Carnegie Mellon University
Language Technologies Institute

Knowledge-Based AI for Common Sense and Language Understanding

Scott E. Fahlman

Professor Emeritus
Language Technologies Institute & Computer Science Department
School of Computer Science
Carnegie Mellon University

September 23, 2019

About Me…

- Undergrad, MS, and PhD at MIT, 1965-1977
  - MS thesis in flexible hierarchical planning.
  - PhD thesis on common-sense knowledge and inference, using an (imaginary) massively parallel machine.
  - Advisors: Patrick Winston, Gerald Sussman, Marvin Minsky
- On CMU faculty since 1978.
  - Worked with Allen Newell, Herb Simon, Raj Reddy
  - From about 1981 to 1996, worked on artificial neural nets, collaborating at times with Geoff Hinton. Funding collapsed in US.
- Ran Justsystem Pittsburgh Research Center, 1996-2000.
- Back at CMU, working on knowledge-based AI ever since.
- Now Professor Emeritus (formally retired). Writing a book.

Four Splits in the AI Community

1. Super-Human vs. Normal-Human Problems and Abilities.
2. Creeping Incrementalism vs. “What will it really take?”
3. ML-Based (“Empirical”) vs. Symbolic KRR
4. Neat (Mathematical, Elegant) vs. Scruffy

What is “Knowledge-Based AI”?

- AI used to be about knowledge and reasoning.
  - Entities (nodes), relations/statements (links), rules for inference.
  - Some assumed this meant formal logic and theorem-proving, and many still do.
    - BIG mistake, in my opinion. Slowed down progress.
- Mid-1980’s: People got impatient, and realized that ML methods could get 80% solutions with little human effort. So symbolic AI “failed”.
- Symbolic, non-statistical AI is now a small minority of the AI world, but a few of us think it will be essential for human-level intelligence.
  - What to call it?
    - “Good old-fashioned AI” or GOFAI
    - “Neoclassical AI” (My idea, but it hasn’t caught on.)
    - “Knowledge-Based AI”
    - Note: Not “Logic-Based” or “Rule-Based” AI, unless that’s what you are really doing.

Big Data to Big Knowledge

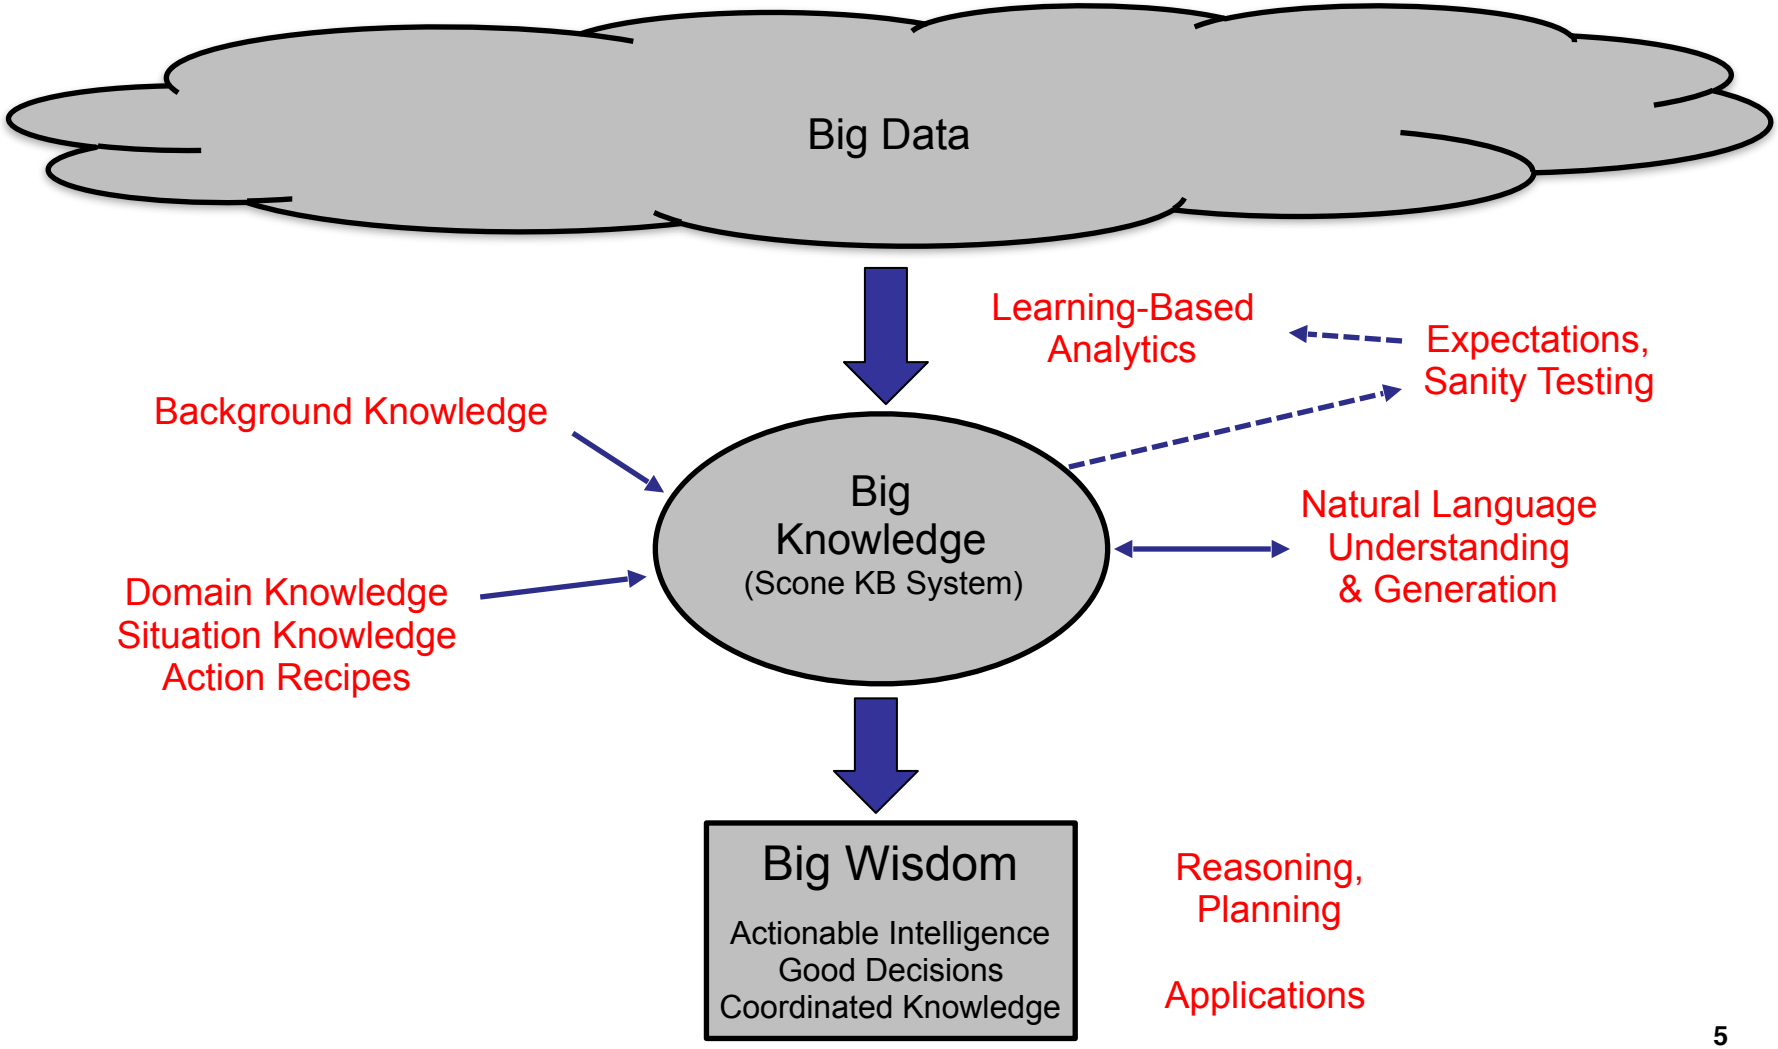

What do we mean by Common Sense?

What is the most important open question in AI?

"Why is the Big Bad Wolf hiding behind the tree, and what is likely to happen when the little girl reaches that point in the path?"

- There's a LOT going on here.
- But any reasonably bright four-year-old can answer this, without any obvious mental effort.

The Scone Knowledge-Base System

Open Source KB System, developed at LTI.

Scone is a “smart memory” module, supporting fast search and reasoning.

- Not a stand-alone AI system.
- Supports efficient search and inference.
- Human-like common-sense reasoning, not theorem proving.
- Scalable to millions of entities and facts on a single-core workstation.

Key Ideas in Scone

- Clean separation between {concepts} and “lexical items”.
- Semantic net with nodes (entities) and links (relations).
- Multiple inheritance with exceptions.
- Expressiveness beyond first-order logic.
  - Allows general statements with exceptions, and higher-order logic (statements about statements).
- Very efficient search and inference.
- Scale up to millions of entities and statements on a serial workstation.

We must give up the idea that inference is theorem-proving. Scone uses “common sense” reasoning, implemented via marker-passing.

Key Concept: Virtual Copy

Suppose I tell you **"Clyde is an elephant".** You now know a LOT about Clyde. I didn't tell you any of that – you're using background knowledge and inference.

- Concepts are mostly not defined, but based on typical-member description.
- A concept node like {elephant} represents an *entity*, and is also the anchor for a *description:* Roles and fillers, statements, IS-A links…
- {Clyde} IS-A {elephant} says that Clyde inherits the typical-elephant description, which inherits {mammal}, {animal}, {physical object}…

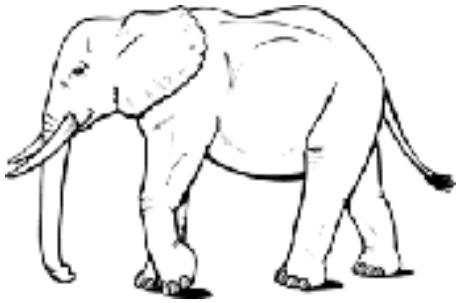

Key Concept: Virtual Copy

- We don't actually copy the inherited descriptions. The Scone inference engine makes it behave as if we did.
- Only create a *real* node for entities that have some new info.
- Each concept can add some new info, specializing what it inherits from its superior.
- And it can subtract some things that would otherwise be inherited.
- This key piece of inference is highly optimized, parallelizable.

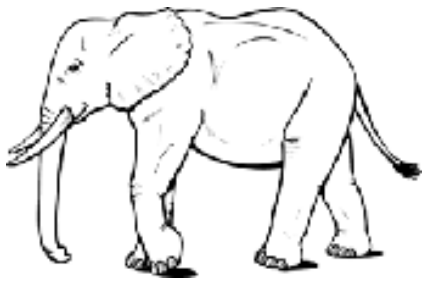
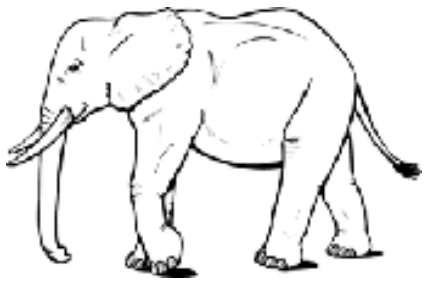
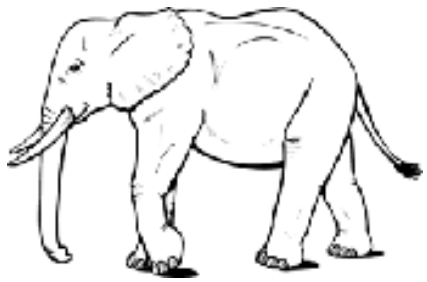

Multiple Contexts

- A *context* is a world view or model within the KB, represented as a regular entity in the KB. {Pittsburgh} is a {city}, but also acts as a context.
- Every entity exists in some context; every statement is active in some context.
- All KB changes, queries, inferences, are with respect to some context.
- "In …" means I'm about to name a context, and that is the active one until further notice. In Pittsburgh, in the 1980s, in my opinion…
- Each context starts off as a clone of some other one, from which it inherits all info. No need to copy.
- Once the context is created, we can add new info and cancel things that would otherwise be inherited.
- Efficient and lightweight, so we can use these freely.

Multiple Contexts

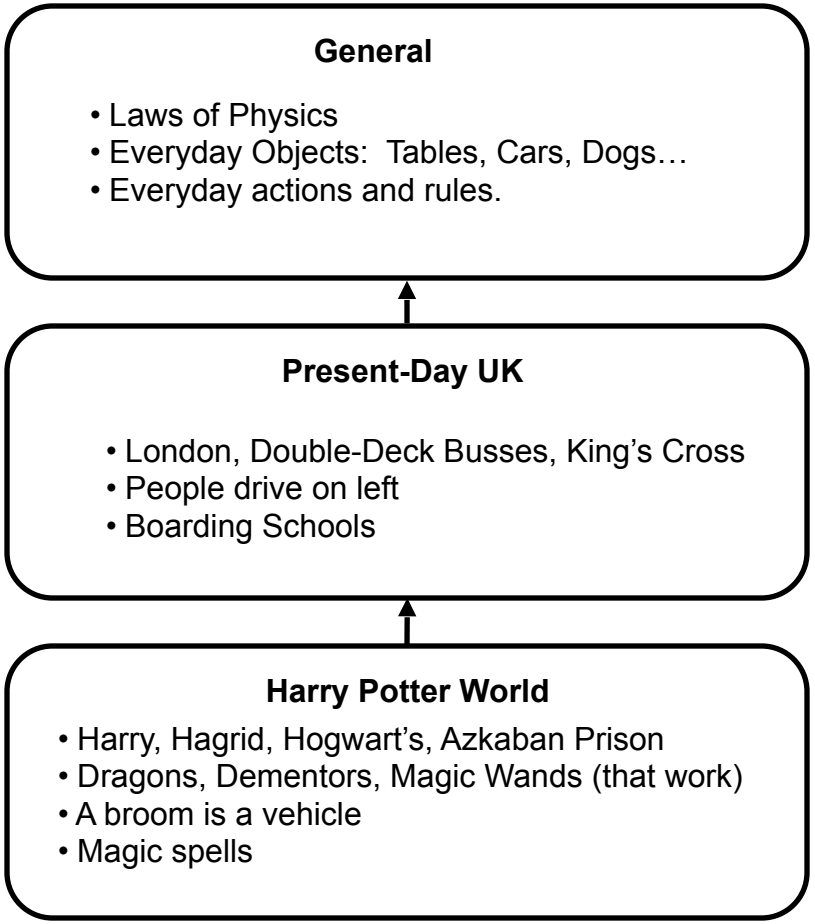

Uses of Multiple Contexts

- Different belief sets and states of mind.
- Model “what if” scenarios.
- Alternatives being considered in planning or analysis.
- “Quarantine” inputs or inferences that are not yet accepted.
- Things that are true at different points in time, or different places (real or fictional).
- Model the world-state after each step in a plan.
- Model knowledge or belief states of different people.
- Attach meta-knowledge to the entire context.

Natural Language: It's All About Meaning

Consider how much background knowledge is required to understand a sentence like this:

> *"In response to last Tuesday's café bombing in Tel Aviv, Israeli tanks today crossed into the Gaza strip and demolished the houses of several suspected Hamas leaders before retreating under a hail of rocks and Molotov cocktails."*

Queries:

- Articles about "revenge"…
- Articles about "angry Palestinian mob"…

Bag-of-words search won't find this, even with synonyms.

Can "distributional semantics" or word-vectors do this?

We Need NLU, not Meaning-Free NLP

- In the past 30 years there has been a gold rush in meaning-free statistical NLP. (Google, text mining…)
  - The amazing thing is that a bag-of-words model works at all.
  - These meaning-free techniques can do valuable tasks, but they will always be limited.
- To take the next big step – real natural-language *understanding* – we must return to a meaning-based approach.
  - That requires a really good knowledge base.
  - Yes, this is difficult, but the problem isn't going away.

Current Scone Activities

- **Reduce Barriers to entry:**
  - Writing a tutorial book: *Knowledge-Based AI: Common Sense, Language Understanding, and the Meaning of Meaning.*
  - Code cleanups & extensions.
  - Extend the core knowledge bases via mining.
- **Natural Language Understanding and Generation in Scone**
  - Inputs, queries, commands, responses, explanations…
  - We won‘t have enough general knowledge until we can read human language.
- **Episodic Representation & Reasoning**
  - Actions, events, plans, time durations, concurrency…
  - Recipe-based planning
  - Add language about actions and plans

Other Scone-Related Applications

- Use Scone as a front-end or reference librarian to access other knowledge bases, data bases, tables on the web, etc.
- Model complex hardware-software systems and maintenance procedures.
- Situation awareness for large-scale disaster management.
- Ontologies for biology, medicine, world events…
- Games.

I believe that Scone can be used in more ways than DBs are today.

The End

To learn more, until the book is done, look at my Knowledge Nuggets blog, hosted on Quora.com:

Just Google “Fahlman knowledge nuggets”