



What Kinds of Knowledge Are Needed for Genuine Understanding?

*Lenhart Schubert
University of Rochester*

NLP in the 60's to mid-80s

KEY INSIGHT:

Need large amounts of knowledge for genuine understanding/ inference/ problem solving.

*"John had lunch at Mario's.
The hamburger he ordered was tasteless,
so he left a small tip."*

*Was John given a hamburger? Did he eat it? Where?
Who gave it to him? Was he satisfied with the food?
Why a small tip? For whom? Did he leave Mario's?*

PERCEIVED PROBLEMS

- **PARSING ACCURACY & COVERAGE**

*"John observed the bird with {binoculars/ yellow tail feathers/
delight/ fellow birders/ its brood of chicks/ ...}"*

- **KNOWLEDGE ACQUISITION BOTTLENECK**

- **ontological knowledge** (Binoculars are a viewing instrument)
- **well-known facts** (Microsoft is a computer software/hardware company)
- **general world knowledge** (Restaurants serve meals to customers for pay).

- **SR/KR EXPRESSIVITY, CONGRUITY WITH LANGUAGE, AND REASONING POWER**

- *in the 80's and 90's this faded more and more from the NLP community's purview*

THE (WIDELY) PERCEIVED SOLUTION:

MACHINE LEARNING (INSTEAD OF MANUAL CODING)

- *statistical, corpus-based techniques*
- *neural nets*

Both exploit the distributional properties of language (in large data sets)

"Time flies like an arrow"

***[N time] [V flies] occurs very often,
[N time] [N flies], and [V time] [N flies] do not.***

SUCSESSES: Where shallow methods suffice...

- *speech recognition*
- *learning probabilistic grammars*
- *machine translation*
- *document retrieval*
- *text-based QA and summarization (in a sense)*
- *sentiment and topic analysis, etc.*

The problems of genuine understanding, general dialogue, reasoning, and problem-solving remain (almost) as challenging as ever!

"Peripheral" advances: speech recognition, input parsing (to some extent), corpus-based factoid & rule gathering, evaluation methodologies

KEY CHALLENGES:

- **A SEMANTIC REPRESENTATION/ KNOWLEDGE REPRESENTATION (SR/KR) COMMENSURATE WITH NL**
- **EFFECTIVE, SCALABLE "SOFT" AND "FIRM" INFERENCE METHODS**
- **TENS OF MILLIONS OF PATTERN-LIKE, SCHEMA-LIKE, AND AXIOM-LIKE KNOWLEDGE ITEMS**

Why we can't depend on ML to yield a SR/KR commensurate with NL:

Machine Learning can't "invent" a meaningful representational framework; it can only learn what you provide examples of (or build into latent variables):

- *if you annotate data with class labels, ML can produce class labels;*
- *if you annotate sentences with parse trees, ML can produce parse trees;*
- *if you annotate sentences with translations, ML can produce translations (and latent syntax transduction rules, if we provide a syntax for these);*
- *if you annotate English db queries with answers, ML can produce answers via formal queries (if the syntax of those queries has been supplied);*
- *if you annotate sentences with "AMR"s, ML can produce "AMR"s;*
etc.;

Why we can't depend on ML to lead directly to human-like reasoning:

Reasoning requires knowledge in an interpretable, modular form

An Example of genuine understanding and reasoning

*Alicia drove home from her holiday visit to the US.
What border(s) did she cross in her drive?*



- Visiting a place generally entails having one's current place of residence away from that place, going to it, staying there in the expectation of some immediate rewards, and then leaving it, typically going back to one's place of residence.
- “Home” is a person's permanent place of residence.
- The US is a country, and leaving a country entails crossing its border.
- Driving takes place on land.
- The only land borders of the US are with Canada and Mexico.

Along with lexical and paraphrase knowledge →

*Alicia crossed the border with Canada OR Mexico
(and perhaps others, e.g., from Mexico into Guatemala).*



Also, it's quite possible that *Alicia is a resident of Canada or Mexico*, and that *she is a citizen of one of those countries*.

A (still) more demanding example

When Randy learned that his metastatic cancer was terminal, he decided to stop receiving chemotherapy and enrol in a hospice program designed to provide palliative care.

Why did Randy decide this?

- **When you learn something, you then know it;**
(so Randy knew he had terminal metastatic cancer);
- **When someone learns (or knows) that something is the case, it is indeed the case;**
(so Randy had terminal metastatic cancer).
- **Terminal metastatic cancer is usually fatal within a few months;**
(thus Randy was destined to die within a few months);
- **Chemotherapy is intended to combat cancer and prolong life, but results in suffering and weakness;**
- **No medical treatment will significantly prolong a terminal cancer patient's life; chemotherapy is a medical treatment;**
- **Stopping something entails it was going on;**
(so Randy had been receiving chemotherapy, and therefore had endured suffering);
- **Receiving palliative care reduces pain and discomfort;**
- **Enrolling in a program to provide certain services will lead to receiving those services;**
(so Randy would receive palliative care);
- **If Randy continued receiving chemotherapy, he would endure further suffering without significant life extension, while palliative care would make him feel better;**
- **People act to optimize expected rewards and minimize suffering;**

so if Randy knew all of the above, his choice of palliative care instead of chemotherapy is explained.

Example continued

How did Randy know all this?

- Commonsense inferences that "jump out" at me also "jump out" at others (who possess the same premise knowledge) – *“simulative inference”*
- All the items above are common knowledge (especially to cancer patients), so the indicated inferences were just as obvious to Randy as they are to me;
- So the presumption previously underscored (Randy knew/inferred all those facts) surely holds, and *Randy's choice is explained*.



(Again, we're also presupposing a lot of lexical & paraphrase knowledge: “terminal” illness, “fatal” condition, “medical treatment”, “palliative care”, “optimize”, “life extension”, “suffering”, etc.)

Expressive devices encountered in the two examples:

- Predication, of course (Alicia *visited* the US, the US is a *country*, and so on)
- Temporally related events (driving somewhere, staying, returning) and their implicit consequences (crossing borders, enjoyment)
- Causal relations (chemotherapy causes suffering, possible life prolongation)
- Geographic relationships (perhaps dependent on "mental maps")
- Conjunction and disjunction (the US borders on Canada *and* Mexico; Alicia drove into Canada *or* Mexico)
- Negation (if you leave a place, you are then *not* in that place; *not* receiving significant benefits from chemotherapy)
- Quantification (*almost everyone* knows that chemotherapy is used to treat cancer; *no medical treatment* will significantly prolong a terminal cancer patient's life;

And that's not all ...

Further demands on expressivity

- Genericity -- almost all the general knowledge listed has the character of applying in “*typical*” cases, *but allowing for exceptions*;
- Patterns of behavior (*holiday trips; the usual course of events if you get cancer*); this is of course *generic* knowledge;
- Modal/mental notions (*learning, knowing, expecting, deciding, inferring, intending, ...*)
- Counterfactuals (deciding against something because of the adverse consequences it *would* have);
- Uncertainty (Alicia's *probable* home country, Randy's life *expectancy*)
- Predicate modification (*terminal metastatic* cancer, *palliative* care, *continuing/stopping* receiving chemotherapy, *feel* better)
- Sentence modification (*probably*, Alicia lives in Canada or in Mexico; *possibly* she lives in Guatemala)
- Predicate reification (*driving* takes place on land; *palliative care* reduces *suffering*)
- Sentence reification (Randy learned *that his metastatic cancer was terminal*)

Are any of these devices peculiar to English?

These devices are available in all languages!

A cultural accident, or a reflection of our “mentalese” semantic types?

- Language and thought are conjectured to have arisen concurrently (e.g., Ian Tattersall, Derek Bickerton, Philip Lieberman, and many others);
- Humans occupy the “cognitive niche” (Pinker) because they can learn, store, manipulate for inference, & communicate meaningful symbolic patterns -- *semantic / knowledge representations*
- Richard Montague showed that there is a tight relationship between linguistic structure and meaning.

So, surely it is reasonable to view language as a mirror of mind; i.e., “mentalese” as language-like, and as richly expressive.



Candidates for semantic representation and knowledge representation

(discussed & critiqued in Schubert, AAAI '15; criticisms here are “caricatures”)

- ***FOL, DRT*** (e.g., Allen, Jurafsky & Martin, Kamp, Heim, J. Bos, ...) [expressively inadequate]
- ***Semantic nets*** (wide spectrum: Shapiro, Sowa, Schubert, ConceptNet, ...); [any notation can be cast as SN]
- ***Description logics*** (CLASSIC, LOOM, OWL-DL, KAON, SROIQ, ...) [expressively inadequate]
- ***Conceptual meaning representations*** (Schank's CD, Jackendoff, FrameNet, ...) [expr./infer. inadequate]
- ***Thematic role representations*** (e.g., Palmer, Gildea, & Xue '10) [expr./infer. inadequate]
- ***Abstract meaning representation*** (AMR) (Banarescu et al. '13) [inferentially inadequate]
- ***Hobbs' "flat" representation*** (Hobbs '06) [conflates distinct types]
- ***Structured English*** (e.g., MacCartney & Manning '09, Dagan et al. '08, ...); [ambiguous, inferentially inadequate]
- ***Montague Grammar*** (English as logic) (e.g., Dowty '79, Chierchia & McConnell-Ginet '00); [unnecessarily complex, higher-order]
- ***Extensional Montague fragments*** (e.g., McAllester & Divan '92, Artzi & Zettlemoyer '13); [expressively inadequate]
- ***DCS trees*** (Liang et al. '11); [expressively inadequate]
- ***Situation semantics*** (Reichenbach '47, Barwise & Perry '83) [abstruse, inferentially inadequate]
- ***Episodic Logic*** (EL) [remedies many of these flaws; still brittle, lacks adequate KB]

Episodic Logic (EL) – Montague-inspired, first-order, situational, intensional

A language-like, first-order SR/KR with a small number of types (but abstract individuals);
 LFs are obtained compositionally from phrase structure trees;

Handles most semantic phenomena shared by NLS; all types of referents;
 allows complex situations/events, with temporal/causal relations

John donated blood to the Red Cross.

[John1 <past donate.v> (K blood.n) Red_Cross] {initial representation}

(some e: [e before Now3] (some x: [x blood.n] [[John1 donate.v x Red_Cross1] ** e]	Skolemize, split	[E1.sk before Now3], [Blood1.sk blood.n], [[John1 donate.v Blood1.sk Red_Cross1] ** E1.sk]
--	------------------	---

Very few people still debate the fact that the earth is heating up {final representation}:

[Fact4.sk fact.n],
 [Fact4.sk = (that (some e0: [e0 at-about Now0]
 [(The z [z earth.n] [z heat_up.v]) ** e0])),
 ((fquan (very.adv few.a)) x: [x (plur person.n)] (still.adv (I v [v debate.v Fact4.sk])))]

Unfortunately, faulty/underspecified syntactic analyses lead to faulty LFs.

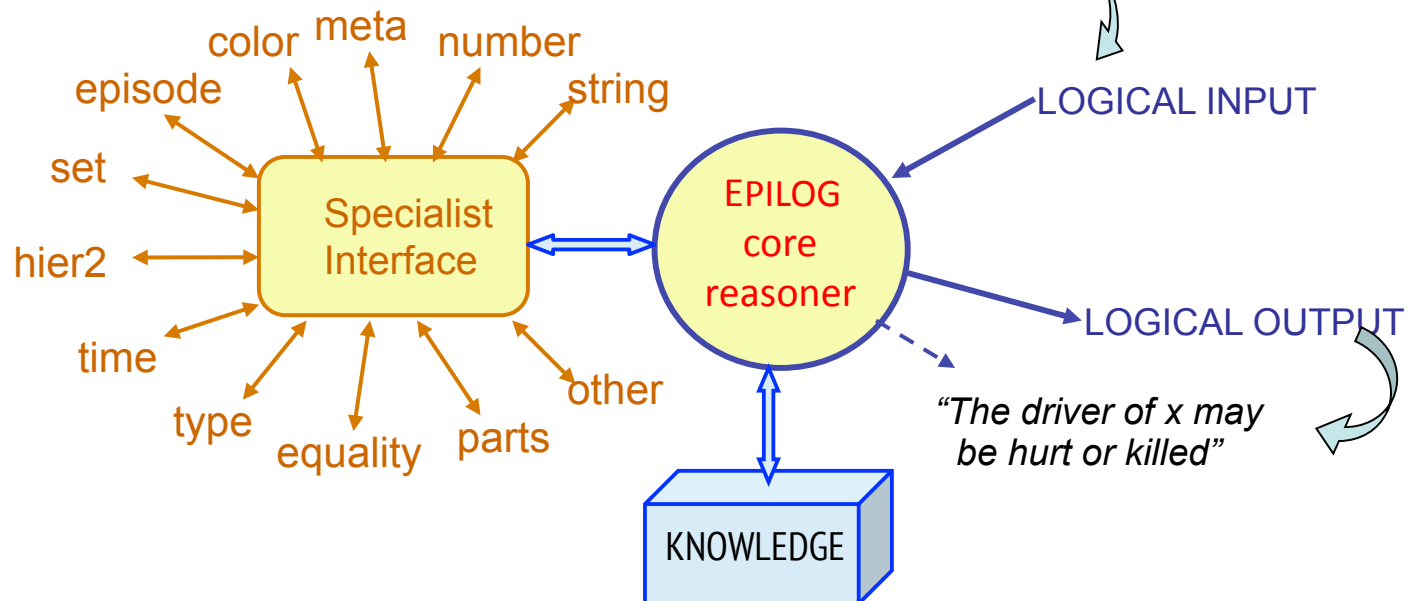
We don't have enough (reliable) semantic pattern knowledge.

The EPILOG system for Episodic Logic

(L. Schubert, C.-H. Hwang, S. Schaeffer, F. Morbini, Purtee, ...)

“A car crashed into a tree. ...”

(some e: [e before Now34]
(some x: [x car] (some y: [y tree]
[[x crash-into y] ** e]]))



Expressive richness does not impede inference:

EPILOG 2 holds its own on large FOL problems (Morbini & Schubert '09);

but inference remains brittle; uncertainty handling is heuristic; KB remains inadequate

Some simple EPILOG Inference Examples

Example using “most” (from J.F. Allen’s ‘Monroe emergency’ domain):

Most front loaders are currently in use;
Whatever equipment is in use is unavailable;
Front loaders are equipment →



Therefore, most front loaders are currently unavailable.

(most x: [x front-loader] [[x in-use] @ Now3]);
(all x: [x equipment] (all e [[x in-use] @ e] => [(not [x available]) @ e]]));
(all x: [x front-loader] [x equipment]) →

(most x: [x front-loader] [(not [x available]) @ Now3])


Example involving attitudes (English glosses of EL formulas)

Alice found out that Mark Twain is the same as Samuel Clemens.
When someone finds out something, s/he doesn't know it to be true at the beginning (of the finding-out event) but knows it to be true at the end.



Therefore:

- Alice didn't know (before finding out) that Mark Twain is the same as Samuel Clemens;
- Alice knew afterwards that Mark Twain is the same as Samuel Clemens;
- Mark Twain is the same as Samuel Clemens (from knowing-axioms).



***EPILOG inference resembles Natural Logic (Nlog) but is more general.
An example beyond the scope of Nlog (again from Allen's Monroe domain):***

*Every available crane can be used to hoist rubble onto a truck.
The small crane, which is on Clinton Ave, is not in use.*

*Therefore: the small crane can be used to hoist rubble from
the collapsed building on Penfield Rd onto a truck.*

- *Every available crane can be used to hoist rubble onto a truck*
(s '(all x (x ((attr available) crane))
(all r (r rubble)
(that (some y (y person)
(some z (z truck) (y (adv-a (for-purpose (Ka (adv-a (onto z) (hoist r))))
(use x)))))) possible))))))
- *The small crane, on Clinton Ave., is not in use.*
(s '(the x (x ((attr small) crane)) ((x on Clinton-Ave) and (not (x in-use))))))
- *Every crane is a device*
(s '(all x (x crane) (x device)))
- *Every device that is not in use is available*
(s '(all x ((x device) and (not (x in-use))) (x available)))



***Can the small crane be used to hoist rubble from the collapsed building
on Penfield Rd onto a truck? (Answered affirmatively by EPILOG in .127 sec)***

(q (p '(the x (x ((attr small) crane))
(some r ((r rubble) and
(the s ((s (attr collapsed building)) and (s on Penfield-Rd))
(r from s)))
(that (some y (y person)
(some z (z truck)
(y (adv-a (for-purpose (Ka (adv-a (onto z) (hoist r))))
(use x)))))) possible))))))

Types of knowledge required

Patterns of predication: for parser guidance

e.g., person see bird; see with viewing instruments;
bird with feathers; tail feathers; colored feathers

Disambiguate

“He saw a bird with {binoculars | yellow tail feathers}”



Patterns of actions/events and relationships: for expanding /connecting sentences

e.g., *Dyn-schema* [Person x visits foreign-country y] (episode e):

Init-conds: [x loc-at u:<some (place in (home-country-of y))>] (episode e1),

Co-conds: [x have v:<some (travel paraphernalia for y e)>] (episode e2),

Steps: [x travel from u to w:<some (place in y)>] (episode e3),

[x do some activities in y] (episode e4),

[x travel from z:<some (place in y)> to u] (episode e5)

Effects: [x obtain g:<some (gratification-from e4)>] (episode e6)

Constr: [e = (join-of e3 e4 e5 e6)], [e1 starts e], [e2 same-time e],

[e3 consec e4 e5], [e6 during e4]

Similarly, possible patterns of events for cancer patient diagnosis, surgery, chemo, ...)

Similarly, “object utility” patterns (what can you do with an apple? a pen? a car? ...)

Similarly, “dispositional” patterns (dogs bark; fragile object tend to break on impact)

Inference based on patterns and schemas seems to have the character

match certain parts → tentatively infer the whole

Types of knowledge, cont'd

Lexical and paraphrase knowledge

e.g., dogs are land mammals; pens are writing instruments;
trees are plants with a tall wooden trunk and a branched canopy ...;
to kill is to render dead; to walk is to advance on foot;
managing to do x entails doing x;
x sells y to z \Leftrightarrow z buys y from z;
make good money \Leftrightarrow earn well \Leftrightarrow be well compensated \Leftrightarrow pull in the bucks;



Conditional/generic world knowledge

e.g., if you drop an object, it will hit whatever is directly beneath it;
most dogs are someone's pet; dogs are generally friendly;
chemotherapy tends to cause nausea and hair loss; ...



Specialist knowledge:

taxonomies, paronomies, temporal relations, arithmetic & scales,
geometric/imagistic representations, sets, symbolic expressions (incl. language), ...



Some attempts to address the “Knowledge Acquisition Bottleneck”




We have plenty of textual “knowledge”, but little

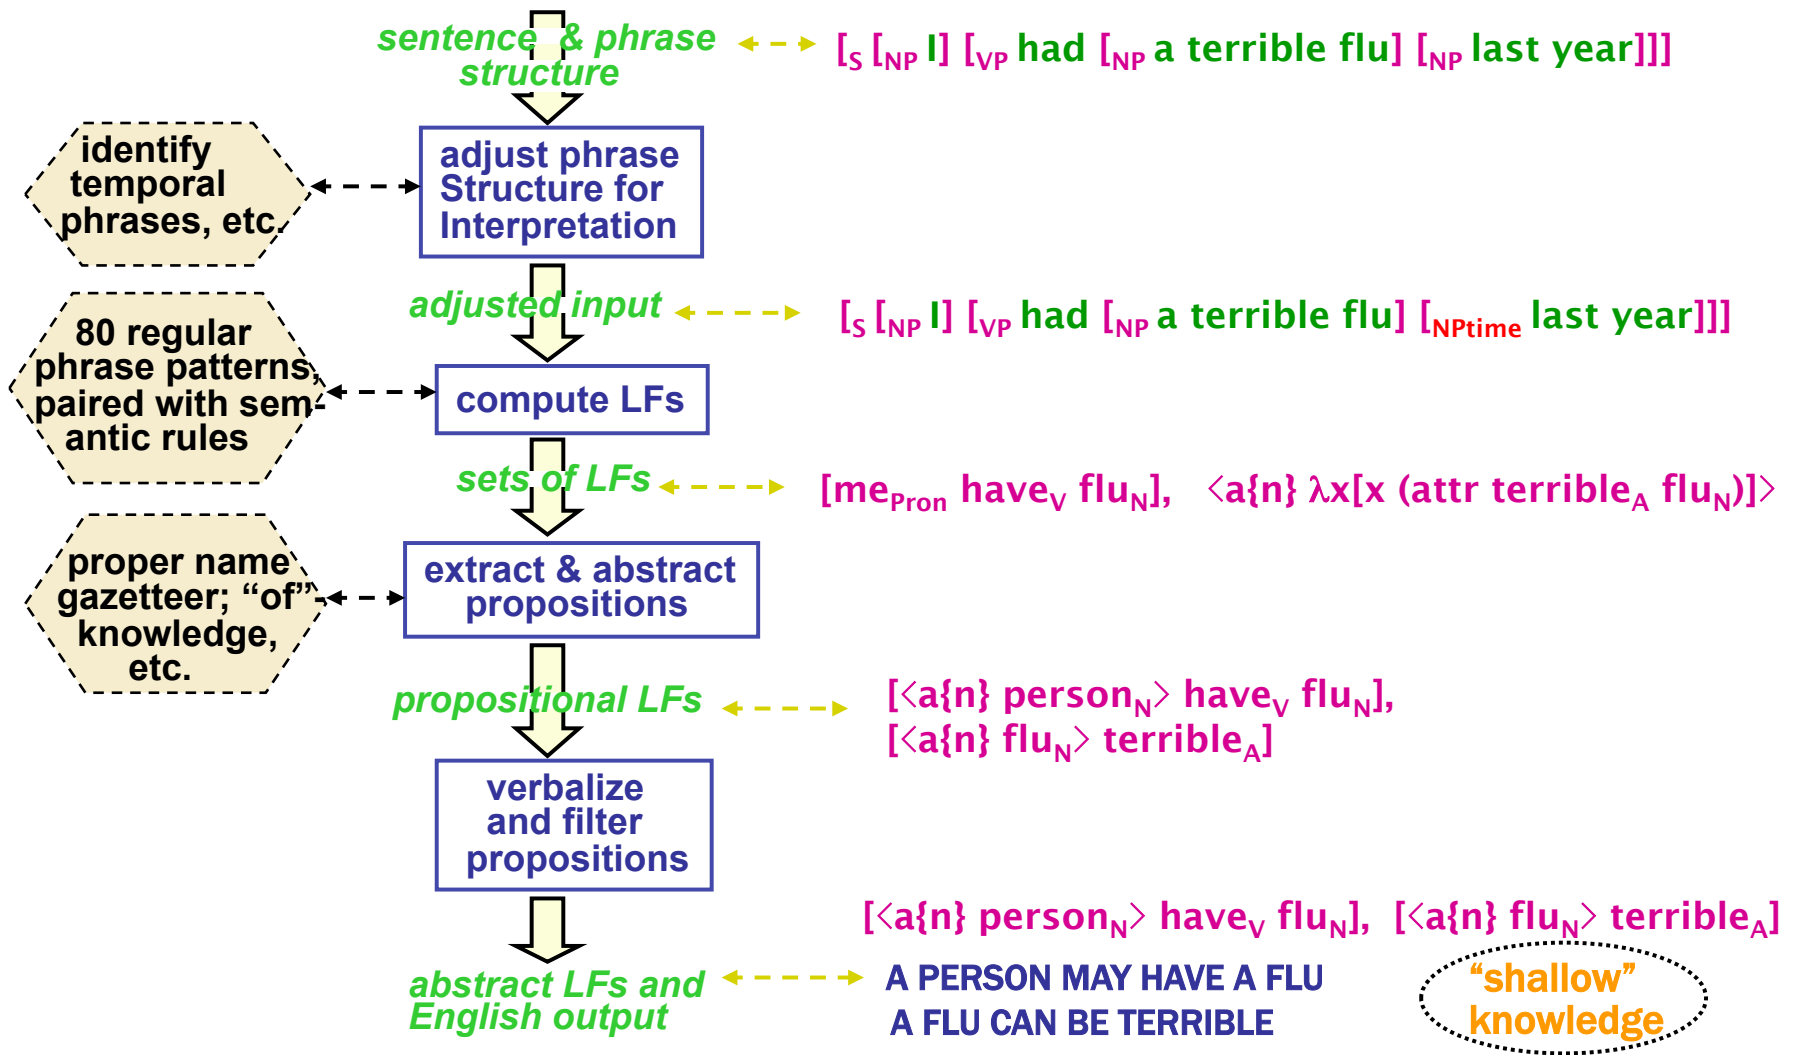
- schematic pattern knowledge for understanding (and behavior);
- “if-then” lexical and world knowledge for reasoning!

Past and ongoing efforts in my Rochester group:

- **KNEXT – knowledge extraction from text**, aimed at general factoids (typically, patterns of predication)
- “**Sharpening**” and abstracting of **KNEXT factoids** into quantified generalizations (**Lore**, successor to KNEXT)
- **Lexical knowledge engineering** (partially automated), esp. for frequent and “primitive” verbs, several VerbNet classes, implicatives, attitudinal verbs
- **Automatic WordNet gloss interpretation** (in progress)
- **Computer-graphics-based object/scene representation & inference**
- **Experimental schema engineering**, to assess syntactic/semantic requirements
- **Schema learning** (future work)



KNEXT architecture



KNEXT Website for World Knowledge: www.cs.rochester.edu/research/knext/browse

See also Jonathan Gordon's <http://www.cs.rochester.edu/research/lore/>

The screenshot shows a Mozilla Firefox browser window titled "KNEXT Factoids: Car". The address bar shows the URL www.cs.rochester.edu/research/knext/browse/prop-query.php?q=car. The page content includes a search bar with the text "KNEXT— Query Results" and "Car". Below the search bar, it states "58,806 results in the active KB. This search was truncated for speed; more results may exist than are being returned." There are two columns of results: "Unsharpened Factoids" and "Sharpened Factoids". The "Sharpened Factoids" column shows a result: "All or most persons occasionally drive a car." Below this result is a box containing the logical form: (all-or-most x: [x person.n] (occasional e (some y: [y car.n] [[x drive.v y] ** e]])) and the text "Sharpened from". At the bottom of the box, it says "Click on a source factoid to see the original English sentences it was learned from."

“Sharpening” KNEXT factoids
(using transduction patterns, semantic info from WordNet, VerbNet, etc.)

E.g., A person may have hair
→ All or most persons permanently have some hair as part

(all-or-most x: [x person]
(some e: [(x . e) permanent]
(some y: [y hair]
[[x have-as-part y] ** e]]))

Sample inferences:

Dana is a person →
Probably, Dana permanently has some hair as part.

(probably
(some e: [(Dana . e) permanent]
(some y: [y hair]
[[Dana has-as-part y] ** e]]))

ACME is a company →
Probably, ACME occasionally announces a product.

(probably
(occasional e
(some y: [y product]
[[ACME announce y] ** e]]))

Accumulating Lexical Knowledge: Primitives, VerbNet, and WordNet

~150 “primitive” verbal concepts, with axioms

(**MOVE, GRASP, SEE, LEARN, MAKE, ASK-OF, CONVEY-INFO-TO, WANT-TBT, ...**)

Several hundred axioms for VerbNet classes

(e.g., state-change concepts like **BREAK, REPAIR, MELT, etc.**) in terms of “primitives” and “predicate parameters”, inserted into 2-4 *axiom schemas* per class)

WordNet: 77,000 formal axioms derived from WN nominal hierarchy

(used mass/count distinction and various other features to distinguish the relation between pairs like **<seawater, water>** → *all seawater is water*, and like **<gold, noble_metal>** → *gold (the kind of stuff) is a noble_metal*)


Gloss interpretation (in progress): E.g., *x **slam2.v** y in event e* → *x **violently strike1.v** y, and x is probably a person and y is probably a thing12.n (physical entity)*. Problems: Many glosses are quasi-circular (*admire* → *feel admiration for*), don’t apply to all synset members, and don’t provide major entailments (like, *dying results in being dead!*)

Graphics-based scene modeling in first-reader stories:

Rosy and Frank see the nest in the apple tree.

Why can’t they tell whether there are eggs in the nest?

Visual occlusion by the nest itself!



Current preliminary schema work: experimenting with schema syntax (like the “**visiting a foreign country**” example), to be used as target formalism for learning schemas from text.

Concluding comments

- Reliable semantic parsing requires reliable syntactic parsing; need guidance by patterns of predication, modification, etc. (KNEXT-like)
- EL provides an expressively adequate representation close to language
- Inference methods beyond those of FOL are implemented
- Probabilistic inference remains ad hoc
- KA bottleneck is under attack, but far from vanquished; schemas (containing EL formulas) seem to be the kind of “soft” representation that will support genuine understanding, reasoning

Publications (from my web page, <http://www.cs.rochester.edu/u/schubert/>)

- L.K. Schubert, "[Semantic representation](#)", *29th AAAI Conference (AAAI15)*, Jan. 25-30, 2015, Austin, TX. [Slides of the talk](#).
- E. Bigelow, D. Scarafoni, L. Schubert, and A. Wilson, "On the need for imagistic modeling in story understanding", *Biologically Inspired Cognitive Architectures 11*, Jan. 2015, 22--28. (Presented at Ann. Int. Conf. on Biologically Inspired Computing Architectures (BICA 2014), Nov. 7-9, MIT, Cambridge, MA, 2014.)
- L.K. Schubert, "[Computational linguistics](#)", in Edward N. Zalta (Principal Editor), *Stanford Encyclopedia of Philosophy*, CSLI, Stanford, CA.(100pp), 2014. L.K. Schubert, "[From Treebank parses to Episodic Logic and commonsense inference](#)", Proc. of the ACL 2014 Workshop on Semantic Parsing, Baltimore, MD, June 26, ACL, 55-60.
- L.K. Schubert, "[NLog-like inference and commonsense reasoning](#)", in A. Zaenen, V. de Paiva, & C. Condoravdi (eds.), *Semantics for Textual Inference*, special issue of Linguistic Issues in Language Technology (LiLT) 9(9), 2013. (The pdf file is a preliminary version, Dept. of Computer Science, University of Rochester, Oct. 2012.)
- J. Gordon and L.K. Schubert, "[WordNet hierarchy axiomatization and the mass-count distinction](#)", 7th IEEE Int. Conf. on Semantic Computing (ICSC 2013), September 16-18, Irvine, CA, 2013.
- A. Purtee and L.K. Schubert, "[TTT: A tree transduction language for syntactic and semantic processing](#)", EACL 2012 Workshop on Applications of Tree Automata Techniques in Natural Language Processing (ATANLP 2012), Avignon, France, Apr. 24, 2012.
- K. Stratos, L.K. Schubert, & J. Gordon, "[Episodic Logic: Natural Logic + reasoning](#)", Int. Conf. on Knowledge Engineering and Ontology Development, (KEOD'11), Oct 26-29, Paris, 2011.
- Jonathan Gordon and Lenhart Schubert, "[Quantificational sharpening of commonsense knowledge](#)", Proc. of the AAAI Fall Symposium on Commonsense Knowledge, Arlington, VA, Nov. 11-13, 2010. L.K. Schubert, "[From generic sentences to scripts](#)", IJCAI'09 Workshop W22, Logic and the Simulation of Interaction and Reasoning (LSIR 2), Pasadena, CA, July 12, 2009.
- F. Morbini and L.K. Schubert, "[Evaluation of Epilog: A reasoner for Episodic Logic](#)", Commonsense'09, June 1-3, 2009, Toronto, Canada. Also available at <http://commonsensereasoning.org/2009/papers.html>
- Benjamin Van Durme and Lenhart K. Schubert, "[Open knowledge extraction using compositional language processing](#)", Symposium on Semantics in Systems for Text Processing (STEP 2008), September 22-24, 2008 - Venice, Italy.
- A.N. Kaplan and L.K. Schubert, "[A computational model of belief](#)", *Artificial Intelligence* 120(1), June 2000, 119-160.
- L.K. Schubert and C.H. Hwang, "[Episodic Logic meets Little Red Riding Hood: A comprehensive, natural representation for language understanding](#)", in L. Iwanska and S.C. Shapiro (eds.), *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*, MIT/AAAI Press, Menlo Park, CA, and Cambridge, MA, 2000, 111-174.