

Meeting the interlocking needs of LF-computation, deindexing, and inference: An organic approach to general NLU*

Chung Hee Hwang & Lenhart K. Schubert

Department of Computer Science, University of Rochester
Rochester, New York 14627, U. S. A.

Abstract

We argue that a “theory bottleneck” encountered in the 70’s and early 80’s in attempts to build comprehensive NLU systems led to a fragmentation of NLU research, which still persists. To some extent, this fragmentation represents an appropriate response to the variety and subtlety of remaining problems; but at this point, it also represents a loss of nerve: NLU is an *organic* phenomenon, and enough has been learned about the vexing problems of the 80’s to try to integrate these insights and build more comprehensive theories and extensible implementations. On that premise, we have been building such a comprehensive framework. Its centerpiece is *Episodic Logic* (EL), a highly expressive knowledge representation well-adapted to the interpretive and inferential needs of general NLU. The logic has been successfully implemented in the EPILOG system, which has been applied to several domains proving EL’s practicality: reasoning about small excerpts from the *Little Red Riding Hood* story, as a natural language interface for the TRAINS planning system at the University of Rochester, and for a message processing application at the Boeing Co.

1 The Organism Awaits: A Manifesto

We report here on an approach to theoretical and practical NLU that aspires to be complete and comprehensive, with respect to all the major syntactic, semantic, and pragmatic phenomena encountered in NL texts. Before launching into technical particulars, we will devote some space to motivating our endeavor, since comprehensive approaches appear to be somewhat unfashionable. For instance, the overwhelming majority of papers at recent ACL conferences have focused on specific issues in grammar formalisms, parsing, discourse, corpus-based analysis, etc., with scant attention to the role of the proposed theories or techniques within a larger NLU framework.

This trend toward specialization represents a striking shift from the heady days of the 70’s and early 80’s, when (at least to a casual observer) each year appeared to bring a more complete picture of the essential structures and processes underlying language understanding and a more impressive set of working systems demonstrating the power of the new ideas and mechanisms. In retrospect, these efforts were fated to encounter two

formidable barriers: a “knowledge bottleneck” and a “theory bottleneck.” The knowledge bottleneck lies in the dependence of genuine understanding on a vast substrate of general and specific knowledge, and the difficulty of imparting all but trivial amounts of such knowledge to machines “by hand.” The theory bottleneck lies in the myriad problems of syntax, semantics, pragmatics and inference that any general NLU or dialog system needs to address. AI-oriented researchers perhaps tended to underestimate the latter bottleneck, armed as they were with new insights and techniques from AI. But the brittleness and inextensibility of the NLU systems of the 70’s and 80’s was ultimately seen to lie as much in their reliance on programmer intuitions and domain-specific hacks to bridge gaps in theoretical foundations, as in the meagerness of their knowledge.

From this perspective, it is not surprising that the field has shifted into “divide and conquer” mode, characterized by fragmentation and specialization. However, we believe that at this point, the dearth of comprehensive theorizing and system-building also represents something of a loss of nerve in the AI/CL community. Divide-and-conquer is tenable up to a point, but ultimately we need to confront language as an *organic* phenomenon, whose facets have no clearly discernible shape independently of each other. Even grammar, the most accessible facet, is moot (witness the various competing syntactic formalisms, and competing analyses of many specific phenomena); and logical form, context, discourse structure, knowledge representation, and inference are utterly hypothetical. But of course they are tightly interlocked, through the coupling between surface form and logical form, between logical form and the ultimate meaning representation (mediated by discourse context), and between the ultimate meaning representation and subsequent conversational (and other) behavior; as well as through the crucial role played by world knowledge in all of these transductions.

Compared to the 70’s and early 80’s, present prospects for principled, integrated NLU are greatly improved. Considerable strides have been taken in our understanding of all aspects of language processing: e.g., grammar, parsing, theories of discourse structure and intention, etc. Most importantly from our perspective, new logical frameworks such as DRT, situation semantics and type and property theories provide new insights into various long-standing semantic conundrums, such as the semantics of events, anaphora, properties, and attitudes.

How, then, do we propose to go about “putting it all together”? We find that the various “building blocks”

*This paper appeared in *Proc. IJCAI-93*, 1297–1302. The authors are grateful to James Allen and Philip Harrison for helpful comments during EL development, and Stephanie Schaeffer for staying with program development for EPILOG. This research was supported in part by the Boeing Co. in Seattle, NSERC Operating Grant A8818, ONR/DARPA Research Contracts No. N00014-82-K-0193 and No. N00014-92-J-1512, and NSF Research Grant IRI-9013160.

offered by various authors and schools of thought simply do not fit together in any obvious way, and leave numerous gaps. What we have therefore been attempting is a *creative synthesis* of a modularized NLU framework, with progressive refinement of the overall abstract architecture and gradual “debugging” and theoretical deepening of the entire account. Lest we lose sight of our computational objectives, we have at all times maintained a tight link between theoretical development and computer implementation (the latter culminating in the EPILOG reasoning system).

We feel that what we have accomplished so far vindicates our methodology. The current framework consists of a GPSG-like approach to grammar with a Montague-like coupling of phrase structure to logical form; “preference”-based methods of parsing and scope disambiguation; a new theory of tense and aspect “deindexing”; a combined semantic representation/knowledge representation; and methods of making deductive and probabilistic inferences in both input-driven and goal-driven modes in this representation. For good reason, the semantic/knowledge representation, called *Episodic Logic (EL)*, has turned out to be the centerpiece of our effort. After all, it is the choice of representation which determines how easily we can derive content from surface form, how fully we can capture the semantic nuances of NL text, and how readily we can perform needed inferences.

2 Representation: The Central Problem

In speaking of a semantic representation or knowledge representation, we take for granted that the representation should admit not only straightforward intuitive interpretation but also a formal denotational semantics. As long as we remain unclear about what sorts of things in the world our symbols can stand for, or how a putative knowledge base can conform with (or deviate from) how things actually are, we risk having the system lapse into total inconsistency and incoherence for all but trivial knowledge bases. At the same time we should not from the outset demand absolute hygiene and mathematical perfection from a representational logic which attempts to encompass all the semantic phenomena of NL which have occupied linguists and philosophers for decades, and for which no agreed-upon unified framework exists. In fact, EL initially had little more than a tentative syntax, ontology, and type structure, but we have now gained an understanding of semantic entailment, finding soundness proofs for some of the inference methods. In the meantime, however, we have never compromised expressive adequacy for theoretical thoroughness. Instead, all the while, the expressive completeness of EL has allowed us to experiment freely with the syntax/logical form interface, tense and aspect deindexing, and inferences based on stories and conversations.

We recognize that attempts at comprehensive theorizing and system-design have not been completely abandoned. For example, Hobbs *et al.* [8; 9] and Charniak *et al.* [3; 4] have kept their sights on full understanding, and in particular, have made very intriguing proposals for full integration of all types of disambiguation,

based on *abduction*. However, the overwhelming majority of NLU researchers have in recent years shown a curious reserve toward full understanding, most strikingly in their approach to the core problem of semantic representation. The existing approaches fall by and large into either of two categories, theory-free approaches or FOL approaches. In a theory-free approach, a representation with no known formal semantics is employed. Consequently, the user of the representation is free to claim that a rich variety of concepts and ideas can be expressed in the representation — for instance, beliefs, actions, goals, habitual behavior, etc. — without ever having to confront technical criticisms about internal consistency of the representation, or what follows from given knowledge. We are not saying that such work is necessarily bad — it is often quite appealing and persuasive with respect to the issues it addresses; but the avoidance of the semantic issue, in any serious sense, leaves the overall framework excessively ill-defined.

FOL approaches limit the fragments of language considered to those which appear to be expressible in FOL, at least in rough-and-ready fashion. This has the advantage that FOL is well-understood syntactically and semantically, but also has the disadvantage that very little real language is easily expressible in it. Consider, for instance, the first sentence of the dialog which was processed by the initial TRAINS-implementation [1]:

“We have to make orange juice”

This simple sentence exemplifies the following interesting semantic phenomena: it expresses an obligation (and puts that obligation, in part, on the hearer); it expresses joint agency; it is tensed (thereby expressing that the joint obligation is upon the hearer and speaker at the time of speech); it involves a mass noun phrase, “orange juice,” which on most modern accounts denotes an abstract *kind*; it involves a verb of creation, “make,” which in combination with its kind-denoting object implies the coming-into-being of some *quantity* of orange juice (this coming-into-being cannot be expressed by existential quantification); and it involves an infinitive whose interpretation is arguably a reified property. While with practice one can become quite good at inventing *ad hoc* FOL-approximations to given sentences, it is quite implausible that any reasonably simple, systematic transduction from syntax to semantics would give an FOL translation of this sentence. We find this situation quite typical for real spoken and written language. So the FOL approach factors out most of language — at least for any *algorithmic* (as opposed to *ad hoc*) mapping from syntax to semantic representation. Thus, we suggest it is time to break away from the restrictiveness of FOL, and settle for nothing less than NL-like expressiveness (without retreating to theory-free representations).

We now briefly review EL and its role in a comprehensive, modular approach to NLU. EL is a first-order logic that is very expressive, formally interpretable, and easily derived from surface utterances, yet allows efficient inference. As mentioned, it is based on a Montague-style coupling between syntactic form and logical form, while incorporating from situation semantics the idea that sentences describe situations (events, states, episodes, even-

tualities, etc.).¹ Moreover, all of this is implemented in at least a preliminary way. The EPILOG system [17], the computer implementation of our logic, makes quite complex inferences, e.g., with utterances from the TRAINS domain [1], telex reports for aircraft mechanical problems in the ARMS (the Boeing Commercial Airplane Reliability and Maintainability Project) domain [16], and excerpts from *Little Red Riding Hood* story [19]. The inferences in these experiments are made quite straightforwardly, despite the richness of our logic, or—we would argue—because of it, and the knowledge they employ is uncontrived (it corresponds quite directly to English sentences, and each individual piece of knowledge is formulated at a maximally general level, rather than being particularized to the needs of a specific story).

In the rest of this paper, we first describe EL syntax and inference rules, using examples from the ARMS application domain, and then discuss how one can get episodic logical form from English input.

3 The Episodic Logical Form

The most important feature of EL is its NL-like *expressiveness*. This makes it easy to derive EL-translations of English sentences, while also providing a basis for concise, easily understood inferences. Its syntax allows lambda abstraction, restricted quantifiers, propositional attitudes, predicate and sentence modifiers, kind forming operators, nominalization operators, action abstraction, DRT-like anaphoric variables, generic conditionals, and other non-standard constructs. Most importantly, however, it makes use of episodic variables in the representation of episodic sentences, making implicit temporal and causal relationships between situations explicit.

To give an idea of the syntax, we show below an ELF representation of the sentence “A large bird bumped into the left wing of (airplane) VB7, causing it to get a crack” (with certain simplifications).²

$$\begin{aligned} & [(\exists e_1: [e_1 \text{ before } u_1] \\ & \quad [(\text{The } x_1: [[x_1 \text{ ((attr left) wing)]} \wedge [x_1 \text{ part-of VB7]} \\ & \quad (\exists y_1: [y_1 \text{ ((attr large) bird)] } [y_1 \text{ bump-into } x_1])]] \\ & \quad ** e_1)] \wedge \\ & (\exists e_2: [e_1 \text{ cause-of } e_2] [(\exists z_1: [z_1 \text{ crack}] [x_1 \text{ get } z_1]) ** e_2])] \end{aligned}$$

This sentence introduces three episodes: u_1 , the utterance event of the above sentence itself; e_1 , an episode of “large bird y_1 bumping into the left wing x_1 ”; and e_2 , an episode of “wing x_1 getting crack z_1 .” Note the clause [e_1 cause-of e_2] that shows the causal relationship between the two episodes. The occurrence of e_1 outside its quantifier scope is allowed thanks to the *parameter mechanism* of EL (more on this later). ‘**’ is an episodic modal operator, such that $[\Phi ** \eta]$ means formula Φ characterizes (or, completely describes) episode η . A weaker

¹We use the term “episode” to emphasize the transient nature of the situations that commonly occur in narratives.

²Below, attr is an operator that transforms a predicate into an *attributive* predicate modifier. Note that we use infix notation for readability and restricted quantification of the form $(Q \alpha: \Phi \Psi)$, where $Q \in \{V, \exists, \text{The, Few, Most, } \dots\}$, α is a variable, and restriction Φ and matrix Ψ are formulas. Readers are referred to [10; 12] for details of EL syntax.

form of this modal operator is ‘*’. $[\Phi * \eta]$ means that Φ partially describes (or, is true in) η . This is similar to the \models (“support”) relation of situation semantics [2; 6], except that we are relating sentence intensions (partial mappings from situations to truth values), rather than “infons,” to situations.

The following examples illustrate attitude predicates as well as kind abstraction, action abstraction, and event abstraction. That, K, Ka, and Ke below are nominalization operators: That maps sentence intensions to propositions,³ K maps predicates into “kinds” of *property*, Ka maps predicates into kinds of *action* or *attribute*, and Ke maps sentence intensions into kinds of *event*.

- *The mechanic doubted that the crack indicated danger.*
 $(\exists e_3: [e_3 \text{ before } u_3]$
 $\quad [(\text{The } x_2: [x_2 \text{ mechanic}]$
 $\quad (\text{The } y_2: [y_2 \text{ crack}]$
 $\quad [x_2 \text{ doubt } (\text{That } (\exists e_4: [e_4 \text{ same-time } e_3]$
 $\quad \quad [[y_2 \text{ indicate } (K \text{ danger})] ** e_4])])])])$
 $** e_3])$
- *The mechanic tried to repair the crack.*
 $(\exists e_5: [e_5 \text{ before } u_5]$
 $\quad [(\text{The } x_3: [x_3 \text{ mechanic}]$
 $\quad (\text{The } y_3: [y_3 \text{ crack}] [x_3 \text{ try } (Ka \text{ (repair } y_3))])])$
 $** e_5])$
- *The inspector suggested that an expert examine the hydraulic seam.*
 $(\exists e_6: [e_6 \text{ before } u_6]$
 $\quad [(\text{The } x_4: [x_4 \text{ inspector}]$
 $\quad (\text{The } y_4: [y_4 \text{ ((attr hydraulic) seam)]}$
 $\quad [x_4 \text{ suggest}$
 $\quad \quad (Ke (\exists z_4: [z_4 \text{ expert}] [z_4 \text{ examine } y_4])])])])$
 $** e_6])$

The DRT-like treatment of indefinites is illustrated by the donkey sentence, “Every inspector who found a crack longer than 2 inches immediately reported it”:

$$\begin{aligned} & (\exists e_7: [e_7 \text{ before } u_7] \\ & \quad [(\forall x_5: [[x_5 \text{ inspector}] \wedge \\ & \quad (\exists e_8: [[e_8 \text{ before } u_7] \wedge [e_8 \text{ during } e_7]] \\ & \quad \quad [(\exists y_5: [[y_5 \text{ crack}] \wedge [(length-of } y_5 \text{ inch)} > 2]] \\ & \quad \quad \quad [x_5 \text{ find } y_5]) ** e_8])]] \\ & \quad (\exists e_9: [e_9 \text{ imm-after } e_8] [[x_5 \text{ report } y_5] ** e_9])]) \\ & ** e_7)]. \end{aligned}$$

Note existential variables y_5 and e_8 that occur outside their quantifier scope. The interpretation of such free variables is much like that of *parameters* in DRT. Such anaphoric variables also play a role in generic conditionals, which, when used with probabilistic inference rules, allow us to express unreliable generalizations such as the following (with some simplifications).

$$\begin{aligned} & [(\exists e [(\exists x: [x \text{ crack}] [(length-of } x \text{ inch)} < .25]) ** e)] \\ & \rightarrow .75, e, x [(\neg [x \text{ dangerous}]) * e]] \end{aligned}$$

A crack shorter than 1/4 inch usually is not dangerous.

The subscript ‘.75’ in the above conditional indicates a lower bound on statistical probability, and e, x are controlled variables. The rule says, roughly, that in at

³In EL ontology, propositions subsume possible facts (see [11; 12; 13]).

least 75% of the cases in which the antecedent is true, the consequent will also be true.

We need to remark here that unlike Davidsonian [5] event variables, our approach allows for episodes involving quantification, such as “Most planes were more than 10 years old, *which* indicated that the airline company had been in financial trouble,” or episodes involving negation, such as “The mechanic did not report the crack; *as a result*, he was fired.” Note that such complex episodes frequently occur as causal antecedents, or are anaphorically referred to or quantified over, in the ARMS (and other) applications.

Our treatment of temporal and locative adverbials views these as providing conjunctive information about the described episode, much as in [7]. For instance, “Airplane VB7 executed an emergency landing in Chicago yesterday” is represented as

$$(\exists e_{10}: [e_{10} \text{ before } u_{10}] \wedge [[e_{10} \text{ during (yesterday-rel-to } u_{10})] \wedge [e_{10} \text{ in-loc Chicago}] \wedge (\exists x_6: [x_6 \text{ ((nn emergency) landing)}] [VB7 \text{ execute } x_6])] ** e_{10})).$$

Above, *nn* (standing for *noun-noun* modification) is an operator that transforms a predicate to a predicate modifier (cf., the *nn* operator of Hobbs *et al.* [9]).

Space limitations prevent a discussion of formal semantics (but see [10; 12]), but we should remark that unlike situation semantics, EL is based on an ontology that allows *possible situations*. These are much like “partial possible worlds,” in that symbols are assigned partial extensions (and antiextensions) relative to them.

4 Inference Rules in Episodic Logic

The main inference rules we have developed are *Rule Instantiation* (RI) and its dual *Goal Chaining* (GC), which resemble forward and backward chaining rules in expert systems. There is also another class of goal-directed methods that consists of standard natural deduction rules such as proving a conditional by assuming the antecedent and deriving the consequent, or proving a universal by proving an arbitrary instance of it. Here we only discuss RI and GC.

Rule Instantiation (RI)

Rule Instantiation, which is heavily used in input driven inference, allows arbitrarily many minor premises to be matched against arbitrarily deeply embedded subformulas of a rule. It subsumes *modus ponens* and *modus tollens*, but can also instantiate generic conditionals. In the unit probability version, with just one minor premise (“fact”), the RI rules are:⁴

$$\frac{R^-(\Phi), F^+(\Psi)}{R^-_\sigma(\neg(F^+(\perp)))} \quad \frac{R^-(\Phi), F^+(\Psi)}{F^+_\sigma(R^-(\top))}$$

where σ unifies Φ, Ψ . R stands for “Rule,” and F for “Fact.” \top and \perp are truth and falsity respectively. The

⁴R. Waldinger (personal communication) pointed out to us that these rules are essentially the same as those of Traugott [21], except that we entirely avoid skolem functions.

+ and – signs are intended to indicate positive and negative occurrence of the embedded Φ, Ψ formulas being unified.⁵ Unification is defined in a way that allows substitution for explicitly quantified, “matchable” variables. A variable in a rule or fact is matchable if it is bound by a positively occurring universal quantifier or negatively occurring existential quantifier. For instance, substitution of w for x in a positively embedded subformula $(\forall x: [x P] [x Q])$ yields $[[w P] \rightarrow [w Q]]$, and the same substitution in a negatively embedded subformula $(\exists x: [x P] [x Q])$ yields $[[w P] \wedge [w Q]]$. The first rule is *sound* if Ψ contains no unmatchable free variables which are bound in F as a whole. The second rule is sound if Φ contains no unmatchable free variables which are bound in R as a whole. So in particular, the first rule is sound if F contains only constants and top-level universal (hence matchable) variables. (See [18] for soundness proof.)

As mentioned, the rules also work for generic conditionals, and we illustrate this case. Consider the following generic conditional:

$$((\exists e [(\exists x: [[x \text{ aircraft}] \wedge [(age-of } x \text{ year}) < 3]] (\exists y: [y \text{ crack}] [y \text{ located-on } x])]) ** e]) \rightarrow .8, e, x (\neg [[y \text{ due-to (K corrosion)}] * e]))$$

If an aircraft less than 3 years old has a crack, usually the crack is not due to corrosion.

Suppose now the translation of a telex contains the following:

$$[[[VB7 \text{ aircraft}] \wedge [(age-of VB7 year) = 2] \wedge [C4 \text{ crack}] \wedge [C4 \text{ located-on VB7}]] ** E1]$$

The two year old aircraft VB7 has a crack.

Then, RI matches these formulas against the antecedent of the conditional, unifying $VB7/x$, $C4/y$ and $E1/e$, and derives (aided by “specialists” in EPILOG):

$$(\neg [[C4 \text{ due-to (K corrosion)}] * E1])^{.8},$$

where the superscripted number, .8, is a lower bound on epistemic probability. (That is, the system concludes that the crack C4 is not due to corrosion, with minimum degree of confidence .8.) For this simple example, either version of RI yields the correct inference. This kind of inference, based on unreliable generalizations, allows evidence for explanations or predictions to be weighted, much as is done in expert systems.

Goal Chaining (GC)

Goal chaining, which dominates goal-driven inference, is a pair of very general chaining rules. Chaining from rule consequents to antecedents is a special case. The following are the GC rules in the unit probability case:

$$\frac{R^+(\Phi), ?G^+(\Psi)}{?G^+_\sigma(\neg(R^+_\sigma(\perp)))} \quad \frac{R^+(\Phi), ?G^+(\Psi)}{?\neg(R^+_\sigma(\neg(G^+_\sigma(\top))))}$$

where σ' “antiunifies” Φ, Ψ (i.e., with positive existen-

⁵ Φ is positively (negatively) embedded in another formula if the only operators embedding Φ are among $\{\neg, \rightarrow, \wedge, \vee, \forall, \exists\}$, and the number of occurrences of \neg , the antecedent scope of \rightarrow , and the restriction scope of \forall among the operators embedding Φ is even (odd).

tials and negative universals in G regarded as matchable). R stands for “Rule,” and G for “Goal.” The first rule is *sound* if Φ contains no unmatchable free variables which are bound in R as a whole. The second rule is *sound* if Ψ contains no unmatchable (e.g., top-level universal) variables which are bound in G as a whole.

For example, consider a rule

$$[(\exists x_1:[x_1 \text{ ((attr hydraulic) joint)}]) \\ (\exists e_1 [(\exists x_2:[x_2 \text{ crack}] [x_2 \text{ located-on } x_1]) ** e_1])] \\ \rightarrow [(\exists e_2:[e_2 \text{ same-time } e_1] [(\text{age-of } x_1 \text{ year}) > 5] ** e_2)] \vee \\ (\exists e_3:[e_3 \text{ before } e_1] \\ [x_1 \text{ receive (K ((attr excessive) pressure))] ** e_3)]]$$

If a hydraulic joint has a crack, either it is more than 5 years old, or received excessive pressure.

Suppose, now, we want to know if a particular hydraulic joint, Joint6, received excessive pressure. Then the goal can be posed as

$$? (\exists e:[e \text{ before } E1] \\ [[\text{Joint6 receive (K ((attr excessive) pressure))] ** e]),$$

with a constraint that $E1$ is at the same time as the utterance time. Since the goal matches the second disjunct of the consequent of the rule, with unification $\text{Joint6}/x_1$, $E1/e_1$, and e/e_3 , we replace the matching part of the rule with \perp , and negate the whole formula. Then we get a new goal

$$? (\neg [(\exists x_2:[x_2 \text{ crack}] [x_2 \text{ located-on } \text{Joint6}]) ** E1] \\ \rightarrow [(\exists e_2:[e_2 \text{ same-time } E1] \\ [(\text{age-of } \text{Joint6} \text{ year}) > 5] ** e_2)] \vee \perp]),$$

which, after simplification, becomes

$$? [(\exists x_2:[x_2 \text{ crack}] [x_2 \text{ located-on } \text{Joint6}]) ** E1] \wedge \\ (\neg [(\exists e_2:[e_2 \text{ same-time } E1] \\ [(\text{age-of } \text{Joint6} \text{ year}) > 5] ** e_2)]).$$

The new goal says, “Joint6 has a crack and is not older than 5 years.” This makes sense because one way of proving “Joint6 received excessive pressure” is to prove that “Joint6 is not older than 5 years and has a crack.” As with the previous RI example, either version of GC gives us the same new subgoal.

The general version of GC allows arbitrarily many subsidiary knowledge base facts to be invoked in the process of chaining from the given goal to a subgoal. For instance, in the above example, knowledge base fact $[\text{Joint6} \text{ ((attr hydraulic) joint)}]$ led to immediate success.

For both RI and GC, the second rule remains to be implemented. Yet, the kinds of EL inferences described so far are carried out efficiently by the EPILOG system [17], a hybrid reasoning system combining efficient storage and access mechanism, forward and backward chaining, an agenda-driven control structure, and multiple “specialists” for taxonomies, temporal reasoning, etc.

5 Computing Episodic Logical Form

As we have emphasized, a crucial feature of EL with respect to the goal of building general NLU systems is the ease with which EL-representations are derived from surface syntax. This is made possible by its NL-like expressiveness.

The initial translation from phrase structure to the

preliminary indexical logical form (LF) is accomplished with GPSG-like syntactic and semantic rules; the final nonindexical episodic logical form (ELF) is obtained by simple recursive deindexing rules. Such a transformation is essential because, to be useful for *inference*, a situational logic need to be nonindexical. Our deindexing algorithm systematically handles tense, aspect, and many temporal adverbials and their interaction, and brings the context information into the logical form, removing context dependency.

For example, the logical form of sentence

“A mechanic repaired the crack yesterday”

is easily computed using the following (somewhat simplified) lexical and phrase structure rules, annotated with corresponding semantic rules. (We use a feature system hierarchically organized as a tree; see [10].)

$$\text{DET}[\text{indef}] \leftarrow a; \exists \\ \text{DET}[\text{def}] \leftarrow \text{the}; \text{The} \\ \text{N} \leftarrow \text{mechanic}; \text{mechanic} \\ \text{N} \leftarrow \text{crack}; \text{crack} \\ \text{NP} \leftarrow \text{DET N}; <\text{DET}' \text{ N}' > \\ \text{V}[\text{past}, \text{NP}] \leftarrow \text{repaired}; <\text{past repair}> \\ \text{VP} \leftarrow \text{V}[\text{NP}] \text{ NP}; (\text{V}' \text{ NP}') \\ \text{NP}[\text{def-time}] \leftarrow \text{yesterday}; \text{Yesterday} \\ \text{ADV}[\text{ep-mod}] \leftarrow \text{NP}[\text{def-time}]; (\text{during NP}') \\ \text{ADVL}[\text{post-mod}] \leftarrow \text{ADV}[\text{ep-mod}]; \\ \lambda P \lambda x ((\text{adv-e ADV}') [x P]) \\ \text{VP} \leftarrow \text{VP ADVL}[\text{post-mod}]; (\text{ADVL}' \text{ VP}') \\ \text{S} \leftarrow \text{NP VP}; [\text{NP}' \text{ VP}']$$

Applying these rules gives us the initial, unscoped logical form ULF (1) shown below. Angle brackets in (1) indicate unscoped expressions. Notice the indexical term *Yesterday* and the indexical sentence operator *past* in (1). *adv-e* (standing for “episode modifying adverbial”) is an operator that maps 1-place predicates over situations into sentential operators. (2) is a scoped LF; *past* has been raised so that it has a scope over the whole formula. The deindexing step from LF (2) to ELF (3) is carried out by the tense-aspect deindexing rules of EL that use *tense trees* as components of a discourse structure. Note that the deindexing step introduces an episodic variable through tense operator *past* into ELF (3). Also in (3), the indexical constant *Yesterday* is disambiguated to *yesterday relative to the utterance episode*, and crack x_1 is resolved to Crack8 (the latter kind of reference resolution is not currently dealt with in our system). ‘ \vee ’ is an extension operator that maps 1-place predicates over situations into sentence intensions. Intuitively, $\vee \pi$ means π is true of the current situation, which is e_1 in this case. Thus, (3) is identical to (4), which may be obtained via meaning postulates.

- (1) $((\text{adv-e} (\text{during } \text{Yesterday})) \\ [<\exists \text{ mechanic}> <\text{past repair}> <\text{The crack}>])$
- (2) $(\text{past} ((\text{adv-e} (\text{during } \text{Yesterday})) \\ (\text{The } x_1:[x_1 \text{ crack}] \\ (\exists y_1:[y_1 \text{ mechanic}] [y_1 \text{ repair } x_1])))$
- (3) $(\exists e_1:[e_1 \text{ before } u_1] \\ [(\vee (\text{during} (\text{yesterday-rel-to } u_1)) \wedge \\ (\exists y_1:[y_1 \text{ mechanic}] [y_1 \text{ repair } \text{Crack8}])) \\ ** e_1])$

- (4) $(\exists e_1:[e_1 \text{ before } u_1]$
 $[[[e_1 \text{ during (yesterday-rel-to } u_1)] \wedge$
 $(\exists y_1:[y_1 \text{ mechanic}][y_1 \text{ repair Crack8}]]$
 $** e_1])$

Our deindexing mechanism is compositional in that operators *past*, *futr*, *perf*, etc., contribute separately and uniformly to the meanings of their operand formulas, driving the generation and traversal of tense trees in deindexing. As an example, we show the Past-rule below.

$$\text{Past: } (\text{past } \Phi)_{\mathbf{T}}$$

$$\leftrightarrow (\exists e_{\mathbf{T}}:[e_{\mathbf{T}} \text{ bef}_{\mathbf{T}} \text{ Emb}_{\mathbf{T}}] \wedge [\text{Last}_{\mathbf{T}} \swarrow \mathbf{T} \text{ orients } e_{\mathbf{T}}]]$$

$$[\Phi_{\circ \swarrow \mathbf{T}} ** e_{\mathbf{T}}])$$

$$\text{Tree transform: } (\text{past } \Phi) \cdot \mathbf{T} = \uparrow (\Phi \cdot (\circ \swarrow \mathbf{T}))$$

\mathbf{T} denotes a tense tree, $e_{\mathbf{T}}$ a “new” episode symbol (i.e., not yet used in \mathbf{T}). $\text{Emb}_{\mathbf{T}}$, $\text{Last}_{\mathbf{T}}$, $\swarrow \mathbf{T}$, $\circ \mathbf{T}$, etc., are easily computed functions on the tense tree (see [14] for details). The recursively deindexed Φ is taken to characterize the new episode $e_{\mathbf{T}}$, which is predicated to be *at* or *before* (“ $\text{bef}_{\mathbf{T}}$ ”) the embedding episode $\text{Emb}_{\mathbf{T}}$, e.g., the utterance event (depending on the aspectual class of Φ and other factors). Tense trees also provide the “points of orientation” (cf., [15; 22]), such as the reference point for a perfect episode or the immediately preceding past episode in a succession of simple past-tensed sentences, and the “*orients*” predication captures this.

As mentioned, the information in the *past* operator and in the time adverbial are essentially interpreted *conjunctively* as in [7]. For detailed discussion of the interpretation of temporal and other adverbials, see [10; 14]. Although we do not yet have grammars and LF rules for large subsets of English, the compositional approach of EL has been proven practical by a start-up implementation that has successfully handled dialogues in a planning system in the TRAINS domain [1]. See [10] for some of our grammar fragments. The ARMS application used a grammar and interpretation mechanism internal to the Boeing Co.

6 Concluding Remarks

We think there is cause for optimism about the possibility of constructing theoretical and computational frameworks for full NLU. Our own efforts in that direction have led to a rather well-integrated conception of syntax, LF, knowledge representation, context, and inference, and of the interfaces linking these into an organic whole. The conception is not yet complete or fully “debugged,” but it is sufficiently far along to have provided a basis for diverse start-up implementations. Unlike most past implemented NLU and inference systems, these implementations strenuously avoid cutting corners in syntax, LF computation, and most of all, knowledge representation and inference. Thus, we have reason to regard the theoretical framework and the implementations as a solid and extensible basis for further work toward the ultimate goal of general NLU.

References

- [1] J. Allen & L. K. Schubert. The TRAINS Project. TR 382, U. of Rochester, 1991.
- [2] J. Barwise. *The Situation in Logic*. CSLI, Stanford, 1989.
- [3] E. Charniak. Motivation analysis, abductive unification, and nonmonotonic equality. *Artificial Intelligence*, 34:275–295, 1988.
- [4] E. Charniak & R. Goldman. A logic for semantic interpretation. In *Proc. 26th Annual Meeting of the ACL*, 87–94, Buffalo, NY, 1988.
- [5] D. Davidson. The logical form of action sentences. In N. Rescher, editor, *The Logic of Decision and Action*. U. of Pittsburgh Press, 1967.
- [6] K. Devlin. *Logic and Information*. Cambridge U. Press, UK, 1991.
- [7] D. Dowty. Tense, time adverbs and compositional semantic theory. *Linguistics & Philosophy*, 5:23–55, 1982.
- [8] J. R. Hobbs, W. Croft, T. Davies, D. Edwards, & K. Laws. Commonsense metaphysics and lexical semantics. *Computational Linguistics*, 13:241–250, 1987.
- [9] J. R. Hobbs, M. Stickel, P. Martin, & D. Edwards. Interpretation as abduction. In *Proc. 26th Annual Meeting of the ACL*, 95–103, Buffalo, NY, 1988.
- [10] C. H. Hwang. *A Logical Approach to Narrative Understanding*. PhD thesis, U. of Alberta, Canada, 1992.
- [11] C. H. Hwang & L. K. Schubert. EL: A formal, yet natural, comprehensive knowledge representation. In *Proc. AAAI-93*, 676–682. Washington, DC, 1993.
- [12] C. H. Hwang & L. K. Schubert. Episodic Logic: A situational logic for natural language processing. In *Situation Theory and its Applications, V. 3*, CSLI, Stanford, 303–338.
- [13] C. H. Hwang & L. K. Schubert. Episodic Logic: A comprehensive, natural representation for language understanding. *Minds and Machines, V. 3*, 381–419.
- [14] C. H. Hwang & L. K. Schubert. Interpreting tense, aspect and time adverbials: A compositional, unified approach. In *Proc. 1st Internat. Conf. on Temporal Logic (ICTL '94)*, 238–264. Bonn, Germany, 1994.
- [15] G. Leech. *Meaning and the English Verb (2nd ed)*. Longman, London, UK, 1987.
- [16] A. Namioka, C. H. Hwang, & S. Schaeffer. Using the inference tool EPLOG for a message processing application. *Internat. J. of Expert Systems*, In print.
- [17] S. Schaeffer, C. H. Hwang, J. de Haan, & L. K. Schubert. *The User's Guide to EPLOG*. Edmonton, Canada, 1991.
- [18] L. K. Schubert. Formal foundations of Episodic Logic. In preparation.
- [19] L. K. Schubert & C. H. Hwang. An Episodic knowledge representation for narrative texts. In *Proc. 1st Int. Conf. Principles of Knowledge Representation & Reasoning (KR '89)*, 444–458, Toronto, Canada, 1989.
- [20] L. K. Schubert & C. H. Hwang. Picking reference events from tense trees: A formal, implementable theory of English tense-aspect semantics. In *Proc. Speech and Natural Language, DARPA Workshop*, 34–41, Hidden Valley, PA, 1990.
- [21] J. Traugott. Nested resolution. In *Proc. 8th International Conference on Automated Deduction (CADE-8)*, 394–402. Springer-Verlag, Oxford, UK, 1986.
- [22] B. L. Webber. Tense as discourse anaphor. *Computational Linguistics*, 14:61–73, 1988.