Language Models as a Knowledge Source for Cognitive Agents

Robert Wray, James Kirk, John Laird

Center for Integrated Cognition (CIC)

Nov 2021

This work is supported by the Office of Naval Research, contract N00014-21-1-2369. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Department of Defense or Office of Naval Research.

CIC @ IQMRI

Autonomous Agents Often Need to Acquire New Knowledge

Effective behavior in a changing environment often requires acquisition of new knowledge

- New and changing tasks
- Changing environment
- Changing constraints and expectations

Agent can "extract" knowledge from various sources to help it adapt and learn

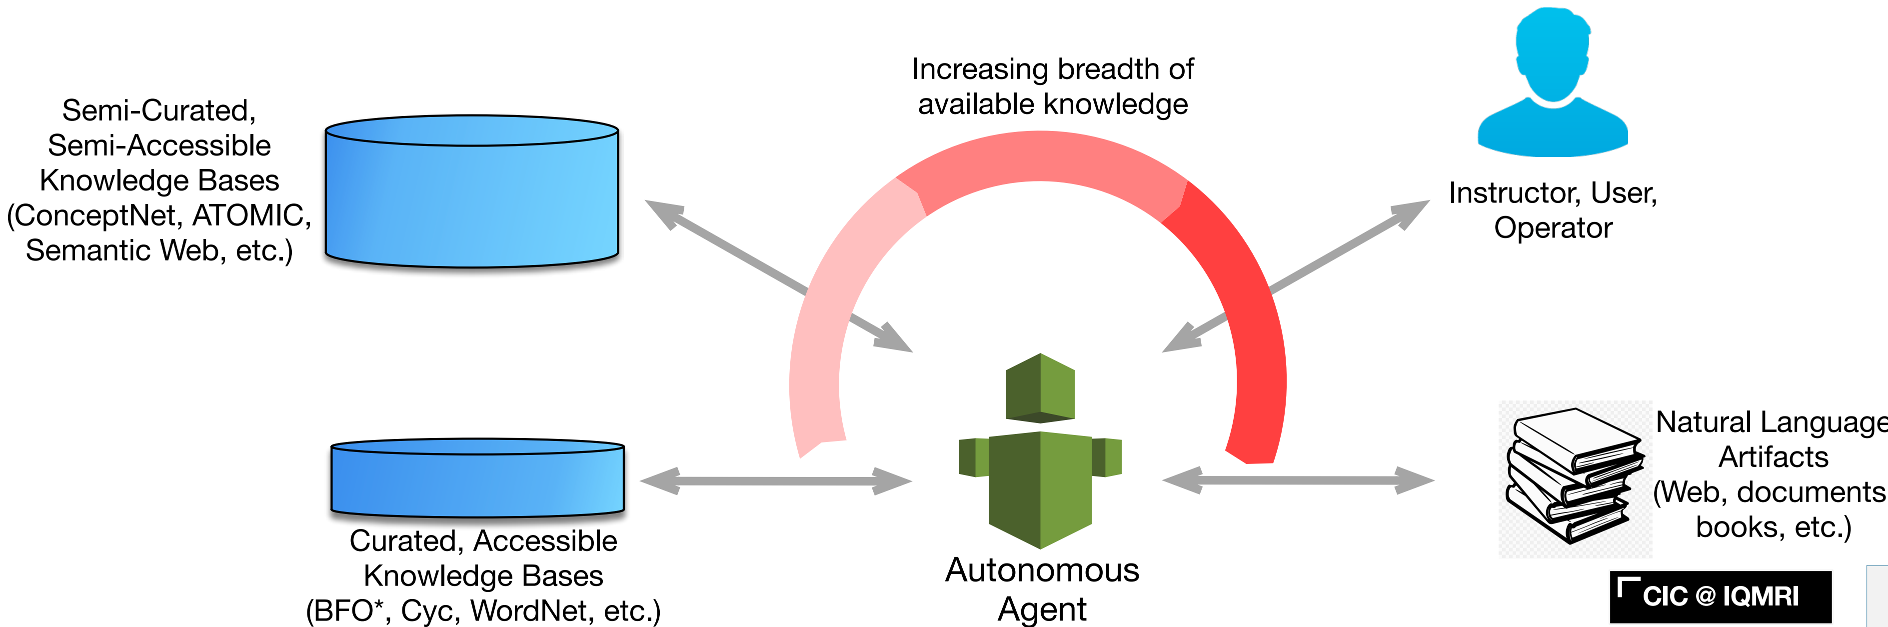

CIC @ IQMRI

Illustration: Work/Task-Support Robot

- Illustration inspired by work on Interactive Task Learning (ITL)/Rosie
- Support humans in household, office, or warehouse environment
- Robots could be produced with general task knowledge but…
- Any specific context may require learning tasks / variations on tasks
  - Office: “Tidy up” the library in the law office or the wet lab
  - Household: Make yogurt; order and replace HVAC filter
  - Warehouse: Escort unbadged personnel


Image: *Multichannel Merchant*


Image: *New Scientist*


Image: *Dreamstimes*

Illustration: Work/Task-Support Robot

- An experienced human instructor could likely “teach” a robot to do these tasks (state of art ITL)
- Current instruction process is tedious and time intensive
  - Explain concepts relevant to new task: “library” “yogurt” “escort”
  - Guide agent in transfer of knowledge: “escort is similar to delivery task, but you are delivering a person”
- Human instructor may not always be available

→Look to other knowledge sources to supplement human instruction

Today**: What is the potential (and challenge) of large language models (LLMs) as an knowledge effective source?**

Image: *Multichannel Merchant*

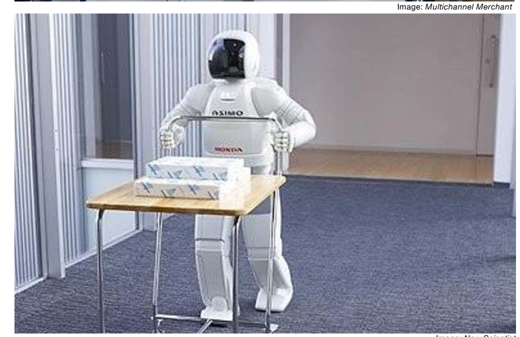

Image: *New Scientist*

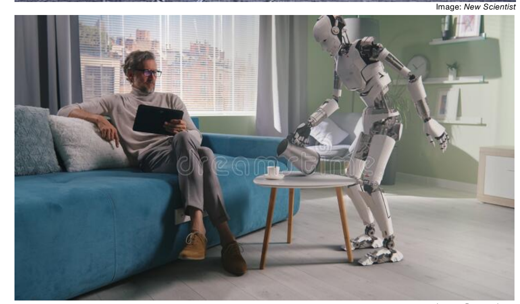

Image: *Dreamstimes*

CIC @ IQMRI

Learning to Make Yogurt (Use Case Example)

Human: Tomorrow, I'm going to teach you to make yogurt.

Robot: Ok, I'll prepare.

…

What equipment do I need to make yogurt?

1. *A yogurt maker. This is the most important piece of equipment. I have the Euro Cuisine YM80 Yogurt Maker, which is a little pricey, but it works great. You can also use a crock pot, but it will take a lot longer to make yogurt.*
2. *A thermometer.* [GPT-3]

What is yogurt? *Yogurt is a dairy product that is made by fermenting milk with a bacteria called Lactobacillus bulgaricus and Streptococcus thermophilus. The fermentation process produces lactic acid, which acts as a preservative. Yogurt is also high in calcium, protein, and vitamins.* [GPT-3]

What are the steps to make yogurt?
1. *Pour milk into the pot.*
2. Warm it until it reaches 85 degrees.
3. Stir in the yogurt starter.
4. Wrap the pot with the blanket to keep the temperature between 85-95 degrees for 6-8 hours.
5. Cool, then refrigerate. [GPT-3]

For all examples in the presentation:
Input to LLM *Response from LLM* [LLM] (Temperature: 0/Deterministic response)

CIC @ IQMRI

A Potential New Knowledge Source: Large Language Models

- Based on sequence to sequence architecture
- Produce a token based on presented tokens
  - Pre-trained on large corpora of text
  - “Large:” M/B/Ts of parameters/weights
- Self-supervised: Learn to produce/predict a token in place within a string of tokens:
  - “The washing machine box is (heavy).” → heavy
- Language models perform very well on a large range of NL benchmark tasks:
  - Text classification (sentiment analysis)
  - Question answering
  - Various kinds of natural language inference
  - …

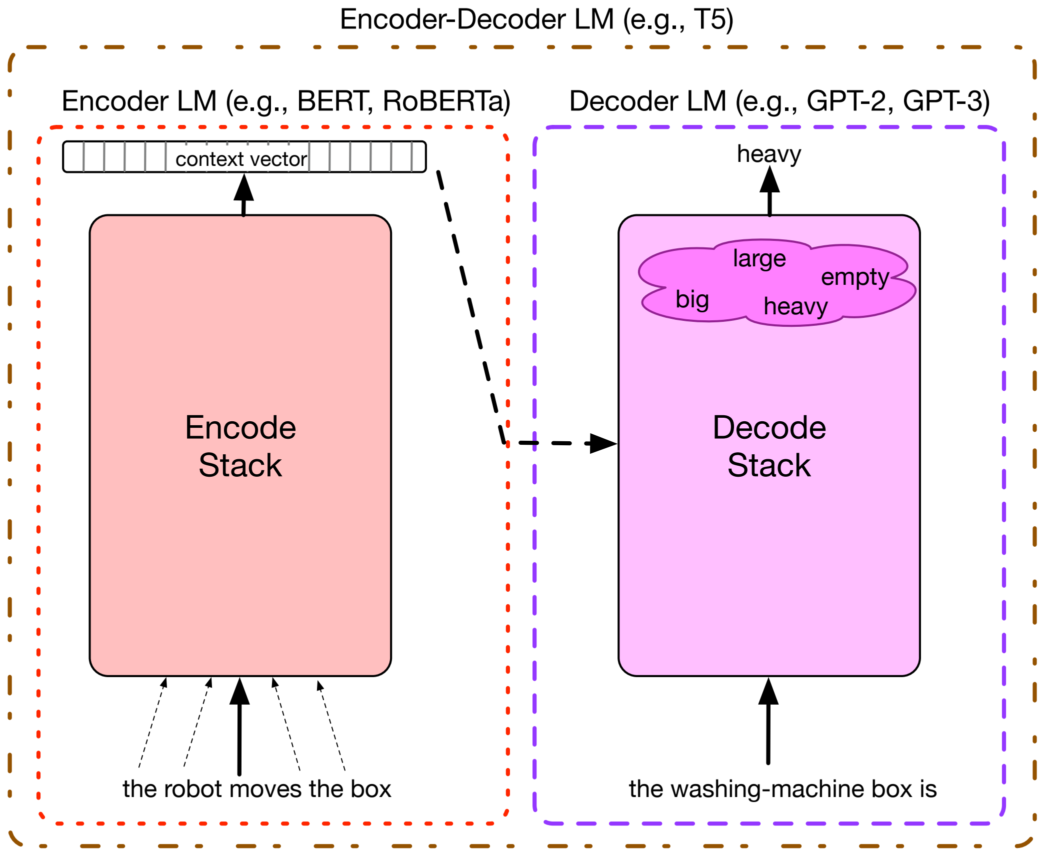

How well can LLMs support agent task learning (as a source of knowledge)?

Suggestive evidence of potential:

- Active, promising research extracting general knowledge from large language models (LLMs)
  - Taxonomic knowledge (parent/child, HAS-A, PART-OF, etc.)
  - Commonsense knowledge
  - Human social knowledge
  - Causal knowledge (e.g., narrative understanding)
  - Planning knowledge
- LLMs are improving quickly (at least on common NL benchmarks)

Current analysis/initial research:

1. Define extraction requirements from a cognitive-agent perspective ("actionable extraction")
2. Sketch a *general* extraction process (based on extraction patterns)
3. Identify core challenges

CIC @ IQMRI

LLM Characteristics Relevant to Task Learning in Cognitive Agents

- Breadth and Depth of Knowledge
  - LLMs *likely* have greater breadth and depth of captured knowledge than existing knowledge stores.

What is yogurt? *Yogurt is a dairy product that is made by fermenting milk with a bacteria called Lactobacillus bulgaricus and Streptococcus thermophilus. The fermentation process produces lactic acid, which acts as a preservative. Yogurt is also high in calcium, protein, and vitamins.* [GPT-3]

Steps to tidy up an office library include:
1. *Remove books that are not used. 2. Sort the remaining books by subject. 3. Place books on shelves in alphabetical order by author. 4. Place books on shelves in chronological order by publication date. 5. Place books on shelves in order of use. 6. Place books on shelves* [GPT-3]

When a visitor to the warehouse lacks a badge, the visitor must be escorted. Duties of an escort include:
1. *Escort visitors to the warehouse. 2. Ensure that visitors do not remove any items from the warehouse. 3. Ensure that visitors do not damage any items in the warehouse. 4. Ensure that visitors do not interfere with the work of warehouse employees. 5. Ensure that visitors do not enter any restricted areas of the warehouse.* [GPT-3]

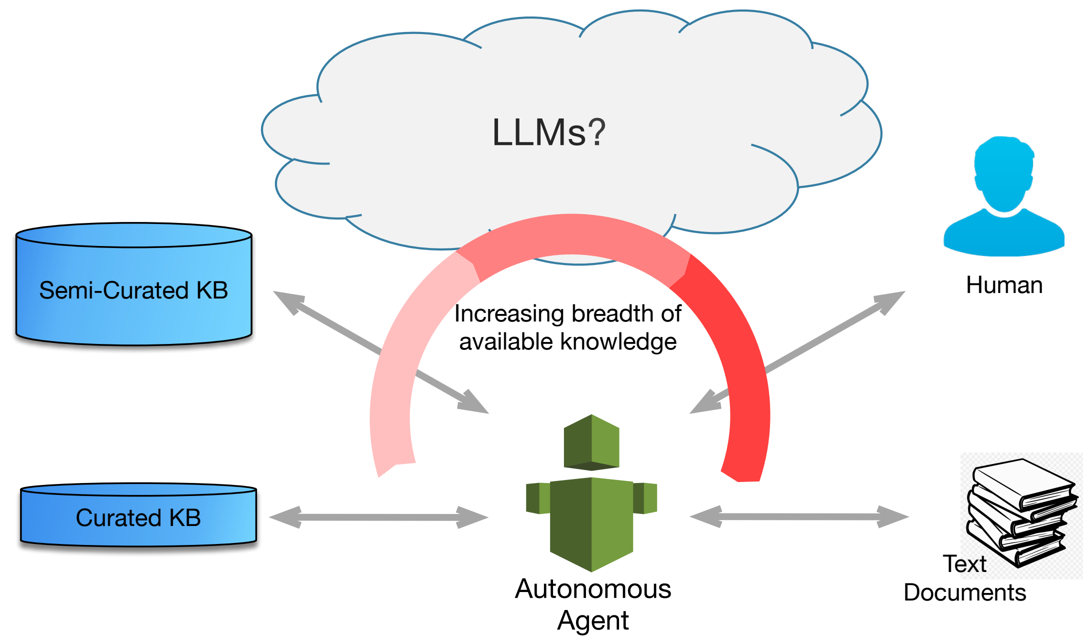

LLM Characteristics Relevant to Task Learning in Cognitive Agents

- Breadth and Depth of Knowledge
  - LLMs *likely* have greater breadth and depth of captured knowledge than existing knowledge stores.
- Provenance and Accuracy of Knowledge
  - Extractable knowledge is conditioned on the source/corpora (which can be erroneous, dated, etc.).

How long can I leave milk on the counter? *The answer to this question depends on the type of milk you have.*
Whole milk: 3-5 days
Skim milk: 5-7 days
Soy milk: 7-10 days [GPT-3]

Who is the current leader of Canada? *The current leader of Canada is Justin Trudeau.* [GPT-3]

Who is the current leader of the UK? *Theresa May is the current Prime Minister of the United Kingdom. She is the leader of the Conservative Party* [GPT-3]

LLM Characteristics Relevant to Task Learning in Cognitive Agents

- Breadth and Depth of Knowledge
  - LLMs *likely* have greater breadth and depth of captured knowledge than existing knowledge stores.
- Provenance and Accuracy of Knowledge
  - Extractable knowledge is conditioned on the source/corpora (which can be erroneous, dated, etc.).
- Relevance of Knowledge
  - LLMs have no awareness of the agent's context (but KBs do not either).

How long can I leave milk on the counter? *The answer to this question depends on the type of milk you have.*
Whole milk: 3-5 days
Skim milk: 5-7 days
Soy milk: 7-10 days [GPT-3]

How long can I leave milk sitting on the counter? *The answer to this question depends on the type of milk you have. Pasteurized milk can be left out for up to two hours at room temperature*. [GPT-3]

CIC @ IQMRI

LLM Characteristics Relevant to Task Learning in Cognitive Agents

- Breadth and Depth of Knowledge
  - LLMs *likely* have greater breadth and depth of captured knowledge than existing knowledge stores.
- Provenance and Accuracy of Knowledge
  - Extractable knowledge is conditioned on the source/corpora (which can be erroneous, dated, etc.).
- Relevance of Knowledge
  - LLMs have no awareness of the agent's context (but KBs do not either).

What are the steps to make yogurt?
1. *Pour milk into the pot.*
2. Warm it until it reaches 85 degrees.
3. Stir in the yogurt starter.
4. Wrap the pot with the blanket to keep the temperature between 85-95 degrees for 6-8 hours.
5. Cool, then refrigerate. [GPT-3]

What are the steps to make yoghurt?
1. *Heat milk to 180 degrees F.*
2. Cool milk to 110 degrees F.
3. Add starter culture.
4. Incubate for 4-8 hours.
5. Add rennet. [GPT-3]

LLM Characteristics Relevant to Task Learning in Cognitive Agents

- Breadth and Depth of Knowledge
  - LLMs *likely* have greater breadth and depth of captured knowledge than existing knowledge stores.
- Provenance and Accuracy of Knowledge
  - Extractable knowledge is conditioned on the source/corpora (which can be erroneous, dated, etc.).
- Relevance of Knowledge
  - LLMs have no awareness of the agent's context (but KBs do not either).
- Situatedness of Knowledge
  - LLMs can only provide general knowledge, not knowledge specific to the agent's exact circumstance.

Does this house have a yogurt maker? Where is it stored?
What size of HVAC filters are needed in this house?
What is/are the organizational method used for storage in this library?
…

CIC @ IQMRI

LLM Characteristics Relevant to Task Learning in Cognitive Agents

- Breadth and Depth of Knowledge
  - LLMs *likely* have greater breadth and depth of captured knowledge than existing knowledge stores.
- Provenance and Accuracy of Knowledge
  - Extractable knowledge is conditioned on the source/corpora (which can be erroneous, dated, etc.).
- Relevance of Knowledge
  - LLMs have no awareness of the agent's context (but KBs do not either).
- Situatedness of Knowledge
  - LLMs can only provide general knowledge, not knowledge specific to the agent's exact circumstance.
- (Lack of) Model of Knowledge
  - LLMs do not specify what knowledge is available, what form it is in, how accurate or dependable it is, etc.
- Accessibility of Knowledge
  - LLMs can produce tokens/text as complex as natural language.
- Structural Integration
  - Computation and latency are both potential issues for online use by agents.

Goals, Implications, Extraction Requirements

- Long-term Goal: *Actionable* extraction
  - To what extent does the extracted knowledge enable new task performance?
  - Actionable need not imply completeness (many sources of knowledge)
- Implications from characteristics
  - Potentially poor accuracy and relevance
    - Agent must assume/consider potential inaccuracy
    - Agent must provide sufficient context (and limit sensitivity) to extract relevant knowledge
  - Lack of model of knowledge
    - Agent needs model of what the LLM is good for, (un)/successful patterns of use, etc.
  - Potential inaccessibility
    - Agent must direct LLM production to match its interpretation capability

Steps in Acquiring Agent Knowledge

- Patterns of acquiring/extracting knowledge in support of task performance recur across different sources

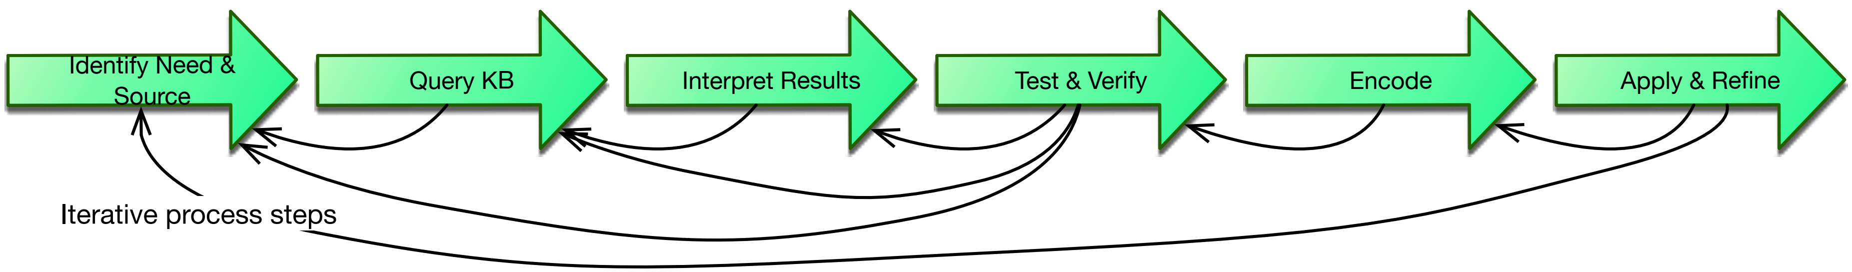

Acquiring Agent Task Knowledge from an LLM

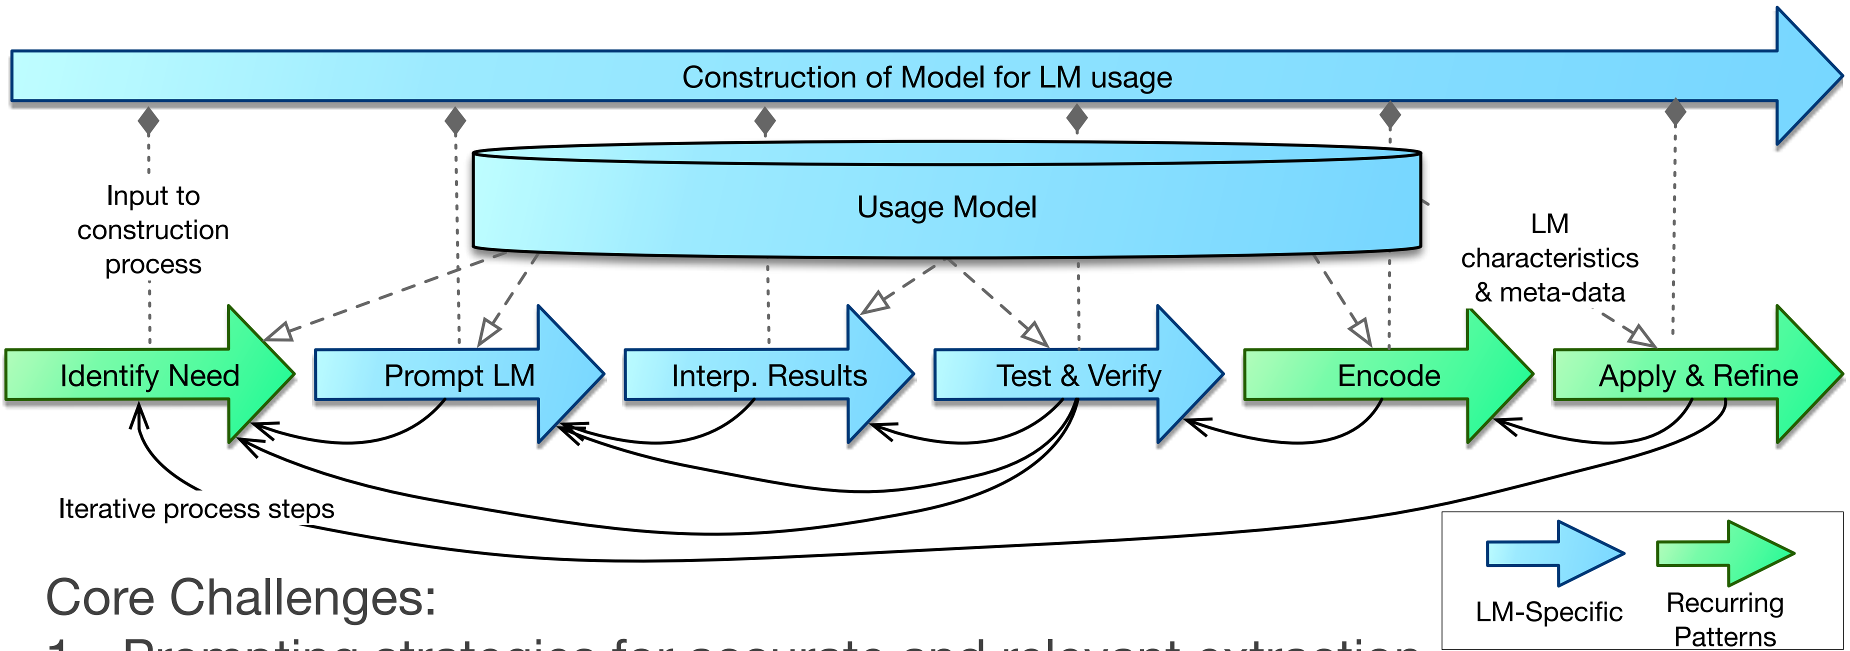

Core Challenges:

1. Prompting strategies for accurate and relevant extraction
2. Constructing a model of LLM usage
3. Guiding/constraining LLM responses (achieving accessible responses)
4. Assessing and verifying results

Methodological challenge: testbed, benchmarks, scaling evaluation

Potential Methods of Extraction

Numerous strategies for extracting knowledge from LLMs

- Fine-tuning LLMs for specific extraction tasks
  - Example: COMET (Bosselut et al., 2019)
  - Potentially expensive (additional training) but possibly apt for general, recurring tasks

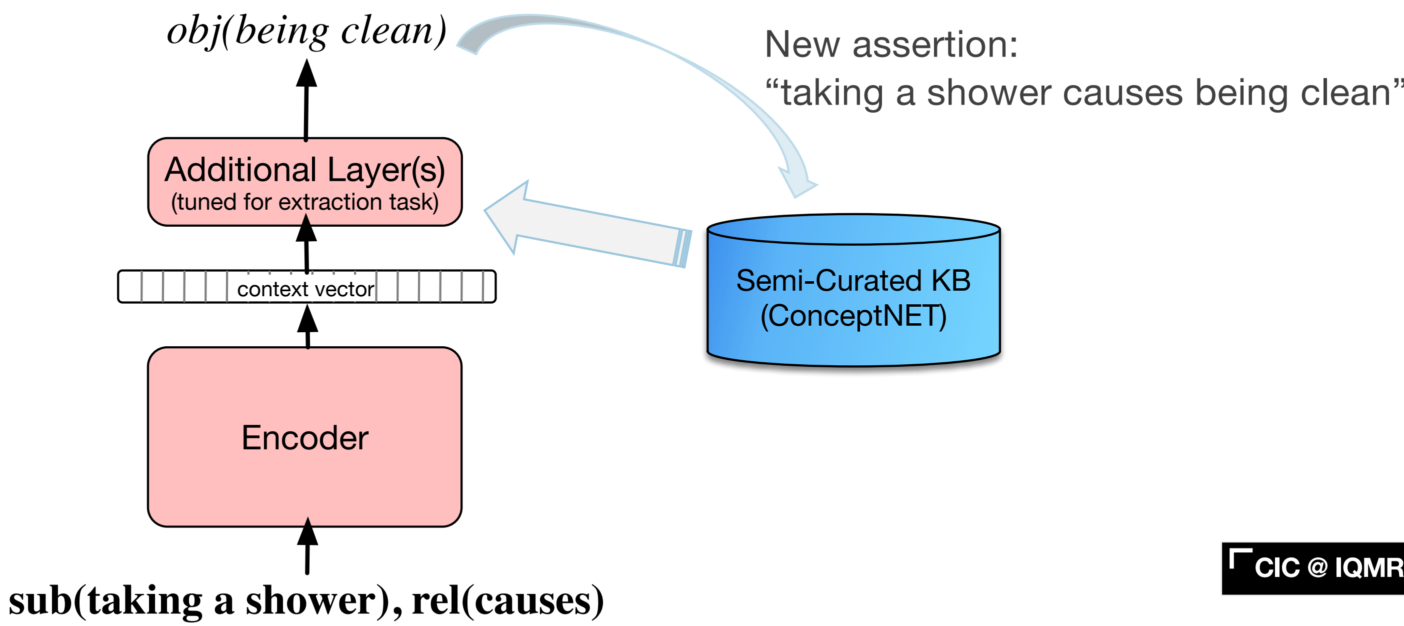

Potential Methods of Extraction

Numerous strategies for extracting knowledge from LLMs

- Fine-tuning LLMs for specific extraction tasks
- Simple prompting using masked language models (MLM)
  - Leverage existing semi-formal understanding of components of tasks (e.g., Problem Space Computational Model/PSCM) to guide/direct extraction of task knowledge.

PSCM Functions	Illustrative Prompt
Problem space description	How do you… Explain …
Goal definition	What is the goal of ?task …
State definition (lexicon)	**?object is also known as a <mask>**
State definition (taxonomic relations)	?object is a type of <mask> ?object has a/is part of a <mask>
Operator definition (lexicon)	?actor can <mask> an ?object
Operator definition (preconditions)	?object is used for <mask>
Operator definition (actions)	?action causes…

Similar in spirit to other work using structured knowledge for extraction (Mostafazadeh et al., 2020; Bian et al., 2021; Kambhampati; et al 2021)

Examples of Template-based Prompting of LLMs

Instantiated Template	LLM Response
Explain how to patrol a warehouse.	*Walk around the warehouse to inspect it and make sure that it is free of people. Search the warehouse for dangerous items like weapons or hazardous materials.* [GPT-3]
The goal of patrolling a warehouse is	*to identify any hazards that may be present.* [GPT-3]
A shelf is also known as a <mask>.	*cabinet, rack, shelf, table, stand, drawer, platform, closet, trunk, box* [RoBERTa]
A wall is part of a <mask>.	*house, building, structure, home, wall, room, castle, prison, fortress, fence* [RoBERTa]
A robot can <mask> a package.	*open, deliver, inspect, carry, assemble, sort, transport, handle, check, sniff* [RoBERTa]
A bay is used for <mask>.	*navigation, fishing, shipping, boats, ships, sailing, transportation, storage, commerce, docking* [RoBERTa]
warehouse robot. A bay is used for <mask>.	*storage, loading, maintenance, equipment, docking, cargo, parking, recycling, deliveries, cleaning* [RoBERTa]
Pushing a box causes	*the box to move and the object inside the box to move with the box.* [GPT-3]

Potential Methods of Extraction

Numerous strategies for extracting knowledge from LLMs

- Fine-tuning LLMs for specific extraction tasks
- Simple prompting using masked language models (MLM)
- Contextual prompting
  - Construct prompts with semantic similarity to desired target

LM Prompt	LM Response
Many household furniture items are stored in the warehouse including	household towels [GPT-2]
	beds, tables, chairs, wardrobes, and more [GPT-3]
Many furniture items are stored in the warehouse including	tables, chairs, and other items [GPT-3]

CIC @ IQMRI

Examples of Template-based Prompting of LLMs

Instantiated Template	LLM Response
Explain how to patrol a warehouse.	*Walk around the warehouse to inspect it and make sure that it is free of people. Search the warehouse for dangerous items like weapons or hazardous materials.* [GPT-3]
The goal of patrolling a warehouse is	*to identify any hazards that may be present.* [GPT-3]
A shelf is also known as a <mask>.	*cabinet, rack, shelf, table, stand, drawer, platform, closet, trunk, box* [RoBERTa]
A wall is part of a <mask>.	*house, building, structure, home, wall, room, castle, prison, fortress, fence* [RoBERTa]
A robot can <mask> a package.	*open, deliver, inspect, carry, assemble, sort, transport, handle, check, sniff* [RoBERTa]
A bay is used for <mask>.	*navigation, fishing, shipping, boats, ships, sailing, transportation, storage, commerce, docking* [RoBERTa]
warehouse robot. A bay is used for <mask>.	*storage, loading, maintenance, equipment, docking, cargo, parking, recycling, deliveries, cleaning* [RoBERTa]
Pushing a box causes	*the box to move and the object inside the box to move with the box.* [GPT-3]

Potential Methods of Extraction

Numerous strategies for extracting knowledge from LLMs

- Fine-tuning LLMs for specific extraction tasks
- Simple prompting using masked language models (MLM)
- Contextual prompting
- Analogical or case-based prompting

Prompt	Response
The household robot charges in the garage. The office robot charges in the maintenance closet. The warehouse robot charges in the	*storage room*
When a room is occupied, the robot should not adjust the temperature or lights. When an office building is unoccupied, the robot should turn the heat to 60 degrees. When the warehouse is unoccupied, the robot should	*turn off the lights*
Q: Where should a robot find a package in an office? A: the mail room Q: Where should a robot find a package in a warehouse?	*A: the shipping department*

[GPT-3]

Guiding/constraining LLM responses (accessible responses)

Concept: Use examples of successful dialogues to shape LLM responses

	Without prior dialogue	Including prior dialogue
Prompt	**Move the package into the cabinet. What is the next goal or subtask of move?**	<u>**Move the box onto the table.**</u> <u>**What is the next goal or subtask of move?**</u> <u>**Pick up the box.**</u> <u>**Put the box onto the table.**</u> <u>**You are done.**</u> **Move the package into the cabinet.** **What is the next goal or subtask of move?**
GPT-3 response	*The next goal or subtask is to move the package into the cabinet. Apply these steps to a goal or subtask until the lowest level of goal or subtask is reached.*	*Pick up the package.* *Put the package into the cabinet.* *You are done.*

Summary & Conclusions

- Cognitive agents routinely reach out to external stores of knowledge to inform their tasks/task learning
- Our goal is to research how effective LLMs can be as such a knowledge source
  - Potential: Encouraging evidence of utility from NL research; fast-improving capabilities
  - Concerns: Cost/Benefit, ROI
    - Accuracy, relevance, and accessibility of results; lack of meta-data (model of knowledge)
- Analysis, design, and research agenda builds on cognitive-systems advancements and understanding
  - Recurring patterns of extraction (across source types)
  - Components of task knowledge (PSCM)
  - Testing and verifying agent knowledge (experiential feedback)

Supplemental

Potential Methods of Extraction

Numerous strategies for extracting knowledge from LLMs

- Fine-tuning LLMs for specific extraction tasks
- Simple prompting using masked language models (MLM)
  - Example: Completion of cloze statements (Petroni et al., 2019)
  - **A robot can <mask> a package**
    open, deliver, inspect, carry, assemble [RoBERTa]
  - Cheap, accessible but potentially inaccurate
  - Positive results possibly attributable to pre-defined prompts (Cao et al., 2021)

Potential Methods of Extraction

Numerous strategies for extracting knowledge from LLMs

- Fine-tuning LLMs for specific extraction tasks
- Simple prompting using masked language models (MLM)
- Contextual prompting
  - Learn the prompt rather than using pre-defined prompts
  - Example: AutoPrompt (Shin et al., 2020)
  - MLM prompt: **<subject> plays in <mask> position**
  - Learned prompt: **<subject> ediatric striker ice baseman defensive <mask>**
  - Greater autonomy and adaptability but not transparent