

Answering Count Queries with Explanatory Evidence

Shrestha Ghosh, Simon Razniewski, Gerhard Weikum

Aggregation of context and instance evidence over multiple texts increases answer coverage and user comprehension



How many songs did John Lennon write for the Beatles?



- > type: song
- > composed by: John Lennon
- > performed by: **THE BEATLES**

Instances counted have certain desired properties

Songs/The Beatles/John Lennon → 51
 (?x, composed by, John Lennon) &
 (?x, performed by, the Beatles) → 168
<https://w.wiki/5E4e>

Incomplete instances, inadequate semantic qualifiers → **wrong KB aggregates**

.. approximately 180 jointly ..
 .. 160 Lennon/McCartney songs ..
 .. 10 No. 1 Beatles songs ..

Counts are scattered over search engine result snippets

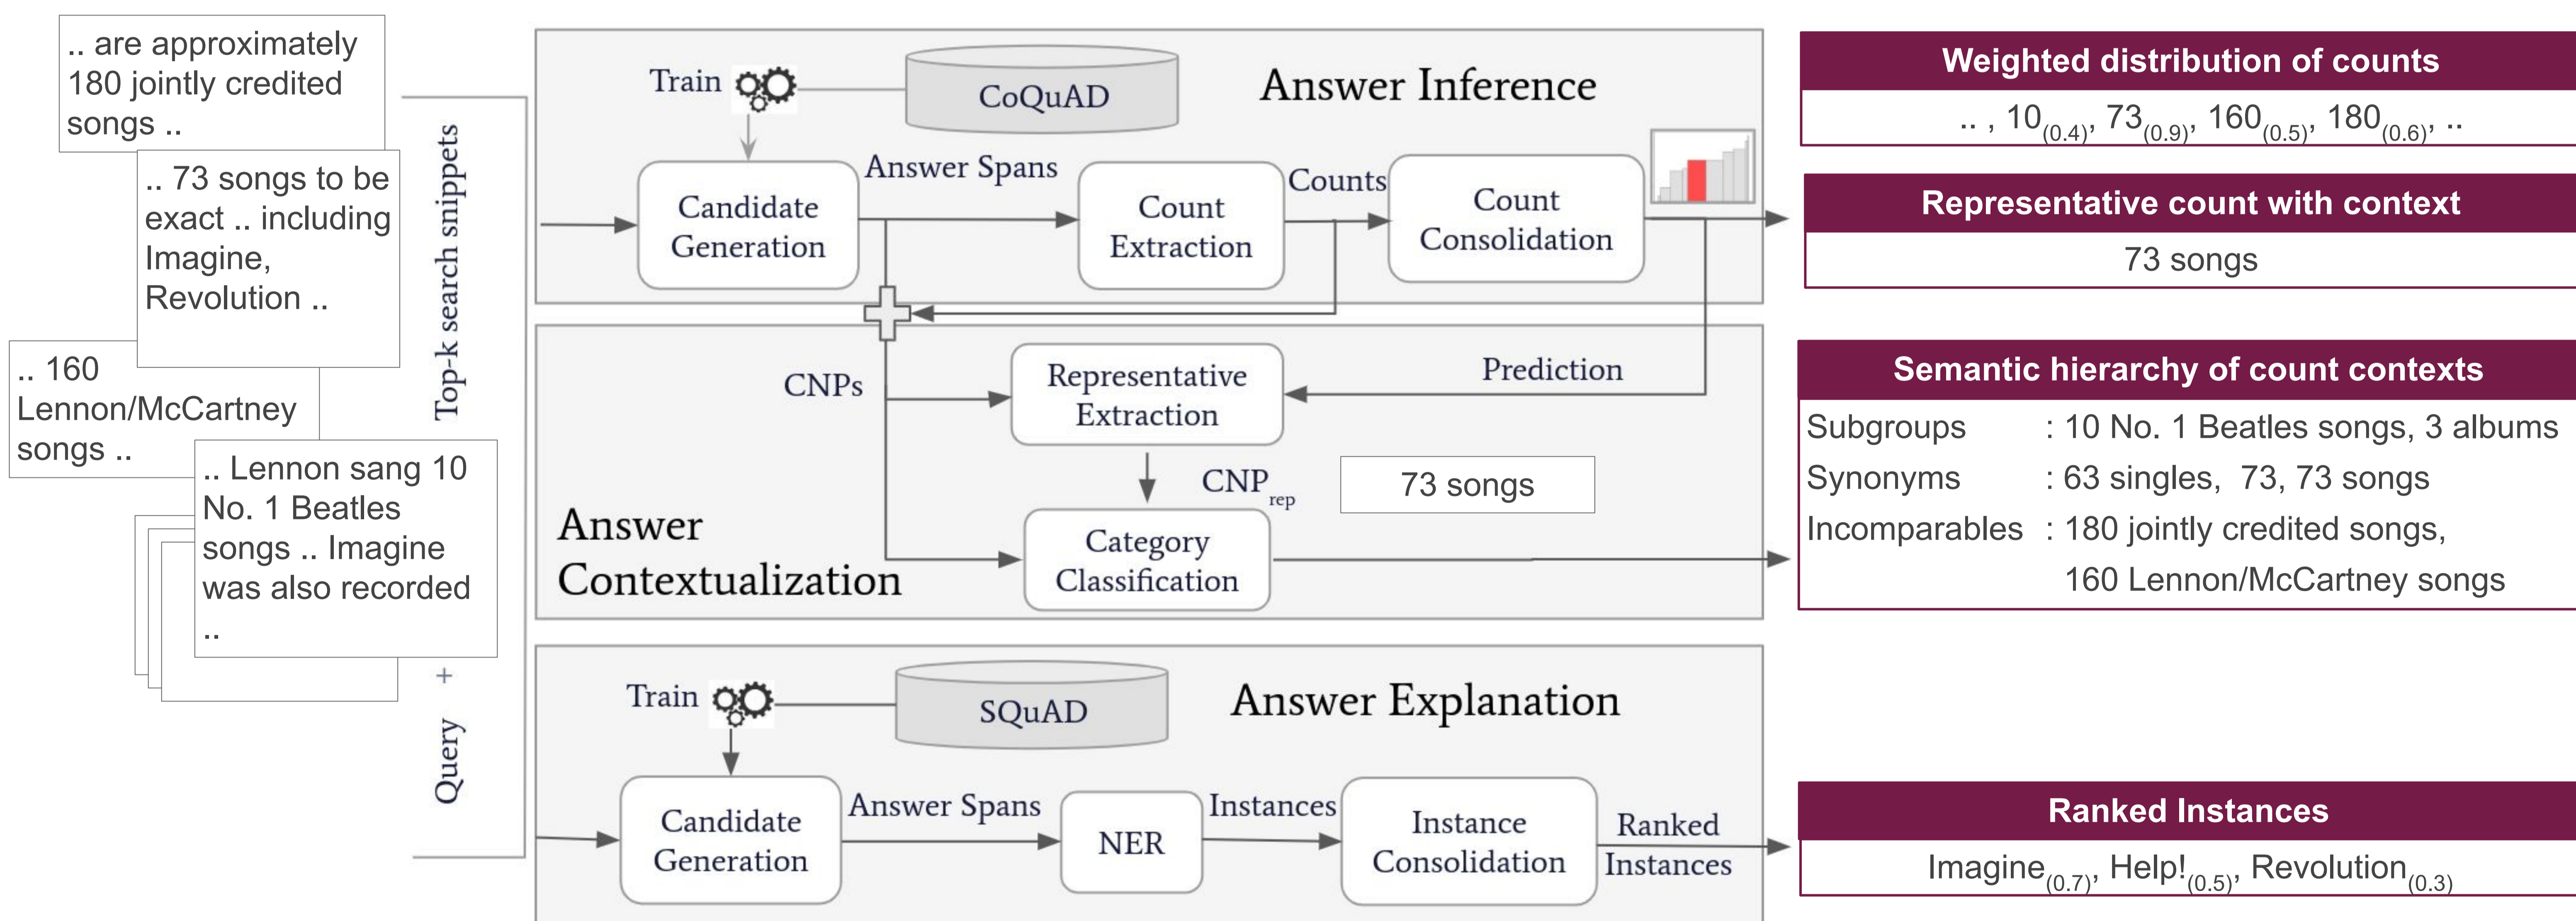
CoQEx System

Main Components

- > **Infers a representative count** from the numeric distribution over snippets
- > **Interprets the count context** by ordering contexts into a semantic hierarchy
- > **Explains through instances** by ranking instances from the snippets

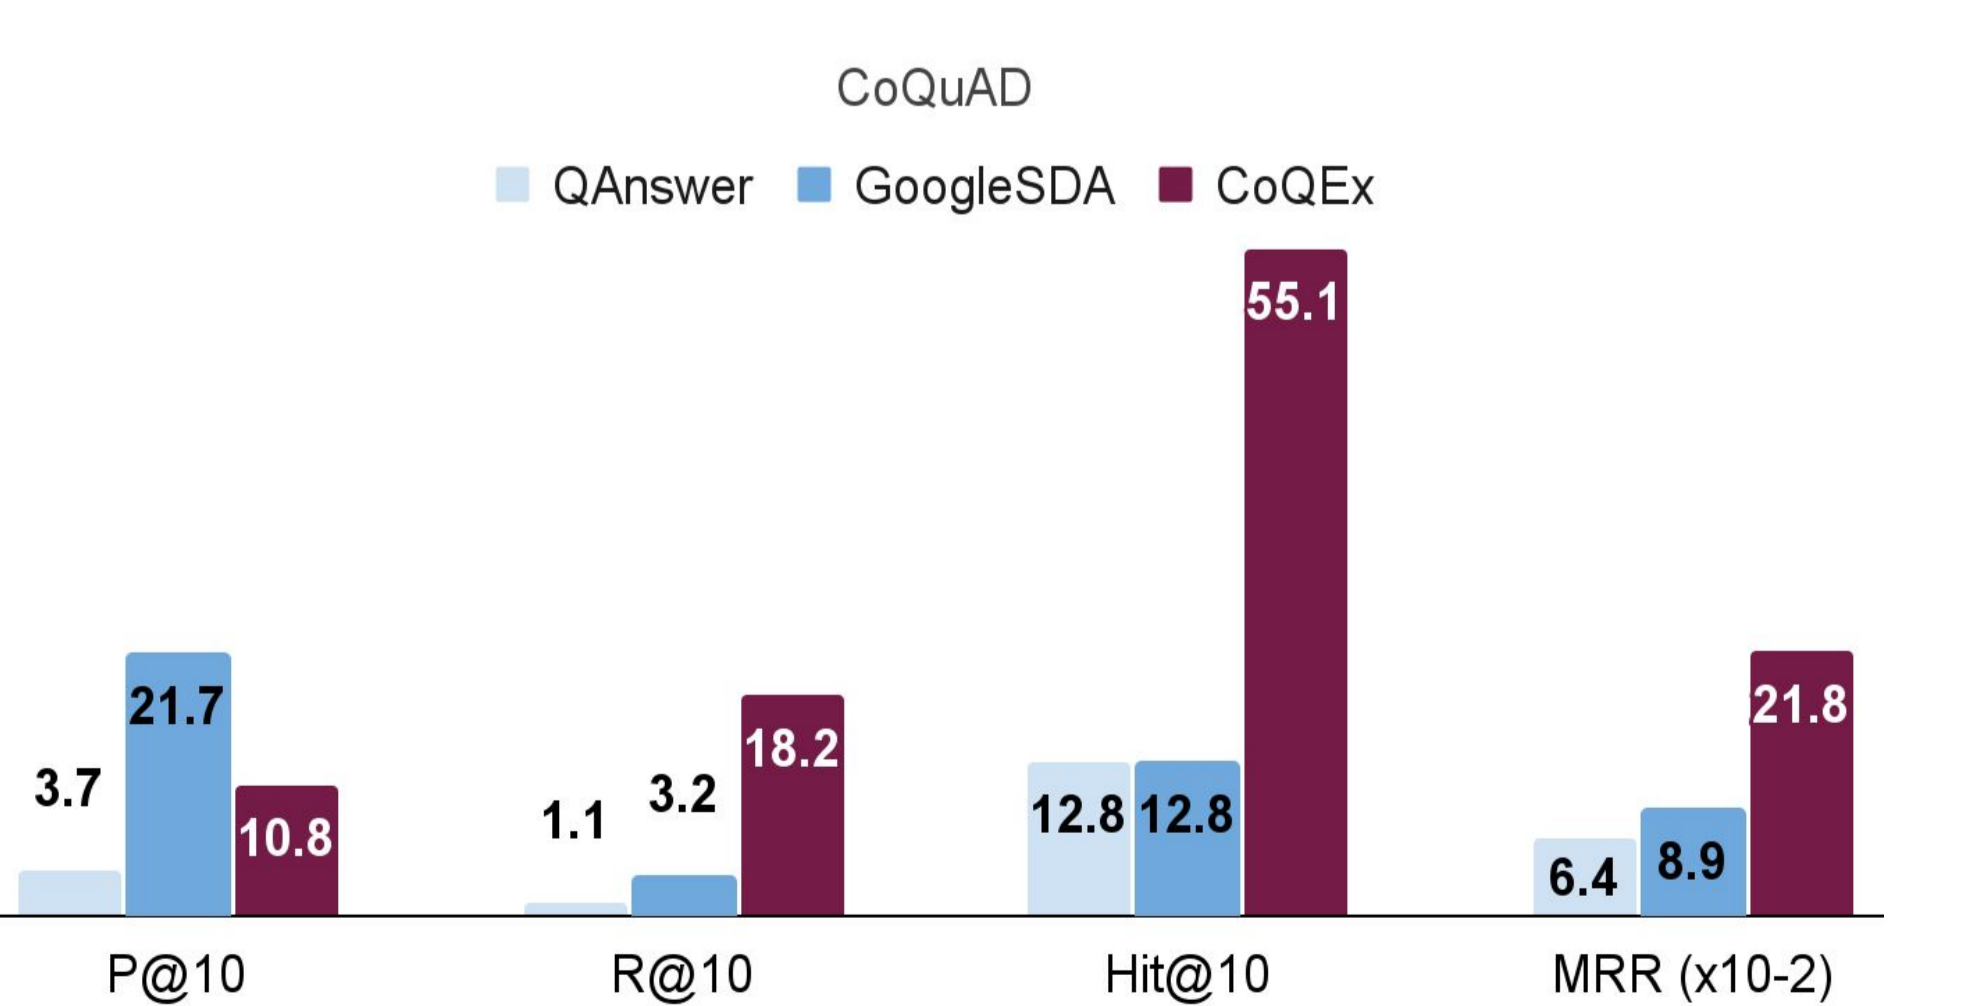
Advantages of CoQEx

- > Aggregates over multiple text snippets
- > Contexts as Synonyms or Subgroups
- > Identifies instances counted in the query

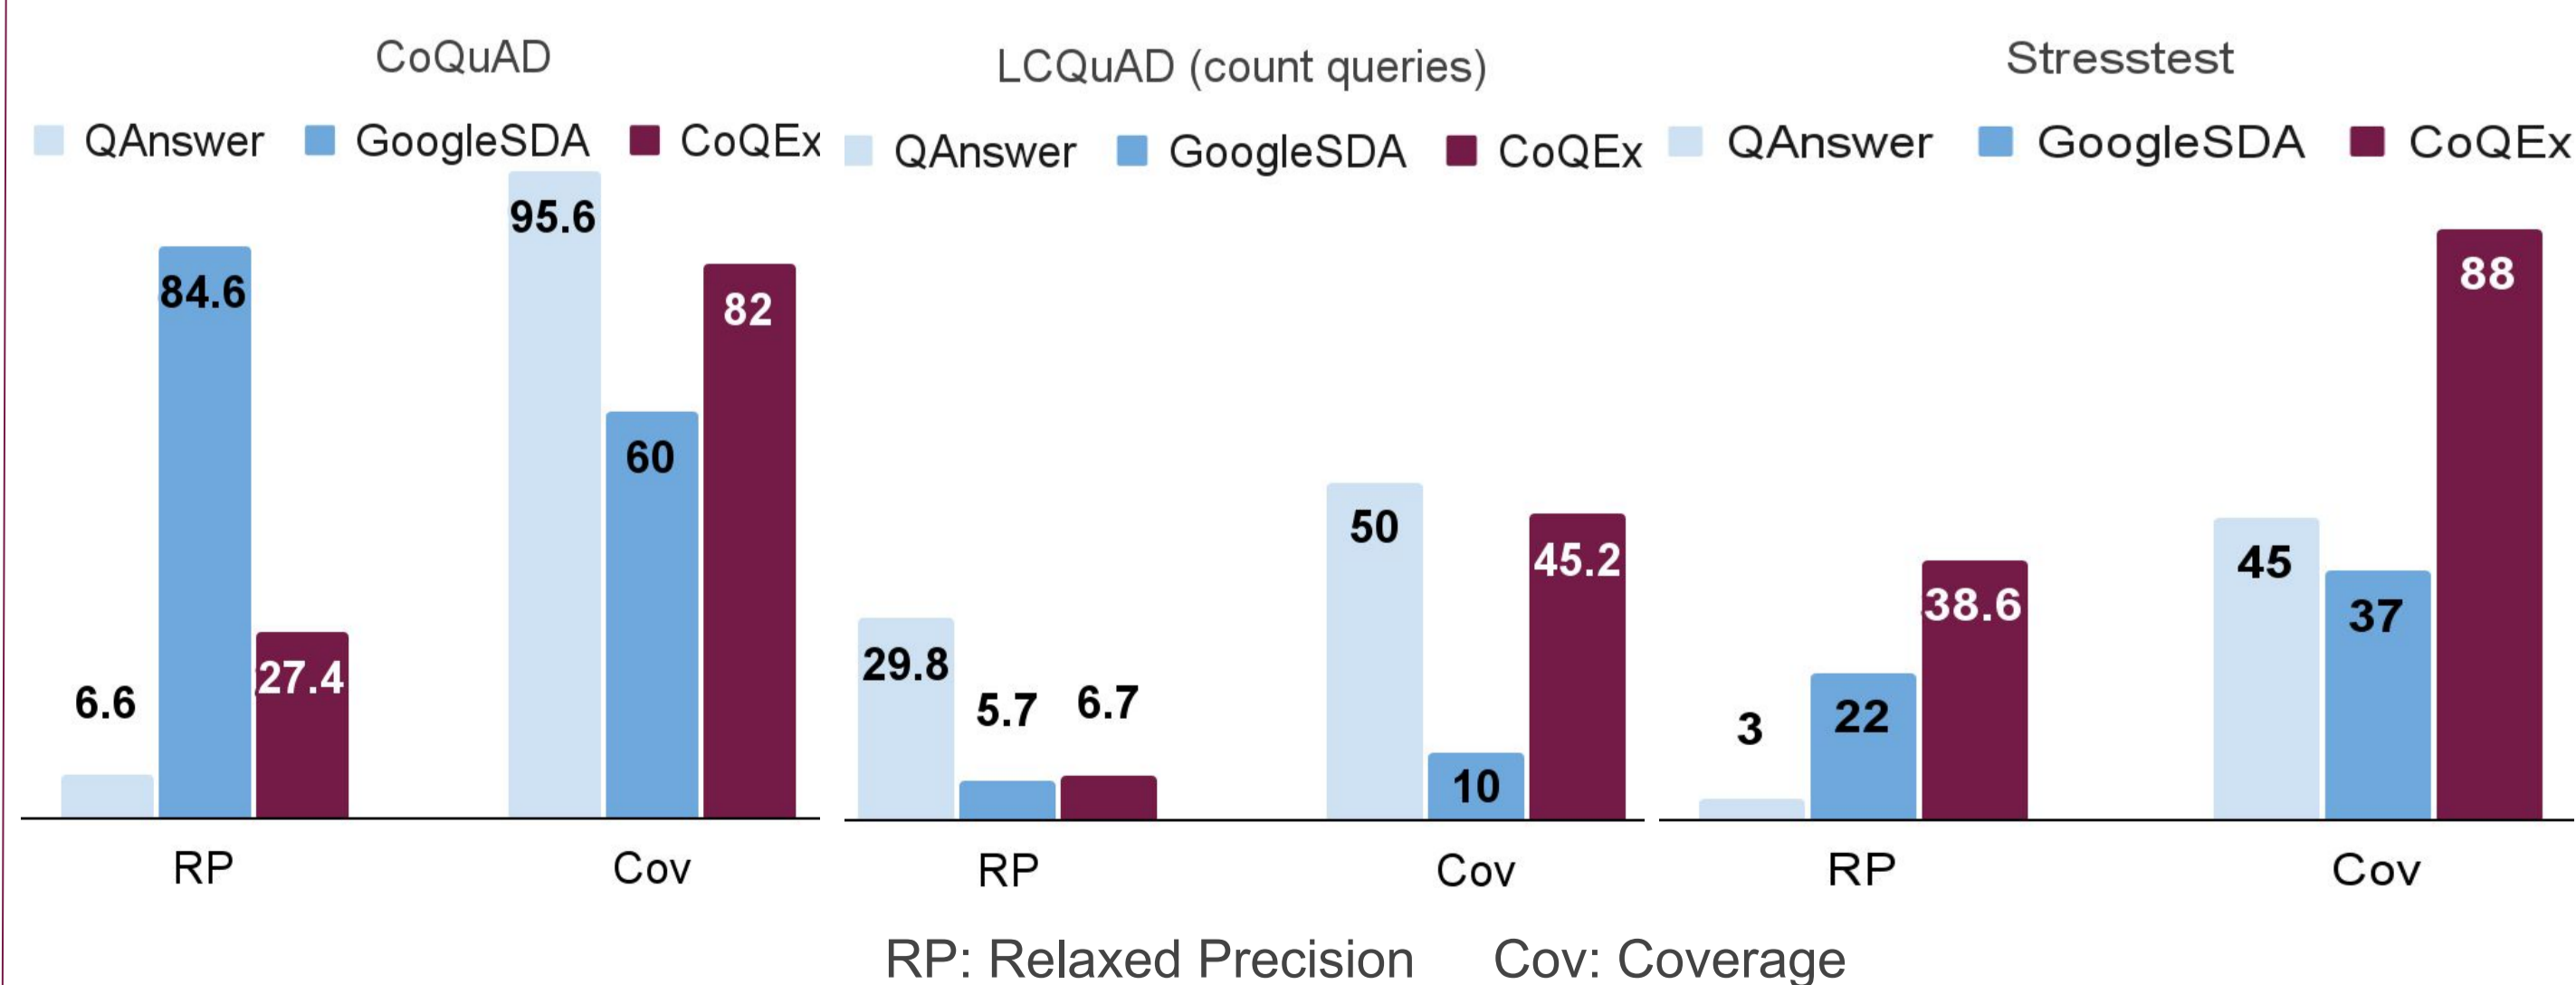


Results

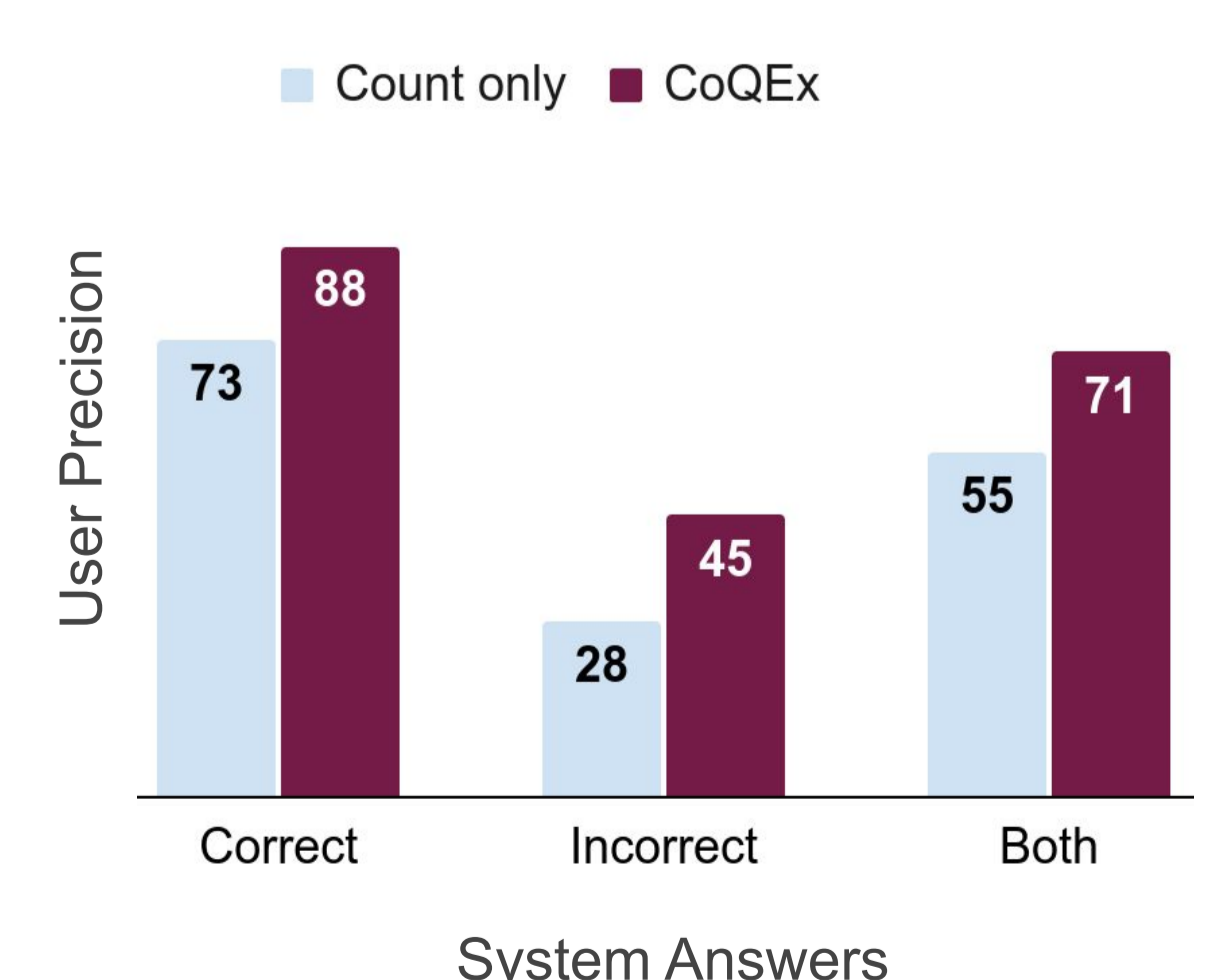
Instances Ranking



Count Inference



User Comprehension



Contact ghoshs@mpi-inf.mpg.de

