

1999 IRCS TECHNICAL REPORT ABSTRACTS

A Formalism for Resource-Oriented Planning Universal Grammar and the Acquisition of Semantic Knowledge (Ph.D. Dissertation work completed at the University of Maryland)

Julien Musolino
IRCS-99-01

This dissertation explores the way in which English-speaking children acquire the meaning of sentences containing negation and quantified noun phrases (QNPs). This investigation is based on a series of psycholinguistic experiments designed to assess children's comprehension of sentences like 'Every horse didn't jump over the fence' or 'Cookie Monster didn't eat two slices of pizza' among others. The major finding is that children around the age of 5 do not interpret these sentences the way adult speakers of English do. This finding raises the following questions (a) How and why do children's interpretations of sentences containing negation and quantified noun phrases differ from those of adults? (b) How do children manage to converge onto the adult system of interpretation? Regarding the first question, it appears that children's non-adult interpretations are nevertheless systematic, i.e. governed by principle. Specifically, children (unlike adults) are found to map overt syntactic relations between QNPs and negation and their relative semantic interpretation isomorphically. This, however, is just a descriptive generalization. The observation of isomorphism is treated as an epiphenomenon, derived from the interplay between a universally encoded dichotomy splitting the class of QNPs and learnability considerations. Regarding the second question, I show that children can move from their system of interpretation to the adult system solely on the basis of positive evidence and thus, that the observed difference does not create a learnability problem. In summary, this dissertation uncovers a new area where the linguistic behavior of children and adults diverge: the comprehension of sentences containing negation and quantified noun phrases. The rest of the dissertation is a methodological statement, namely that it is not only desirable but also possible to account for the observed difference between children and adults without invoking any differences between the two groups beyond minimal conceptual necessity. To the extent that this goal is achieved, the present investigation emphasizes the role played by the theory of Universal Grammar and language learnability in helping us understand language development and its biological basis.

Where to Look? Automating Certain Visual Attending Behaviors of Human Characteristics (Ph.D. Dissertation)

Sonu Chopra-Khullar
IRCS-99-02

This thesis proposes a computational framework for generating visual attending behavior in an embodied simulated human agent. Such behaviors directly control eye and head motions, and guide other actions such as locomotion and reach. The implementation of these concepts, referred to as the AVA, draws on empirical

and qualitative observations known from psychology, human factors and computer vision. Deliberate behaviors, the analogs of scanpaths in visual psychology, compete with involuntary attention capture and lapses into idling or free viewing. For efficiency, the embodied agent is assumed to have access to certain properties of the 3D world (scene graph) stored in the graphical environment. When information about a task is known, the scene graph is queried. When an agent lapses into free viewing or idling, no task constraints are active so a simplified image analysis technique is employed to select potential directions of interest. Insights provided by implementing this framework are: a defined set of parameters that impact the observable effects of attention, a defined vocabulary of looking behaviors for certain motor and cognitive activity, a defined hierarchy of three levels of eye behavior (endogenous, exogenous and idling) and a proposed method of how these types interact, a technique of modifying motor activity based on visual inputs, and a technique that allows for anticipation and interleaving of eye behaviors for sequential motor actions. AVA generated behavior is emergent and responds to environment context and dynamics. Further, this method animates behavior at interactive rates. Experiments supporting several combinations of environment and attending conditions are demonstrated, followed by a discussion of an evaluation of AVA effectiveness.

Lifetime of Human Visual Sensory Memory: Properties and Neural Substrate (Ph.D. Dissertation work completed at New York University)

Wei Yang
IRCS-99-03

The classic partial-report procedure was modified to optimize the condition to measure the transient decay of visual sensory memory (VSM, also known as iconic memory). A model was developed to isolate the VSM and visual working memory (VWM) underlying the partial-report performance. The decay of VSM in each subject was well characterized by a single exponential function, thus a lifetime could be defined for VSM decay in individual subjects. It was found that intensive practice with partial-report task prolonged VSM lifetime. This practice effect shows an unexpected adaptive property of VSM and reveals VSM lifetime as a specific dimension for perceptual learning. Of the stimulus parameters, a change of the mean luminance of the stimuli from that of the background shortened the VSM lifetime. Such a "luminance effect" is consistent with the temporal properties of the spatial frequency channels in the visual pathway, most likely revealing the differences in the time course of the decay of the memory traces in these channels. To identify the neural substrate of VSM, the lifetime for the decay of the neural activation trace in the human primary visual cortex (area V1) were deduced from the visually evoked potential (VEP) recordings for each subject. There was a precise match between the V1 lifetime and the VSM lifetime for each subject. The match even held when both psychophysical and physiological lifetimes were prolonged by practice, or shortened by enhancing the mean

luminance of the stimuli above the background. This precise match indicates that the cortical location for VSM is V1.

Path Constraints for Databases with or without Schemas (Ph.D. Dissertation)

**Wenfei Fan
IRCS-99-04**

This dissertation introduces a path constraint language and investigates its associated implication and finite implication problems.

This path constraint language has proven useful in a variety of database contexts, ranging from semistructured data as found for instance on the Web, to structured data such as data in object-oriented databases. It is capable of expressing natural integrity constraints that are not only a fundamental part of the semantics of the data, but are also important in query optimization.

Path constraint implication is investigated for two models for semistructured data: the semistructured data model and the deterministic data model. Databases in these models are unconstrained by any type system or schema. For the semistructured data model, it is shown that, despite the simple syntax of the constraint language, its associated implication problem is r.e. complete and its finite implication problem is co-r.e. complete. However, in light of these undecidability results, several decidable fragments of the constraint language are identified. These fragments suffice to express many important integrity constraints such as referential integrity, inverse relationships and local database constraints. For the deterministic data model, it is shown that the implication and finite implication problems for the path constraint language are finitely axiomatizable and decidable in cubic-time.

Path constraint implication is also studied for structured data, i.e., data constrained by a schema. In the context of three practical object-oriented data models, a number of complexity results on the implication and finite implication problems for the path constraint language are established. In addition, the interaction between path constraints and type systems is investigated. It is demonstrated that adding a type to the data may in some cases simplify the analysis of path constraint implication, and in other cases make it harder. More specifically, it is shown that there is a path constraint implication problem that is decidable in PTIME in the untyped context, but that becomes undecidable when a type system is added. On the other hand, there is an implication problem that is undecidable in the untyped context, but that becomes not only decidable in cubic-time but also finitely axiomatizable when a type system is imposed.

Compositional Semantics for Unification-based Linguistic Formalisms

**Shuly Wintner
IRCS-99-05**

Contemporary linguistic formalisms have become so rigorous that it is now possible to view them as very high level declarative programming languages. Consequently, grammars for natural languages can be viewed as programs; this view enables the application of various methods and techniques that were proved useful for programming languages to the study of natural languages. This paper adapts the notion of program composition, well developed in the context of logic programming languages, to the domain of linguistic formalisms. We study alternative definitions for the semantics of such formalisms, suggesting a denotational semantics that we show to be compositional and fully-abstract. This facilitates a clear, mathematically sound way for defining grammar modularity.

A Motion Control Scheme for Animating Expressive Arm Movement (Ph.D. Dissertation)

**Diane Chi
IRCS-99-06**

Current methods for figure animation involve a tradeoff between the level of realism captured in the movements and the ease of generating the animations. We introduce a motion control paradigm that circumvents this tradeoff--it provides the ability to generate a wide range of natural-looking movements with minimal user labor.

Effort, which is one part of Rudolf Laban's system for observing and analyzing movement, describes the qualitative aspects of movement. Our motion control paradigm simplifies the generation of expressive movements by proceduralizing these qualitative aspects to hide the non-intuitive, quantitative aspects of movement. We build a model of Effort using a set of kinematic movement parameters that defines how a figure moves between goal keypoints. Our motion control scheme provides control through Effort's four-dimensional system of textual descriptors, providing a level of control thus far missing from behavioral animation systems and offering novel specification and editing capabilities on top of traditional keyframing and inverse kinematics methods. Since our Effort model is inexpensive computationally, Effort-based motion control systems can work in real-time.

We demonstrate our motion control scheme by implementing EMOTE (Expressive MOTion Engine), a character animation module for expressive arm movements. EMOTE works with inverse kinematics to control the qualitative aspects of end-effector specified movements. The user specifies general movements by entering a sequence of goal positions for each hand. The user then expresses the essence of the movement by adjusting sliders for the Effort motion factors: Space, Weight, Time, and Flow. EMOTE produces a wide range of expressive

movements, provides an easy-to-use interface (that is more intuitive than joint angle interpolation curves or physical parameters), features interactive editing, and real-time motion generation.

A Computational Analysis of Information Structure Using Parallel Expository Texts in English and Japanese (Ph.D. Dissertation)
Nobo Komagata
IRCS-99-07

This thesis concerns the notion of ‘information structure’: informally, organization of information in an utterance with respect to the context. Information structure has been recognized as a critical element in a number of computer applications: e.g., selection of contextually appropriate forms in machine translation and speech generation, and analysis of text readability in computer-assisted writing systems.

One of the problems involved in these applications is how to identify information structure in extended texts. This problem is often ignored, assumed to be trivial, or reduced to a sub-problem that does not correspond to the complexity of realistic texts. A handful of computational proposals face the problem directly, but they are generally limited in coverage and all suffer from lack of evaluation. To fully demonstrate the usefulness of information structure, it is essential to apply a theory of information structure to the identification problem and to provide an evaluation method.

This thesis adopts a classic theory of information structure as binomial partition between theme and rheme, and captures the property of theme as a requirement of the contextual-link status. The notion of ‘contextual link’ is further specified in terms of discourse status, domain-specific knowledge, and linguistic marking. The relation between theme and rheme is identified as the semantic composition of the two, and linked to surface syntactic structure using Combinatory Categorical Grammar. The identification process can then be specified as analysis of contextual-link status along the linguistic structure.

The implemented system identifies information structure in real texts in English. Building on the analysis of Japanese presented in the thesis, the system automatically predicts contextually-appropriate use of certain particles in the corresponding texts in Japanese. The machine prediction is then compared with human translations. The evaluation results demonstrate that the prediction of the theory is an improvement over alternative hypotheses. We then conclude that information structure can in fact be used to improve the quality of computational applications in practical settings.

Tree Description Grammars and Underspecified Representations (Ph.D. Dissertation work completed at the University of Tuebingen)
Laura Kallmeyer
IRCS-99-08

In this thesis, a new grammar formalism called (local) Tree Description Grammar (TDG) is presented that generates tree descriptions. This grammar formalism brings together some of the central ideas in the context of Tree Adjoining Grammars (TAG) on the one hand, and approaches to underspecified semantics for scope ambiguities on the other hand.

First a general definition of TDGs is presented, and afterwards a restricted variant called local TDGs is proposed. Since the elements of a local TDG are tree descriptions, an extended domain of locality as in TAGs is provided by this formalism. The tree descriptions generated by local TDGs are such that the dominance relation (i.e. the reflexive and transitive closure of the parent relation) need not be fully specified. Therefore the generation of suitable underspecified representations for scope ambiguities is possible.

The generative capacity of local TDGs is greater than the one of TAGs. Local TDGs are even more powerful than set-local multicomponent TAGs (MC-TAG). However, the generative capacity of local TDGs is restricted in such a way that only semilinear languages are generated. Therefore these languages are of constant growth, a property generally ascribed to natural languages. Local TDGs of different rank can be distinguished depending on the form of derivation steps that are possible in these grammars. This leads to a hierarchy of local TDGs. For the classes of this hierarchy, a pumping lemma is proven.

In order to describe the relation between two languages, synchronous local TDGs are introduced. The synchronization with a second local TDG does not increase the generative power of the grammar in the sense that each language generated by a local TDG that is part of a synchronous pair of local TDGs, also can be generated by a single local TDG. This formalism of synchronous local TDGs is used to describe a syntax-semantics interface for a fragment of French which illustrates the derivation of underspecified representations for scope ambiguities with local TDGs. In this framework, island constraints for quantifier scope ambiguities arise as a natural consequence of the locality of the grammar.