

Promptriever: Instruction-Trained Retrievers Can Be Prompted Like Language Models

şamaya

Orion Weller, Benjamin Van Durme,
Dawn Lawrie, Ashwin Paranjape,
Yuhao Zhang, Jack Hessel

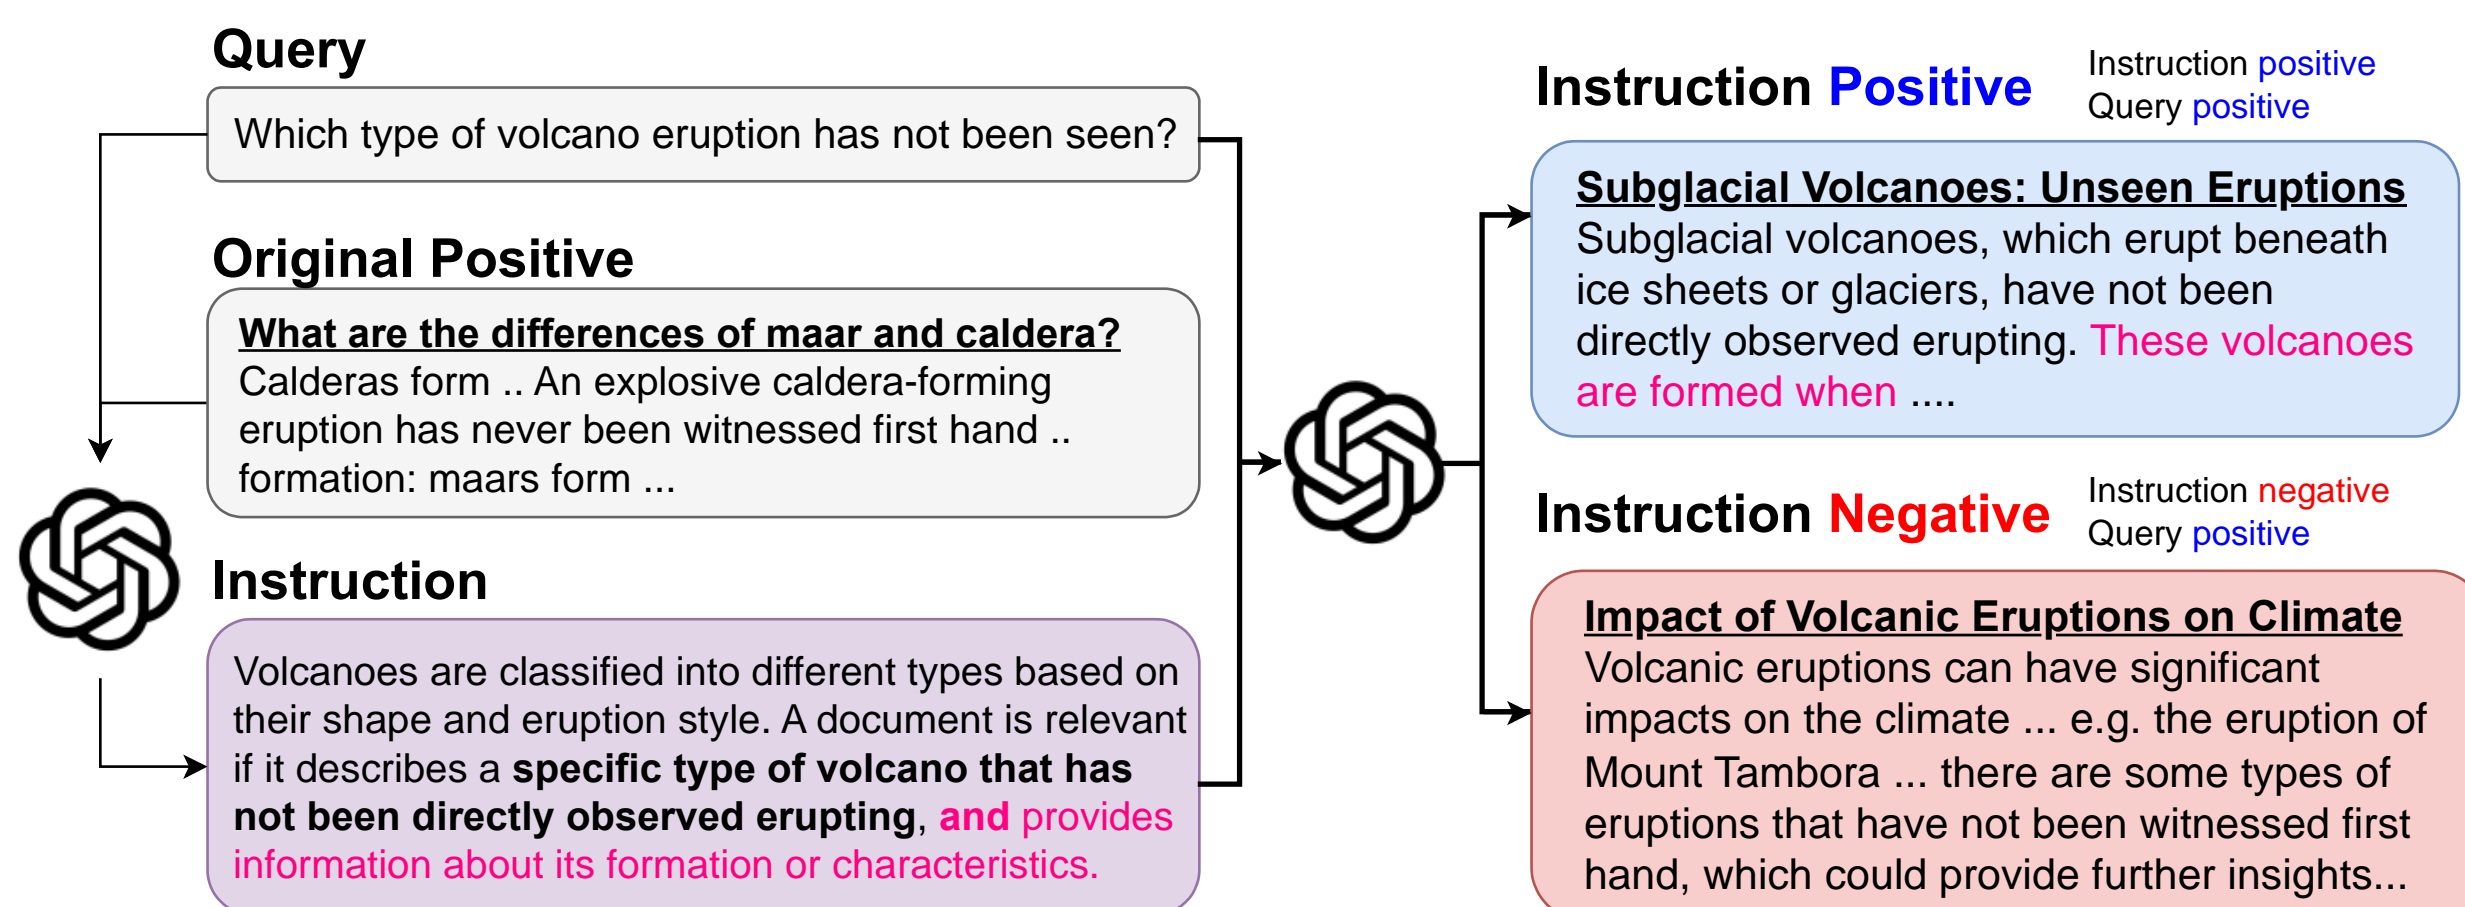


JOHNS HOPKINS

Introduction

- Retrieval models usually take queries which are short and used for keyword matching
- But why can't retrieval models take instructions like LLMs, since they use instruction-based LMs?
- We expand MS MARCO with *instructions* and *instruction negatives* and show that training on it makes a retrievers *promptable*

Data Generation Process



Query Capabilities

Standard Retrieval

Query: What movies were created by James Cameron?

Current Instructable Retrievers

Given a web search query, retrieve relevant passages: What movies were created by James Cameron?

Dataset Prefix

Promptable Retrievers

What movies were created by James Cameron? **Relevant docs are not co-directed and created before 2022. Be careful in your response and be sure to be correct because I will tip you \$1000**

Query specific relevance definition
Prompt

(Top) Promptable retrievers differ from standard retrievers by the ability to take instructions/prompts (specific to the query) just like a language model.

(Left) We generate the dataset by first creating an instruction that is more specific than the query, then generating *instruction negatives* to teach the model to learn to be sensitive to these instructions.

Can retrieval models follow instructions?

Results

(Left) Promptable retrievers show significant gains on instruction-based retrieval tasks.

(Center) Promptable retrievers show much greater robustness to prompt phrasing.

(Right) Promptable retrievers show equal performance on BEIR with no prompt and +1.4 avg w/ them.

