Competency Questions for ontologies

C. Maria Keet

University of Cape Town, South Africa, mkeet@cs.uct.ac.za

Educational Series on Applied Ontology (ESAO)
December 17, 2024, online

Outline

Outline

CQs in ontology development – Examples

- What is assessed in the walking task?
- What level of expertise is required to use TexShop?

CQs in ontology development – Examples

- What is assessed in the walking task?
- What level of expertise is required to use TexShop?
- Is [this animal] a herbivore?
- Which animals eat [these animals]?
  - 'templated' CQs
  - Assumes subclasses of *animal* to instantiate the CQ;
  - e.g.:
    Is an elephant a herbivore?
    Is a lion a herbivore?
    Which animals eat impalas?
    Which animals eat mice?

CQs in ontology development – Usage

- To describe and demarcate the scope and requirements
- Queries for obtaining information from the ontology (e.g., to find an ontology for reuse)
- Validation of ontologies on coverage
- Broadly: serving purposes in knowledge acquisition, organisation, and validation for ontologies

Usage in TDD – example [Davies et al.(2019)]

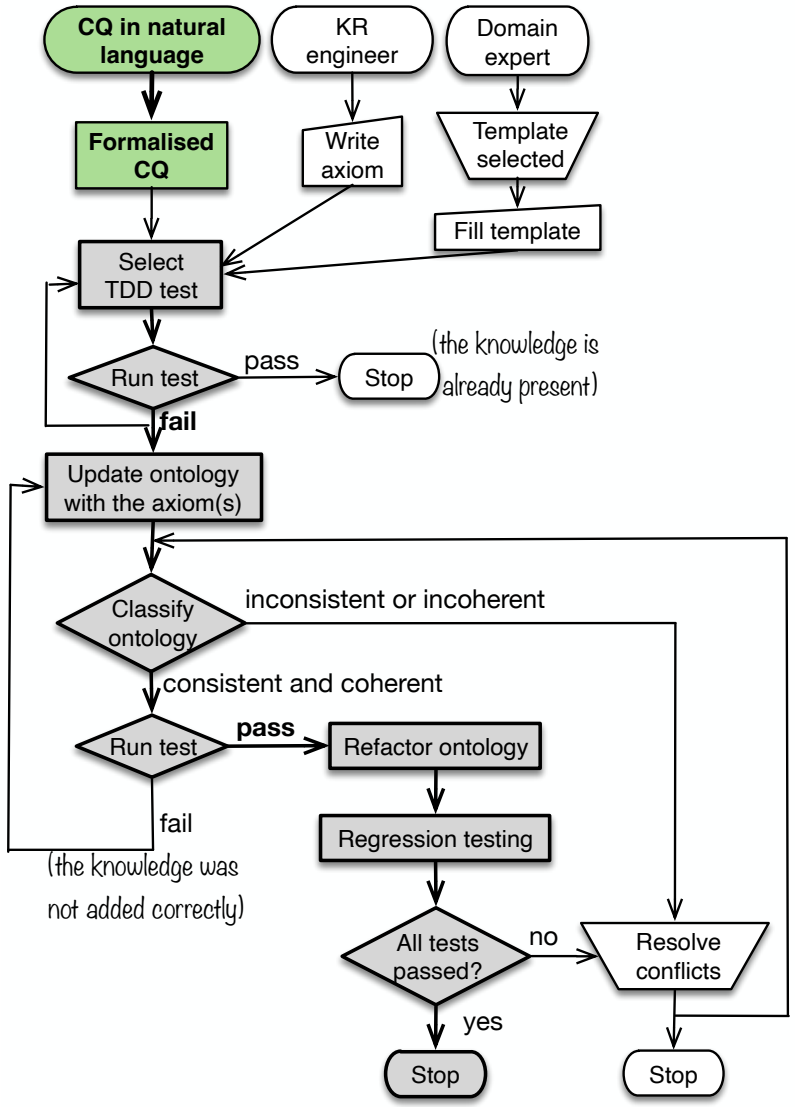

Usage in TDD – naive illustration and reality check

Ontology authoring

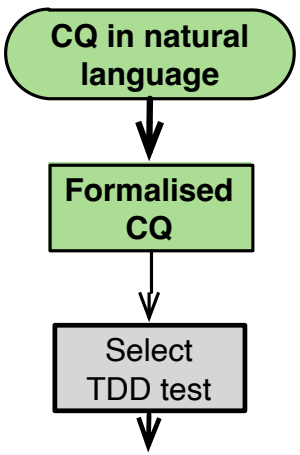

What is assessed in the walking task?

$\top \sqsubseteq \exists assessedIn.WalkingTask$

CQ from dem@care;

Usage in TDD – naive illustration and reality check

Ontology authoring

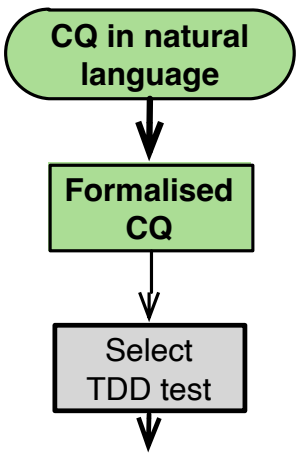

What is assessed in the walking task?

~~$\top \sqsubseteq \exists assessedIn.WalkingTask$~~

$WalkingTask \sqsubseteq \exists measuredData.measuredData$

CQ from dem@care;

Usage in TDD – naive illustration

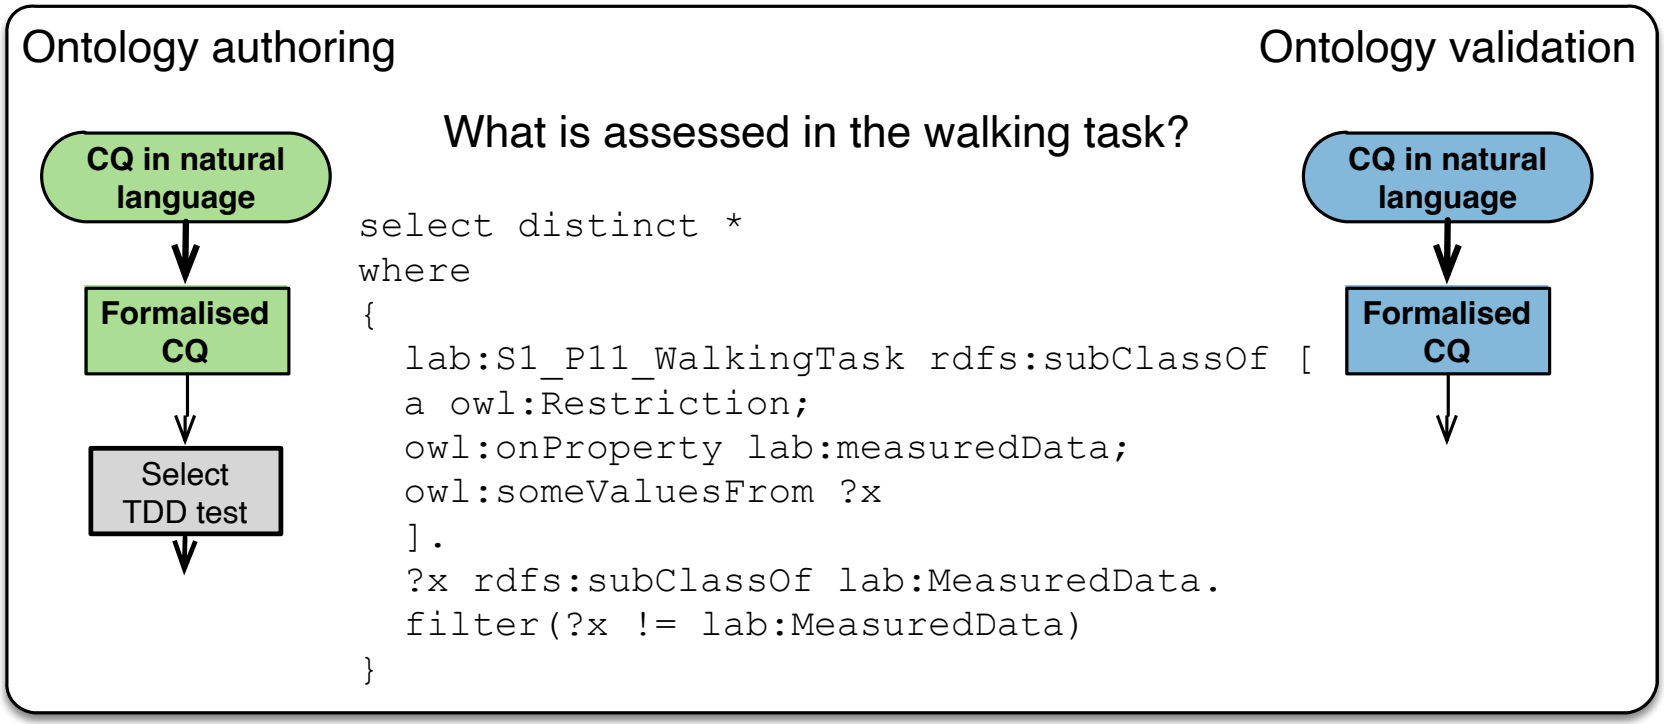

CQ from dem@care; SPARQL-OWL query by [Potoniec et al.(2020)]

It still might look obvious, but...

1. Who has to write those CQs?
2. What about assistance in authoring CQs?
3. How to translate CQ into OWL to develop the ontology?
4. How to translate CQ into SPARQL-OWL to query the ontology?
5. What are good CQs for ontology development?
6. Which CQs are bad or wrong and why?
7. Does that hold for all tasks in ontology development?
8. What is a competency question?
9. (What is a question?)

Outline

Manual guidance [Keet et al.(2019)]

- Avail of a novel dataset of CQs that had been analysed into 'patterns' [Wiśniewski et al.(2019)]
- Analyse those patterns
- Specify the CNL
- Evaluate on coverage
- Proof-of-concept tool (which then may feed into a next cycle)

Those patterns

- Wisniewski' et al's [Wiśniewski et al.(2019)] patterns from 234 type-level CQs from 5 ontologies; linguistic analysis with chunking
  Which country do I have to visit to see elephants?
  Which EC1 PC1 I PC1 PC2 EC2?

Those patterns

- Wisniewski' et al's [Wiśniewski et al.(2019)] patterns from 234 type-level CQs from 5 ontologies; linguistic analysis with chunking
  Which country do I have to visit to see elephants?
  Which EC1 PC1 I PC1 PC2 EC2?
- Analyse & harmonise those 106 patterns: their shape (e.g., min and max ECs, number of slots), synonyms ('type'/'kind'), singular/plural, redundant words ("or not"), impersonal or not ("can we"), negation

Those patterns

- Wisniewski' et al's [Wiśniewski et al.(2019)] patterns from 234 type-level CQs from 5 ontologies; linguistic analysis with chunking
  Which country do I have to visit to see elephants?
  Which EC1 PC1 I PC1 PC2 EC2?
- Analyse & harmonise those 106 patterns: their shape (e.g., min and max ECs, number of slots), synonyms ('type'/'kind'), singular/plural, redundant words ("or not"), impersonal or not ("can we"), negation
- Total: 92 default templates and 40 variants

  Competency question **La**nguage for specifying **R**equirements for an **O**ntology, model, or specification (CLaRO)

CLaRO editor

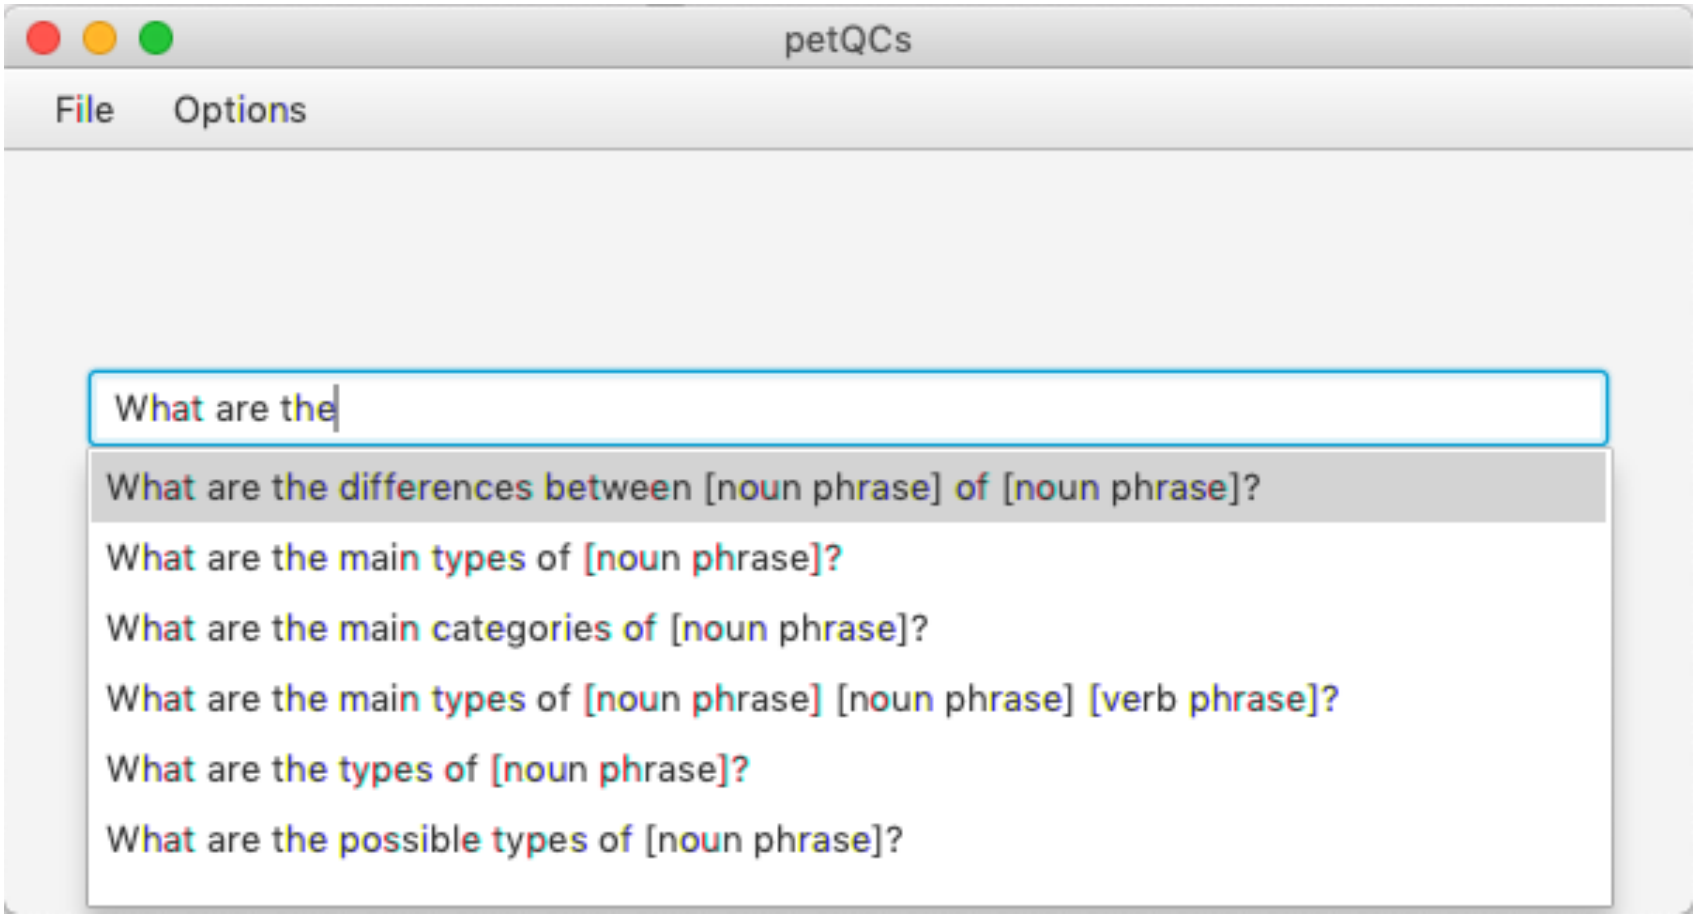

https://github.com/mkeet/CLaRO

CLaRO editor

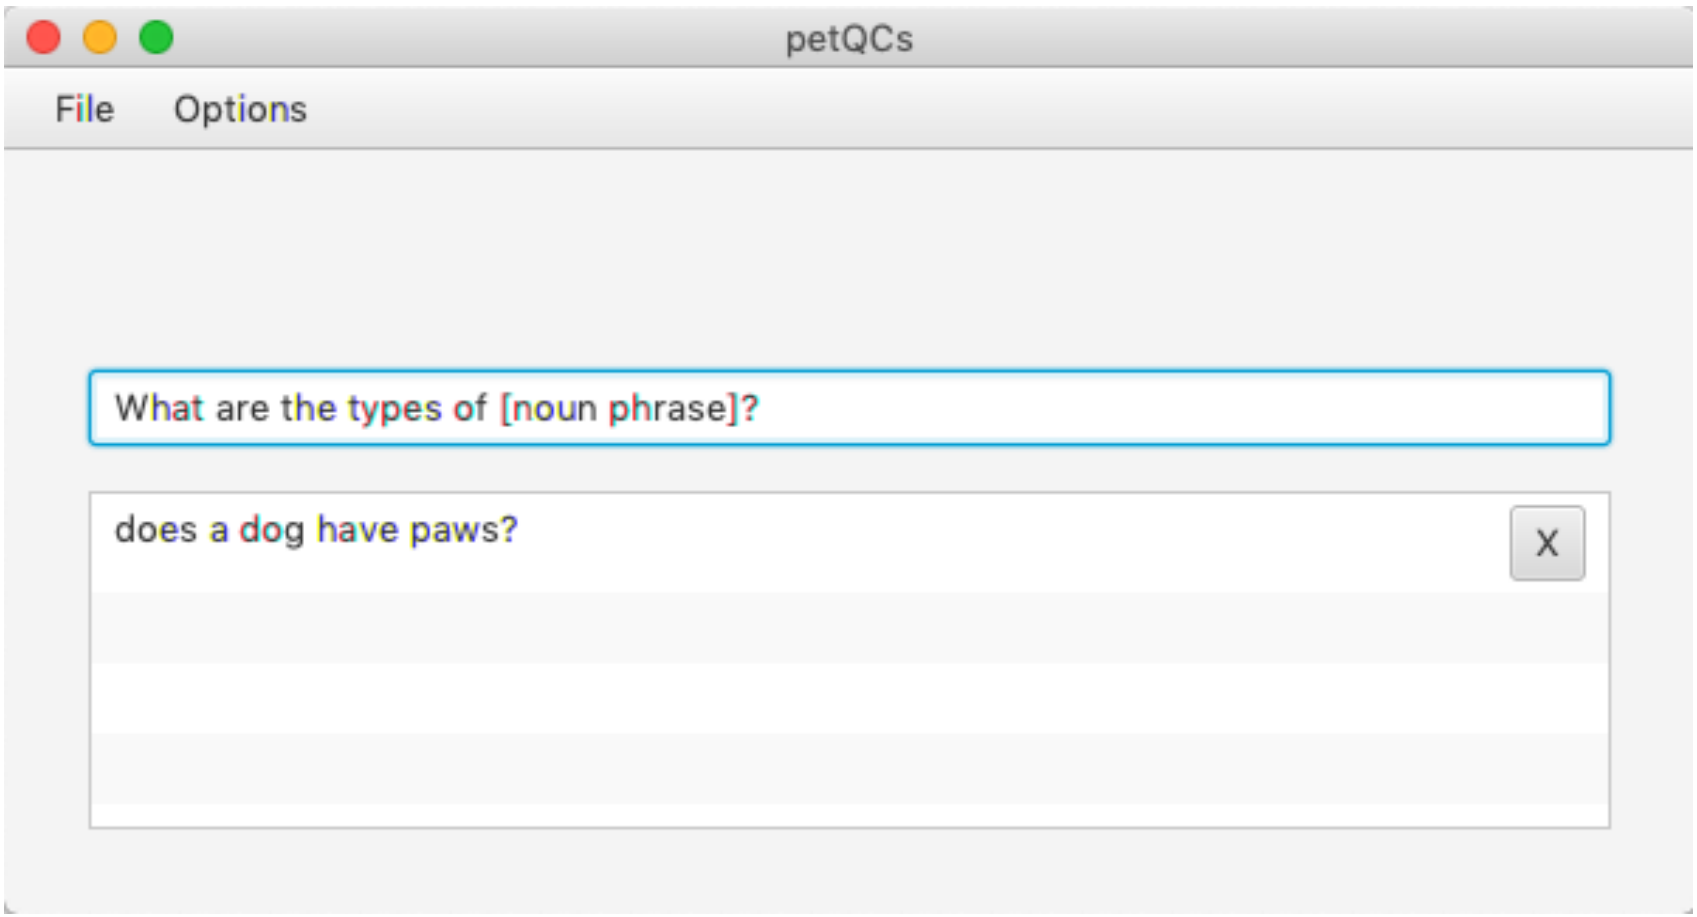

https://github.com/mkeet/CLaRO

CLaRO editor

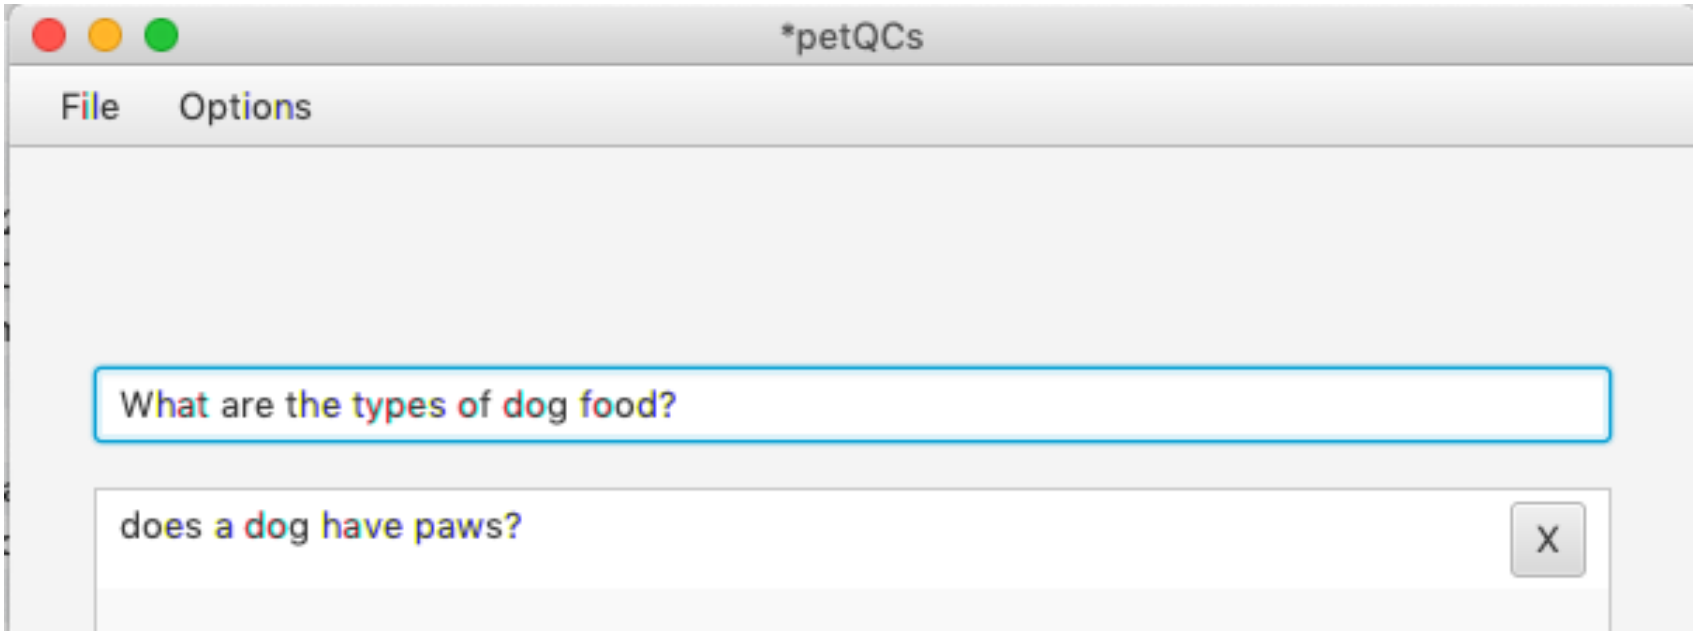

https://github.com/mkeet/CLaRO

Limitations

- There's assistance, but you still need to write something
- Data-driven: any errors in the CQ data set and in chunking propagates to the CNL
- Coverage ok-ish for unseen CQs [Keet et al.(2019)]
- Coverage improved by cleaning up the 234 CQs and training on more CQs [Antia and Keet(2021)]
- Coverage unlikely to ever be 100% even for English

Automating CQ authoring

- From scratch with a corpus-based approach and an LLM; e.g., [Antia and Keet(2023)]:
- Retrofitting to an existing ontology with an LLM-based approach; e.g., [Alharbi et al.(2024)]:
- See also the very recent review in [Alharbi et al.(2024)]

Automating CQ authoring

- From scratch with a corpus-based approach and an LLM; e.g., [Antia and Keet(2023)]:
  - create/get text corpus, determine important statements
  - convert into questions (T5)
  - filter questions with CLaRO to extract 'probably more suitable as CQ' ones
- Retrofitting to an existing ontology with an LLM-based approach; e.g., [Alharbi et al.(2024)]:
- See also the very recent review in [Alharbi et al.(2024)]

Automating CQ authoring

- From scratch with a corpus-based approach and an LLM; e.g., [Antia and Keet(2023)]:
  - create/get text corpus, determine important statements
  - convert into questions (T5)
  - filter questions with CLaRO to extract 'probably more suitable as CQ' ones
- Retrofitting to an existing ontology with an LLM-based approach; e.g., [Alharbi et al.(2024)]:
  - Extract triples from ontologies
  - LLM prompting; e.g.: "Based on <statement>, generate a list of relevant question" + statement
  - Filter output (remove redundancies, questions that a clearly not CQs for ontologies)
- See also the very recent review in [Alharbi et al.(2024)]

Outline

Some first considerations

- An easy CQ: Which plants eat animals?
- Mapped into SPARQL-OWL:

```
SELECT DISTINCT ?eats
   WHERE {
        ?eats rdfs:subClassOf awo:plant, [
        a owl:Restriction ;
        owl:onProperty awo:eats;
        owl:someValuesFrom awo:animal
        ] .
   FILTER(?eats != owl:Nothing)
}
```

example from [Potoniec et al.(2020)]

Some first considerations

- An easy CQ: Is the output format of [it] proprietary?
- Mapped into a not-so-simple query:

```
ASK WHERE { $PPx1$ rdfs:subClassOf [
  a owl:Restriction ; owl:onProperty swo:SWO_0000087 ;
  owl:someValuesFrom [ a owl:Restriction ;
      owl:onProperty swo:SWO_0004002 ;
      owl:someValuesFrom ?format ]  ] .
?format rdfs:subClassOf obo:IAO_0000098 .
filter(not exists { ?subformat rdfs:subClassOf ?format .
filter(isURI(?subformat) && ?subformat!=?format && ?subformat!=owl:Nothing) })
filter(not exists {?os_sw rdfs:subClassOf swo:SWO_0000001,
     [ a owl:Restriction ; owl:onProperty swo:SWO_0000087 ;
        owl:someValuesFrom [a owl:Restriction ;
            owl:onProperty swo:SWO_0004002 ;
            owl:someValuesFrom ?format] ] ,
            [a owl:Restriction ; owl:onProperty swo:has_license ;
             owl:someValuesFrom [a owl:Restriction ;
               owl:onProperty swo:SWO_9001002 ;
               owl:someValuesFrom swo:SWO_9000020 ] ].
      filter(isURI(?os_sw) && ?os_sw!=owl:Nothing) }) }
```

example from [Potoniec et al.(2020)]

Issues (1/2)

- Only 131 CQs of the 234 could be converted as such into SPARQL-OWL
- Four main issues of *untranslatability* of the CQ into SPARQL-OWL [Wiśniewski et al.(2019)]
  - Lack of required vocabulary in the ontology (n=58)
    - A gap in the ontology? Scope drift (be it intentionally or not)? Other?
  - Ambiguity of the CQ (n=26)
  - The CQ is stated against ABox instead of TBox and out of scope for the dataset (n=19)
    - TBox CQs are assumed applicable for ontologies
  - Inability to express the query (n=8)

Issues (2/2)

- Recent manual re-assessment of those 234, just the CQs [Keet and Khan(2024)]:
  - 53 problematic questions (23%)
  - 17 are easily solvable grammar issues
  - 9 were about 'can I do x'/'how to do x' rather than about content of the ontology
  - Rest had a range of issues; e.g., "fastest" software, imprecision with, e.g., "possibly problematic" behaviour

Issues (2/2)

- Recent manual re-assessment of those 234, just the CQs [Keet and Khan(2024)]:
  - 53 problematic questions (23%)
  - 17 are easily solvable grammar issues
  - 9 were about 'can I do x'/'how to do x' rather than about content of the ontology
  - Rest had a range of issues; e.g., "fastest" software, imprecision with, e.g., "possibly problematic" behaviour
- **Relative** issues: wrt the ontology's content, or representation or query language limitations
- **Absolute** issues: syntax issues, questions no ontology will answer, vague statements

Semantically faulty questions?

- "semantically faulty" questions [Wiśniewski(2015)]

Q: "Which natural number is smaller than 0?"

'A' 'none', but 'none' is not a natural number that was asked for

Semantically faulty questions?

- "semantically faulty" questions [Wiśniewski(2015)]

Q: "Which natural number is smaller than 0?"

'A' 'none', but 'none' is not a natural number that was asked for

- Applying that to CQs for ontologies:
  - A CQ should not have always the empty set as answer
  - And maybe also that at least the category of the entity in the answer is among the intended one
  - e.g.: for What is assessed in the walking task?, that the domain expert would expect, a.o., gait and latency data

Digging deeper

- Question complexity, implicit assumptions [Cohen(1929)]
- Its *function* as an "information-seeking act" [Watson(2021)]
- With knowledge acquisition vs. organisation *goals* [Ram(1991)]
- *Motivations* behind asking questions; e.g., to "expose" a colleague for non-performance of duties [Watson(2021)]
- Types: e.g., Hypothesis-scanning (e.g., "is it a melon?") vs. constraint-seeking (e.g., "is it edible?") questions

  [Bertolazzi et al.(2023), Ruggeri and Lombrozo(2014)]
- Questions as sets of declaratives [Knuth(2003), Wiśniewski(2015)]

Digging deeper

- Competency questions elsewhere to, e.g., assess fitness of a witness in a trial
  - Test whether the human is competent in some way

Digging deeper

- Competency questions elsewhere to, e.g., assess fitness of a witness in a trial
  - Test whether the human is competent in some way
- Competency questions for ontologies: competency of *what*?
  - The ontology as a whole
  - Some modelling style or modelling pattern used in the ontology
  - A small piece of knowledge asked for in the CQ (class, property, axiom, vocabulary presence)
  - And the annotations, too?

Preliminary hierarchy [Keet and Khan(2024)]

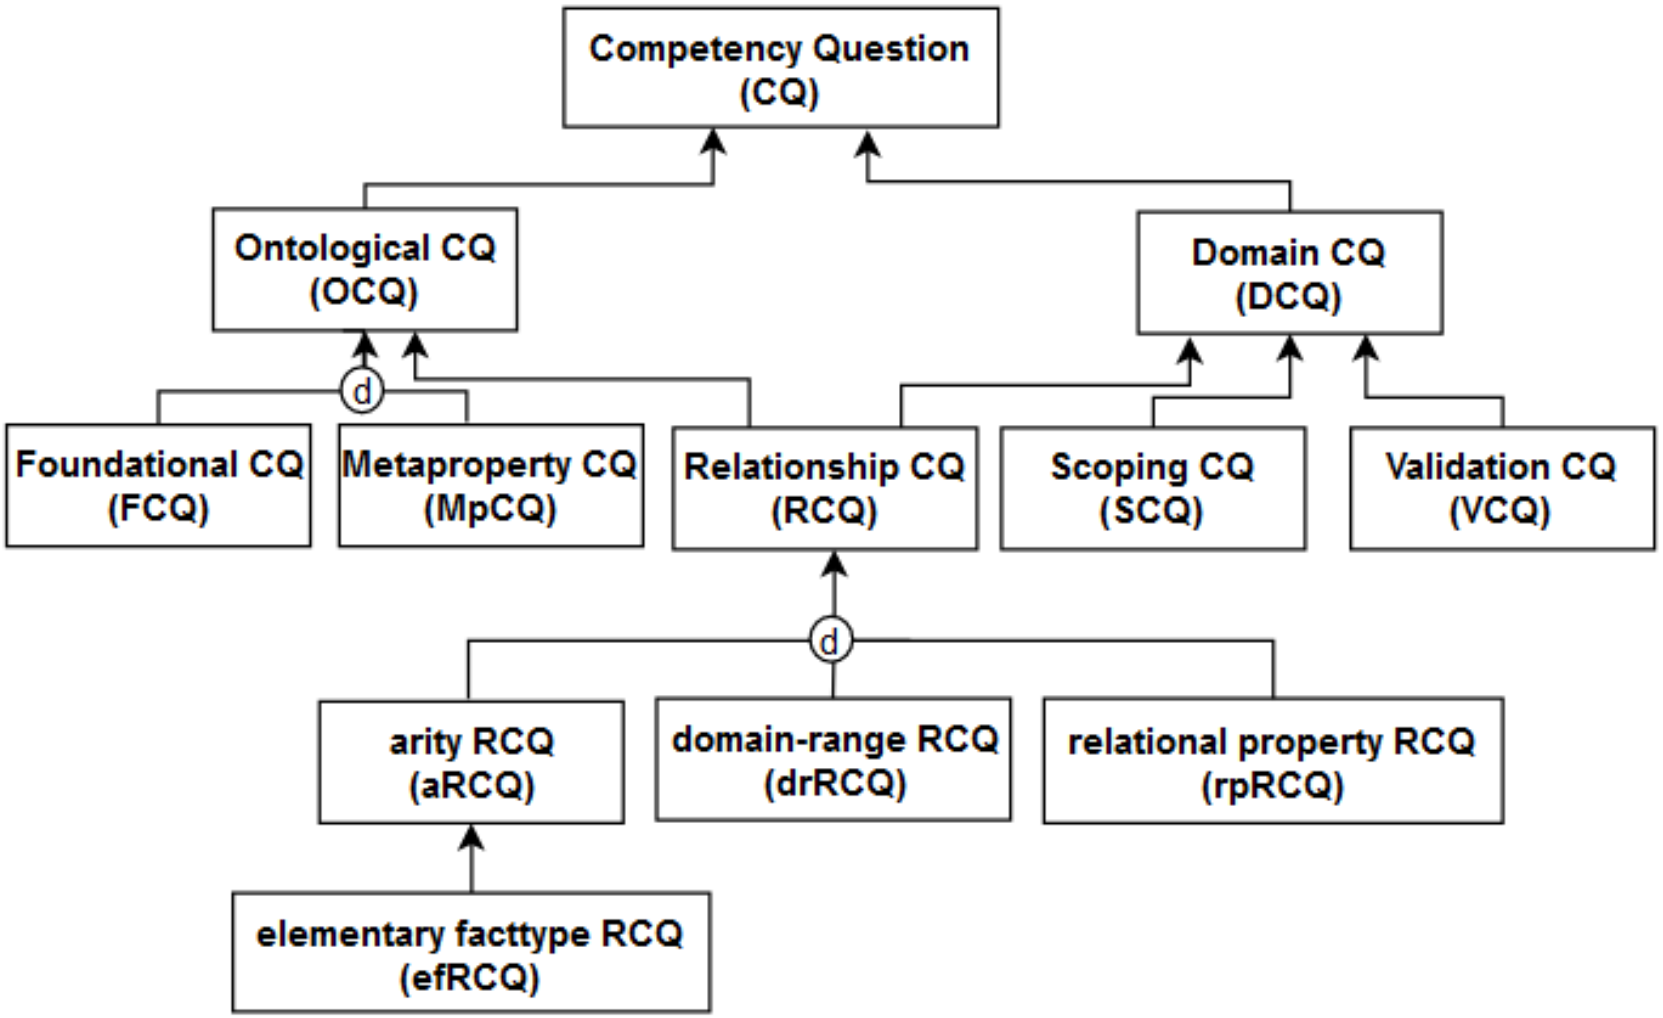

Selection of examples

- FCQs fit the notion of Ram's [Ram(1991)] "knowledge organisation" goal—where precisely to link it to that ontology
- MpCQs and RCQs are information-seeking, "knowledge acquisition" [Ram(1991)]
- VCQs align with question as "a system of assertions that answers that question" [Knuth(2003)]

Selection of examples

- FCQs fit the notion of Ram's [Ram(1991)] "knowledge organisation" goal—where precisely to link it to that ontology
  - Is [X] something that is happening or occurring?
  - Is [X] a collection of disjoint self-standing single objects?
- MpCQs and RCQs are information-seeking, "knowledge acquisition" [Ram(1991)]

- VCQs align with question as "a system of assertions that answers that question" [Knuth(2003)]

Selection of examples

- FCQs fit the notion of Ram's [Ram(1991)] "knowledge organisation" goal—where precisely to link it to that ontology
  - Is [X] something that is happening or occurring?
  - Is [X] a collection of disjoint self-standing single objects?
- MpCQs and RCQs are information-seeking, "knowledge acquisition" [Ram(1991)]
  - Is each instance of a coffee bean necessarily (at all times of its existence) an instance of a coffee bean? (rigid)
  - Does the "marriage" relationship always involve two individuals? (arity)
  - If a human loves, must it also love itself? ((ir)reflexivity)
- VCQs align with question as "a system of assertions that answers that question" [Knuth(2003)]

Repositories of CQs: collecting to analyse (and reuse?)

- CQ data set [Wiśniewski et al.(2019)] (n=234); no substantive metadata other than ontology it was developed for,
  `https://github.com/CQ2SPARQLOWL/Dataset`
- Repositories of Ontology CQs (ROCQS) [Keet and Khan(2024)]; (n=438, with 38 FCQs, 33 VCQs, 323 SCQ, 27 RCQs, and 17 MpCQs; includes the CQ dataset), with more metadata (also which FO/feature, templated or not, type, and more)
  `http://www.meteck.org/files/ROCQS/ROCQS.htm`
- Tamma's list with links (last week):
  `https://github.com/KE-UniLiv/CQ-benchmark/tree/main`

Outline

Discussion: plenty TBD (1/2)

- Automated CQ generation efforts, CQ quality issues
- Quality metrics of CQs
- CQ formalization:
  - CQ to SPARQL-OWL [Wisniewski et al.(2021)]
  - CQ to OWL (or another suitable logic)
  - Challenges: CQ presuppositions [Potoniec et al.(2021)] and modelling styles [Fillottrani and Keet(2019)]
- Methods and techniques for more effective use of CQs in ontology development (where and how)

Discussion: plenty TBD (2/2)

- Philosophical and logical works served to analyse, identify, and identify types CQs; there may be more, more refined types
- Collection of more CQs for ontologies for better ontology authoring, possibly also benchmarking
- Integration of CQs with existing ODEs
- User studies on use, effectiveness, relevance etc.

Conclusions

- CQs as complex acts with diverse purposes, used throughout the ontology engineering lifecycle
- Towards a framework for understanding and applying CQs in ontology development
- Need for CQ quality, CQ authoring assistance, deployment/use of CQs in ontology engineering

References I

Reham Alharbi, Valentina Tamma, Floriana Grasso, and Terry Payne. An experiment in retrofitting competency questions for existing ontologies. In *Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing*, SAC '24, page 1650?1658, New York, NY, USA, 2024. Association for Computing Machinery. doi: 10.1145/3605098.3636053.

Reham Alharbi, Valentin Tamma, Floriana Grasso, and Terry R. Payne. A review and comparison of competency question engineering approaches. In Mehwish Alam, Marco Rospocher, Marieke van Erp, Laura Hollink, and Genet Asefa Gesese, editors, *Knowledge Engineering and Knowledge Management*, pages 271–290, Cham, 2024. Springer.

Mary-Jane Antia and C. Maria Keet. Assessing and enhancing bottom-up CNL design for competency questions for ontologies. In *Proceedings of the Seventh International Workshop on Controlled Natural Language (CNL 2020/21)*. ACL, 2021. URL https://aclanthology.org/2021.cnl-1.11.

Mary-Jane Antia and C. Maria Keet. Automating the generation of competency questions for ontologies with agocqs. In Fernando Ortiz-Rodriguez, Boris Villazón-Terrazas, Sanju Tiwari, and Carlos Bobed, editors, *Knowledge Graphs and Semantic Web*, pages 213–227, Cham, 2023. Springer.

Leonardo Bertolazzi, Davide Mazzaccara, Filippo Merlo, and Raffaella Bernardi. ChatGPT's Information Seeking Strategy: Insights from the 20-Questions Game. In C. Maria Keet, Hung-Yi Lee, and Sina Zarrieß, editors, *Proceedings of the 16th International Natural Language Generation Conference, INLG 2023, Prague, Czechia, September 11 - 15, 2023*, pages 153–162. Association for Computational Linguistics, 2023. URL https://aclanthology.org/2023.inlg-main.11.

Felix S Cohen. What is a question? *The monist*, pages 350–364, 1929.

References II

Kieren Davies, C. Maria Keet, and Agnieszka Lawrynowicz. More effective ontology authoring with test-driven development and the TDDonto2 tool. *International Journal on Artificial Intelligence Tools*, 28 (7):1950023, 2019.

P. R. Fillottrani and C. M. Keet. Dimensions affecting representation styles in ontologies. In *1st Iberoamerican conference on Knowledge Graphs and Semantic Web (KGSWC'19)*, volume 1029 of *CCIS*, pages 186–200. Springer, 2019. 24-28 June 2019, Villa Clara, Cuba.

C. M. Keet, Z. Mahlaza, and M.-J. Antia. CLaRO: a controlled language for authoring competency questions. In E. Garoufallou et al., editors, *13th Metadata and Semantics Research Conference (MTSR'19)*, volume 1057 of *CCIS*, pages 3–15. Springer, 2019. 28-31 Oct 2019, Rome, Italy.

C. Maria Keet and Zubeida Khan. On the roles of competency questions in ontology engineering. In *24th International Conference on Knowledge Engineering and Knowledge Management (EKAW'24)*, volume 15370 of *LNAI*, pages 123–132. Springer, 2024. November 26-28, Amsterdam, The Netherlands.

Kevin H. Knuth. What is a question? In *AIP Conference Proceedings*, volume 659, pages 227–242. AIP, 2003. doi: 10.1063/1.1570546.

Jedrzej Potoniec, Dawid Wisniewski, Agnieszka Lawrynowicz, and C. Maria Keet. Dataset of ontology competency questions to SPARQL-OWL queries translations. *Data in Brief*, 29:105098, 2020. doi: 10.1016/j.dib.2019.105098.

References III

Jedrzej Potoniec, Dawid Wisniewski, and Agnieszka Lawrynowicz. Incorporating presuppositions of competency questions into test-driven development of ontologies (S). In *The 33rd International Conference on Software Engineering and Knowledge Engineering, SEKE 2021, USA, July 1 - July 10, 2021*, pages 437–440. KSI Research Inc., 2021.

Ashwin Ram. A theory of questions and question asking. *Journal of Learning Sciences*, 1(3/4):273–318, 1991.

Azzurra Ruggeri and Tania Lombrozo. Learning by asking: How children ask questions to achieve efficient search. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*, pages 1335–1340, 2014. URL `https://escholarship.org/uc/item/32c734r9`.

Lani Watson. What is a question. *Royal Institute of Philosophy Supplement*, 89:273–297, 2021. doi: 10.1017/s1358246121000114.

Andrzej Wiśniewski. Semantics of questions. *The handbook of contemporary semantic theory*, pages 271–313, 2015.

Dawid Wisniewski, Jedrzej Potoniec, and Agnieszka Lawrynowicz. SeeQuery: An automatic method for recommending translations of ontology competency questions into SPARQL-OWL. In *CIKM '21: The 30th ACM International Conference on Information and Knowledge Management*, pages 2119–2128. ACM, 2021.

Dawid Wiśniewski, Jedrzej Potoniec, Agnieszka Ławrynowicz, and C. Maria Keet. Analysis of ontology competency questions and their formalizations in SPARQL-OWL. *Journal of Web Semantics*, 59:100534, 2019. doi: 10.1016/j.websem.2019.100534.

Thanks to collaborators and students: Agnieszka Lawrynowicz, Zubeida Khan, Dawid Wisniewski, Jedrzej Potoniec, Mary-Jane Antia, and Zola Mahlaza

Questions?