Roles, Co-Descriptors, and the Formal Representation of Quantified English Expressions

William A. Martin [1]

Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

A scheme is proposed for representing the logical form of English sentences, wherein the meaning of a network node depends on how it is related through "role-in" links to nodes representing more aggregate entities. It is argued that roles are a natural device for capturing many linguistic and philosophic distinctions, and that they are convenient for computational processing. In particular, it is shown how role-in links may be advantageously used in lieu of quantifier scope to represent quantificational dependencies.

1. Introduction

In representing the semantics of English sentences, it is traditional to distinguish "logical form" from semantic content. Chomsky [2], for instance, introduces LF (his version of logical form) as a linguistic level of representation between syntax (phrase structure) and semantics.

Representations of logical form, which are typically based on predicate calculus or lambda calculus, must be carefully chosen for empirical adequacy as well as computational convenience. There are at least three general ways in which predicate or lambda calculus might be carried over into representations of logical form:

(1a) directly (as by a theorem prover);

(1b) by replacing quantification with dynamically scoped iteration procedures (as in Woods [30]); or

(1c) by translating into a semantic network formalism (as in Hendrix [10]).

The second and third are more process-oriented; that is, they more closely relate elements of the representation and states in the interpretation process. In my view, a process orientation is needed if issues of efficiency are to be discussed at all.[2] One key to efficiency is the delaying of decisions as long as possible. This is easier to achieve when differences between alternative interpretations are minimized.

The two process-oriented approaches noted above differ in their treatment of scope dependencies. For example, we can model "each bottle" in

(2) Each cork is fastened to each bottle by a small wire basket.

with a quantifier like ($\forall$x: bottle), and then translate that into an explicit FOR loop clause. Alternatively, we can employ Skolem functions, with "wire" and "basket" modeled as functions of a variable ranging over "bottle". These functions and variables can then be represented in a semantic network, along the lines discussed in "What's in a Link?" [29].

I choose to adopt a particular semantic network approach wherein it is more appropriate to view meanings of nodes *de*compositionally (i.e., as dependent upon what they are constituents of) rather than compositionally (i.e., as fully determined by the meanings

[1] This paper was substantially revised after Professor Martin's death by a group of his colleagues to satisfy length constraints and to take advantage of the reviewers' extensive and thoughtful comments. We wish to thank the editor and reviewers for their cooperation under these unfortunate circumstances. Prof. Martin's original acknowledgment follows.

Ken Church, Lowell Hawkinson, Mitchell Marcus, Peter Szolovits, and Lucia Vaina read an earlier version of this manuscript and made many helpful comments. Ellen Lewis and Anne Schmitt did an excellent job of preparing the manuscript and figures. This research was supported by the Defense Advance Research Projects Agency and monitored by the Office of Naval Research under contracts N00014-75-C-0661 and N00014-80-C-0505.

[2] It is unclear how to pose efficiency questions in a non-process-oriented framework, where even effectiveness may be rejected as a desideratum.

0362-613x/81/030137-12$01.00

of their constituents). More specifically, the meaning of a network node depends (decompositionally) on how it is related though "role-in links"[3] to nodes representing more aggregate entities. The decompositional view is preferable for explaining a number of linguistic phenomena. I will show, in this paper, that many important linguistic and philosophical distinctions can be analyzed as differences in role-in linkages. Among these are the intensional/extensional, referential/attributive, and distributive/collective distinctions.

My approach turns out to be more complex than some, but I will argue that it captures subtle cases in a more natural way and that it is superior for computational purposes. Indeed it has been my experience (as in the development of MACSYMA [13]) that a representational formalism with more constructs often permits more efficient processing than a formalism based on a smaller set of more primitive operators. I suspect that this experience will strike home with many readers familiar with the difficulties encountered by predicate-calculus-based theorem provers.

2. Roles

Figures 1a-1c illustrate the sort of semantic network I use. (In these and subsequent figures, only those nodes and links are shown that are relevant to the issues under discussion. By convention, links without arrowheads are to be taken as directed upwards.)

The networks shown in this figure have two kinds of nodes, *generic* and *individual*, and three kinds of links, *inherits-from*, *may-be-described-as*, and *role-in*. Generic nodes (e.g. ARCH) represent generic objects; individual nodes (e.g. ARCH-1 and ARCH-2) represent instances of generic objects. Inherits-from links are used to say that the individuals ARCH-1 and ARCH-2 inherit properties from the generic ARCH. May-be-described-as links are used to indicate that an ARCH is a STRUCTURE.[4] Finally, role-in links are used to indicate that R1 and R2 are parts of the generic ARCH. The "meaning" of a node depends primarily upon role-in links that connect it to nodes representing more aggregate entities.

My representation is similar to several in the literature. Generic nodes correspond to Minsky frames [16] and roles (parts of generic objects) correspond to slots (parts of frames). One can also draw analogies with set theory. May-be-described-as links and inherits-from are similar to set inclusion and set membership.

[3] Role-in links could be viewed logically as Skolem dependencies.

[4] May-be-described-as and inherits-from are similar to "is-a" links.

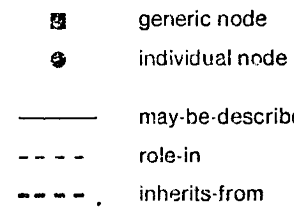

Figure 1a. Notation used in networks.

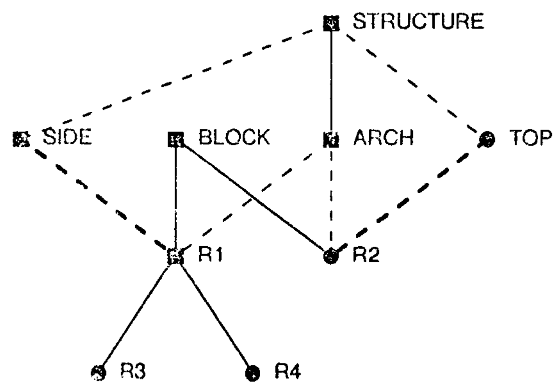

Figure 1b. Simplified structural description of an arch.

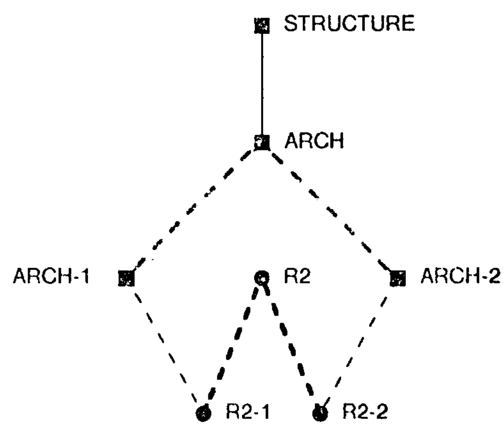

Figure 1c. An example of individual nodes.

That is,

(3a) ARCH may-be-described-as STRUCTURE

(3b) ARCH-1 inherits-from ARCH

are analogous to

(4a) the set of all ARCHes is a subset of the set of all STRUCTUREs

(4b) ARCH-1 is a member of the set of all ARCHes

However, there is an important difference in the semantic interpretation: may-be-described-as links are interpreted as relating *intensions* expressed by generic nodes, whereas subset relations operate on sets of individuals. I consider the intension/extension distinction extremely important, because intensional representations appear to be more natural for certain sentences.[5]

My representation does not deal with contextual problems by assuming a set of objects called contexts (or partitions) and then somehow stipulating a context link for each node. Instead, I argue that roles allow a more natural modeling of the context phenomena. For example, Figure 2 shows how Hendrix [10] represents the fact that legal persons can own physical objects. Rectangles (and squares) represent partitions (contexts). An individual or set is in a partition if its representing circle is in the rectangle representing that partition. In this figure, the individual I is an implication that any individual owning X has an individual agent Y, which is a legal person; an individual object Z, which is a physical object; an individual start-time t_1, which is a time; and an individual end-time t_2, which is also a time.

This same information is shown in Figure 3, using the notation presented here. Figure 3 eliminates the individuals I and X, the implication node, the context surrounding I, and the context surrounding Y, Z, t_1, and t_2. Some of these are replaced in Figure 3 with a richer variety of node types; I believe these types have better linguistic and philosophical motivations than Hendrix's abstract individuals, contexts, and implications.

3. Reference and Definite Descriptions

Russell's [22] analysis of a singular definite description required the existence of a unique object satisfying the description in order for the expression to denote anything, and hence it fails to account for the successful reference of a noun phrase like "the clock" in

(5) Did you wind the clock?

Strawson [25] views these definite descriptions in relativistic terms: even when the referent cannot be determined in an absolute sense, as a particular clock for example, the description still has some meaning relative to whatever the context turns out to be.

(6) When shall we say that a hearer knows what particular is being referred to by a speaker? Consider first the following case. A speaker tells a story which he claims to be factual. It begins, "A man and a boy were standing by a fountain," and it continues: "The man had a drink". Shall we say that the hearer knows which or what particular is being referred to by the subject expression in the second sentence? We might say so. For, of a certain range of two particulars the words "the man" serve to distinguish the one being referred to, by means of a description which applies only to him. But though this is, in a weak sense, a case of identification, I shall call it only a *story-relative* or, for short, a *relative* identification. For it is identification only relative to a range of particulars (a range of two members) which is itself identified only as the range of particulars being talked about by the speaker. That is to say, the hearer, hearing the second sentence, knows *which* particular creature is being referred to *of the two particular creatures being talked about by the speaker*; but he does not, without this qualification, know what particular creature is being referred to. The identification is within a certain story told by a certain speaker. It is identification within his story but not identification within history.

Following this line of thought, sentence (5) might be represented as follows, where clock-1 describes a particular clock and the-clock-1 describes the clock relative to the discourse.

(7)
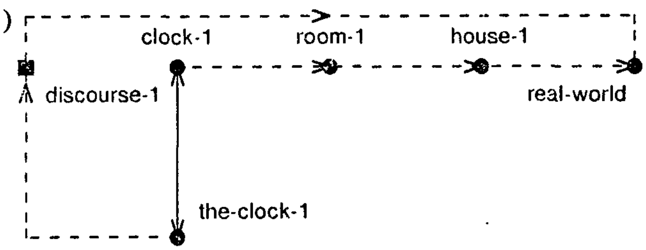

Clock and Clock-1 are *co-descriptors* because they both describe the same individual. (The term co-description is chosen over co-reference or co-denotation to avoid implications of intensionality and existence.)

Co-description is appropriate in (7), but surely not every occurrence of the-clock will refer to clock-1. For example, Webber [28] points out that in

(8a) Wendy bought a yellow T-shirt that Bruce had liked.

(8b) It cost twenty dollars.

an appropriate description of the entity referred to by "it" is not "the yellow T-shirt that Bruce had liked," since (8a) is true even if Bruce had liked several T-shirts (and both the speaker and the listener were aware of the fact). Nor is it "the yellow T-shirt that Bruce had liked and Wendy bought," since (8b) can be true even if Wendy had bought several such T-shirts. An appropriate description is something like "the yellow T-shirt that Bruce had liked and that

[5] I choose to avoid the question of formal adequacy because almost all systems, intensional as well as extensional, are formally adequate in the sense that they are capable of representing all (effective) functions. A much more interesting question is whether they are capable of representing all functions *in a natural way*.

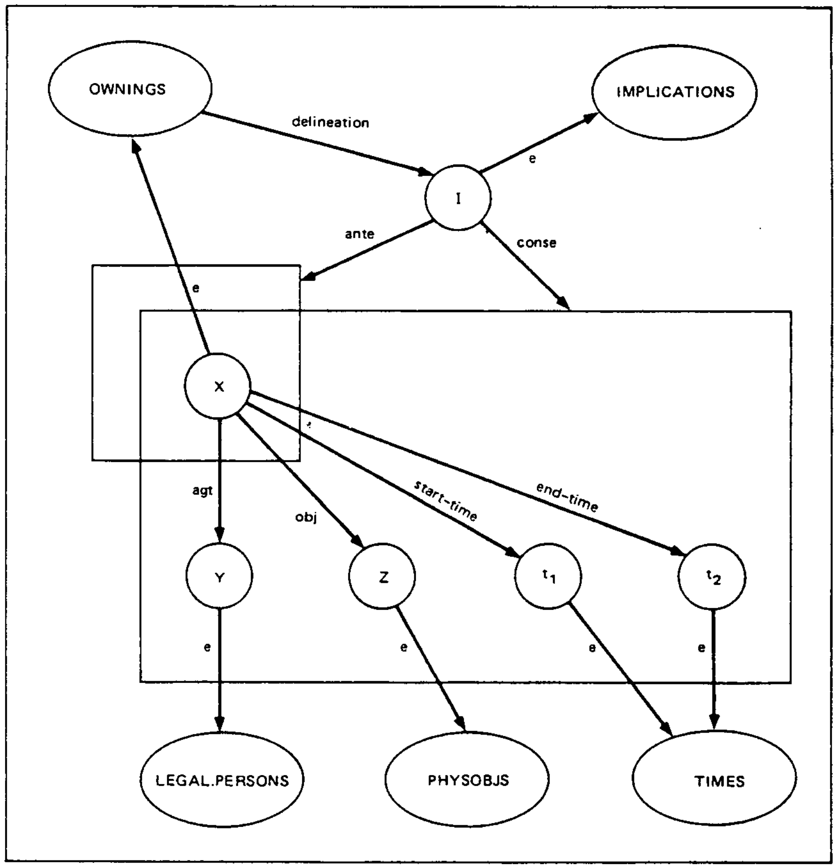

Figure 2. The delineation theorem of ownings (Hendrix).

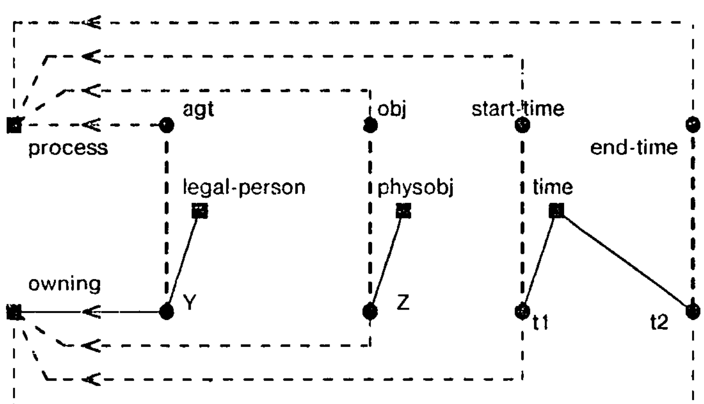

Figure 3. The information of Figure 2 in the notation of this paper.

Wendy bought and that was mentioned in (8b)". Accordingly, (7) needs to be amended so that the-clock-1 is relative to the sentence "did you wind the clock".

(9)

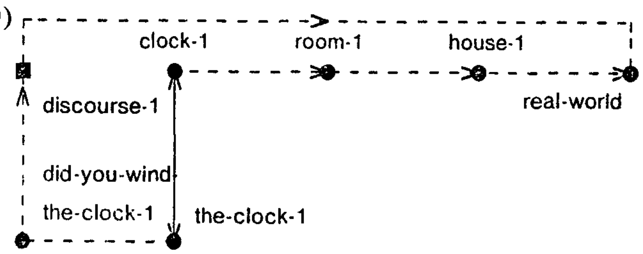

3.1 Referential and Attributive Use of Definite Descriptions

Donnellan [3] has suggested that definite descriptions like "the clock" have two uses: referential and attributive. Attributive use asserts something about "whoever or whatever is the so-and-so"; referential use is intended to allow the hearer to identify at once, in his own terms, who or what is being talked about.

(10) I will call the two uses of definite descriptions I have in mind the attributive use and the referential use. A speaker who uses a definite description attributively in an assertion states something about whoever or whatever is the so-and-so. A speaker who uses a definite description referentially in an assertion, on the other hand, uses the description to enable his audience to pick out whom or what he is talking about and states something about that person or thing. In the first case the definite description might be said to occur essentially, for the speaker wishes to assert something about whatever or whoever fits that description; but in the referential use the definite description is merely one tool for doing a certain job— calling attention to a person or thing— and in general any other device for doing the same job, another description or a name, would do as well. In the attributive use, the attribute of being the so-and-so is all important, while it is not in the referential use.

To illustrate this distinction, in the case of a single sentence, consider the sentence, "Smith's murderer is insane". Suppose first that we come upon poor Smith foully murdered. From the brutal manner of the killing and the fact that Smith was the most lovable person in the world, we might exclaim, "Smith's murderer is insane". I will assume, to make it a simpler case, that in a quite ordinary sense we do not know who murdered Smith (though this is not in the end essential to the case). This I shall say, is an attributive use of the definite description.

The contrast with such a use of the sentence is one of those situations in which we expect and intend our audience to realize whom we have in mind when we speak of Smith's murderer and, most importantly, to know that it is this person about whom we are going to say something.

For example, suppose that Jones has been charged with Smith's murder and has been placed on trial. Imagine that there is a discussion of Jones's odd behavior at his trial. We might sum up our impression of his behavior by saying, "Smith's murderer is insane". If someone asks to whom we are referring, by using this description, the answer here is "Jones". This, I shall say, is a referential use of the definite description.

The distinction is brought out nicely by the following pair of sentences from Moore [17] and their corresponding semantic net representations (Figures 4 and 5).

(11a) The President has been married since 1945. *referential*

(11b) The President has lived in the White House since 1800. *attributive*

Sentence (11a) refers to the person who is currently President, while sentence (11b) refers to whoever was president at each point in time since 1800. The truth of (11b) is not dependent on who is filling the role of President. In fact, we can go even farther. In the past the President has been killed and it has taken a while to swear in the new President. During this time there is no President, yet (11b) is still true. In (11a), by contrast, the referent of "the President", president-1, is not a role in "the President is married", as shown in Figure 4. Therefore, we are not at liberty to have the referent change with time.[6]

The referential/attributive distinction might have been ascribed to differences of scope in quantified expressions.

(12a) (the p : president) (∀ t since 1945) p has been married at time t. *referential*

(12b) (∀ t after 1800) (the p : president) p has lived in the white house at time t. *attributive*

(13a) ((λ (p) (true-at-time (married p)
(every (t) (after t 1945))))
(the (y) (pres y))) *referential*

(13b) (true-at-time (live-in (the (y)
(pres y)) white-house)
(every (t) (after t 1800))) *attributive*

---

[6] A method for the sequential binding of quantified variables in determining the truth of an expression was introduced by Hintikka [11]. Although he was working with predicate calculus, he obviously had in mind the same general strategy proposed here.

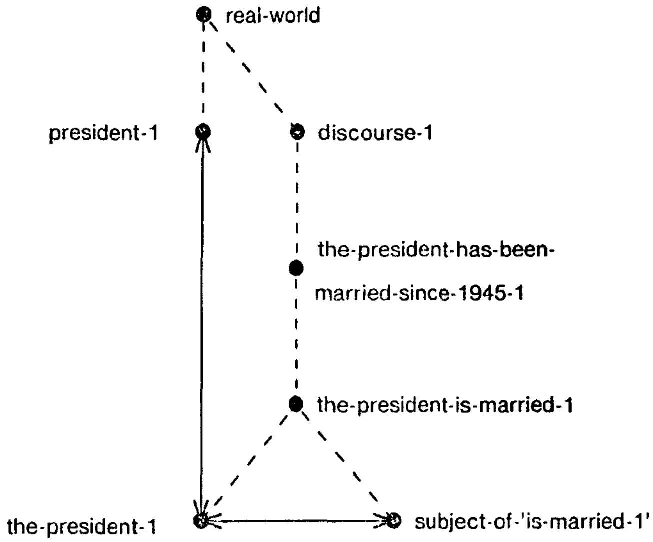

Figure 4. Referential use of *the president.*

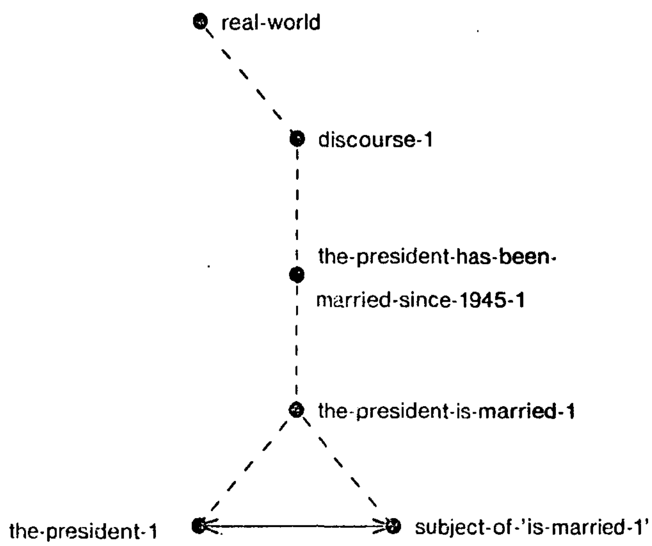

Figure 5. Attributive use of *the president.*

My representation, in contrast, makes it easy for the interpreter to delay the decision because it minimizes the difference between the two cases. It is also possible for the interpreter never to decide which reading is intended. For example, in a sentence such as, "The president has owned a terrier since February," the listener will not be able to distinguish between the two readings. If a later sentence requires the listener to make a distinction, only a minimal alteration to the representation is required. Partee [19] makes an argument that further supports this approach.

(14) ... having a particular individual in mind (the "referential" case) and knowing nothing about an individual other than some descriptive phrase (the "attributive" case) may be just two extremes on a continuum of "vividness". One may consider, for instance, the case of a detective tracking down a criminal and obtaining more and more clues, including fingerprints, voice recordings, photographs of varying clarity, etc. It is not at all clear at what point the detective, who may be described as "looking for the man who did so-and-so" stops looking for "whoever it is that did so-and-so" and starts looking for a particular individual.

3.2 Subject Co-Descriptors (Verb Phrase Deletion)

Both Figures 4 and 5 employ co-description between *the-President* and *subject-of-x*, where *x* is the sentential predicate. Sentences (15a) and (15b) from Sag [23] offer further support for this convention.

(15a) The chickens are ready to eat, and the children are too.

(15b) John likes flying planes, and Bill does too.

In these examples, the subject plays the same role in both conjuncts. That is, if the chickens are ready to be eaten, then the children are also ready to be eaten, and if the chickens are ready to do the eating, then the children are ready to do the eating. However, (15a) cannot mean that the chickens are ready to be eaten and the children are ready to do the eating. Sentence (15b) illustrates the same point. It is ambiguous whether John likes to fly planes himself or whether he would prefer someone else do the flying, but whichever way he likes the flying done, Bill will want it done the same way.

The following sentence from Partee [20] has the same sort of ambiguity as (15a) and (15b) above; I will use it to illustrate how such ambiguities are dealt with in my scheme of representation.

(16) The prosecutor believed that he would win the case, and so did the defense attorney.

The missing verb phrase can be understood in two ways, either that the defense attorney would win or that the prosecutor would win. This can be explained in terms of my representation, as shown in (17a) and (17b) below. Note that he-1 is co-descriptive with the subject role in (17a) and with the-prosecutor-1 in (17b). Thus (17a) represents the sloppy attributive interpretation where he-1 refers to whatever the subject happens to be, whereas (17b) represents the non-sloppy referential interpretation where he-1 refers to the-prosecutor-1 in both conjuncts.

(17a) *sloppy*

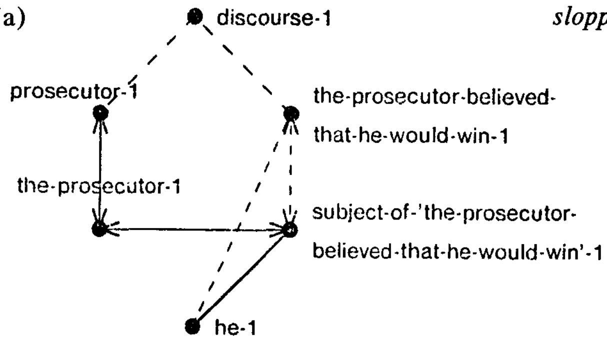

(17b) *nonsloppy*

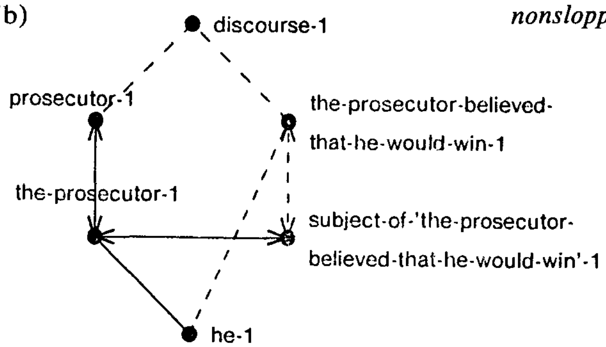

3.3 Restrictive and Non-Restrictive Modifiers

Another opposition which can be captured with the attributive/referential distinction is that between restrictive and non-restrictive modifiers.

(18a) My uncle, who is 70, is bald. *non-restrictive*

(18b) My uncle who is 70 is bald. *restrictive*

Sentence (18a) has a non-restrictive relative clause, "who is 70", modifying, "my uncle". That is, "who is 70" is not used to pick out the uncle who is bald, but just to give extra information about him. This sentence is equivalent to the conjunction, "My uncle is bald & my uncle is 70." By contrast, in the restrictive reading (18b), "who is 70" picks out a particular uncle. The distinction between (18a) and (18b) can be captured as follows.

(19a) *non-restrictive*

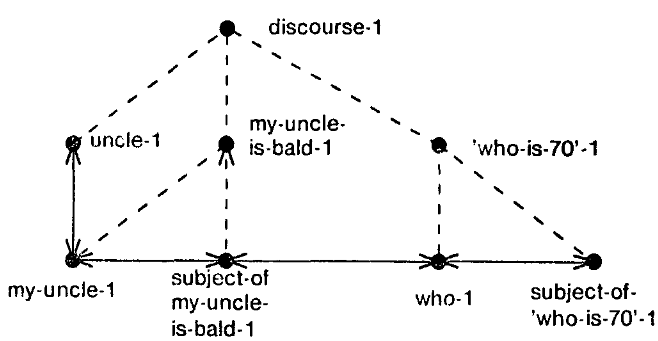

(19b) *restrictive*

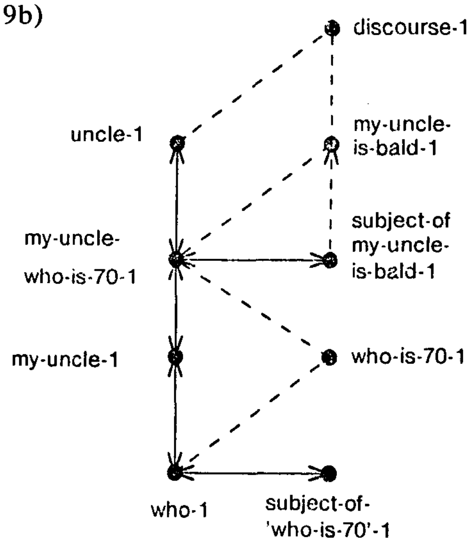

Note that the head of the relative clause must be either attributive or referential for both main and relative clauses. It is not possible to have (20), for example. This argues that the relative pronoun "who" should be treated as a co-descriptor of a sentence role of the main clause, not a discourse role.

(20) *The President, who has been married since 1945, has lived in the White House since 1800.

4. Discourse Iteration

We have seen that explicit nesting of structural descriptions can explain the referential/attributive distinction, sloppy binding, and restrictive/descriptive relative clauses. This section will show that quantifier interpretation can also be explained in a similar manner. Consider sentence (21), which is three ways ambiguous, as shown in (22a)-(22c).

(21) Every boy wants a lion.

(22a) *They will take any lion.*

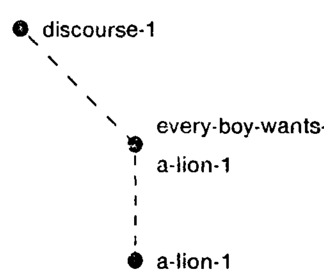

(22b) *Each wants his own specific lion.*

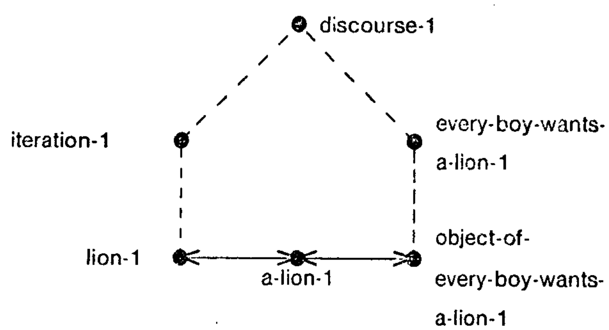

(22c) *They all want the same specific lion.*

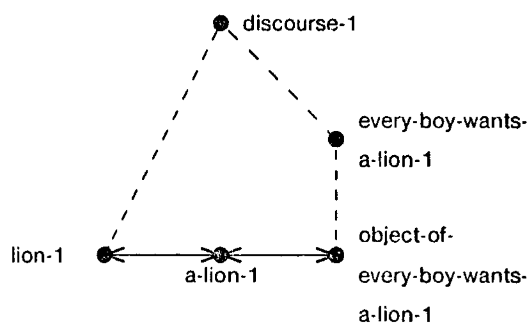

The three readings (22a)-(22c) can be interpreted as follows. In each case, want-a-lion must be true for each boy. In (22a), a-lion-1 is attributive: it can be picked as needed for each wanting. In (22b), a-lion-1 is referential, but it refers to a role in an iteration of wanting at the discourse level. (Not surprisingly, given my account, this is a difficult reading for people to construct.) Finally, (22c) gives the by-now-familiar referential case where every boy is constrained to wanting the same lion.[7] Let us now turn to a discussion of the discourse iteration construct used in (22b).

4.1 Procedural Representation of Knowledge

By *procedural representation of knowledge*, I mean knowledge in the form of procedures. For example, the knowledge that every boy wants a lion could be cast in the form "if one were to check every boy in question and count those who want lions, then the count of those who want lions would be equal to the count of the boys checked". Woods [30] introduced a *FOR* iteration construct for representing knowledge of quantified propositions procedurally.[8] Examples of the use of this construct are:

(23a) (FOR EVERY X / CLASS : (P X) ; (Q X))
Every X in CLASS that satisfies P
also satisfies Q.

(23b) (FOR SOME X / CLASS : (P X) ; (Q X))
Some X in CLASS that satisfies P
also satisfies Q.

(23c) (FOR GEN X / CLASS : (P X) ; (Q X))
A generic X in CLASS that satisfies P
also satisfies Q.

(23d) (FOR THE X / CLASS : (P X) ; (Q X))
The single X in CLASS that satisfies P
also satisfies Q.

The FOR statement applies the filter P(x) in turn to each element in the CLASS, and then applies Q(x) to those elements that pass the filter. The key words EVERY, SOME, GEN, THE, and so on, specify a particular enumeration function.

I prefer to represent quantified phrases in terms of role-in links (interpretable as Skolem dependencies) rather than iteration procedures, and to use predicates on intensions and sets in addition to predicates on individuals. By converting Woods' FOR construct to my notation, the difference between the two referential readings of "every boy wants a lion" can be spelled out in more detail as shown in (24a) and (24b) below. (Note that (24a) is a slight elaboration of (22b), and that (24b) is identical to (22c)). (24a) captures both the EVERY and GENeric options. In both cases, lions are individuated by the boys who want them. The individuating is represented by the role-in link (Skolem modification) from lion-1 to want-boy-lion-1.

(24a) *Every boy wants his own lion.*

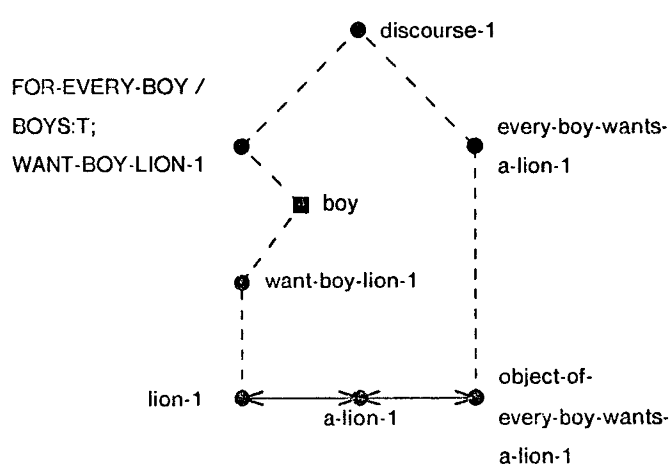

[7] Partee [19] points out that this type of sentence and those of the previous section pose a problem for the analysis of the indefinite article in terms of the features +specific [7]. Something more is needed to produce all the readings. Features can be used as follows: if a hearer has decided a node is definitely attributive it can be marked +attributive, if it is referential the co-descriptor is shown. Otherwise, it is unmarked and has no co-descriptor.

[8] Readers familiar with iteration macros and programming languages will see that Woods' FOR is quite limited, e.g. iteration is limited to a single variable. I use it here because it is simple, is well explained by Woods, and will facilitate comparison of my suggestions with Woods' earlier work.

(24b) *Every boy wants the same specific lion.*

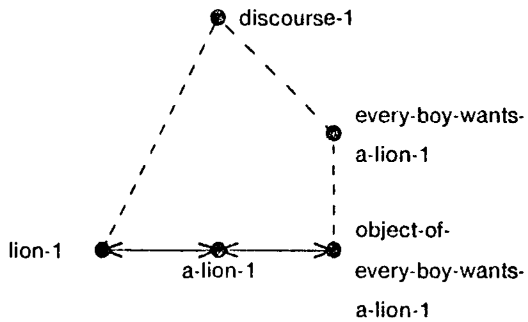

In contrast to (24a), (24b) contains no iteration at all. In this case, all the boys want the same lion. Since lions are not individuated with respect to boys, there is no motivation to introduce an iteration (or Skolem dependency). On this reading, "every boy wants a lion" can be conceptualized just like "a boy wants a lion".

4.2 Collectives, Distributives, and Pluralities

In applying predicates to pluralities, we must be careful not to confuse a predicate on a plurality with a predicate on a collective. If a process, like *wanting*, is done by a plurality, then a plurality of wantings occurs. This contrasts with a collective, which acts as a single unit. Fauconnier [5] demonstrates this distinction with (25).

(25a) *The men* {gathered, united, quarreled}. *collective*

(25b) *The men* took off their hats. *distributive*

(25c) *The men* carried the couch. *plurality*

As a further example, observe that in

(26) Everybody gave $1000 to many of the men.

we must decide whether the men receiving money receive $1000 apiece or just participate in a group receiving $1000. We must also decide whether "everybody" acted individually or collectively. Note that one cannot personally receive $1000 if $1000 is given to a group he is in, but one can be personally told the news if the news is told to a group he is in. From this we see that the distributive reading may or may not require a separate instance of some state or process for each individual.

A plural noun group may be understood in any of the three ways. Vendler [27] gives examples intended to show that the choice between "each" and "every" influences our preferences.

(27) Suppose I show you a basket of apples and I tell you, "Take all of them". If you started to pick them one by one I should be surprised. My offer was sweeping: you should take the apples, if possible, "en bloc". Had I said, "Take every one of them," I should not care how you took them, provided you do not *leave* any behind. If I say, "Take each of them," one feels that the sentence is unfinished. Something like, "Take each of them and examine them in turn," is expected. Thus I expect you to take them one after the other not *missing* any.

"All" favors collective interpretation, but permits distributive interpretation. "Every" favors distributive interpretation. "Each" strongly favors distributive interpretation and favors distinct actions for each individual.

4.3 Multiple Iterations

The previous sections suggested that an iteration need be constructed only when two descriptions in the sentence are related in a special way, e.g., where one is individuated by the other. Otherwise, predicates on collectives or pluralities provide an adequate description. In this way, I avoid iteration, which is relatively expensive from a computational perspective. There are other ways to eliminate iterations. For example, the apparent double iteration in sentence (28a) can be reduced to a single iteration. Formula (28b) represents the natural reading where each cork is associated with one bottle. It might appear that it requires two nested FOR loops as in (28c), but in fact, it can be implemented with a single loop (28d).

(28a) Each cork is fastened to each bottle by a small wire basket.

(28b) ($\forall$x) ($\forall$y) if x is a cork & y is a bottle
& x is the cork in bottle y
then x is fastened to y by a small
wire basket.

(28c) Forall x in corks do *two nested loops*
Forall y in bottles
if x is the cork in bottle y then
x is fastened to y by a small wire basket.

(28d) Forall y in bottles do *single loop*
let x = the cork in bottle y
x is fastened to y by a small wire basket.

Most people understand this sentence by iterating over bottles and individuating the corks by the bottles. The single loop implementation can be derived straightforwardly from my representation. In some other representation, it may require a very clever compiler optimization, such as "loop jamming".

Sentences (29a) and (29b) are similar to (28a) in this respect.

(29a) All the boys kissed all the girls.

(29b) Each man and each woman will be joined in marriage here tonight.

These sentences provide evidence that a quantifier like "each" doesn't necessarily set up an iteration. It usually does, and therefore these sentences are a bit odd. But these sentences are understood when each distributed quantified expression has a discourse co-description which *participates* in an iteration. Setting up an iteration in (29a) and (29b) expresses the specific dependence of boys on girls, and of each man and woman on each other.

4.4 Human Processing of Quantified Expressions

In the previous example we could have had each man married to each woman! Indeed, it is not always obvious how an iteration should be constructed. When someone hears the statement

(30) A requirement for the course is the carving of a block of wood into each of the 12 designs.

his reasoning might be as follows: "Well, let's see. We take the wood and carve the first design. (He pursues the distributive referential reading with one block.) Oh! Oh! Now how do we carve a second design, the block is used up. Well, maybe we could fit the twelve designs on one block, or we could cut the block into twelve pieces. Or maybe I should abandon the referential reading and use twelve blocks." There can be no doubt that world knowledge (pragmatics) is required to choose between such readings. Consider now the following pair of sentences.

(31a) Everybody at MIT knows a dialect of LISP.

(31b) Everybody at IJCAI knows a dialect of LISP.

Everyone at a university might conceivably know the same specific dialect, while everyone at an international conference might not.

Van Lehn [26] reports that when people are given a sentence like

(32) A quick test confirmed that every drug was psychoactive.

they claim they understand it, but are then unable to state whether there was one test per drug or only one test for all. This ambiguity can be expressed in predicate calculus using quantifiers.

(33a) $(\exists\ \text{test})\ (\forall\ \text{drug})\ C(\text{test, drug})$

(33b) $(\forall\ \text{drug})\ (\exists\ \text{test})\ C(\text{test, drug})$

But in my representation, the ambiguity is represented by the presence or absence of a role-in link, which makes it easy to delay disambiguation. There is no shuffling of quantifiers to shift between the two interpretations; instead, disambiguation is achieved by simply inserting a role-in link, when and if appropriate.

4.5 Ambiguity and Generality

The problem with (32) arises in part because of the difficulty of distinguishing between a *general* expression and an *ambiguous* one. For example,

(34)

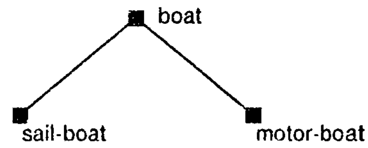

illustrates that both a sail boat and a motor boat can be described by the general expression boat. This contrasts with

(35)

where there is no generalization of the *pronoun-mine* and *mineral-mine* senses of "mine". Finally,

(36)

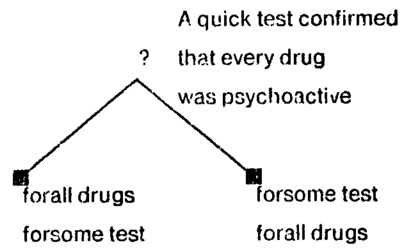

points out that it is uncertain whether "a quick test confirmed that every drug was psychoactive" has a general sense. The general sense of "boat" is resolved semantically — by choosing between more specific concepts such as "sailboat" or "motorboat". If (35) does have a general sense it might be resolved semantically, but it might also be resolved pragmatically — by binding the concept into the discourse structure in different ways. Philosophers have largely ignored this issue since they have been interested in formal languages whose terms are logically unambiguous. Van Lehn's results would tend to indicate that people do have a general sense of this expression.

The advantage of the representational scheme proposed here is that it allows the resolution of ambiguity by incrementally adding to what is already present. We have discussed many reasons why scoping decisions should be delayed, ranging from computational efficiency to Van Lehn's empirical observations. Role-in links provide a simple way to accomplish this.

5. Conclusion: The Asymmetry of Roles

Roles introduce a very interesting asymmetry. In my representation, the meaning of a node depends on its *roles* in more aggregate entities, not just on its *constituents*. For example, I use role-in links as Skolem modifiers where others have employed quantifier scope. In this way, I can represent quantificational dependencies as role-in links to nodes representing more aggregate entities. I have argued that roles are a natural device for explaining many linguistic and philosophical distinctions, and they are convenient for computational processing.

Section 3 showed how roles can help represent referring expressions. The referential/attributive distinction was formulated in terms of roles so that the attributive use asserts something as a role in such-and-such whereas the referential use asserts something as being the such-and-such without a role dependency. I choose to use roles where others (e.g. Moore) have employed scope dependencies, because the ambiguity persists even when the scoping context is unavailable, as in "The president has owned a terrier since February". In this way, I can delay the binding decision because I can represent the dependencies (or lack thereof) without hypothesizing an outer context.

Furthermore, I have argued that verb phrase deletion provides additional support for the claim that role dependencies are associated with references, not with contexts. That is, in a sentence like "The prosecutor believed that he would win the case, and so did the defense attorney", there are two ways to interpret "he": as a co-descriptor of the subject or as a co-descriptor of the prosecutor. Note that whichever way it is taken in the first conjunct, it will be taken the same way in the second. There is a natural explanation for this if the co-description dependency is associated with the reference "he", whereas in a system where this dependency was associated with the context, an ad hoc stipulation would be required.

Section 4 argued that quantifier dependencies should be associated with references as opposed to contexts. My approach leads to a natural representation for collectives, distributives, and pluralities. By minimizing the differences in the representation of these three cases, an interpreter is in a better position to delay binding decisions. This is consistent with Van Lehn's empirical observations. Furthermore collectives and pluralities can be interpreted without iteration, saving considerable processing effort.

References and Related Work

1. Brachman, R.J. (1978) "On the Epistomological Status of Semantic Networks", in N.V. Findler (ed) *Associative Networks — The Representation and Use of Knowledge in Computers*, Academic Press, New York.
2. Chomsky, N. (1980) "On Binding", *Linguistic Inquiry* 11, 1-46.
3. Donnellan, K. (1966) "Reference and Definite Descriptions", *The Philosophical Review* 75, 281-304.
4. Fahlman, S.E. (1979) *NETL: A System for Representing and Using Real World Knowledge*, MIT Press, Cambridge, Mass.
5. Fauconnier, G. (1975) "Do Quantifiers Branch" *Linguistic Inquiry* 6.4.
6. Fauconnier, G. (1978) "Is There a Linguistic Level of Logical Representation?" *Theoretical Linguistics*, Vol. 5, No. 1, Walter De Gruyter, Berlin.
7. Fillmore, C.J. (1967) "On the Syntax of Preverbs" *Glossa* 1, 91-125.
8. Goldstein, I.P. and Roberts, B.R. (1977) "NUDGE — A Knowledge Based Scheduling Program" *Proceedings of 5th IJCAI*, available from Dept. of Computer Science, Carnegie-Mellon Univ., Pittsburgh, Pa.
9. Hayes, P.J. (1979) "On Semantic Nets, Frames and Associations", *Proceedings of 5th IJCAI*, available from Dept. of Computer Science, Carnegie-Mellon Univ., Pittsburgh, PA.
10. Hendrix, G. (1978) "Encoding Knowledge in Partitioned Networks", N.V. Findler (ed) *Associative Networks — The Representation and Use of Knowledge in Computers*, Academic Press, New York.
11. Hintikka, K.J.J. (1976) "Quantifiers in Logic and Quantifiers in Natural Languages", in S. Korner (ed) *Philosophy of Logic*, Oxford: Basil Blackwell, 208-232.
12. Jackendoff (1972) *Semantic Interpretation in Generative Grammar*, MIT Press, Cambridge, Mass.
13. Mathlab Group (1977) "Macsyma Reference Manual", Laboratory for Computer Science, MIT.
14. Martin, W.A. (1979) "Descriptions and the Specialization of Concepts" in P. Winston (ed) *Artificial Intelligence, An MIT Perspective*, MIT Press, Cambridge, Mass.
15. Martin, W.A. (1979) "Roles, Co-descriptors, and the Formal Representation of Quantified English Expressions", MIT/LCS/TM-139.
16. Minsky, M. (1975) "A Framework for Representing Knowledge" in *The Psychology of Computer Vision*, P.H. Winston, ed. McGraw-Hill, New York.
17. Moore, R.C. (1973) "D-SCRIPT: A Computational Theory of Descriptions" *Advance Papers of the 3rd IJCAI*, 223-229.
18. Moore, R.C. (1979) "Reasoning About Knowledge and Action" MIT Computer Science PhD thesis.
19. Partee, B.H. (1972) "Opacity, Co-Reference, and Pronouns" in *Semantics of Natural Language*, Davidson and Harmon (eds) D. Reidel Pub. Co. Dordrecht-Holland.
20. Partee, B.H. (1978) "Bound Variables and Other Anaphors" in *Proceedings of Theoretical Issues in Natural Language Processing — 2*, available from ACM.
21. Quine, W.V.O. (1960) *Word and Object*, MIT Press, Cambridge, Mass.
22. Russell, B. (1905) "On Denoting" in H. Feigl and W. Sellars (eds) *Readings in Philosophical Analysis*, 85-102, Appleton-Century-Crofts, Inc., New York, 1949.
23. Sag, I. (1976) "Deletion and Logical Form" MIT Linguistics PhD Thesis, available through Indiana University Linguistics Club.
24. Sidner, C. (1979) "Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse", MIT-AI-TR 537, MIT Computer Science PhD Thesis.
25. Strawson, P.F. (1959) *Individuals: An Essay in Descriptive Metaphysics*, Anchor Books Edition, (1963).
26. Van Lehn, K.A. (1978) "Determining the Scope of English Quantifiers" MIT Artificial Intelligence Laboratory Report AI-TR-483.
27. Vendler, Z. (1967) *Linguistics in Philosophy*, Cornell University Press, Ithaca, New York.
28. Webber, B.L. (1978) "A Formal Approach to Discourse Anaphora", Bolt Beranek and Newman Report 3761.
29. Woods, W.A. (1975) "What's in a link? Foundations for semantic networks.", in *Representation and Understanding*, D. G. Bobrow and A. M. Collins (eds.), Academic Press, New York, pp. 35-82.
30. Woods, W. (1977) "Semantics and Quantification in Natural Language Question Answering", Bolt Beranek and Newman Report 3687.

William A. Martin

William A. Martin, Professor in the Department of Electrical Engineering and Computer Science and in the Sloan School of Management at MIT and a long-time member of the Association for Computational linguistics, died Tuesday, June 2, 1981, after an extended illness. He was 43.

Bill Martin managed to crowd into an all-too-short life three distinct research careers. After receiving his doctorate from MIT, he, in association with Joel Moses, led the MACSYMA development project. MACSYMA, a computer system for applied mathematicians, has become one of the few large applied-AI systems that have ever-growing user communities. He then switched into the field of automatic programming, where he was responsible for developing the PROTOSYSTEM automatic programming system. This effort culminated in the development of HIBOL, a high-level business-oriented language. Changing fields once again, Bill became an expert in computational linguistics, concentrating on problems in parsing and knowledge representation.

Bill's involvement with computational linguistics was triggered by the need he perceived for a natural language interface to PROTOSYSTEM. His early efforts in this area began with an investigation of the structure of dialogues. This led to the development of first MAPL and later OWL as formal languages for representing knowledge of English, of the application domain, and of PROTOSYSTEM's procedures. Subsequently, the problems of English language processing came to be the focal point of his work, motivating his full concentration on computational linguistics during the last few years.

Bill's foremost talents were his ability to acquire an immense knowledge of both the most significant and the most intricate problems of any domain that he chose to focus on and his ability to forge and maintain a large coherent research project. His attention has alternated between these two somewhat dissimilar motifs, as individual researcher and as project head. He was a first-rate engineer throughout, one who believed and put into practice the notion that good engineering principles could be used to resolve all scientific and application problems. Believing in the notion of a "small infinity," Bill demonstrated, on many occasions, that seemingly unbounded problems will succumb to the careful analysis of their hundreds of cases, when the "infinity" of their complexity is indeed only on the order of hundreds.

Bill was a native of Oklahoma City, graduating from Northwest Classen High School, where he made his mark as a wrestler as well as a student. He was an Oklahoma state wrestling champion, and continued his active interest in athletics as an undergraduate at MIT, where he was a member of Beta Theta Pi fraternity. He received the bachelor's, master's, and PhD degrees from MIT in 1960, 1962, and 1967, respectively, all in electrical engineering, and was appointed to the MIT faculty in 1968.

He survived by his wife, Susan Y. Forbes Martin; three children, Jamie, Tad, and Jon; his parents Earl and Barbara Martin of Oklahoma City; and a sister Jane Anne Slane, also of Oklahoma City. A memorial fund has been established in Professor Martin's name by the Department of Electrical Engineering and Computer Science at MIT. It will award an annual prize for the best MIT undergraduate thesis in computer science.