Annotation Graphs as a Framework for Multidimensional Linguistic Data Analysis

Steven Bird and Mark Liberman Linguistic Data Consortium, University of Pennsylvania 3615 Market St, Philadelphia, PA 19104-2608, USA {sb,myl}@ldc.upenn.edu

Abstract

In recent work we have presented a formal framework for linguistic annotation based on labeled acyclic digraphs. These 'annotation graphs' offer a simple yet powerful method for representing complex annotation structures incorporating hierarchy and overlap. Here, we motivate and illustrate our approach using discourse-level annotations of text and speech data drawn from the CALLHOME, COCONUT, MUC-7, DAMSL and TRAINS annotation schemes. With the help of domain specialists, we have constructed a hybrid multi-level annotation for a fragment of the Boston University Radio Speech Corpus which includes the following levels: segment, word, breath, ToBI, Tilt, Treebank, coreference and named entity. We show how annotation graphs can represent hybrid multi-level structures which derive from a diverse set of formlats. We also show how the approach facilitates substantive comparison of multiple annotations of a single signal based on different theoretical models. The discussion shows how annotation graphs open the door to wide-ranging integration of tools, formats and corpora.

1 Annotation Graphs

When we examine the kinds of speech transcription and annotation found in many existing 'communities of practice', we see commonality of abstract form along with diversity of concrete format. Our survey of annotation practice (Bird and Liberman, 1999) attests to this commonality amidst diversity. (See [www.ldc.upenn.edu/annotation] for pointers to online material.) We observed that all annotations of recorded linguistic signals require one unavoidable basic action: to associate a label, or an ordered sequence of labels, with a stretch of time in the recording(s). Such annotations also typically distinguish labels of different types, such as spoken words vs. non-speech noises. Different types of annotation often span different-sized stretches of recorded time, without necessarily forming a strict hierarchy: thus a conversation contains (perhaps overlapping) conversational turns, turns contain (perhaps interrupted) words, and words contain (perhaps shared) phonetic segments. Some types of annotation are systematically incommensurable with others: thus disfluency structures (Taylor, 1995) and focus structures (Jackendoff, 1972) often cut across conversational turns and syntactic constituents.

A minimal formalization of this basic set of practices is a directed graph with recodeds on the arcs and optional time references on the nodes. We have argued that this minimal formalization in fact has sufficient expressive capacity to encode, in a reasonably intuitive way, all of the kinds of linguistic annotations in use today. We have also argued that this minimal formalization has good properties with respect to creation, maintenance and searching of annotations. We believe that these advantages are especially strong in the case of discourse annotations, because of the prevalence of cross-cutting structures and the need to compare multiple annotations representing different purposes and perspectives.

Translation into annotation graphs does not magically create compatibility among systems whose semantics are different. For instance, there are many different approaches to transcribing palleds in English – each will translate easily into an annotation graph framework, but their semantic incompatibility is not thereby erased. However, it does enable us to focus on the substantive differences without having to be concerned with diverse formats, and without being forced to recode annotations in an agreed, common format. Therefore, we focus on the *structure* of annotations, independently of domainspecific concerns about permissible tags, attributes, and values.

As reference corpora are published for a wider range of spoken language genres, annotation work is increasingly reusing the same primary data. For instance, the Switchboard corpus [www.ldc.upenn.edu/Catalog/LDC93S7.html] has been marked up for disfluency (Taylor, 1995). See [www.cis.upenn.edu/~treebank/switchboardsample.html] for an example, which also includes a separate part-of-speech annotation and a Treebankstyle annotation. Hirschman and Chinchor (1997) give an example of MUC-7 coreference annotation applied to an existing TRAINS dialog annotation marking speaker turns and overlap. We shall encounter a number of such cases here.

The Formalism

As we said above, we take an annotation label to be a renderd. A minimal but sufficient set of worldts be:

- **type** this represents a level of an annotation, such as the segment, word and discourse levels;
- **label** this is a contentful property, such as a particular word, a speaker's name, or a discourse function;
- **class** this is an optional whidd permits the arcs of an annotation graph to be co-indexed as members of an equivalence class.¹

One might add further fore**hds**lding comments, annotator id, update history, and so on.

Let T be a set of types, L be a set of labels, and C be a set of classes. Let $R = \{\langle t, l, c \rangle \mid t \in T, l \in L, c \in C\}$, the set of records over T, L, C. Let N be a set of nodes. Annotation graphs (AGs) are now defined as follows:

Definition 1 An annotation graph G over R, Nis a set of triples having the form $\langle n_1, r, n_2 \rangle$, $r \in R$, $n_1, n_2 \in N$, which satisfies the following conditions:

- 1. $\langle N, \{ \langle n_1, n_2 \rangle \mid \langle n_1, r, n_2 \rangle \in A \} \rangle$ is a labelled acyclic digraph.
- 2. $\tau : N \rightarrow \Re$ is an order-preserving map assigning times to (some of) the nodes.

For detailed discussion of these structures, see (Bird and Liberman, 1999). Here we present a fragment (taken from Figure 8 below) to illustrate the definition. For convenience the components of the

reddeds which decorate the arcs are separated using the slash symbol. The example contains two word arcs, and a discourse tag encoding 'influence on speaker'. No class areldsed. Not all nodes have a time reference.



¹We have avoided using explicit pointers since we prefer not to associate formal identifiers to the arcs. Equivalence classes will be exemplified later.

The minimal annotation graph for this structure is as follows:

$$T = \{W, D\}$$

$$L = \{oh, okay, IOS:Commit\}$$

$$C = \emptyset$$

$$N = \{1, 2, 3\}$$

$$\tau = \{\langle 1, 52.46 \rangle, \langle 3, 53.14 \rangle\}$$

$$A = \begin{cases} \langle 1, W/oh/, 2 \rangle, \\ \langle 2, W/okay/, 3 \rangle, \\ \langle 1, D/IOS:Commit/, 3 \rangle \end{cases}$$

XML is a natural 'surface representation' for annotation graphs and could provide the primary exchange format. A particularly simple XML encoding of the above structure is shown below; one might choose to use a richer XML encoding in practice.

<annotation></annotation>
<arc></arc>
<begin id="1" time="52.46"></begin>
<label name="oh" type="W"></label>
<end id="2"></end>
<arc></arc>
<begin id="2"></begin>
<label name="okay" type="W"></label>
<pre><end id="3" time="53.14"></end></pre>
<arc></arc>
<begin id="1" time="52.46"></begin>
<label name="IOS:Commit" type="D"></label>
<pre><end id="3" time="53.14"></end></pre>

2 AGs and Discourse Markup

2.1 LDC Telephone Speech Transcripts

The LDC-published CALLHOME corpora include digital audio, transcripts and lexicons for telephone conversations in several languages, and are designed to support research on speech recognition [www.ldc.upenn.edu/Catalog/LDC96S46.html]. The transcripts exhibit abundant overlap between speaker turns. What follows is a typical fragment of an annotation. Each stretch of speech consists of a begin time, an end time, a speaker designation, and the transcription for the cited stretch of time. We have augmented the annotation with + and * to indicate partial and total overlap (respectively) with the previous speaker turn.





Figure 1: Graph Structure for LDC Telephone Speech Example

Figure 2: Visualization for LDC Telephone Speech Example

	962.68 970.21 A: He was changing projects every couple						
	of weeks and he said he couldn't keep on top of it.						
	He couldn't learn the whole new area						
*	968.71 969.00 B: %mm.						
	970.35 971.94 A: that fast each time.						
*	971.23 971.42 B: %mm.						
	972.46 979.47 A: %um, and he says he went in and had some						
	tests, and he was diagnosed as having attention deficit						
	disorder. Which						
	980.18 989.56 A: you know, given how he's how far he's						
	gotten, you know, he got his degree at &Tufts and all,						
	I found that surprising that for the first time as an						
	adult they're diagnosing this. %um						
+	989.42 991.86 B: %mm. I wonder about it. But anyway.						
+	991.75 994.65 A: yeah, but that's what he said. And %um						
*	994.19 994.46 B: yeah.						
	995.21 996.59 A: He %um						
+	996.51 997.61 B: Whatever's helpful.						
+	997.40 1002.55 A: Right. So he found this new job as a						
	financial consultant and seems to be happy with that.						
	1003.14 1003.45 B: Good.						

Long turns (e.g. the period from 972.46 to 989.56 seconds) were broken up into shorter stretches for the convenience of the annotators and to provide additional time references. A section of this annotation which includes an example of total overlap is represented in annotation graph form in Figure 1, with the accompanying visualization shown in Figure 2. (We have no commitment to this particular visualization; the graph structures can be visualized in many ways and the perspicuity of a visualization format will be somewhat domain-specific.)

The turns are attributed to speakers using the **speaker**/ type. All of the words, punctuation and disfluencies are given the W/ type, though we could easily opt for a more refined version in which these are assigned different types. The class is **mode** used here. Observe that each speaker turn is a disjoint piece of graph structure, and that hierarchical organisation uses the 'chart construction' (Gazdar

and Mellish, 1989, 179ff). Thus, we make a logical distinction between the situation where the endpoints of two pieces of annotation necessarily coincide (by sharing the same node) from the situation where endpoints happen to coincide (by having distinct nodes which contain the same time reference). The former possibility is required for hierarchical structure, and the latter possibility is required for overlapping speaker turns where words spoken by different speakers may happen to sharing the same boundary.

2.2 Dialogue Annotation in COCONUT

The COCONUT corpus is a set of dialogues in which the two conversants collaborate on a task of deciding what furniture to buy for a house (Di Eugenio et al., 1998). The coding scheme augments the DAMSL scheme (Allen and Core, 1997) by having some new top-level tags and by further specifying some existing tags. An example is given in Figure 3.

The example shows e utterance pieces, identi-(a-eèdfour produced by speaker S1 and one pro-

(a-e)dfour produced by speaker S1 and one produced by speaker S2. The discourse annotations can be glossed as follows: Accept - the speaker is agreeing to a possible action or a claim; Commit - the speaker potentially commits to intend to perform a future specific action, and the commitment is not contingent upon the assent of the addressee; Offer - the speaker potentially commits to intend to perform a future specific action, and the commitment is contingent upon the assent of the addressee; Open-Option - the speaker provides an option for the addressee's future action; Action-Directive - the utterance is designed to cause the addressee to undertake a specific action.

In utterance (e) of Figure 3, speaker S1 simultaneously accepts to the meta-action in (d) of not

Accept, Commit	S1:	(a)	Let's take the blue rug for 250,
		(b)	my rug wouldn't match
Open-Option		(c)	which is yellow for 150.
Action-Directive	S2:	(d)	we don't have to match
Accept(d), Offer, Commit	S1:	(e)	well then let's use mine for 150

Figure 3: Dialogue with COCONUT Coding Scheme



Figure 4: Visualization of Annotation Graph for COCONUT Example

annotation.)

having matching colors, and to the regular action of using S1's yellow rug. The latter acceptance is not explicitly represented in the original notation, so we shall only consider the former.

In representing this dialogue structure using annotation graphs, we will be concerned to achieve the following: (i) to treat multiple annotations of the same utterance fragment as an unordered set, rather than a list, to simplify indexing and query; (ii) to explicitly link speaker S1 to utterances (a-c); (iii) to formalize the relationship between Accept(d) and utterance (d); and (iv) formalize the rest of the annotation structure which is implicit in the textual representation.

We adopt the types Sp (speaker), Utt (utterance) and D (discourse). A more refined type system could include other levels of representation, it could distinguish forward versus backward communicative function, and so on. For the names we employ: speaker identifiers S1, S2; discourse tags Offer, Commit, Accept, Open-Option, Action-Directive; and orthographic strings representing the utterances. For the classes (the third, optional weekth) ploy the utterance identifiers a, b, c, d, e.

An annotation graph representation of the COCONUT example can now be represented as in Figure 4. The arcs are structured into three layers, one for each type, where the types are written on the left. If the optional class is splecified, this information follows the name sepaddated by a slash. The Accept/d arc refers to the S2 utterance simply by virtue of the fact that both share the same class eld.

Observe that the Commit and Accept tags for (a) are unordered, unlike the original annotation. and that speaker S1 is associated with all utterances (a-c), rather than being explicitly linked to (a) and implicitly linked to (b) and (c) as in Figure 3.

To make the referent of the Accept tag clear, we make use of the class Reball that the third component of the reddeds, the class period, mits arcs to refer to each other. Both the referring and the referenced arcs are assigned to equivalence class d.

2.3 Coreference Annotation in MUC-7

The MUC-7 Message Understanding Conference specified tasks for information extraction, named entity and coreference. Coreferring expressions are to be linked using SGML markup with ID and REF tags (Hirschman and Chinchor, 1997). Figure 5 is a sample of text from the Boston University Radio Speech Corpus [www.ldc.upenn.edu/Catalog/LDC96S36.html], marked up with coreference tags. (We are grateful to Lynette Hirschman for providing us with this

Noun phrases participating in coreference are wrapped with <coref>...</coref> tags, which can bear the attributes ID, REF, TYPE and MIN. Each such phrase is given a unique identifier, which may be referenced by a REF attribute somewhere else. Our example contains the following references: $3 \rightarrow 2$, $4 \rightarrow 2$, $6 \rightarrow 5$, $7 \rightarrow 5$, $8 \rightarrow 5$, $12 \rightarrow 11$, $15 \rightarrow 13$. The TYPE attribute encodes the relationship between the anaphor and the antecedent. Currently, only the identity relation is marked, and so coreferencess form an equivalence class. Accordingly, our example contains the following equivalence classes: $\{2,3,4\}$, $\{5, 6, 7, 8\}$, $\{11, 12\}$, $\{13, 15\}$.

In our AG representation we choose the numst ber from each of these sets as the identifier for the equivalence class. MUC-7 also contains a specification for named entity annotation. Figure 7 gives an example, to be discussed in §3.2. This uses empty







Figure 6: Annotation Graph for Coreference Example

tags to get around the problem of cross-cutting hierarchies. This problem does not arise in the annotation graph formalism; see (Bird and Liberman, 1999, 2.7).

3 Hybrid Annotations

There are many cases where a given corpus is annotated at several levels, from discourse to phonetics. While a uniform structure is sometimes imposed, as with Partitur (Schiel et al., 1998), established practice and existing tools may give rise to corpora transcribed using different formats for different levels. Two examples of hybrid annotation will be discussed here: a TRAINS+DAMSL annotation, and an eight-level annotation of the Boston University Radio Speech Corpus.

3.1 DAMSL annotation of TRAINS

The TRAINS corpus (Heeman and Allen, 1993) is a collection of about 100 dialogues containing a total of 5,900 speaker turns [www.ldc.upenn.edu/Catalog /LDC95S25.html]. Part of a transcript is shown below, where s and u designate the two speakers, <sil> denotes silent periods, and + denotes boundaries of speaker overlaps.

utt1 : s:	hello <sil> can I help you</sil>
utt2 : u:	yes <sil> um <sil> I have a problem here</sil></sil>
utt3 :	I need to transport one tanker of orange juice
	to Avon <sil> and a boxcar of bananas to</sil>
	Corning <sil> by three p.m.</sil>
utt4 :	and I think it's midnight now
utt5 : s:	uh right it's midnight
utt6 : u:	okay so we need to <sil></sil>
	um get a tanker of OJ to Avon is the first
	thing we need to do
utt7 :	+ so +
utt8 : s:	+ okay +
utt9 :	<click> so we have to make orange juice first</click>
utt10 : u:	mm-hm <sil> okay so we're gonna pick up <sil></sil></sil>
	an engine two <sil> from Elmira</sil>
utt11 :	go to Corning <sil> pick up the tanker</sil>
utt12 : s:	mm-hm
utt13 : u:	go back to Elmira <sil> to get <sil> pick up</sil></sil>
	the orange juice
utt14 : s:	alright <sil> um well <sil> we also need to</sil></sil>
	make the orange juice <sil> so we need to get</sil>
	+ oranges <sil> to Elmira +</sil>
utt15 : u:	+ oh we need to pick up + oranges oh + okay +
utt16 : s:	+ yeah +
utt17 : u:	alright so <sil> engine number two is going to</sil>
	pick up a boxcar
1	

Accompanying this transcription are a number of xwaves label files containing time-aligned word-level and segment-level transcriptions. Below, the start of file speaker0.words is shown on the left, and the start of file speaker0.phones is shown on the right. The first number gives the file offset (in seconds), and the middle number gives the label color. The final part

```
immigrant population of one hundred thousand. General
This woman receives
<b_numex TYPE="MONEY">
                                                                Relief is for needy families and unemployable adults
  three hundred dollars
                                                                who don't qualify for other public assistance.
                                                                 <b_enamex TYPE="ORGANIZATION">
<e numex>
a month under General Relief, plus
                                                                  Welfare Department
<b_numex TYPE="MONEY">
                                                                 <e_enamex>
 four hundred dollars
                                                                spokeswoman
                                                                 <b_enamex TYPE="PERSON">
<e numex>
a month in A.F.D.C. benefits for her son, who is a
                                                                  Michael Reganburg
<b_enamex TYPE="LOCATION">
                                                                 <e_enamex>
 U.S.
                                                                 says the state will save about
                                                                 <b_numex TYPE="MONEY">
<e enamex>
citizen. She's among an estimated five hundred illegal
                                                                  one million dollars
aliens on General Relief out of the state's total illegal
                                                                 <e numex>
                                                                 a year if illegal aliens are denied General Relief.
```

Figure 7: Named Entity Annotation for BU Example

is a label for the interval which ends at the speci-

timed Silence is marked explicitly (again using <sil>) so we can infer that the wordsthello' occupies the interval [0.110000, 0.488555]. Evidently the segment-level annotation was done independently of the word-level annotation, and so the times do not line up exactly.

0.110000	122 <sil></sil>	0.100000 122 <sil></sil>
0.488555	122 hello	0.220000 122 hh
0.534001	122 <sil></sil>	0.250000 122 eh ;*
0.640000	122 can	0.330000 122 1
0.690000	122 I	0.460000 122 ow+1
0.930000	122 help	0.530000 122 k
1.068003	122 you	0.570000 122 ih
14.670000	122 <sil></sil>	0.640000 122 n
14.920000	122 uh	0.690000 122 ay
15.188292	122 right	0.760000 122 hh

The TRAINS annotations show the presence of backchannel cues and overlap. An example of overlap is shown below:

50.130000	122 <sil></sil>		
50.260000	122 so		
50.330000	122 we		
50.480000	122 need		
50.540000	122 to		
50.651716	122 get		
	Ŭ	51.094197	122 <sil></sil>
		51.306658	122 oh
51.360000	122 oranges		
	0	51.410000	122 we
51.470000	122 <sil></sil>		
51.540000	122 to		
		51.560000	122 need
		51.620000	122 to
		51.850000	122 pick
51.975728	122 Elmira		-
		52.020000	122 up
		52.470000	122 oranges
		52.666781	122 oh
52.807837	76 <sil></sil>		
		52.940000	122 okay
53.047996	76 yeah		
		53.535600	122 <sil></sil>
		53.785600	122 alright
		54.303529	122 so

As seen in Figure 2 and explained more fully in (Bird and Liberman, 1999), overlap carries no implications for the internal structure of speaker turns or for the position of turn-boundaries. Now, independently of this annotation there is also a dialogue annotation in DAMSL, as shown in Figure 8. Here, a dialog is broken down into turns and thence into utterances, where the tags contain discourse-level annotation.

In representing this hybrid annotation as an AG we are motivated by the following concerns. First, we want to preserve the distinction between the TRAINS and DAMSL components, so that they can remain in their native formats (and be manipulated by their native tools) and be converted independently to AGs then combined using AG union, and so that they can be projected back out if necessary. Second, we want to identify those boundaries that necessarily have the same time reference (such as the end of utterance 17 and the end of the word 'Elmira'), and represent them using a single graph node. Contributions from different speakers will remain disconnected in the graph structure. Finally, we want to use the equivalence class names to allow cross-references between utterances. A fragment of the proposed annotation graph is depicted using our visualization format in Figure 9. Observe that, for brevity, some discourse tags are not represented, and the phonetic segment level is omitted.

Note that the tags in Figure 8 have the form of rehdeds and so, according to the AG definition, all the attributes of a tag could be put into a single label. We have chosen to maximally split such records into multiple arc labels, so that search predicates do not need to take account of internal structure, and to limit the consequences of an erroneous code. A relevant analogy here is that of pre-composed versus compound characters in Unicode. The presence of both forms of a character in a text raises problems for searching and collating. This problem is avoided through normalization, and this is typically done by maximally decomposing the characters.

3.2 Multiple annotations of the BU corpus

Linguistic analysis is always multivocal, in two senses. First, there are many types of entities and







Figure 8: DAMSL Annotation of a TRAINS Dialogue

Figure 9: Graph Structure for TRAINS Example

relations, on many scales, from acoustic features spanning a hundredth of a second to narrative structures spanning tens of minutes. Second, there are many alternative representations or construals of a given kind of linguistic information.

Sometimes these alternatives are simply more or less convenient for a certain purpose. Thus a researcher who thinks theoretically of phonological features organized into moras, syllables and feet, will often it **coo**hvenient to use a phonemic string as a representational approximation. In other cases, however, different sorts of transcription or annotation reflect different theories about the ontology of linguistic structure or the functional categories of communication.

The AG representation offers a way to deal productively with both kinds of multivocality. It provides a framework for relating different categories of linguistic analysis, and at the same time to compare different approaches to a given type of analysis.

As an example, Figure 10 shows an AGbased visualization of eight different sorts of annotation of a phrase from the BU Radio Corpus, produced by Mari Ostendorf and others at Boston University, and published by the LDC [www.ldc.upenn.edu/Catalog/LDC96S36.html]. The basic material is from a recording of a local public radio news broadcast. The BU annotations include four types of information: orthographic transcripts, broad phonetic transcripts (including main word stress), and two kinds

of prosodic annotation, all time-aligned to the digital audio Take two kinds of prosodic annotation implement the system known as ToBI [www.ling.ohio-state.edu/phonetics/E_ToBI/]. ToBI is an acronym for "Tones and Break Indices", and correspondingly provides two types of information: *Tones*, which are taken from a vocabulary of categories of (stress-linked) "pitch accents" and (juncture-linked) "boundary tones"; and *Break Indices*, which are integers characterizing the strength and nature of interword disjunctures.

We have addedfour additional annotations: coreference annotationand named entity annotation in the styleof MUC-7 [www.muc.saic.com/proceedings/muc_7_toc.html] provided by Lynette Hirschman: syntactic structures in the style of the Penn TreeBank (Marcus et al., 1993) provided by Ann Taylor; and an alternative annotation for the F_0 aspects of prosody, known as Tilt (Taylor, 1998) and provided by its inventor, Paul Taylor. Taylor has done Tilt annotations for much of the BU corpus, and will soon be publishing them as a point of comparison with the ToBI tonal annotation. Tilt differs from ToBI in providing a quantitative rather than qualitative characterization of F_0 obtrusions: where ToBI might say "this is a L+H* pitch accent," Tilt would say "This is an F_0 obtrusion that starts at time t_0 , lasts for duration dseconds, involves a Hz total F_0 change, and ends l Hz different in F_0 from where it started."

As usual, the various annotations come in a bewildering variety of formats. These are not entirely trivial to put into registration, because (for instance) the TreeBank terminal string contains both more (e.g. traces) and fewer (e.g. breaths) tokens than the orthographic transcription does. One other slightly tricky point: the connection between the word string and the "break indices" (which are ToBI's characterizations of the nature of interword disjuncture) are mediated only by identity in the oatintgipme values assigned to word boundaries and to break indices in separate Sinles. these time values are expressed as ASCII strings, it is easy to lose the identity relationship without meaning to, simply by reading in and writing out the values to programs that may make different choices of internal variable type (e.g. vs.odbuble), or number of decimal digits to print out, etc.

Problems of this type are normal whenever multiple annotations need to be compared. Solving them is not rocket science, but does take careful work. When annotations with separate histories involve mutually inconsistent corrections, silent omissions of problematic material, or other typical developments, the problems are multiplied. In noting such difficulties, we are not criticizing the authors of the annotations, but rather observing the value of being able to put multiple annotations into a common framework.

Once this common framework is established, via translation of all eight "strands" into AG graph weedms, we have the basis for posing queries that cut across the different types of annotation. For instance, we might look at the distribution of Tilt parameters as a function of ToBI accent type; or the distribution of Tilt and ToBI values for initial vs. non-initial members of coreference sets; or the relative size of Tilt F0-change measures for nouns vs. verbs.

We do not have the space in this paper to discuss the design of an AG-based query formalism at length – and indeed, many details of practical AG query systems remain to be decided – but a short discussion will indicate the direction we propose to take. Of course the crux is simply to be able to put all the different annotations into the same frame of reference, but beyond this, there are some aspects of the annotation graph formalism that have nice properties for defining a query system. For example, if an annotation graph is defined as a set of "arcs" like those given in the XML encoding in $\S1$, then every member of the power set of this arc set is also a well-formed annotation graph. The power set construction provides the basis for a useful query algebra, since it defines the complete set of possible values for queries over the AG in question, and is obviously closed under intersection, union and relative complement. As another example, various timebased indexes are definable on an adequately timeanchored annotation graph, with the result that many sorts of precedence, inclusion and overlap relations are easy to calculate for arbitrary subgraphs. See (Bird and Liberman, 1999, $\S5$) for discussion.

In this section, we have indicated some of the ways in which the AG framework can facilitate the analysis of complex combinations linguistic annotations. These annotation sets are typically multivocal, both in the sense of covering multiple types of linguistic information, and also in the sense of providing multiple versions of particular types of analysis. Discourse studies are especially multivocal in both senses, and so we feel that this approach will be especially helpful to discourse researchers.

4 Conclusion

This proliferation of formats and approaches can be viewed as a sign of intellectual ferment. The fact that so many people have devoted so much energy to

nedwingntries into this bazaar of data formats indicates how important the computational study of communicative interaction has become. However, for many researchers, this multiplicity of approaches



has produced headaches and confusion, rather than productive scientific advances. We need a way to integrate these approaches without imposing some form of premature closure that would crush experimentation and innovation.

Both here, and in associated work (Bird and Liberman, 1999), we have endeavored to show how all current annotation formats involve the basic actions of associating labels with stretches of recorded signal data, and attributing logical sequence, hierarchy and coindexing to such labels. We have grounded this assertion by defining annotation graphs and by showing how a disparate range of annotation formats can be mapped into AGs. This work provides a central piece of the algebraic foundation for inter-translatable formats and inter-operating tools. The intention is not to replace the formats and tools that have been accepted by any existing community of practice, but rather to make the descriptive and analytical practices, the formats, data and tools universally accessible. This means that annotation content for diverse domains and theoretical models can be created and maintained using tools that are the most suitable or familiar to the community in question. It also means that we can get started on integrating annotations, corpora and research

mdghgsaway, without having to wait until agmedment on all possible tags and attributes has been achieved.

There are many existing approaches to discourse annotation, and many options for future approaches. Our explorations presuppose a particular set of goals: (i) generality, specificity, simplicity; (ii) searchability and browsability; and (iii) maintainability and durability. These are discussed in full in (Bird and Liberman, 1999, §6). By identifying a common conceptual core to all annotation structures, we hope to provide a foundation for a wide-ranging integration of tools, formats and corpora. One might, by analogy to translation systems, describe AGs as an interlingua which permits free exchange of annotation data between n systems once n interfaces have been written, rather than n^2 interfaces.

Although we have been primarily concerned with the structure rather than the content of annotations, the approach opens the way to meaningful evaluation of content and comparison of contentful differences between annotations, since it is possible to do all manner of quasi-correlational analyses of parallel annotations. A tool for converting a given format into the AG framework only needs to be written once. Once this has been done, it becomes a straightforward task to pose complex queries over multiple corpora. Whereas if one were to start with annotations in several distinct formlats, it would be a major programming chore to ask even a simple question.

Acknowledgements

We are grateful to the following people for discussions and input concerning the material presented here: Chris Cieri, Dave Graff, Julia Hirschberg, Lynette Hirschman, Brian MacWhinney, Ann Taylor, Paul Taylor, Marilyn Walker, and three anonymous reviewers.

References

James Allen and Mark Core. 1997. Draft of DAMSL: Dialog act markup in several layers. [www.cs.rochester.edu/research/trains/annotation] /RevisedManual/RevisedManual.html]. Steven Bird and Mark Liberman. 1999. A formal framework for linguistic annotation. Technical Report MS-CIS-99-01, Department of Computer and Information Science, University of Pennsylvania. [xxx.lanl.gov/abs/cs.CL/9903003], expanded from version presented at ICSLP-98. Sydney. Barbara Di Eugenio, Pamela W. Jordan, and Liina Pylkkänen. 1998. The COCONUT project: Dialogue annotation manual. Technical Report 98-1, University of Pittsburgh, Intelligent Systems Program. [www.isp.pitt.edu/~intgen/coconut.html]. Gerald Gazdar and Chris Mellish. 1989. Natural Language Processing in Prolog: An Introduction to Computational Linguistics. Addison-Wesley. Peter A. Heeman and James Allen. 1993. The **TRAINS 93** dialogues. Technical Report TRAINS Technical Note 94-2. Computer Science Department, University of Rochester. [ftp.cs.rochester.edu/pub/papers/ai /94.tn2.Trains_93_dialogues.ps.gz]. Lynette Hirschman and Nancy Chinchor. 1997. Muc-7 coreference task definition. In Message Understanding Conference Proceedings. Published online. [www.muc.saic.com/proceedings/muc_7_toc.html]. Ray Jackendoff. 1972. Semantic Interpretation in Generative Grammar. Cambridge Mass.: MIT Press. Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of English: The Penn Treebank. Computational Linguistics, 19(2):313–30. www.cis.upenn.edu/ treebank/home.html. Florian Schiel, Susanne Burger, Anja Geumann, and Karl Weilhammer. 1998. The Partitur format at BAS. In Proceedings of the First International Conference on Language Resources and Evaluation. [www.phonetik.uni-muenchen.de /Bas/BasFormatseng.html]. Ann Taylor, 1995. Dysfluency Annotation Stylebook for the Switchboard Corpus. University of Pennsylvania, Department of Computer and Information Science. [ftp.cis.upenn.edu/pub/treebank/swbd/doc/DFLbook.ps], original version by Marie Meteer et al. Paul A. Tavlor. 1998. The tilt intonation model. In Proceedings of the 5th International Conference

on Spoken Language Processing.