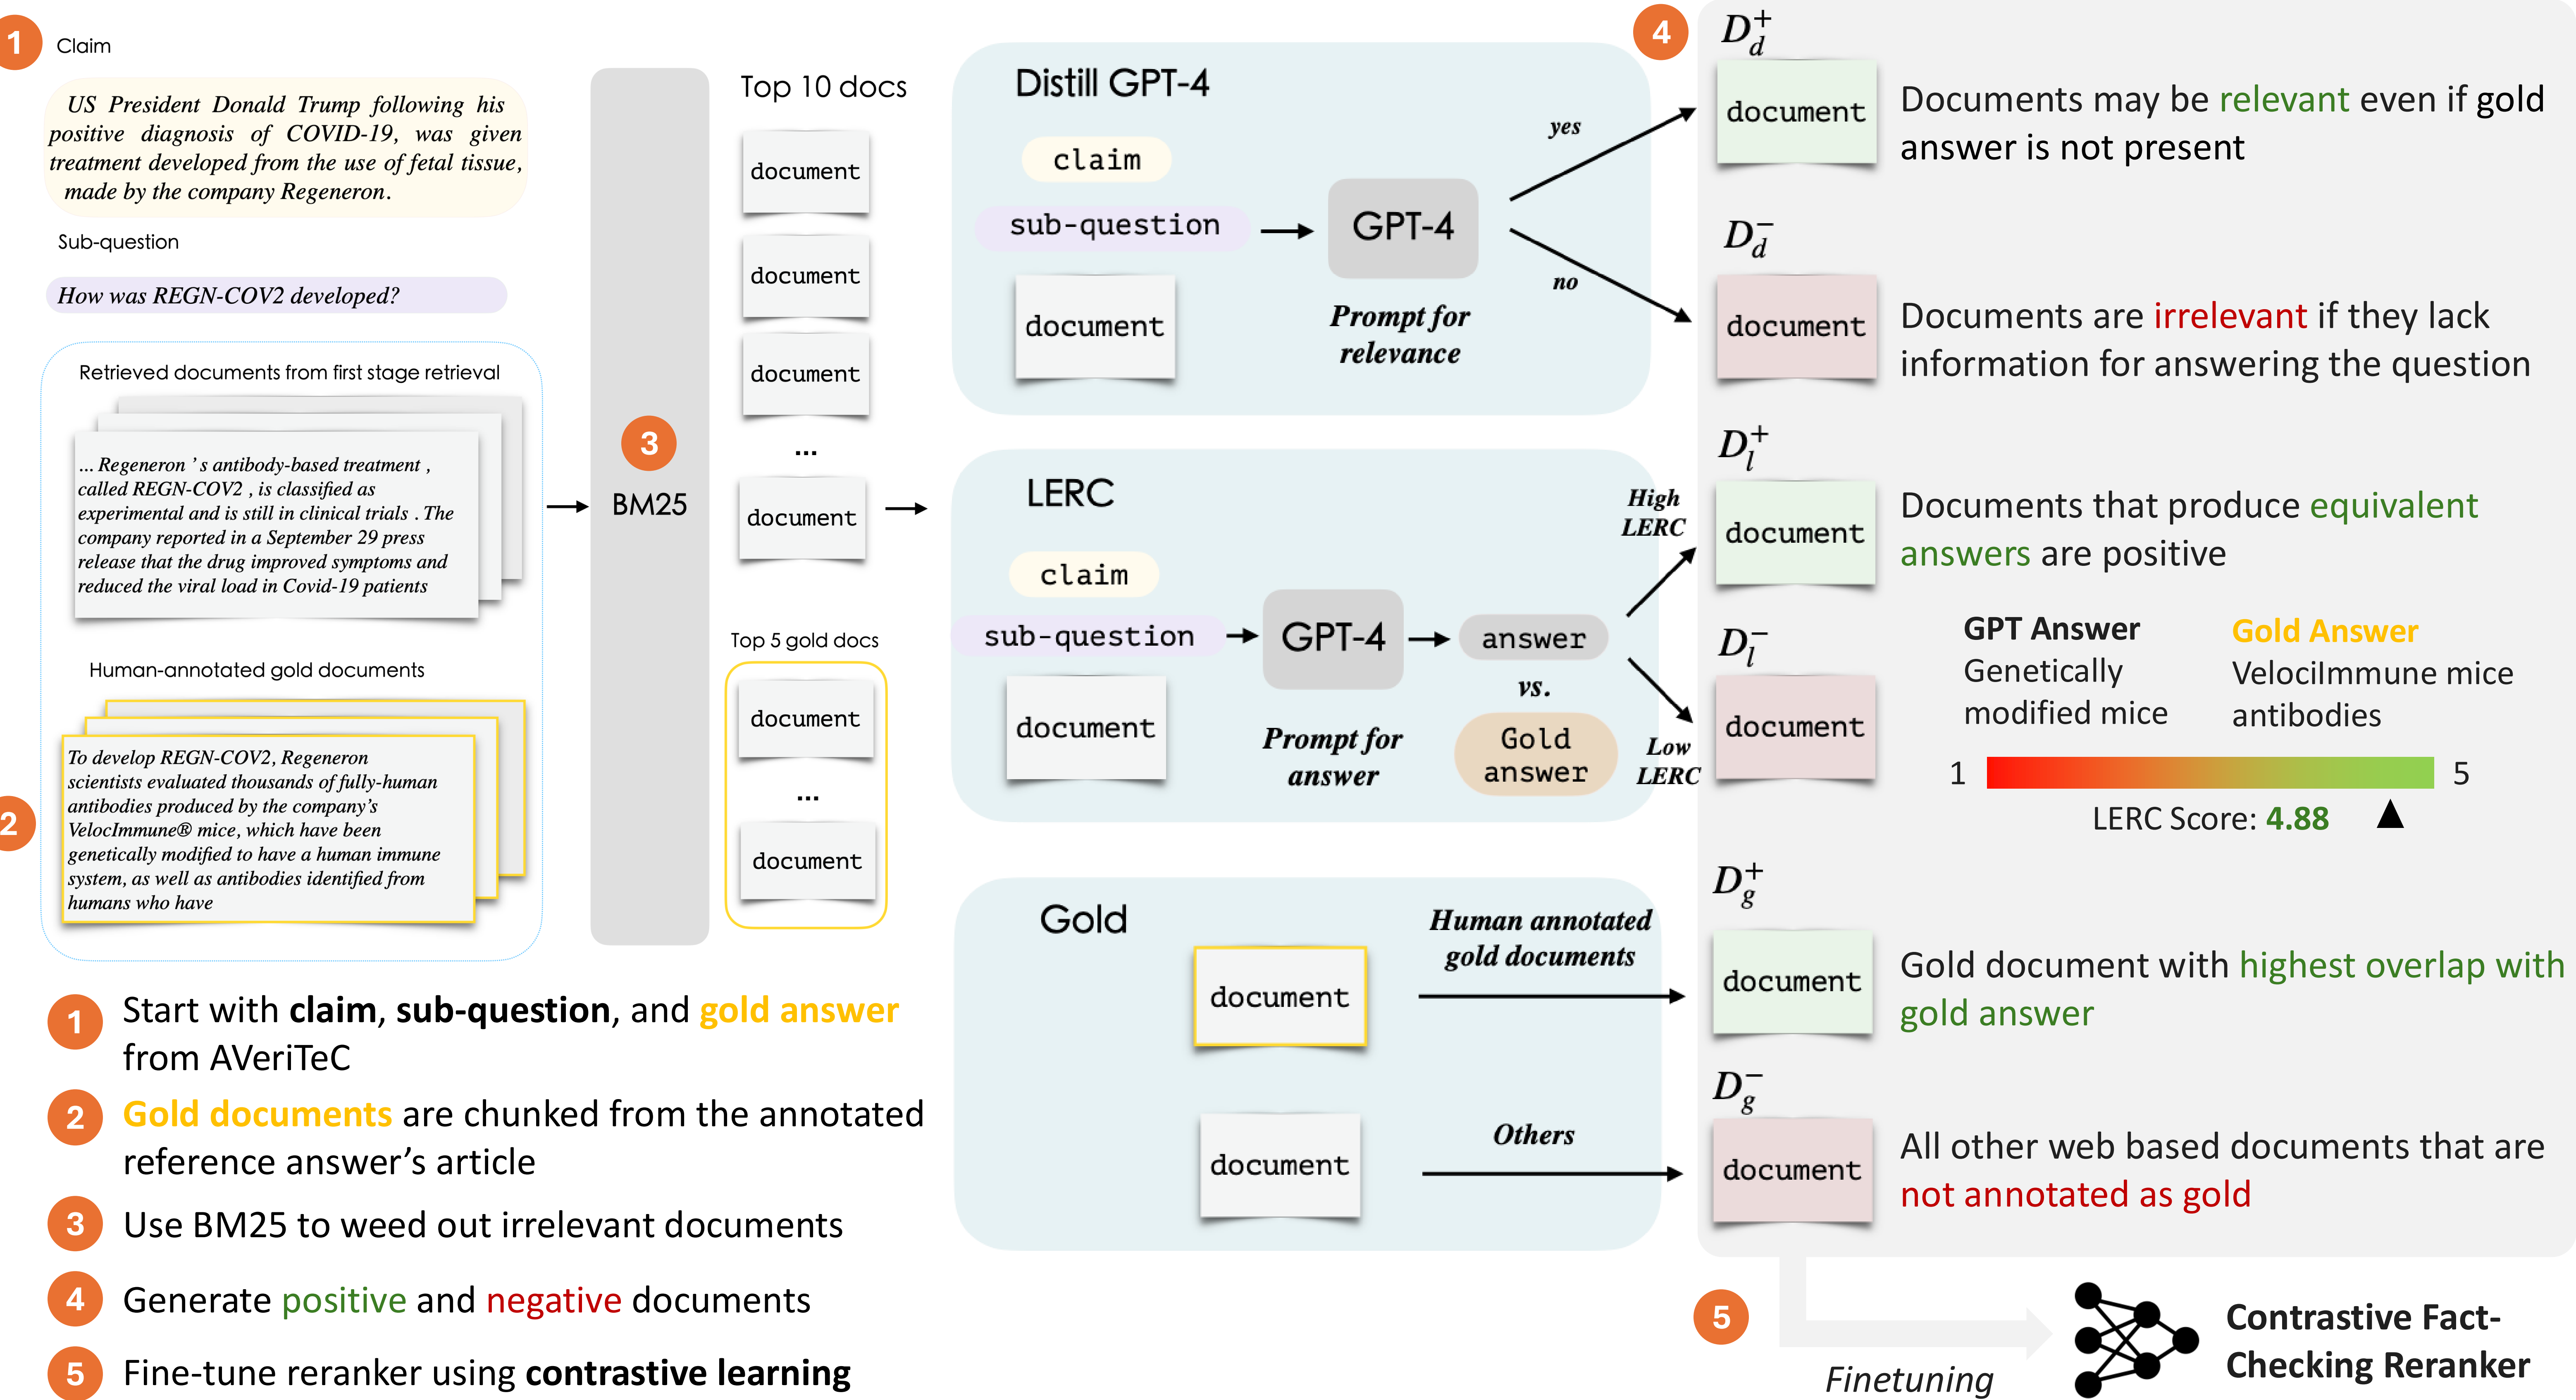
Our Contribution

Novel supervision signals for fine-tuning dense retrievers for complex fact-checking

Why is retrieval for fact-checking hard?

- Claims are often only **subtly true or false**
- Must **infer context** from the claim
- Documents may only obliquely address a false claim and **lack token overlap** with the question

Method: Generate positive and negative documents for contrastive training



Results: Do we see improvement in downstream veracity classification? **Yes!**

