Ontologies

- An ontology → computer interpretable, with rules, relationships, constraints over a vocabulary.
- Used to represent, reason and explore autonomous systems. Using ontologies to represent the decision making, the environment. Used for representing user requirements information, to measure performance
- Nothing has to be created from 0, ontologies can be extended
- How do we use this to evaluate humans? Challenges as humans are coupled with the autonomous systems. Humans can be interviewed to find out what has happened. Evaluate human trust, how trust is built, from small aspects to complicated situations. Inject interesting scenarios and faults for testing. Experiencing autonomy according to degrees of complexity, to learn from simple ones first. Consider different types of people.
- Ontologies → introspection, to understand what systems do and why→ knowledge models pulled out of the system
- Challenges: eliciting what we want the system to do (requirements) when in an interdisciplinary group of people with different expertise, needs, vision; expressibility of the human language and translating this into technical development; understanding systems with many components that are composed but need to be clear and safe, with structure and meaning to us.
- Making robots understand us → natural language, common meaning → to program robots (executable code generation) that are provable correct
- Ontologies as tools for knowledge representation → types of robots, parameters and data types, logics (areas of concern)→ how to reason about errors and ensure fault recovery

Discussion:

In the context of verification, ontologies give us a mapping from robot domain to user domain; does this help us with verification? Like if we have two modules that have been described, will verification problem be hard/simpler/about the same? Depends on how we are recomposing the modules. A “common” baseline of definitions and meaning is needed for these modules. Ontology will let you know why it did at its output; combination will help verification, but the new module’s decision logic will have to be verified.

How defining ontologies locks you in to something to a certain extent? Understanding natural language is a separate problem than that of ontologies. Semantics need to be common to compare performance of systems.

As processes jump from manual to automated, can ontologies be flexible to accommodate the differences? Ontologies do not have to be rigid.

Working ontology is necessary for the composability of systems. Reasoning element needs to be decomposed; what are the challenges about making an ontology over reasoning? Ontology for system model? It is difficult to agree, because every research group develops models and meanings. Invite folks to join standards committee for ontologies.

Confused about ontologies: What do we mean by ontologies? Fragment of first order logic, or complete system specification? Where do you draw the bounds? There is not a consensus in the ontology community. Least common denominator would be the dictionary terms (e.g. first-order logic fragments). Reasoning about the dictionary is the hard problem.

On the need for flexibility in systems, we don't know how to specify basic concepts. Can ontologies help us by specifying weakness of concepts to accommodate the specification of these concepts? Do we satisfy conditions? Possibly a fuzzy logic approach for qualifiers… Not sure if ontologies enable this. Ontologies are best at defining hard concepts, relating them.

Design for verification: where in the design/development pipeline should we start talking about ontological specification? Should begin at the requirements stage. Co-design is far more feasible than adding it on after the fact. Understanding why is tied to what you need about the system → building introspection from the beginning.

What is relationship between variables that are typed? Direct correlation between typed variables and ontological concepts.

Natural language uses statistical parser using dictionaries that are available to understand what the user has uttered; concepts retrieved from existing ontology. Ontologies are static; not dynamic.

Is it possible to learn new semantics in an ontology? No, not at this moment; though work is in progress. There have been `20 questions' based work to try to fit new concepts into hierarchy of existing concepts.

Discussing ontologies for non-natural language applications: Trying to decompose robot's command space into tasks and behaviors, then re-compose into scripts or sequences that can be sent, tested on ground. Ontologies are created ahead of time.

How do you ensure that the ontology has abstracted the right level of detail to accommodate new tasks? Encoded an ontology as an XML schema, take information out of sensing system and treat as instances of ontology, stream to db. Assists with verifying certain properties that should be true based on ontology. Supports reasoning over the perception system e.g. a car should have 4 wheels. Like model-based diagnosis(?) maybe.

How do you verify that the ontology has the right pieces? It's a challenge. Would run the planning system that knows its action set, do reasoning over action set to get from current state to goal. If it doesn't succeed, might imply something is missing but more thought needed.

Needed to standardise ontologies to make them open, trustworthy, available for everyone, to find common ground, for applications to autonomous systems.