Historical Development and Future Directions in Data Resource Development

Workshop Attendees: Martha Palmer, Stephanie Strassel, Randee Tangi
Report Contributors: Martha Palmer, Randee Tangi, Stephanie Strassel,
Christiane Fellbaum, Eduard Hovy

Introduction

On February 25,26 2007, the second of two workshops entitled "Meeting of the MINDS: Future Directions for Human Language Technology," sponsored by the U.S. Government's Disruptive Technology Office (DTO), was held in Marina del Rey, California. "MINDS" is an acronym for **M**achine Translation (MT), **I**nformation Retrieval (IR), **N**atural Language Processing (NLP), **D**ata Resources (Data), and **S**peech Understanding (ASR) which were the 5 areas, each addressed by a number of experienced researchers, at this workshop. The goal of these working groups was to identify and discuss especially promising future research directions that could lead to major paradigm shifts and which are yet to be funded. As a continuation of a prior workshop where each group (except Data Resources) had first reviewed major past developments in their respective fields and the circumstances that led to their success, this workshop began by reviewing and refining the suggestions for future research emanating from the first workshop. The different areas then met in pairs to discuss possible cross-fertilizations and collaborations. Each area is responsible for producing a report proposing 5 to 6 "Grand Challenges."

1. Historically Significant Developments in Data Resource Development

The widespread availability of electronic text and the more recent advent of linguistic annotation have revolutionized the fields of ASR, MT, IR and NLP.

Transcribed Speech Beginning with TI46[1] "the availability of common speech corpora for speech training, development, and evaluation, has been critical in creating systems of increasing capabilities. Speech is a highly variable signal, characterized by many parameters, and thus large corpora are critical in modeling it well enough for automated systems to achieve proficiency. Over the years, these corpora have been created,

---

[1] http://www.ldc.upenn.edu/Catalog/readme_files/ti46.readme.html

annotated, and distributed to the world-wide community by the National Institute of Science and Technology (NIST), the Linguistic Data Consortium (LDC), and others. The character of the recorded speech has progressed from limited, constrained speech materials to masses of progressively more realistic, spontaneous and "found speech." [MINDS07 Speech Understanding Report].

TI46 is "a corpus of isolated spoken words which was designed and collected at Texas Instruments (TI) in 1980." There are 16 speakers, both male and female, speaking 46 words. It was followed by TIMIT, The DARPA TIMIT Acoustic-Phonetic Continuous Speech Corpus. "TIMIT contains a total of 6300 sentences, 10 sentences spoken by each of 630 speakers from 8 major dialect regions of the United States." [2] Equally influential have been the *read* Wall Street Journal, the conversational speech Switchboard corpus and Hub4 for Broadcast News.

"Many labs and researchers have benefited from the availability of common research tools such as HTK, Sphinx, CMU LM toolkit, SRILM toolkit, etc. Extensive research support combined with workshops, task definitions, and system evaluations sponsored by DARPA and others have been essential to today's system developments." MINDS07 Speech Understanding Report].

Publicly available electronic lexical resources Semantic Annotation must be performed against a lexical inventory. The most widely used lexical database of English is WordNet [3](Miller 1995, Fellbaum 1998). In WordNet, nouns, verbs, adjectives and adverbs are grouped into sets of cognitive synonyms (synsets); words and synsets are interrelated by means of lexical and conceptual-semantic relations including super-/subordinate relations, part-whole relations, antonymy, and lexical entailment. The resultant network structure makes it possible not only to uniquely determine all the meanings of a word but also to quantify its semantic similarity to other words. WordNet's format makes it a useful tool for computational linguistics and natural language processing. Because it is freely and publicly available for download, WordNet has become a de facto community standard for the English lexicon, providing a common sense inventory for language processing systems. Richer lexical resources such as FrameNet (Fillmore, et. al., 2001) and VerbNet (Kipper, et al., 2006) all provide links to WordNet synsets. WordNets have been created for dozens of genetically and typologically unrelated languages. As these wordnets are linked to the Princeton WordNet, their potential for use in automatic machine translation is great. WordNet's sense inventory is very fine-grained, and automatic word sense discrimination – the basis for virtually all NLP applications – has been limited when it has relied exclusively on WordNet's structure. Large scale sense tagging efforts, currently underway at Princeton, Colorado and ISI, should provide sufficient data for successful training of automatic taggers.

Document Collections. "The Text REtrieval Conference (TREC), co-sponsored by the National Institute of Standards and Technology (NIST) and U.S. Department of Defense,

---

[2] http://www.ldc.upenn.edu/Catalog/readme_files/timit.readme.html
[3] http://wordnet.princeton.edu

was started in 1992 as part of the TIPSTER Text program. Its purpose was to support research within the information retrieval community by providing the infrastructure necessary for large-scale evaluation of text retrieval methodologies." This effort has been highly effective in fostering major advances in the Information Retrieval and Question Answering and has spawned several similar efforts in Europe and Asia.

"The TREC conference series has produced a series of test collections. Each of these collections consists of a set of documents, a set of topics (questions), and a corresponding set of relevance judgments (right answers). Different parts of the collections are available from different places as described on the data page (http://trec.nist.gov/data.html). In brief, the topics and relevance judgements are available at http://trec.nist.gov/data.html, and the documents are available from either the LDC (Tipster Disks 1 – 3) or NIST (TREC Disks 4 – 5), information on collections other than English can be found at http://trec.nist.gov/data.html."[4]   These collections include: Ad hoc Test Collections, Web Test Collections, Blog Track, Confusion Track, Enterprise Track, Filtering Track, Genomics Track, HARD Track, Interactive Track, Legal Track, Novelty Track, Robust Track, Query Track, Question Answering Track, and the SPAM collection.

Parallel Corpora The rebirth of statistical machine translation in the 1990's, after it was first introduced by Claude Shannon in 1949, is directly attributable to the availability of the Hansards,[5] the official records of the Canadian Parliament. These are kept by law in both French and English, and may be legally reproduced and distributed as long as "it is accurately reproduced and that it does not offend the dignity of the House of Commons or one of its Members." This provided the researchers at IBM with a large parallel French/English corpus of closely aligned, literal translations which is ideal for statistical word alignment techniques. The astonishingly accurate translations their system was able to produce revolutionized the Machine Translation field, and is the basis for the increasingly accurate statistical machine translation of Chinese and Arabic to English which is being currently funded by DARPA-GALE, and which is in daily use in Iraq.

Linguistic Annotation as Training Data The creation of the Penn Treebank (Marcus et al, 1993) and the word sense-annotated SEMCOR (Miller et al., 1994, Fellbaum et al., 1997) have shown how even limited amounts of annotated data can result in major improvements in complex natural language understanding systems. These annotated corpora have led to the training of stochastic natural language processing components which resulted in high-level improvements for parsing and word sense disambiguation (WSD), on the same scale as previously occurred for Part of Speech tagging by the annotation of the Brown corpus and, more recently, the British National Corpus (BNC) (Burnard, 2000). They have also encouraged the development of an increasingly wide

---

[4] http://trec.nist.gov/overview.html

[5]

http://www2.parl.gc.ca/housechamberbusiness/ChamberSittings.aspx?View=H&Language=E&Parl=39&Ses=1

variety of corpora with richer and more diverse annotation. These include the ACE annotations (Named Entity tags, nominal entity tags, coreference, semantic relations and events), semantic annotations such as sense tags (Palmer et al, 2004), semantic role labels as in PropBank (Palmer, et. al., 2005), NomBank (Meyers, et. al., 2004), and FrameNet, and pragmatic annotations such as coreference (Poesio and Vieira, 1998, Poesio 2004), TimeBank (Ferro et al, 2004; Pustejovsky et al 2005) and the Penn Discourse Treebank (Miltisikaki, et al., 2004).

2. Future Directions

1) **Science of annotation.** The success of applying machine learning techniques to annotated training data has had a major impact on the field. This work is limited, however, by the availability of suitably annotated training material. For NLP systems to be trained to produce transformations that add substantially new information to textual input, they have to first be exposed to similar information in context. Thus the current depth of NLP research is defined by the depth of linguistic annotation that is available, with supervised machine learning techniques doing in many cases an admirable job of automatically producing the same types of annotation—as long as the annotation is carried out consistently. The higher the Inter-Annotator agreement is, and the greater the consistency and coherence of the original annotation, the higher probability of acceptable performance of the trained systems. Not surprisingly there is now an insatiable demand for more and more annotated data: the same types of annotations for different genres and different languages; newer, richer annotations for the original languages; parallel annotations of parallel corpora; the merging of annotations that were first done independently; and new formatting for pre-existing annotations that makes them easier to merge. For today's NLP systems, the annotation defines the task, and increasingly rich annotations are the key to more sophisticated systems. Clearly annotation work needs to become much more widely distributed to cope with this need.

What precisely should be annotated, however, remains a matter of discussion. Every annotation effort requires that someone design the annotation scheme—its principal classes, their interrelationships, the conditions or circumstances under which the annotations may be added—and without careful thought there is always a danger that the annotation reflect some hasty decisions not based on solid principles. That is, the best annotation enterprises often rest upon solid theory of other fields, such as linguistics (for language structures), rhetoric (for text structure and argumentation), stylistics (for style and subjectivity), lexicography (for lexical semantics), and so on. Furthermore, the actual procedure of annotation is not simple. A clear understanding of linguistic phenomena does not translate directly into the development of annotated data that is suitable for training machine learning systems. There are issues with respect to consistency, coherence and lack of complexity of the annotation that also need to be well understood.

At the moment successful annotation is more of an art than a science, and there are only three or four sites (in the world) that have the requisite depth of experience and know-how to produce useable annotations is a timely fashion. It is critical that these institutions

codify their knowledge in such a way that it can be transferred readily to other sites.   The field needs a concerted effort on creating an explicit description of a step by step process by which useful annotations can be achieved. Guidelines and methodologies are needed for community wide answers to the following questions:

- The corpus
    - What constitutes a balanced, representative, and timely corpus?
- The annotation
    - What linguistic phenomenon is the target of the annotation and how can that be represented?
    - How can a manual be created that ensures that the choices the annotators are faced with can be decided rapidly, accurately and consistently?
    - How can the usefulness of the resulting annotated corpora as training data be guaranteed and measured?
- The annotators
    - What qualifications, training and supervision do they need?
    - What are realistic productivity estimates?
    - What tools should they be provided with?
    - What are the principles for good annotation interface design?
- Inter-annotator agreements and disagreements
    - How should agreement be measured (kappas, F-measures, confusion matrices, etc?) and what are realistic expectations for different types of tasks?
    - How can different sources for disagreements be identified and addressed, such as carelessness, vague or ambiguous instructions, and obscure or ambiguous data? Which ones are susceptible to good software design of annotation interfaces and which ones are not?

Each different type of annotation requires a stable, language independent methodology based on guidelines and universally accessible tools. The guidelines need to explicate the details of each individual annotation process as well as interactions between the different types of annotation. For instance, the guidelines for the Proposition Bank outline a process that begins with creating a Frame File for each individual verb in the corpus to be annotated. The Frames Files provide invaluable additional direction for the annotation of the individual verbs over and above the general annotation guidelines. The NomBank annotation in turn begins by referencing the verb Frame Files for associated nominalizations whenever possible. The same approach has been used successfully for the Chinese PropBank/NomBank. However, it needs to be tested on at least two or three additional disparate language families before it can be considered a stable part of a comprehensive linguistic annotation process.

The emergence of agreed upon principles and techniques for an "annotation science" is the only way that successful annotation can become widespread and the increasing demand can be met. The creation of an"annotation science" and the resolution of the

issues listed above will require a substantial research effort that would be greatly facilitated by a regular technical forum for discussion such as an ACL SIG.

A measure of the success of this objective would be substantial progress on providing the public domain infrastructure described in Challenge 2. Even in a rosy distant future where unsupervised techniques that do not need training data will prevail, annotated corpora will still be needed for testing and evaluation purposes and to define the desired levels of representation.

2) **Robust, extensible annotation infrastructure.** Along with a better understanding of a methodology for annotation there should be a set of public domain tools and interfaces that can support, and to a certain degree enforce, "best practice" annotation guidelines. The simplest way to ensure that everyone adheres to these guidelines and agreed upon formats is to have everyone using the same tools and APIs. This could also provide a framework for facilitating interaction with automatic taggers as pre-processing components, as described in (3).   There are also similar data tracking needs shared by all annotation projects for which a toolkit could provide a general purpose solution.

The current emphasis on fast, cheap annotation impedes the type of collaboration described in (3) since it does not allow time for experimentation with more sophisticated approaches that would both incorporate already existing technologies and tools, and would generate systems that are generally useful to other researchers, rather than systems that are task specific (See Challenge 3). "Quick and dirty" software development tends to be the norm, dictated by short funding cycles and an emphasis on task-based software. A shift in focus to developing reusable, library-based code with task-specific modules and customizations would greatly improve the overall quality and cohesiveness of annotation systems to the benefit of the entire MINDS community of users.

Development of such an industrial strength toolkit will only be possible with **sustained support** for a robust linguistic annotation infrastructure that can marry the "annotation science" developed in (1) with well understood principles of software engineering. Coherently layered annotation clearly benefits both technology developers and resource developers. The CRI program at NSF now explicitly acknowledges annotation as an important research area contributing to Computing Research Infrastructure, but those funds are limited. The support of other agencies is needed to turn this into a widespread research effort.

Some of the key challenges of this program would be:
- Defining data desiderata to support machine learning techniques
- Defining principles for creating modular annotations which can easily be layered with other annotations. For example, clearly defined standoff XML annotations and compatibility between layers would allow for a standard query builder that knows how to access layered annotations.
- Implementing an annotation infrastructure starter kit for new efforts in linguistic resource creation that embodies these desiderata and principles. This would ideally be centrally localized where it can be downloaded, and where people can

submit annotated data to be archived and distributed. Copywrited data can't be shipped back and forth but standoff annotations can be.

- A regular forum (such as the ACL SIG mentioned in (1)) that includes funders for more public discussion of data resource priorities such as
    - New genres (e-mails, blogs, text messages, meetings, ?)
    - New languages for the resource kit
    - Richer annotations

3) **Closer integration of emergent technology**.   There is considerable evidence that the productivity of manual annotation can be speeded up by pre-processing the data with sufficiently accurate automatic taggers (Chiou, et al., 2001). However, current annotation practices frequently fail to take advantage of this approach, possibly because of the difficulty of integrating these systems into new annotation tasks. Even more benefit could be derived from using sophisticated machine learning techniques to aid in the selection of instances to be tagged, in order to maximize their utility and minimize the total annotation effort. For simple classification tasks like Word Sense Disambiguation there are accepted practices such as active learning techniques, (Chen, et al., 2006). However, for more complex annotations such as syntactic structure, pinpointing novel or unfamiliar items in the data remains an unsolved problem. Fundamental research is needed to develop informed active learning techniques for complex annotation. This will require collaboration between the researchers in the community doing annotation and those developing the machine learning techniques used for system building. Closer ties between annotators and the ASR, NLP, MT, IR and Machine Learning communities are also needed for joint efforts to develop techniques to aid data selection and quick access to modular automatic taggers for preprocessing of data. This will facilitate faster, easier, more complete integration of emergent technology into the human annotation pipeline, which will in turn improve the quality and availability of annotated data for those communities.

Research is needed to explore issues such as **when** in the annotation pipeline technology can be folded in with the maximal benefit, and **what** levels of accuracy are necessary for the automatic taggers to provide a benefit.   A careful study of bootstrapping automatic taggers from small amounts of annotated data to find the minimal amounts of additional annotation required to achieve maximal performance is needed.   Recognizing the importance of the improvement of annotation practices as a valuable research area in its own right will allow the time for progress on this front, and will encourage technology developers to hand over state-of-the-art systems for early integration.   The goal should be, in addition to faster and most useful training data, the identification and selection of hard/rare/special data for annotation that supports better use of limited human annotation resources – bigger "bang for the buck." This will stimulate research in machine learning techniques as well as maximizing the impact of limited amounts of annotation, allowing humans to focus on the parts that really require human judgment.   We need to annotate smarter, not necessarily harder.

4) **Richer annotations**. As described above, progress in natural language processing is being led by the definition of increasingly rich levels of representation. Given sufficient, good quality training data, it is likely automatic taggers can be built to replicate whatever representations they were trained on, to a degree relatively close to human performance. This makes the search for increasingly rich levels of representation that can be successfully annotated the highest priority in the field. This is fundamental, theoretical research requiring expertise in linguistics, annotation and machine learning. Examples of immediate pressing priorities for the following application areas are listed:

- **Machine Translation** – Parallel corpora with parallel annotations, i.e., source documents in multiple languages, translated into multiple languages (all pairings) and sentence aligned; with (complete or partial) coherent annotations across source and translated documents in each language.

- **Spoken Language Understanding** – Rich corpora transcriptions that include phonetic information, term detection, semantic roles and sense tags, discourse structure, dialect information, etc.

- **Natural Language Processing** – Identification of events and relations between them (causal, temporal, modal, etc.), discourse relations, etc. Also parallel corpora (see above) for research on projecting annotations onto (translated data in) another language given reference annotations in one language.

- **Information Retrieval** – annotations for speech corpora that can facilitate the application of search engines.

5) **Language resource kits.** There has long been recognition of the need to have basic language processing resources available for a broad spectrum of languages. When a language unexpectedly becomes of vital strategic importance, quickly having access to resources for either the language in question or for a closely related language could be crucial. These resources would enable rapid ramp-up of a basic news understanding system for text and for speech, and allow information retrieval over document sets. There is an existing effort at LDC to develop *language resource kits* for a small set of languages: Urdu, Thai, Hungarian, Bengali, Punjabi, Tamil, and Yoruba. The resources being developed for these languages include:

- monolingual text, parallel text, part-of-speech taggers, morphological analyzers, and Named Entity annotation.

This level of resource development is at best minimal. Speech recognition needs are completely ignored, as are syntactic and semantic structure. A more complete resource kit would also contain:

- **text** - at least 100K words of parallel text (news domain), tagged for nominal entities & coreference, basic syntactic annotation, basic predicate argument structure, topics + relevance judgments for news articles;
- **speech** - a pronouncing dictionary, a minimum of 100 hours of audio, a minimum of transcription of 10-25 hours, 50% Broadcast News, 50% Broadcast Conversation, possibly interviews.

Ideally these language kits should be made available for 100-200 of the less commonly taught languages.

6) **Broad coverage, empirically grounded lexical resources**. In the quest for improving the portability of supervised stochastic systems, one under-utilized resource is the lexicon. Many supervised approaches depend heavily on lexical cues, and balk when given data with out-of-vocabulary lexical items. One would think that general purpose lexical resources providing similarity links between items would be helpful in alleviating this impasse. So far that has not been the case. Either the lexical resources are not organized in the most effective way, or the stochastic systems are not set up to make use of them, or both.

There is widespread agreement that certain new lexical resources are extremely desirable and could be of great benefit.   For example, in order to address the pressing problem of correctly recognizing references to the same person in different documents which can also be in different languages, the community has been clamoring for a normalized list of references with pointers to instantiations in documents; a cross-document, cross-lingual, Entity Co-reference corpus. From one perspective, this can be viewed as a name lexicon that, in addition to the list of names, includes multilingual spelling variants of each name as well as multiple document instances. It can include nominal entity references as well as named entity references. All of the documents in the set would have all instances of every listed item marked.   There are several other examples of "lexicons" that are effectively specific types of lists of items with corresponding pointers to instances in data. A classic example is a bilingual lexicon, which in this case would also include pointers to word usage instances in documents in a parallel corpus for the two languages. Example-based Machine Translation systems contain a vast amount of information which can be seen as another variant of this type of resource. Again, this resource could potentially be of use for improving the portability capabilities of statistical Machine Translation systems.   Tying the "word lists" to instances in corpora helps bridge the gap between the classic lexicon and the needs of a stochastic system.

Another drawback of existing lexicons is that each application area has its own specially tuned lexicon with information customized to that domain that the other areas do not need and have no interest in providing. Statistical MT systems could be extracting extremely useful bilingual lexicons from aligned corpora, but without phonological information they will be of no use to ASR systems. At the moment the various lexicons are so diverse that different systems cannot ensure that they are referring to the same items. A public domain resource that lists all of the relevant types of information for each lexical

unit, thus enabling the ASR, NLP, MT and IR systems to recognize that they are all dealing with the same lexical items is essential.

The next step is to create monolingual resources that can provide support for automatically detecting semantic similarity. These resources should also be empirically grounded, with pointers to instances that illustrate the desired paraphrases and inferences. For instance, phrases such as "the rising stock prices" and "the surge in the stock market" can both be linked to a similarity set such as "rise, surge," in the same way that a bilingual lexicon links the English/Spanish translations "rising, levantando" and corresponding aligned sentences. The current surge in sense tagged data is making progress in this area, but much more can be done.

Looking even farther ahead, we will need even richer lexical resources (monolingual and bilingual) that describe event hierarchies, and are closely integrated with ontologies associated with rich knowledge bases that can support inferencing (Hovy, et al., 2006). This is the only way we will be able to move forward from our current shallow semantic processing to the deeper levels of comprehension needed for complex question answering and information extraction. Here the empirical grounding might be phrases or parts of sentences that provide the evidence for perceived entailments. These "nuggets" of information correspond closely to the "compositional units" that the MT community is also interested in focusing on.

Bibliography

Chen, J., A. Schein, L. Ungar and M. Palmer. 2006. An Empirical Study of the Behavior of Word Sense Disambiguation. *Proceedings of NAACL-HLT 2006*

Chiou, F., Chiang, D., and Palmer, M., Facilitating Treebank Annotation Using a Statistical Parser, *Proceedings of HLT 2001, First International Conference on Human Language Technology Research,* 4 pp., J. Allan, ed., Morgan , San Francisco, 2001, pp. 117-123.

Fellbaum, C. 1998, ed. *WordNet: An Electronic Lexical Database*. Cambridge: MA: MIT Press.

Fellbaum, C. and J. Grabowski, and S. Landes. 1997. ``Analysis of a hand-tagging task" *In: M Palmer & M. Light (eds.) Tagging text with lexical semantics., ACL*, Washington DC, pp. 34-40.

Fillmore, Charles J. and Collin F. Baker. (2001). Frame semantics for text understanding. *In the Proceedings of NAACL WordNet and Other Lexical Resources Workshop* Pittsburgh, June.

Hovy, E.H., M. Marcus, M. Palmer, S. Pradhan, L. Ramshaw, and R. Weischedel. 2006. OntoNotes: The 90% Solution. Short paper. *Proceedings of the Human Language*

Technology / North American Association of Computational Linguistics conference (HLT-NAACL 2006).
New York, NY.

Kipper, Karin., Korhonen, Anna, Ryant, Neville, and Palmer, Martha, Extending VerbNet with Novel Verb Classes, *Proceedings of the Fifth International Conference on Language Resources and Evaluation – LREC'06,* May, 2006, Genoa, Italy.

Marcus, M.P, B. Santorini, and M.A.Marcinkiewicz. 1993. Building a large annotated corpus of English; The Penn Treebank. *Computational Linguistics, 19(2): 313-330*

Meyers, A., Reeves, R, Macleod, C., Szekely, R., Zielinska, V., Young, B., and Grishman, R., (2004), The NomBank Project: An Interim Report, *Proceedings of the HLT-EACL Workshop on Frontiers in Corpus Annotation*, Boston, Massachussetts.

Miller, George A., Martin Chodorow, Shari Landes, Claudia Leacock and R. G. Thomas. ``Using a semantic concordance for sense identification.'' In: *Proceedings of the ARPA Human Language Technology Workshop,* San Francisco, Morgan Kaufman, 1994.

Miller,   George   A.   1995.   ``WordNet: a   lexical   database   for   English.'' *Communications of the ACM 38 :39-41*.

Miltsakaki,E., Prasad, R., Joshi, A., and Webber, B.. Annotating Discourse Connectives and their Arguments, In *Proceedings of the HLT/NAACL Workshop on Frontiers in Corpus Annotation.* Boston, MA. 2004

Palmer, Martha, Dan Gildea, Paul Kingsbury, (2005) The Proposition Bank: A Corpus Annotated with Semantic Roles, *Computational Linguistics Journal, 31:1.*

Poesio, M., 2004. The MATE/GNOME Scheme for Anaphoric Annotation, Revisited, *Proceedings of SIGDIAL, Boston, MASS. Poesio and Vieira, 1998 A corpus-based investigation of definite description use. Computational Linguistics, 24(2).*