Analogical Question-Answering

Kenneth D. Forbus

Northwestern University

Includes joint work with
Danilo Ribeiro, Max Crouse, Tom Hinrichs, Maria Chang, and Michael Witbrock

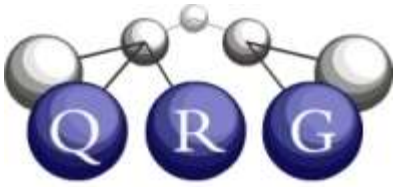

Overview

- Hypothesis: Human reasoning abilities and rapid learning arise from
  - Multiple layers of rich, relational representations
  - Heavy use of analogy to learn and reason
- Analogical Training for Q/A
  - geoquery, identifying processes, and a Kiosk
- Step Semantics: Framework for expressing how processes are described in natural language
- Results on AI2's ProPara dataset
- Conclusions

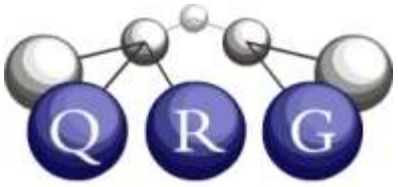

Companion Cognitive Architecture

- Cognitive Architecture = systems that capture broader aspects of cognitive function
  - e.g. ACT-R, SOAR, …
- Big idea for Companions: S*oftware social organisms*
  - Work with people using natural modalities
    - Natural language, sketching
    - Adding speech and vision, via collaboration with MSR
  - Learn and adapt over extended periods of time
  - Maintain themselves

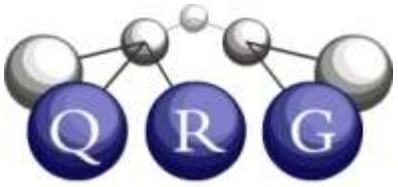

Our Natural Language Approach

- Focus is on deep understanding
- Produces representations that support inference
  - OpenCyc ontology
  - Logical and numerical quantification
  - Quotation
  - Counterfactuals
- Has been used in
  - Learning by reading texts
  - Multimodal learning by reading
  - Moral decision making
  - Conceptual change modeling

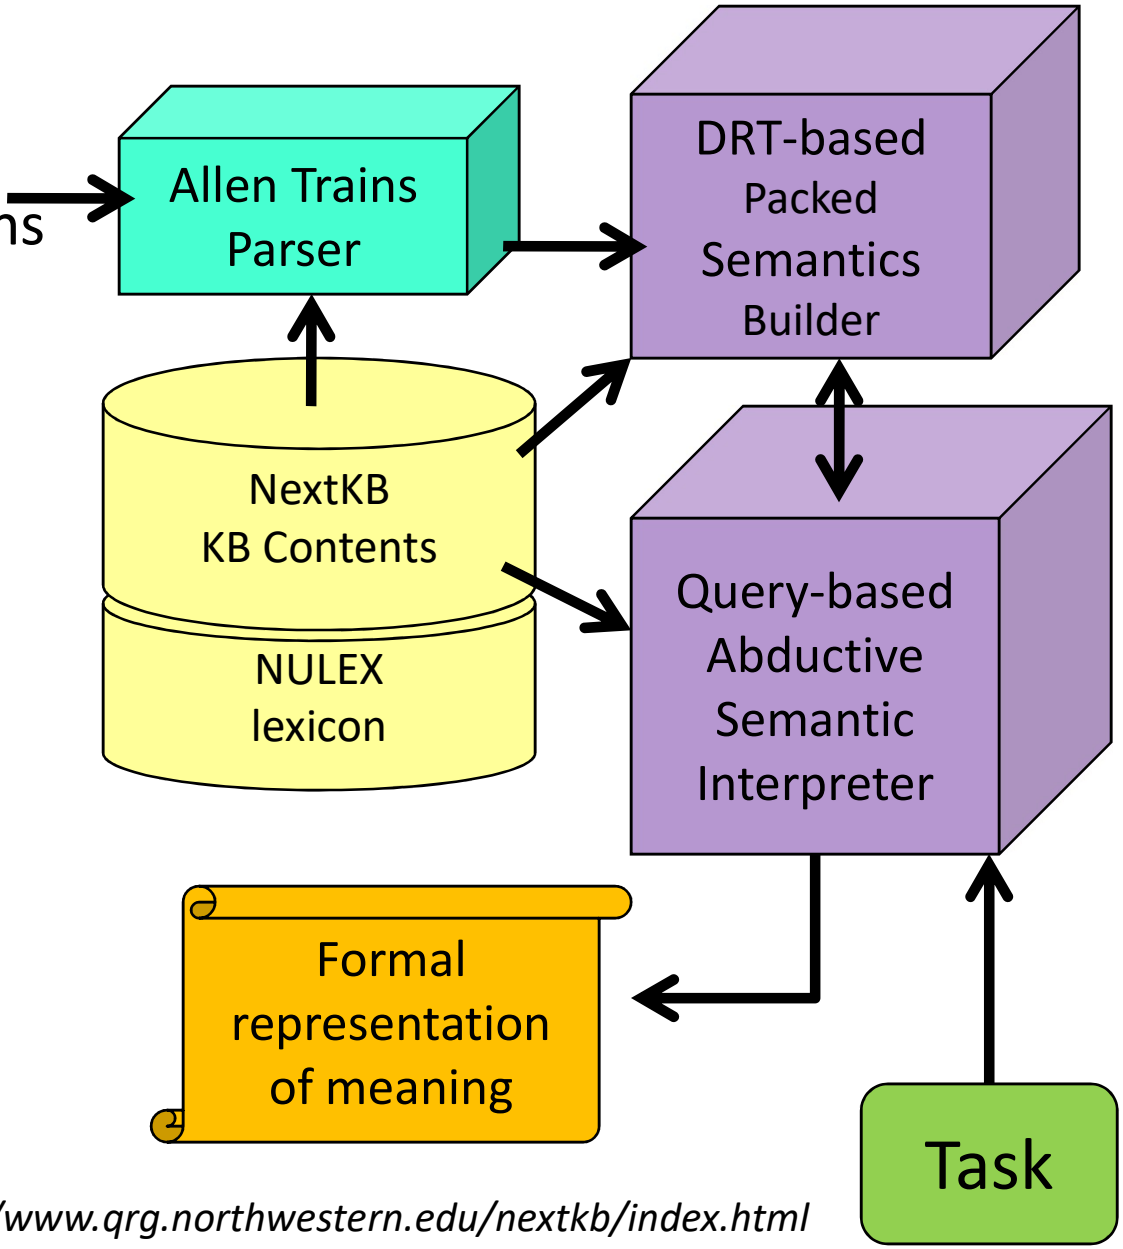

http://www.qrg.northwestern.edu/nextkb/index.html

Layers of Representation in Language

- Allows separation of concerns
  - e.g. Allen's *collaborative problem solving* dialogues
- Higher levels link language to cognitive concerns

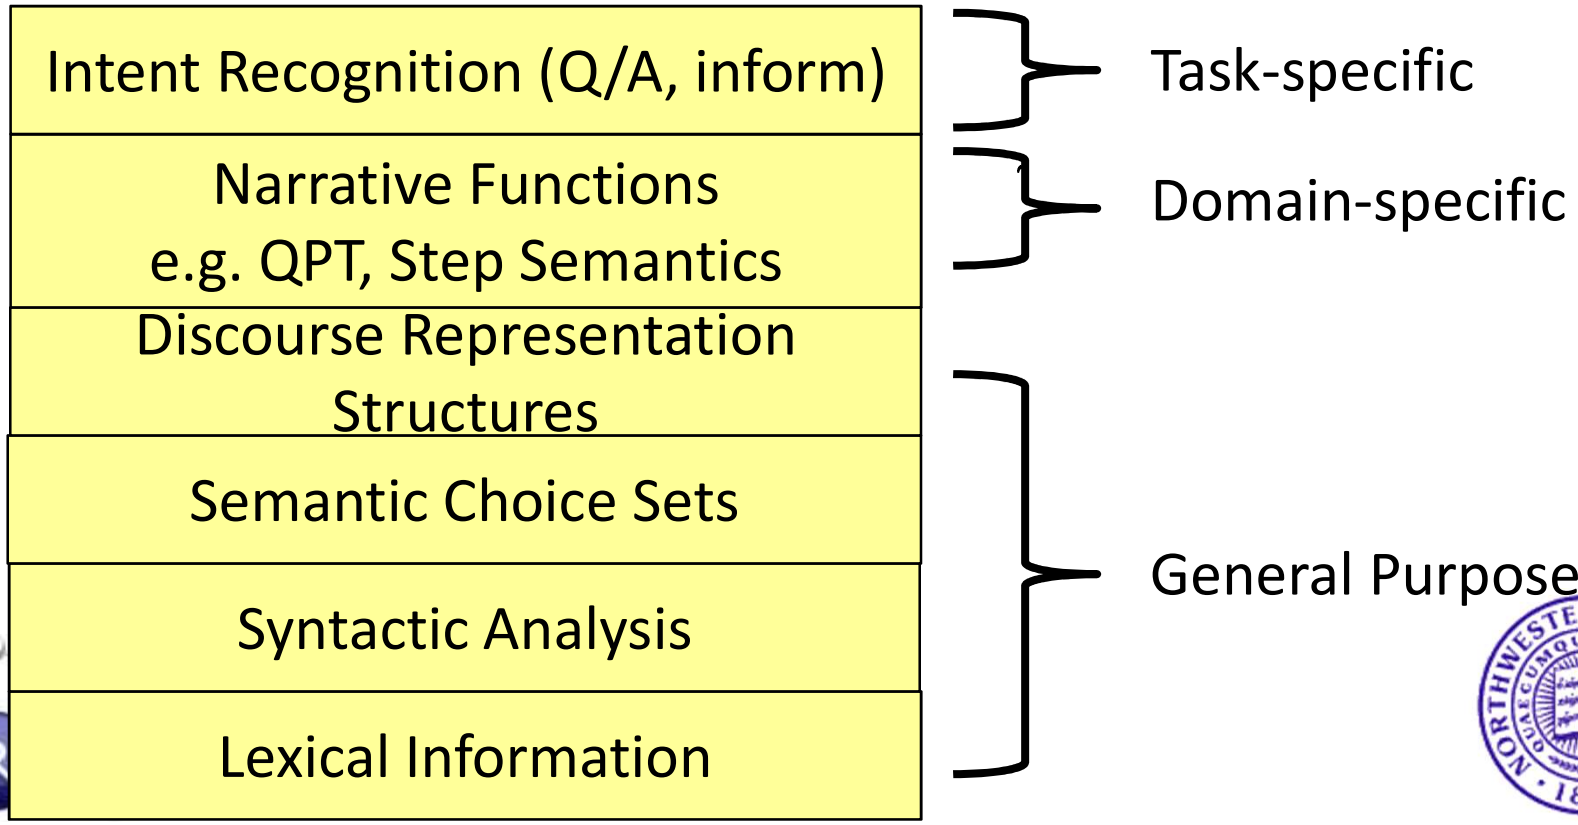

The Analogical Mind

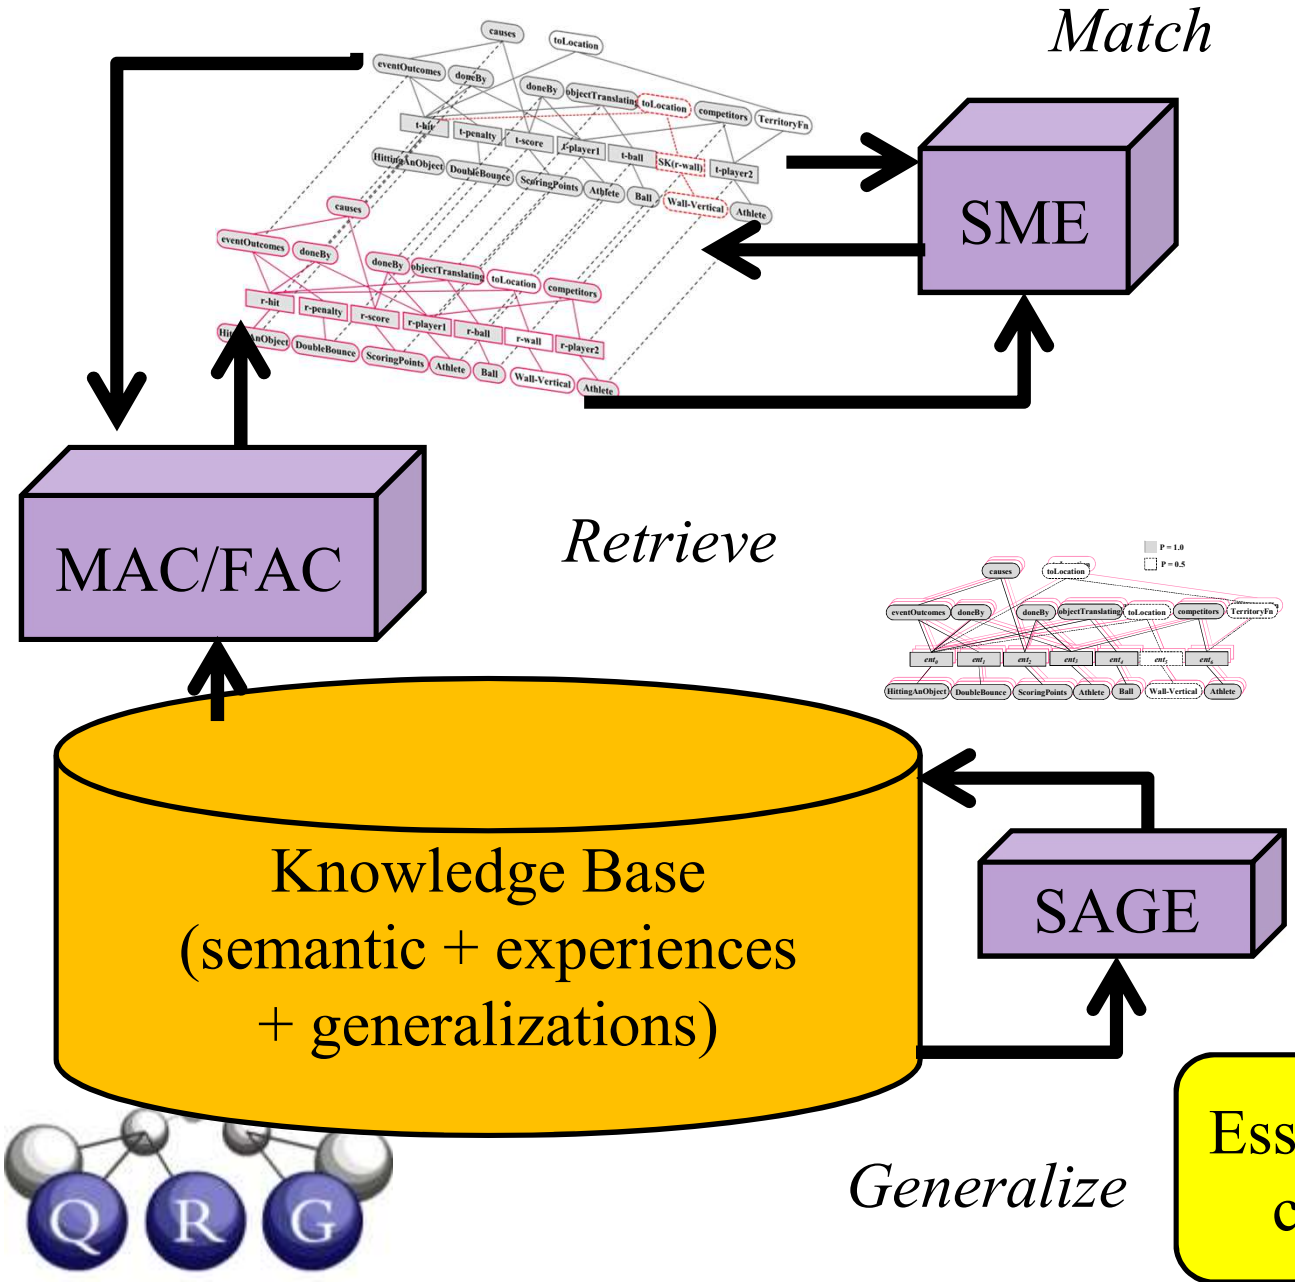

- SME compares examples or generalizations to new situations
- MAC/FAC retrieves experiences and generalizations from LTM
- SAGE incrementally produces generalizations
- Generalizations are probabilistic, partially abstracted, without logical variables

Essence of the Companion cognitive architecture

Analogical Training for Q/A

- Idea: Adapt general-purpose semantic parser by constructing cases from NL question/answer pairs
- Generate a connection graph using the KB to connect predicates in the language analyses of the question and the answers
  - Unannotated data
- Produce *query cases* that suggest components of queries when applied by analogy to new texts
  - Can also produce cases for analogical word sense disambiguation (Barbella, 2013)

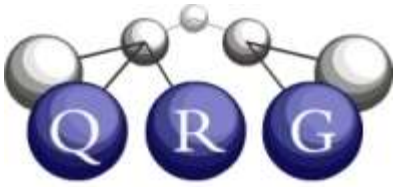

Crouse et al. AAAI 2018

Example: Geoquery

- Train: Q: “Which capitals are in states that border Texas?” A: “Baton Rouge, Little Rock, Oklahoma City, Santa Fe.”
- Train: Q: “What rivers are in Utah?” A: “Colorado, Green, San Juan.”
- Test: “What rivers flow through states that Alabama borders?”
- Test: “What are the cities in states through which the Mississippi flows?”
- Test: “What are the cities in states that border states through which the Mississippi flows?”

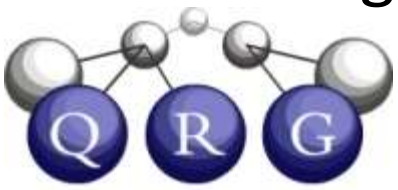

(Crouse et al. AAAI 2018)

Geoquery Results

	2-fold average (880)	10-fold average (880)
Ours-Uncorrected	81.5%	81.9%
Ours-Corrected	83.3%	83.2%

	250/250 split	680/200 split
Zettlemoyer & Collins, 2005	-	79.3%
Kwiatkowski *et al*, 2010	-	88.6%
Liang *et al*, 2011;2013	87.6%	91.4%

- Zetlemoyer & Collins / Kwiatkowski *et al* require annotated logical forms for training
- Liang *et al* 2013 uses a domain-specific lexicon with some manually defined lexical triggers

Learning Curve

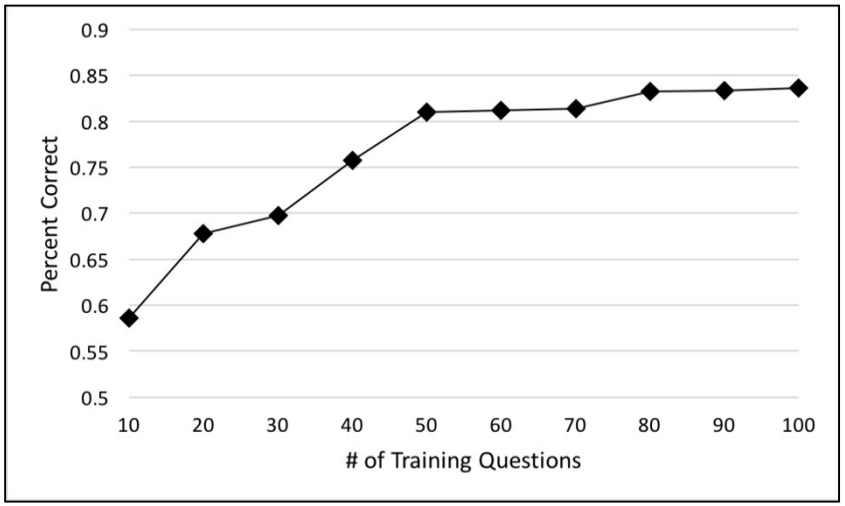

- 58.6% correct with only 10 training examples
- Comparable performance to 10-fold with only 50 questions (81.0%)

Other Tasks for Analogical Q/A

Identifying Processes in Text

- Example:
  Q: A student is investigating changes in the states of matter. The student fills a graduated cylinder with 50 milliliters of packed snow. The graduated cylinder has a mass of 50 grams when empty and 95 grams when filled with the snow. The packed snow changes to liquid water when the snow is put in a warm room. Which statement best describes this process?
  A: melting
- Learns to map from everyday terms to abstract relations in QP model
- 73% accuracy on corpus of science test questions

QR 2018

Answering Visitor Questions

- Back-end for interactive kiosk with vision and speech I/O
- Training takes only 11 Q/A pairs on average per question type

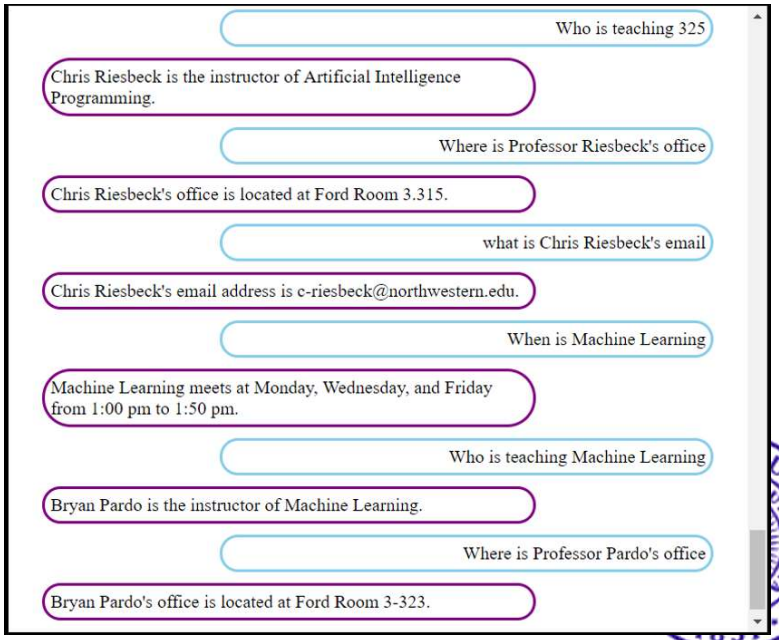

ProPara Dataset (Dalvi et al., 2018)

- Identify state changes in text about processes

  "Chloroplasts in the leaf of the plant traps light from the sun. The roots absorb water and minerals from the soil. This combination of water and minerals flows from the stem into the leaf. Carbon dioxide enters the leaf. Light, water and minerals, and the carbon dioxide all mix together. This mixture forms sugar (glucose) which is what the plant eats. Oxygen goes out of the leaf through the stomata."

- A system ought to be able to answer questions like "Where is sugar produced?" (A: In the leaf)

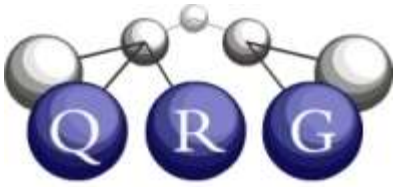

Step Semantics

- Framework for representing processes in text
  - Independent of task and data set
- Provides interface layer between language and conceptual knowledge
  - Needed because of incrementality of language
  - e.g. QP Frames, Kuehne 2004; McFate & Forbus 2016
- Step types are mutually exclusive
  - One sentence can give rise to multiple steps
  - One step can be described in multiple sentences

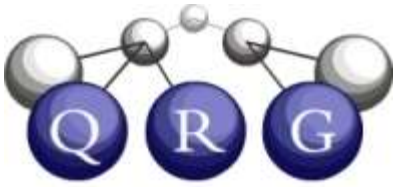

Steps Integrate Continuous and Discrete Representations

- Often intermingled in text
- Both can be used for the same phenomenon
  - "Cook the roast" versus "Wait until the temperature of the roast rises to 160° F"
- Continuous changes captured by QP theory
  - Prior work in QR suggests that people have models encompassing multiple perspectives and time-scales (e.g. Falkenhainer & Forbus, 1991; Rickel & Porter, 1994)
- Discrete changes captured by OpenCyc events, Frame Semantics, STRIPS operators

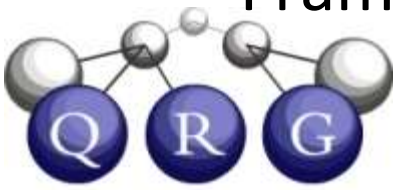

Four Kinds of Steps

1. Change of existence
   - Creating/destroying something
   - Example: When a puddle evaporates, it no longer exists
2. Change of property
   - Discrete relationship or attribute value
   - Example: Movement changes location, painting changes color

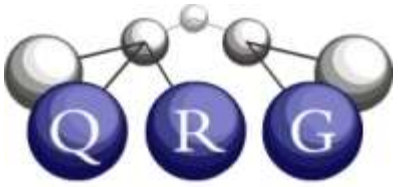

Four Kinds of Steps, continued

3. Change of quantity
   - A quantity changes over an interval
   - Example: The tide is rising
   - The continuous processes causing this are often implicit
4. Subprocess occurrence
   - Processes treated as atomic with respect to the current process being described
   - Example: The water cycle is described in terms of evaporation, condensation, and preciptation

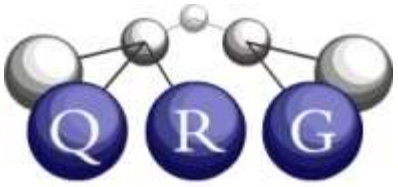

Recognizing Steps in Language

- Partial account, based on FrameNet (FN) and OpenCyc (OC)
- Example: Creation Steps
  - FN_Creating, OC: CreationEvent;
    FN_Giving_birth, OC: BirthEvent
  - Lexemes include create, assemble, form, formation, generate, make, produce, …
- See paper for the rest

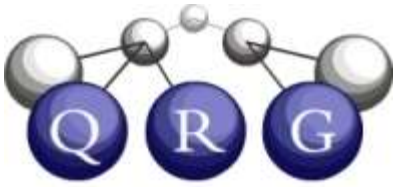

ProPara - State Change Grid

- State changes are hand annotated by Turkers Grid keeps track of the state of each entity (Participant) through time
  - " - " = Entity doesn't exist
  - " ? " = Entity exists but its location is unknown.

		Participants:				
Paragraph (seq. of steps):		**water**	**light**	**CO2**	**mixture**	**sugar**
	state0	**soil**	**sun**	**?**	**-**	**-**
Roots absorb water from soil						
	state1	**roots**	**sun**	**?**	**-**	**-**
The water flows to the leaf.						
	state2	**leaf**	**sun**	**?**	**-**	**-**
Light from the sun and CO2 enter the leaf.						
	state3	**leaf**	**leaf**	**leaf**	**-**	**-**
The light, water, and CO2 combine into a mixture.						
	state4	**-**	**-**	**-**	**leaf**	**-**
Mixture forms sugar.						
	state5	**-**	**-**	**-**	**-**	**leaf**

Time

Image: Dalvi et al., 2018

System Overview

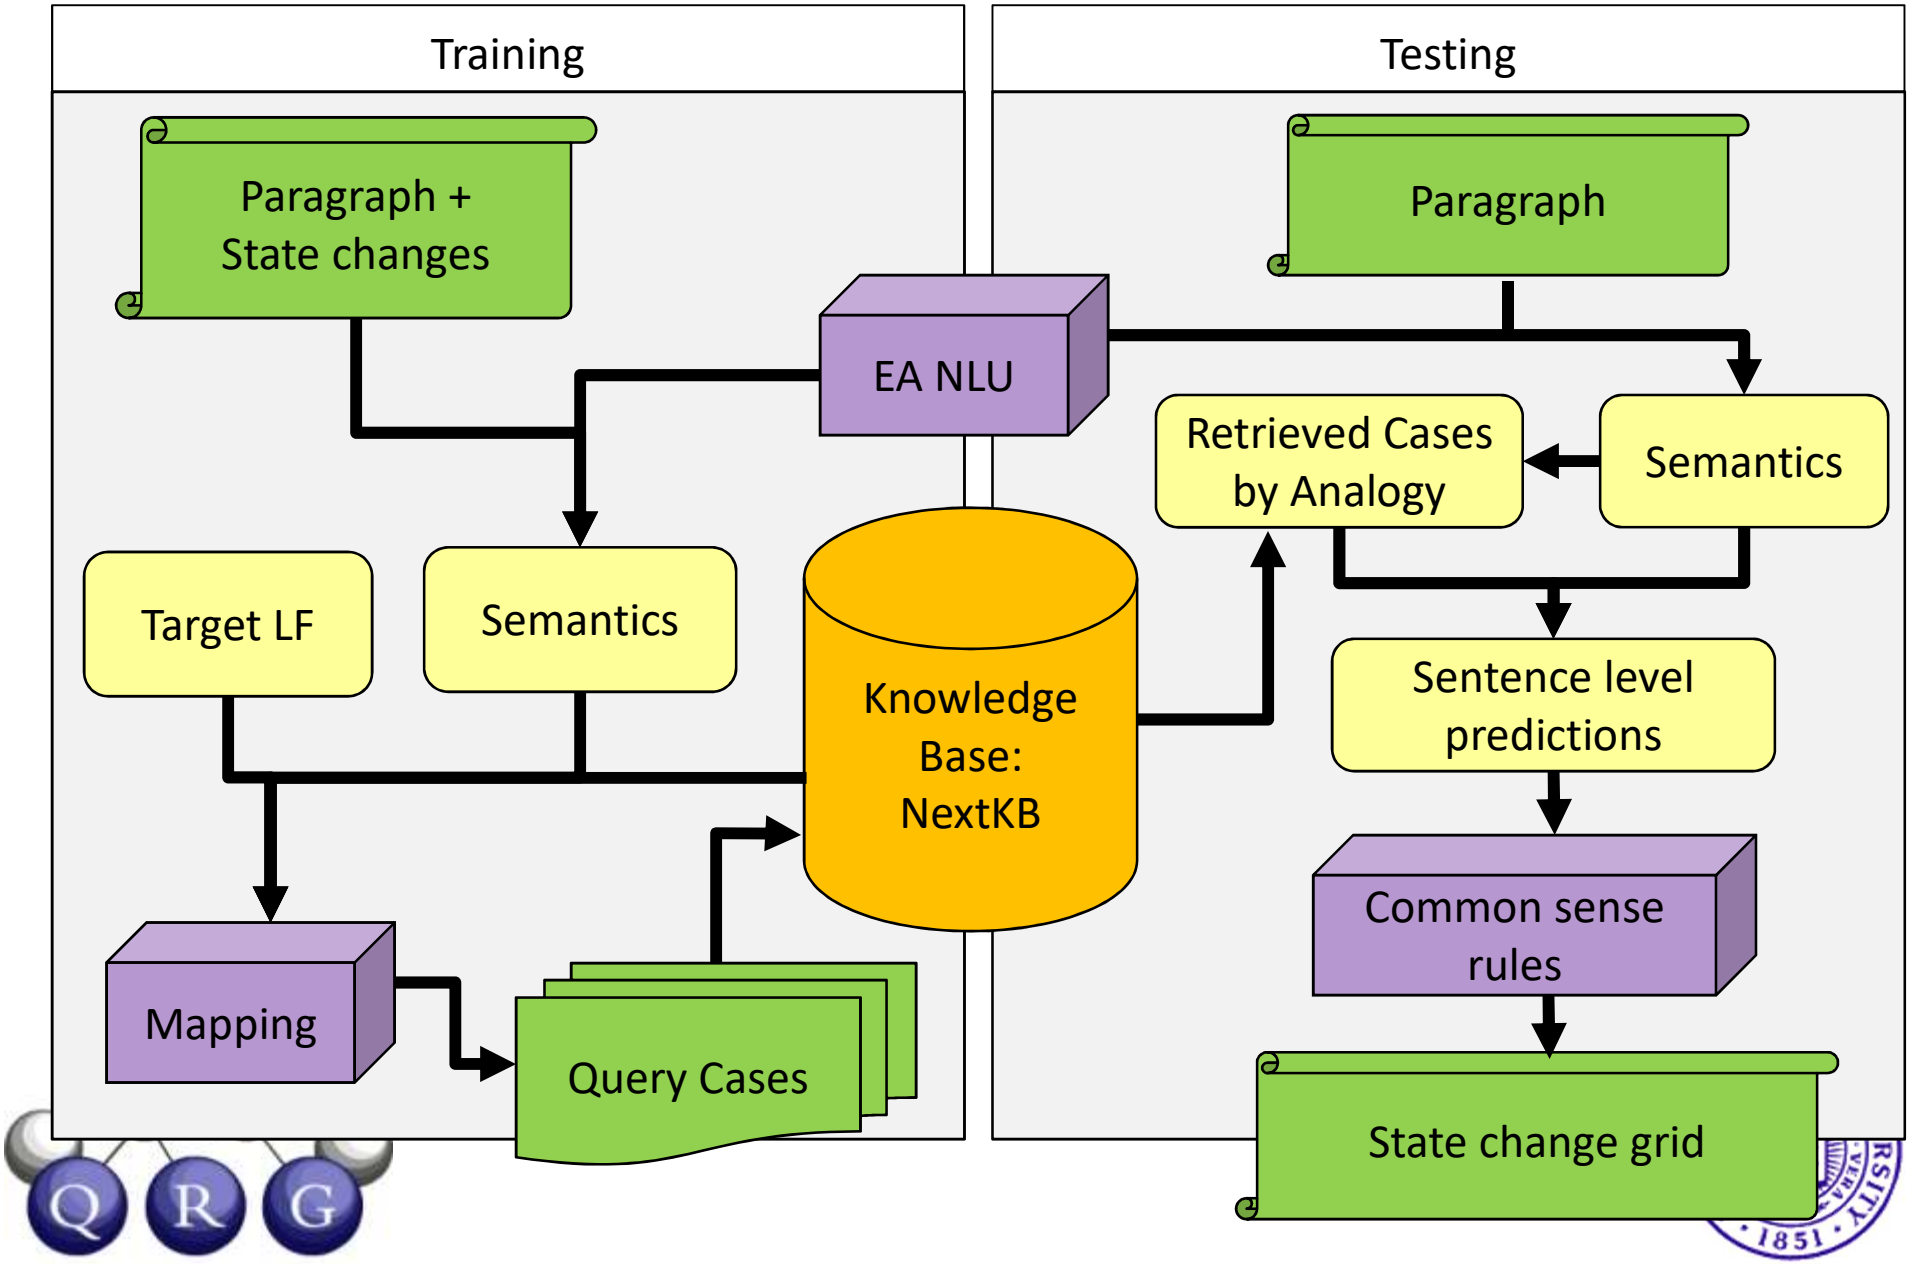

Training

- Create mapping between outputs of NLU system and target logical form
- Add to semantics information about **participants** and **locations**
  - Participants and locations identified by partial string matching, e.g. "recycle bin" = "recycling bins"
- Target logical forms, from Step types

Creation Event

```
(isa ?participant Participant)
(isa ?event CreationEvent)
(outputsCreated ?event
                ?participant)
(outputsCreatedLocation ?event
                        tolocation)
```

Movement Event

```
(isa ?participant Participant)
(isa ?fromlocation Location)
(isa ?tolocation Location)
(isa ?event MovementEvent)
(objectMoving ?event
              ?participant)
(fromLocation ?event
              ?fromlocation)
(toLocation ?event ?tolocation)
```

Destruction Event

```
(isa ?participant Participant)
(isa ?event DestructionEvent)
(inputsDestroyed ?event
                 ?participant)
```

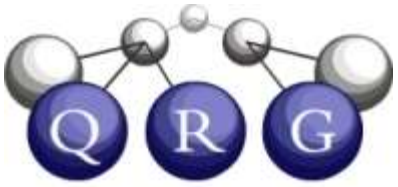

Constructing Mappings

- Uses Structure-Mapping principles with interleaved re-representation
  - Optimized via hill-climbing
  - Steiner tree used to bridge disjointed variables
- Re-representation uses NextKB ontology
  - Mapping between concepts:
    { AbsorptionEvent → TransferIn → GeneralizedTransfer → MovementEvent }
  - Mapping between role relations:
    { stuffUsed → EventOrRoleConcept → objectMoving }

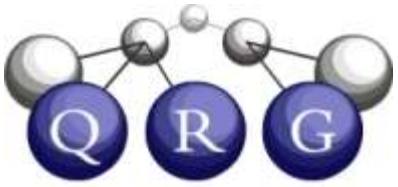

Query Cases

- Provide Step Semantics interpretation of new text
- Applied via analogical retrieval and matching

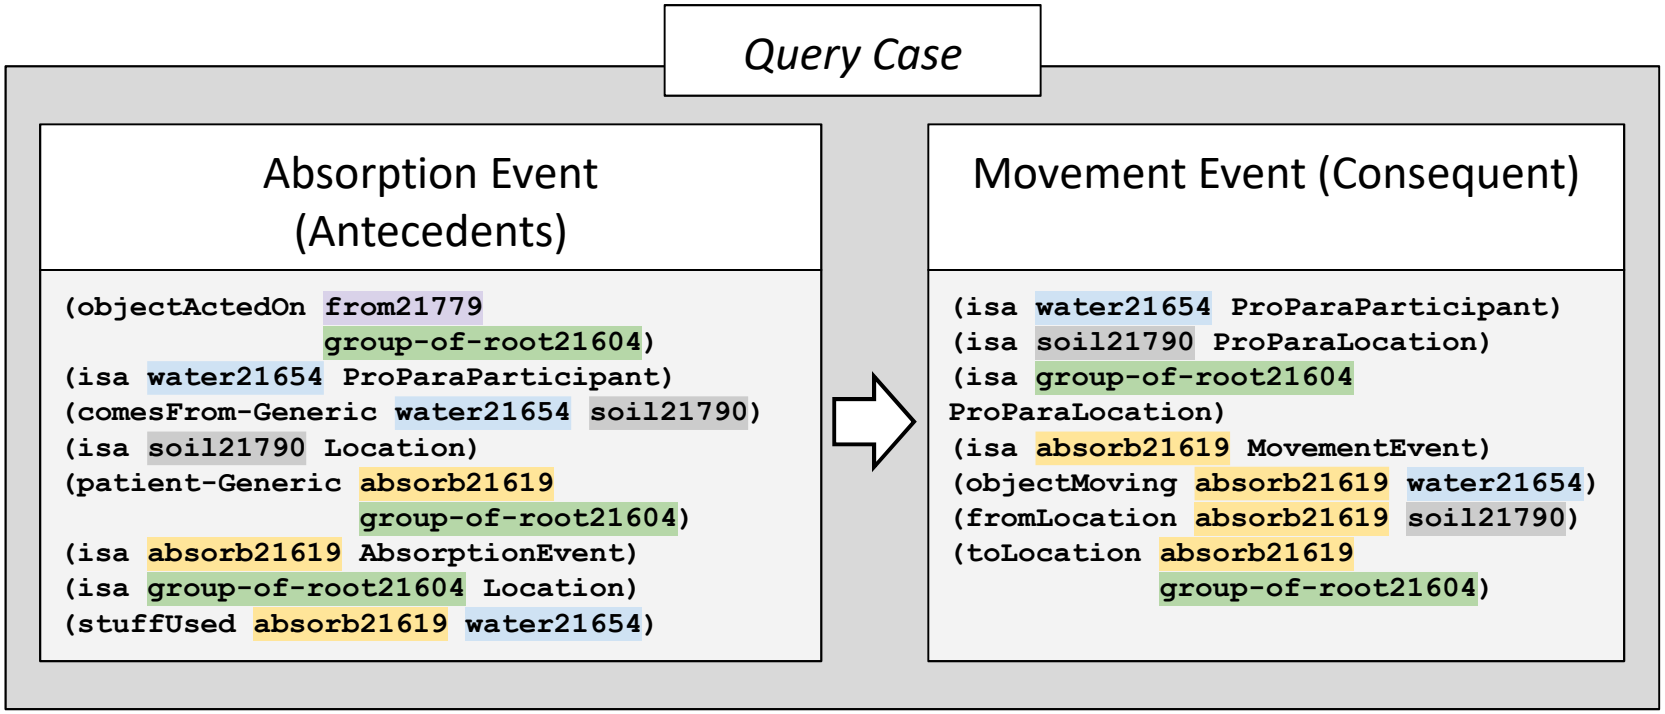

Applying Query Cases

1. Retrieve cases using MAC/FAC
2. Use SME to match query cases to sentence semantics, computing a score for each case
3. Select query case with highest score and generate state change from consequent
4. Use common sense rules to fill out the output state change grid
   - **Inertia:** states are propagated, both forward and backwards, until a new state change occurs
   - **Collocation:** If a participant X is converted to participant Y (X is destroyed when Y is created), & position of Y is unknown, then Y is assigned X's previous position

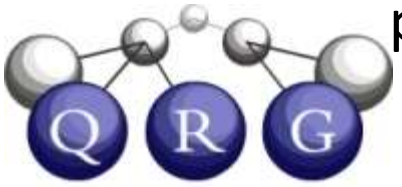

Other ProPara Work

- Rule based:
  - ProComp

NP V NP

EXAMPLE	"Martha carves toys."
SYNTAX	AGENT V PRODUCT
SEMANTICS	NOT(EXIST(START(E), PRODUCT)) EXIST(RESULT(E), PRODUCT)

IF (Agent "carve" Product)
THEN before: not exists(Product)
& after: exists(Product)

- Feature based:
  - Logistic Regression + CRF

- Artificial Neural Network based:
  - QRN
  - EntNet
  - ProLocal
  - ProGlobal
  - ProStruct
  - KM-MRC

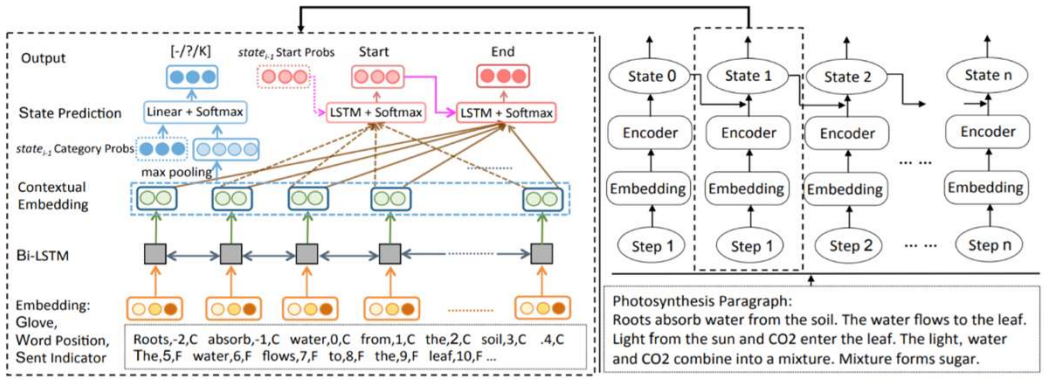

Images: Dalvi et al., 2018 and Clark et al., 2018

Sentence Level Evaluation (Dalvi et al., 2018)

- Cat-1: Is p created (destroyed, moved) in the process?
- Cat-2: When is p created (destroyed, moved)?
- Cat-3: Where is p created (destroyed, moved from/to)

	Model	Cat-1	Cat-2	Cat-3	Macro averaged
Rule Based	ProComp	57.14	20.33	2.40	26.62
Artificial Neural Networks	ProLocal	62.65	30.50	10.35	34.50
	ProGlobal	**62.95**	36.39	35.90	45.08
	KG-MRC	62.86	40.00	**38.23**	**47.03**
Analogical QA	Our Model	49.50	**43.92**	17.13	36.85

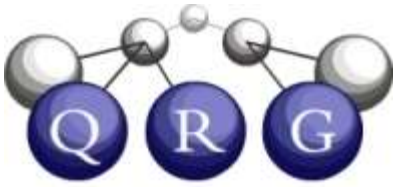

Document Level Evaluation (Tandon et al., 2018)

- Q1: What are the inputs to the process?
- Q2: What are the outputs of the process?
- Q3: What conversions occur, when and where?
- Q4: What movements occur, when and where?

	Model	Precision	Recall	F1 Score
Artificial Neural Networks	ProLocal	**77.4**	22.9	35.3
	ProGlobal	46.7	**52.4**	49.4
	ProStruct	74.2	42.1	53.7
	KG-MRC	64.5	50.6	**56.7**
Analogical QA	Our Model	65.2	43.0	51.9

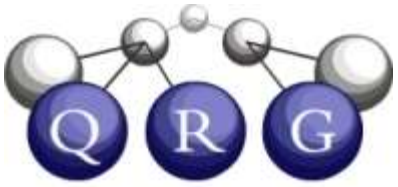

Conclusions

- Analogical Q/A provides a promising approach
  - Competitive performance on Geoquery, some aspects of ProPara
  - Can be quite data-efficient
- Same language system, multiple tasks
  - Evidence for the utility of the cognitive architecture approach to building AI systems

Future Work

- Continue improving grammar
  - Hand-engineering
  - Analogical learning of new constructions (McFate 2018)
- Explore more ways to exploit KB knowledge
  - Including incrementally expanding it via learning through tasks such as ProPara
- Extend the Step Semantics implementation
  - Test on learning by reading science texts
  - Test on other datasets (e.g. Recipes dataset)

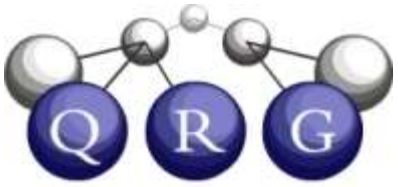

ML Needs to Evolve

- Train, Test method is holding us back
  - Organisms don't start from scratch for every task
  - Organisms learn incrementally, not in train/test cycles
- Understanding human-scale learning requires investigating cognitive architecture
  - People use rich relational representations, our learning systems should too
  - Ultimately frame their own learning problems, gather their own data, evaluate progress

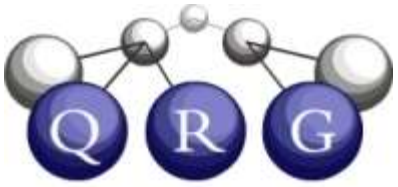

Details and Diversions

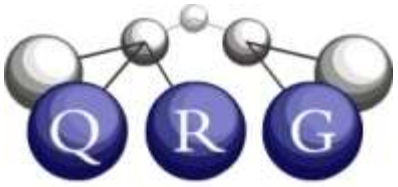

The Other Steps

- Property Change Steps
  - Lots of these, e.g. FN_Cause_changes, FN_Change_of_phase_scenario, …
- Quantity Change Steps
  - FN_Change_position_on_a_scale, FN_Change_of_temperature, …
  - Lexemes: rise, balloon, fluctuate, increase, heat, warm, cool, chill, refrigerate
- Subprocess/Event Steps
  - FN_Motion, FN_Fluidic_Motion, FN_Giving, …
  - Role relations describe changes in participants, e.g. “from” and “to” identify start and end locations

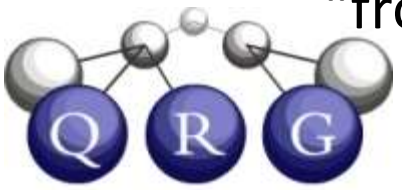

Companion-based Kiosk

- Physical installation in reception area of new Computer Science space in Seely Mudd Building
  - Display, array microphone, depth camera, speakers for multimodal interaction with people
  - September 2018 move-in date
- Goal: Provide useful information for students, visitors, and people working in the building

NextKB: A Resource for Cognitive Systems

- Goals: Broaden coverage, support collaborations and transitions
- Integrates OpenCyc ontology w/FrameNet, Verbnet
- Includes a broad lexicon
- Creative Commons Attribution licensing

	Comlex	Nulex
#Words	38,000	86,680
	Rcyc	**FrameNet**
VerbSemtrans	4,821	6,204
NounSemtrans	4,833	16,145
AdverbSemtrans	119	298
AdjectiveSemtrans	2,158	2,505
PrepSemtrans	194	147