Interacting with Intelligence

Phil Cohen

Chief Scientist, AI and SVP, Advanced Technology

VoiceBox TECHNOLOGIES

Are we stuck in a conversational local minimum?

Why?   How to get out

Voice AI

- Semantic Parsing
- Dialogue/plan-reasoning

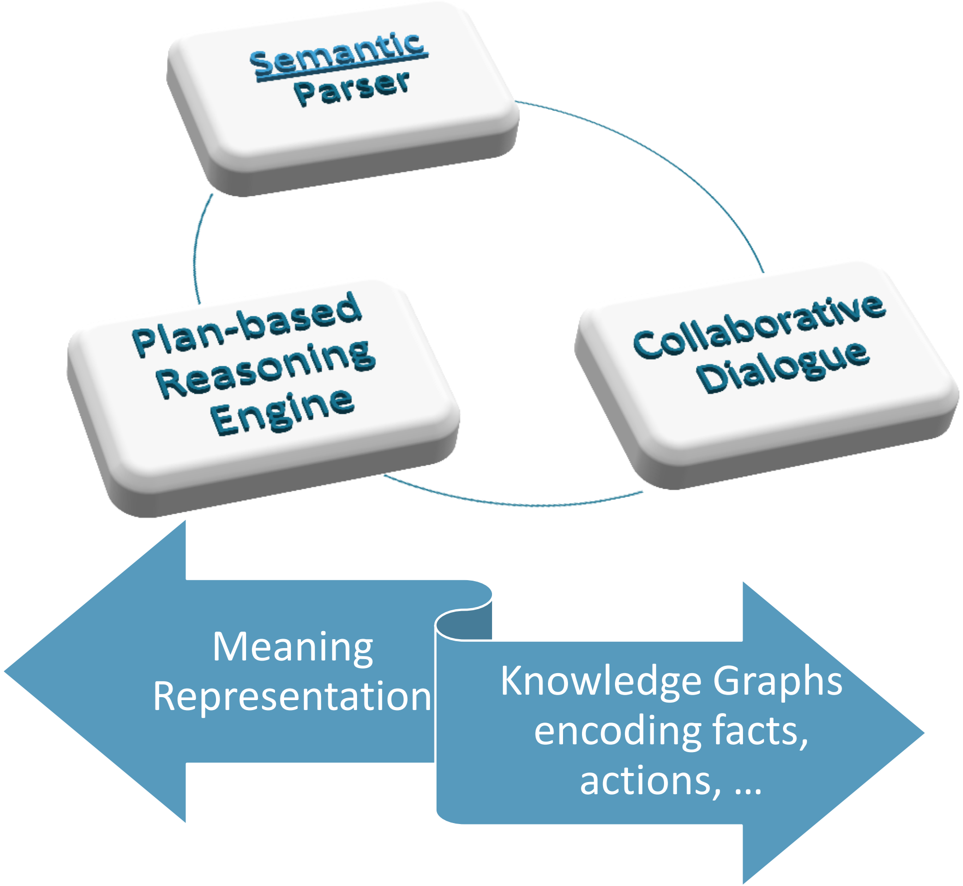

Current State of the Practice -- NLU

Statistical "intent" classification ~ Propositional Content

- Easy to train simple classifiers; nice toolkits
- Limited expressiveness to date
- Challenges
  - » Multi-"intent", multi-domain, multi-utterance turns
  - » Yes/no questions
  - » Anaphora
  - » Complex constructions
  - » Compositionality
  - » Speech act types expressing intentions other than commands/actions
- "Light is better here"

Semantic Parser

An old idea is new again

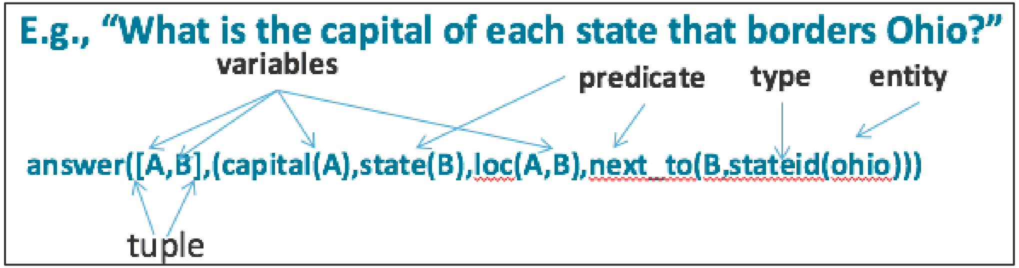

- Process that maps a sentence to a representation of its *semantics*, i.e., its *meaning*
- Meaning is represented as a ***logical form***, a logical language including:
  - » Entities, such as objects in the domain, events, variables, tuples,
  - » Relations such as predicate/argument structures, types, and
  - » Operators, including conjunction, quantification, superlatives, comparatives, aggregation, sequence, conditional, variables etc.

- Logical forms are *compositional* – meaning of whole is a function of meaning of parts
- Logical forms can be
  - » Vague, with pronouns and referential expressions resolved later via context
  - » Fused with LFs from other modalities
  - » Input to learning and inference
  - » Input to dialogue management subsystems
  - » Mapped to backend data sources,
  - » Executed to retrieve data or invoke APIs,

Demonstration

Geography Q&A System

Query

How long are the rivers that flow through each state that borders California?

Answer

Length	River	State
805	gila	arizona
1670	snake	oregon
1953	columbia	oregon
2333	colorado	arizona
2333	colorado	nevada

Logical Form:

answer([A,B,C],
(len(B,A),river(B), traverse(B,C),state(C),
next_to(C, california), state(california)))

Linguistic Constructions

Linguistic construction	Geo
Relative clauses	that border California
Comparatives	taller than Mount Shasta
Superlatives	tallest mountain
Negation	does not border California
Anaphora (pronouns)	their population which one …
Ordinals	What is the capital of the second one.
Ordinals + superlatives	What is the capital of the second largest state?
Ellipsis	Which is taller than shasta?
Quantifiers	the capital of each state that borders California
Yes/No Questions	Does Pennsylvania have a larger population than Ohio?
Conjunction	what is the capital and population of the states that border both Ohio and New York

More complex queries

Query:

How tall is the highest mountain in the state with the smallest population

Elevation	Mountain	State	Population
6194	mckinley	alaska	401800

```
answer([A,B,C,D],
(highest(B,
        (smallest(D,(state(C),
        population(C,D))),
  mountain(B),loc(B,C))),
elevation(B,A)))
```

More question-answering

More Examples

» **What is the capital of each state that borders Ohio?**

- answer([A,B],(capital(A),state(B),loc(A,B), next_to(B,stateid(ohio)), state(stateid(ohio)))

City	State
Charleston	West Virginia
Frankfort	Kentucky
Harrisburg	Pennsylvania
Indianapolis	Indiana
Lansing	Michigan

Yes/No Q's, comparatives

» **Is the tallest mountain in california higher than the tallest mountain in Alaska?**

answer(A,(highest(B,(mountain(B),loc(B,stateid(california)),
state(stateid(california)))),
highest(C,(mountain(C),loc(C,stateid(alaska)),
state(stateid(alaska)))),
higher(B,C)->A=yes;A=no))

YesOrNo	Mountain	Mountain
yes	mckinley	whitney

Other use cases

Complex Commands

Set the temperature to 70 degrees at 8am every other tuesday

Find the latest version of this file and send it to phil and john

Close all the shades except in the kitchen

Standing Orders

When Bob replies to this message, send his reply to my team

Turn off the sprinklers for twenty four hours whenever we get one inch of rain

Prohibitions

Don’t schedule meetings at lunch time

Assertions:

I don’t like text messages when I’m in a meeting

Hybrid Symbolic/Statistical Semantic Parser

	Development / Training	Test	Arbitrated
Top Priority	Symbolic (S) Statistical (St) **N= 3853**	84.2% 91.0% **N=980**	**94.3%**
Other Priority	Statistical **N= 3963**	S: 61.6% St: 81.9% **N=883**	**85.5%**

Challenges

- Where do the predicates come from?
- Relationship to Ontologies/KBs?
- What are the LF constructions?
- How to map various linguistic constructions to LF constructions ?
- How to map those predicates to backend APIs/DBs?
- Tooling
- Deep Learning
- Data collection

Dialogue

- Mostly undeveloped in current virtual assistants
- Finite-state, system driven, speech act sequences, authored, rigid
- Slot-filling (GUS (Bobrow et al., 1977), DARPA's Air Travel Information System (ATIS), 1990+)
- Task-oriented (Grosz, 1977)
- Plan-based (Allen, Cohen, Perrault, 1978-1980), Collaborative (Allen, Cohen, Grosz, Levesque, Sidner, 1990's)
- Most current research and practical dialogue systems are Slot Fillers.
  - » Why?
    - Easy to implement, uses statistical "intent" classification, learn policies for responding
    - "Light is better here",
- Slot-filling is not easily extended to other types of dialogues; should be a special case
- **What do we want, and how do we get there?**

Collaborative dialogue driven by plan recognition

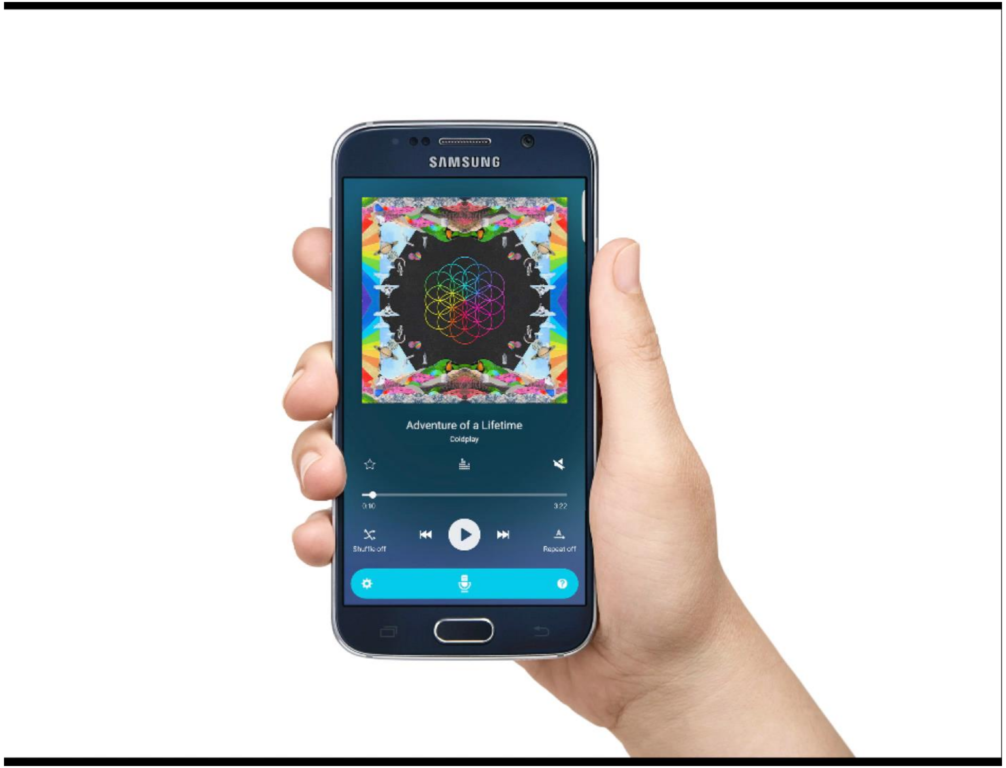

Please note:

- Contextual reference to song
- Referring to group as “their” and “they”
- System infers purpose of question is for user to see next concert in region
- System offers to buy tickets
- System informs user that plan will fail (venue is sold out)
- System finds another method to see concert
- System offers to buy tickets for user
- System proactively provides seating chart to get information needed to purchase ticket

Common Patterns of Plan Inference

Intent / Plan recognition

- Request(Action) → Action_1 → Effect_2 → Action_2… → Effect_n (top goal)
- Req(Time of Next concert) → (Date/Loc) → Go to Loc → Attend Concert → Hear musical group

Precondition failure

Have ticket

- Check preconditions
- On failure, find another plan to achieve goal
  - » Find substitute objects for goals, edit plan, confirm
  - » Other operations possible -- system plans to overcome obstacles; system plans to achieve goals
- Common precondition failures --
  - » business must be open to conduct commerce in person → find one that is open
  - » Item is in stock → find another business selling item, (e.g., concert venue)
  - » Cannot arrive on time for an appointment → make another appointment; change vendor

Use of Knowledge Base

- **Coldplay** is a musical group
  - *"Adventure of a lifetime"* is a song by Coldplay
- Musical groups play multiple live concerts
- Each live concerts takes place at a venue
- Venues sell tickets for concert events
- *Precondition*: To attend event, one must have a ticket
- There is a limited number of tickets
- Concerts are Entertainment Events, which are Events
- Events take place at times and locations
- People like to view events
- *Precondition*: To view an event, the viewer must be at the same place at same time
- Viewer must travel to the venue

…

Need not be first principles reasoning.