Drawing on Linguistic Context to Resolve Ambiguities

OR

How to imrove recongition
in noisy domains

V Kripasundar

A dissertation proposal
submitted to
the Department of Computer Science
at the State University of New York, Buffalo
in partial fulfillment of the requirements for
the degree of Doctor of Philosophy

Dissertation Committee
Dr. Rohini Srihari *Chair*
Dr. William J. Rapaport
Dr. S. N. Srihari

June 15, 1994

To
the countless animals
that have been valued below humans,
and hence experimented upon
in the name of science

ABSTRACT

Since almost every natural source of language is fraught with noise, any attempt to process this input by machine must utilise disambiguating techniques to improve recogniser performance. The purpose of my research is to bring the vast amounts of contextual information present in natural English to bear on the performance of the recogniser. I hope to demonstrate that recognition in a noisy domain can be facilitated by the proper application of information theory to the cues that the language has to offer.

"Show me your friends, and i'll tell you who you are" — *Anon.*

Typesetting notes:

underline	used only for emphasis
sans serif	technical terms
italics	foreign phrases *viz*, stuff like *sans serif*
"quotes"	inline quotations, and/or
"quotes"	informal/casual use of words (some words are casualler than others ...)
'quote'	mention (and not 'use')
[Yeti 94]	citings
i	first person singular (*ie*, i)

Kindly use context to resolve inconsistencies.
And humour the spelling.

Contents

1 Introduction **1**

1.1 Background 1

1.2 Motivations 2

1.3 Corpus techniques 3

2 The Problem **4**

2.1 The domains 4

2.2 The recogniser 5

2.3 The post-processing 5

2.4 The solution 6

3 Information-theoretic underpinnings **7**

3.1 Entropy, perplexity and peculiarity 7

3.2 Measures of contextual information 8

4 My contributions, hopefully **9**

4.1 A context window in the input stream 9

4.2 A context hierarchy 10

4.3 A unified framework for the above 11

5 Results so far **12**

6 What needs to be done **14**

7 Loose ends **15**

8 Conclusion **15**

1 Introduction

In my dissertation, i wish to examine the use of context in improving machine recognition of English input. I use the term context to refer to linguistic context, rather than domain-specific ideas like visual context. In other words, i plan to exploit the constraints that the syntactic, semantic, and lexical contexts impose on a word, so as to help the recogniser make more informed choices among word candidates for the various word positions in the input stream. I'll assume a word-based recogniser, such as those used in the recognition of handwriting.

The focus of this research will be in identifying individual techniques for utilising the different types of context, and in achieving an effective integration of these techniques. This can best be achieved by treating natural language as an information source, and the recognition medium as a noisy channel for this information. The effectiveness of contextual information as a source of disambiguation is directly related to the accuracy of the language model used. Information-theoretic measures such as perplexity and peculiarity can be used to measure this.

1.1 Background

English has long been studied from an information-theoretic point of view as a stochastic process. Most of these studies have modelled English as a stream of characters, and have helped quantify the intuitive notion that natural language is redundant. Shannon's pioneering research into the entropy of English led him to conclude that the English alphabet had an entropy of approximately 1.4 bits per character. Further studies have used larger data samples and more rigorous analyses to demonstrate that the exact value is somewhere between 1.3 and 1.7. This means that the next character in a stream of English text can be predicted with an accuracy of approximately 30%, without any knowledge of (or analysis of the text for) its semantic content. Although this analysis applies at the character level, it is intuitively clear that the redundancy is, if anything, even greater at the word level. Cover [91], Brown [92a,b] and others have studied this to some degree.

Word-level redundancy is demonstrated by several phenomena in human understanding of English. Our ability to communicate effectively through speech even in very noisy environments (such as at a carnival, or in a café), and our ability to infer the meanings of words from context are two immediate examples. In fact, this research is precisely an attempt to model the former. The ease with which we sail through spoonerisms, malapropisms, and puns is further evidence of our aptitude for detecting gaps in communication and filling them. The fundamental idea in the use of context is that the human mind can take in concepts in any order at all, figure out how they fit in with each other, and use background knowledge to fill gaps in.

However, analysis (or synthesis!) of the meanings of words is beyond the purview of this research, and hence i will not talk about the issues involved in tackling the kind of ambiguity (*ie*, semantic ambiguity) illustrated by the title of this thesis*.

In the rest of this section, we will look briefly at the motivation for the use of contextual information,

*If you didn't notice the "titular"† ambiguity until i drew your attention to it, that is but another demonstration of the effortless ease with which we disambiguate input. It also serves to strengthen my (future) point that we should draw from the techniques that humans use so successfully.

†webster says: tit.u.lar /ˈtich-(ə-)lər/ aj. [L *titulus* title] existing in the title only ☺

and for the use of corpus-based language models in utilising this information.

1.2 Motivations

Ever since computers were invented, we have been trying to get them to understand "humanspeak". Thus, the problem of machine recognition of English input has been of relevance since the big bang, almost. The need for practical systems that the lay user can interact with in natural language cannot be overstated. Apart from such obvious advantages as providing a pleasant interface to the user, natural language capability would provide the machine with an (heretofore unthinkable) opportunity to learn from the countless knowledge sources that are available to humans, and would also allow machines to assist humans in tasks like paraphrasing, translation, and précising and thus even facilitate interactions between humans.

To dispense with the foundational "Why?" before it arises, i would like to say here that machine recognition is a "problem" because of noise in the input to the recogniser. When we talk about the recogniser's domain being "noisy", we would naturally like our system to detect (and recover from) noise in generation as well as noise introduced in transmission. Firstly, the input to the recogniser could be ambiguous and/or error-prone (noise in generation). Specifically, in the domains of online handwriting and natural speech, the language used is highly informal, and "ungrammatical" (for example, by the use of sentence fragments, abbreviations and the like). The actual input signal itself will have further noise from the process of online editing that occurs in spontaneous language generation. Additionally, the input mode (such as a scanner or speech recognition hardware) would corrupt the input signal further (noise in transmission). Finally, the recogniser (say, an OCR) might itself add to the noise. It would be the task of post-processing modules (*eg*, the linguistic context module that i advocate here) to get the chaff out, and pass the grain onto the "true" language processor, which would presumably "understand" the input.

The inadequacy of straightforward approaches to this problem has long been recognised, and has resulted in greater respect for, and interest in, human processing of natural language input. Studies in fields ranging from psychology to AI have demonstrated conclusively that humans rely heavily on context to achieve speed and accuracy in processing. Both (*ie*, speed and accuracy) are aspects in which current machine performance trails sorely behind human performance. This awareness has led recently to work that exploits context to speed up the recogniser, and to make it more accurate in its judgements. Much of this research has been in the area of speech processing, and some of it has also seen application in machine reading.

I have elsewhere argued [93] for the inescapable need for the use of context in communication.

> ... It has been observed by Barwise and Perry that one of the important aspects of language, called efficiency, is that much of the semantical contribution of linguistic use arises from circumstances of utterance, and not directly from the structure of the sentence used (deictic usage, for example). Communication is possible in this context only if language use and world knowledge directly interact in the cognitive agent's mind. ... (p. 10)

With a little help from François Rabelais, Samuel [90] expresses the same idea (p. 295): "Perception abhors a vacuum.".

There is also a lot of evidence from cognitive psychological studies that contextual information

affects the activation levels of lexical candidates during the process of perception. Warren [70] describes the phonemic restoration effect, where although a phoneme was substituted with extraneous noise, listeners perceived the utterance as intact. Marslen-Wilson's cohort model [90] gives evidence to show that "[human] perceptual processing is based on competition between simultaneously active candidates" (p. 148). Such results also motivate the notion of context window that i will use extensively in my research.

1.3 Corpus techniques

Automatic machine learning of the degrees of relevance of various contextual cues is achieved by training the system on corpora. Although corpus-based research was frowned upon, as recently as a decade ago, by mainstream computational linguists, it has gained acceptance as perhaps the only method that offers some hope in handling "real" natural language. This change of attitude can be attributed in large measure to the demonstrable success of corpus-based techniques in speech processing, which were in turn motivated by the availability of divers large corpora. At the risk of flogging a dead horse, i would like to point to the work of Garside, Leech and Sampson [87] as evidence of the effectiveness of, and the need for, corpus linguistics. Sampson makes the case for corpora better than i possibly could, and i'll do well to simply quote him on the matter:

> ... If it is correct to describe linguistic behaviour as rule-governed, this is much more like the sense in which car-drivers' is governed by the Highway Code than the sense in which the behaviour of material objects is governed by the laws of physics, which [latter] can be never violated. ... [Sampson 92, p. 429]

His point is that, except in that most excruciatingly formal of settings, namely writing to be published, we never bother to edit our "mistakes" in language generation. People working in speech processing are all too aware of this fact, for they have to deal with the phenomenon of on-the-fly editing that all of human speech is subject to. If we want our machines to handle such (unedited) input, we must be prepared to teach them "raw" English. Since our "Highway code" is inadequate for this task, we must have the mental fortitude to let them learn by themselves.

Here it is important to consider the issue of psychological or cognitive validity of "number-based" approaches. It will be argued that numbers and fractions play no part in our generation or analysis of natural language for communication. Even Sampson and company do not appear to defend the use of corpus techniques as being cognitively valid. I, on the other hand, would like to take the position that corpus linguistics is not just an "anything-that-works" approach to machine recognition of language. Admittedly, attempts to introspect about the exact probability that a given word (say, 'horse') will be followed by another ('rider') are bound to fail, and precisely because such probabilities and weights, if present, must be below the symbol-processing level of consciousness.

However, i would like to argue that statistical techniques are very similar in spirit to neural computation. In fact, it is my thesis that corpus analysis and the computation of statistical measures (transition probabilities and so on) is exactly equivalent to the automatic learning of weights that happens in neural nets. As Aleksander and Morton [93] point out, symbolic approaches (such as generative grammars) are logic-based, and suffer from the inescapable need for a crutch in the form of a programmer. Edward

Sapir [21] identifies[†] the defect in this crutch when he says "All grammars leak." (p. 38). The idea behind this statment is that no amount of planning on the part of the grammarian-programmer can be sufficient to handle all possible natural language utterances. This almost sounds like a paraphrase of Sampson's point! Corpus techniques are based on the logic that bypassing programmer-generated grammars (and, instead, making the system learn from actual utterances) will help avoid the concomitant shortcomings in the plumbing.

2 The Problem

Secure in my assumption that i have convinced you of the need for a language module that exploits contextual information, i would like to talk about the problem itself, and what i hope can be accomplished within a reasonable time-frame. The subsections that follow talk in more detail about the recogniser, the post-processing, and the techniques involved therein.

I would like to aim for a reasonably broad range of applications for the theoretical ideas that i propose and test out, although i will have a relatively narrow focus in the implementation part of the project. The [my-contrib] section, for instance, talks about a context hierarchy that spans several levels, even though my work will largely be confined to the couple of levels at the top of the hierarchy.

2.1 The domains

Rather than specialise our approach to the task of machine reading of natural language input, it will be useful to look at different sources of spontaneous language, and to evolve techniques that are applicable to all or most of them. The following is a summary of the different kinds of input sources for machine recognition of natural language:

- speech recognition
  - discourse
  - narrative *eg*: dictations, recitations
- ascii-based
  - email (short messages addressed to a few persons)
  - news articles (long-ish articles written for a general audience)
- writing (*or*, "visual")
  - machine-print
  - handwriting (cursive, discrete, or mixed)
    - online (*ie*, pen-based human-computer communication)
    - offline (*ie*, images of handwriting)

[†] unbeknownst to himself

Several features of this classification can immediately be noted. Firstly, email and news articles roughly correspond to the discourse-*vs*-narrative distinction in ascii-based input. Secondly, offline handwriting is similar to machine-print in the essentially static nature of the input. Thirdly, online handwriting is similar to speech in that both sources provide signals that are temporal in nature, and represent spontaneous generation of language as well as the signal itself. Finally, it is an empirical fact that ascii-based input and speech input are similar to online handwriting, in terms of the style of language used.

2.2 The recogniser

Since our techniques should not be limited to the realm of machine reading, we need a general formalism for the input recogniser. I would like to target all word-based recognition systems[‡]. A word-based recogniser identifies word-breaks in the input stream, and, for each resulting word-position, attempts to identify the word that the input signal is a token of. Given the potential infinitude of input sentences, it seems inevitable that any practical recognition system must, of necessity, break its input down into tokens. This approach is natural, and is justified by the fact that the generator itself constructs the sentences from word-tokens that are presumably chosen from a lexicon of word-types.

The output of such a recogniser would then be an ordered set of candidate words for each word-position in the input stream. While the ordering is itself a measure of confidence, the recogniser may further associate a confidence-value with each of its word choices to communicate finer resolution in its output. It is these confidence-values that the language module must update to reflect contextual knowledge and expectation.

Given this set of word-confidence pairs for each of the word positions, the task of the language module would be to process these word-candidates, so as to maximise the degree to which the sequence of top choices matches the input stream. In general, we would actually be interested in maximising the probability, for every value of n, that the truth lies within the top n word-choices. We would also like to use linguistic context to actually reject some of the word-candidates that the recogniser computes. The motivation behind such deletions in the choice-list is that a substantial reduction in the average size of the choice-lists would speed up subsequent post-processing significantly.

2.3 The post-processing

The language module can, in principle, perform several types of "post-processing" on the word-candidate lists that the recogniser outputs for the different word-positions. These can broadly be classified into emendments and amendments of each word-list.
Emendings correct or modify the existing word-list, and can be divided into:

- re-ranking the word-list (and assigning new confidence-values to its entries), *and*

[‡]Here, the word 'word' should be interpreted liberally to mean the "basic" unit of the language. Thus, for many non-Indo-European languages, a "word" might be closer to what English calls a clause or a sentence. Several Amerindian languages exhibit this phenomenon known as polysynthesis.

Corollary: my techniques should be applicable to any language, as long as 'word' is suitably defined, and all dependent definitions changed accordingly.

- [stronger] deleting those entries of the word-list that have very low confidence associated with them (as a consequence of applying contextual knowledge to refine the word hypotheses)

Amendings do not modify the word-list *per se*, but augment the word-recogniser output. They, too, can be classified into two types:

- giving feedback to the word recogniser along the lines of:
  - syntactic/semantic categories expected in a given word position (eg: numbers, proper nouns etc)
  - prefix/suffix that has been guessed/ruled out, *and*
- [stronger] proposing new words that the recogniser did not even suggest as candidates

There is, of course, the utopian possibility that the language module may actually detect word segmentation errors on the part of the word recogniser, and give it feedback on the same.

Several researchers view this setup as a "generator-filter" model (*cf.* [Woods 83]), where the recogniser is viewed as a generator of word hypotheses, and the post-processing module is interpreted as "filter"ing out unreliable hypotheses. It seems to me more appropriate to regard each stage as a refiner of hypotheses. This makes it (arguably) more meaningful, in that it enables us to treat the setup as a cascade of refiner modules, and the entire process as successive "credentiation" of word hypotheses. This also allows the language module to be designed and used as a set of closely interacting modules, instead of restricting it to be a black box of techniques that may not logically belong with each other.

2.4 The solution

The techniques that i plan to use for achieving the amending and emending are rooted in information theory. Most of the technical ideas that we brush against now will receive better treatment in the next couple of sections. We know that the recogniser output is a series of word-lists which represent the recogniser's guesses for the word positions. This structure has been called a trellis by [Rohini&Charlotte] for its two-dimensional shape. Our essential technique will be to slide a window across the trellis, updating the word-list in each position depending on the confidence-values littered all over the trellis.

The system will initially be trained to learn transition statistics at different levels of the context hierarchy. This training will be performed offline on the email corpus. The trained system can then perform real-time analysis of the trellis, and use the confidence-values on the trellis in conjunction with its own (internal) learned values of transition probabilities, to update the trellis. It is likely that simple morphological pre-processing will be required to make the process efficient and effective. Once all updates have been performed, the context window can be slid on the input, and the process repeated on the new trellis. The width of the window is likely to be large ($\approx$10), and will influence the speed and the accuracy of the language module.

The work of Zwitserlood, Marslen-Wilson, and others demonstrates that multiple lexical candidates are simultaneously active during perceptual processing, and that these lexical candidates have activation levels associated with them. This gives us motivation for the sliding window technique. We will use

entropy optimisation techniques to integrate the contextual information that is available within the context of the window. This minimises the perplexity of the resultant language model.

We will further need to use good peculiarity indices [Dunning 93] to accurately model "rare" events — events which, in practice, constitute 20-30% of all events! (p. 62). Dunning demonstrates how a Pearson's χ^2 test can give very misleading results for low-frequency events, and develops a measure based on likelihood ratios. These ideas are likely to be very relevant to our system, especially if we try to incorporate a very large, learning lexicon.

3 Information-theoretic underpinnings

Here i present a brief outline of the information-theoretic ideas which contribute to the proper and effective use of contextual information. The central justification for the use of statistical methods in improving recognition is that natural English can be very effectively modelled as a stochastic process. In practice, this means that each word-position can be treated as a random variable that can take any word in the English language as its value. The issue of context then translates into that of determining the type and degree of influence that one or more of these random variables have on others.

3.1 Entropy, perplexity and peculiarity

Researchers in information theory have come up with several related formalisations of the ideas of context and contextual influence. The most commonly used measures among these are mutual information, conditional entropy, and joint entropy. (The concept of entropy is perhaps the most fundamental formal conceptualisation of the idea of "information", and any standard textbook on information theory, such as Cover&Thomas [91], will provide a good introduction to entropy and related notions.) All three measures are defined over pairs of random variables[§], say X and Y, and are closely related to each other. While the mutual information $H(X,Y)$ is a correlation measure that gives us the amount of information that Y gives us about X, the conditional entropy $I(X|Y)$ is a measure of the information that is "still hidden" or "left over" in X when Y is known. Thus, for any Y, we have: $H(X) = H(X,Y) + I(X|Y)$.

The usefulness of a particular language model is directly dependent on its ability to predict the next word in the input stream, based on the history of the stream. This ability is measured by the perplexity of the model. The concept of perplexity is closely related to that of maximum likelihood estimators (MLE's). An MLE is a measure that attempts to predict the most probable value of a random variable given some constraints. In our case, the constraints are the contextual information that the system has deduced. MLE's are derived by trying to optimize perplexity.

We also need a meta-level measure for deciding which definition of perplexity is acceptable at a given point. What we need is a peculiarity index that helps us decide whether we have sufficient instances of the co-occurrence for the (inferred) correlation to be statistically significant. Dunning's work [93] in comparing "traditional" χ^2 tests against likelihood ratio tests demonstrates the need for such self-monitoring.¶

[§]Note that, in the context of modelling English, we will be interested only in discrete random variables. Hence i will discreetly drop the qualifier, and use the phrase 'random variable'.

¶Dunning found that words like 'landscape', 'penetrate', 'specify', and 'winter' occurred at least once, but ⋯ ⟹

3.2 Measures of contextual information

Many researchers [Smadja XX, Rose XX] have suggested that the notion of mutual information directly captures the idea of context. However, our research indicates that non-standard measures are likely to yield better fruit. The motivation for this claim is multi-pronged. Firstly, playing around with the formula (*cf.* the tables in §5) shows how a biased measure performs significantly better than the mutual information score (known as MIS). Secondly, MIS is deficient in its ability to detect one-sided correlations, and it appears that asymmetric influence measures are required to properly handle them. For example, the suggestion that 'age' exerts as much influence on 'Paleolithic' as *vice versa* (in the phrase 'Paleolithic age') seems ridiculous, to say the least. Thus it seems quite unlikely that any symmetric information measure can accurately capture the co-occurrence relationship between the two words. Doyle [1892] holmes in on the issue when he says:

> Singularity is almost invariably a clue. The more featureless and commonplace a crime is, the more difficult is it to bring it home.

Further, it has also been observed that mutual information tends to overstate the significance of unexceptional cases, thus making it unsuitable for measuring contextual information properly.

What is needed here is a one-sided information score (OSIS), something that serves as a measure of influence of one word on another, rather than as a simple, symmetric, "co-existence probability" of two words. The table below illustrates how an OSIS can be effective in detecting lexical and lexico-semantic associations.

letter	$\rightarrow$	to	1.24
the	$\leftarrow$	editor	1.15
financial	$\leftarrow$	institutions	2.438
interest	$\leftarrow$	rates	5.211
interest	$\rightarrow$	rates	4.902
third	$\leftarrow$	world	1.286
third	$\rightarrow$	world	1.928
offended	$\rightarrow$	by	3.185
implicated	$\rightarrow$	in	1.72
added	$\rightarrow$	that	1.681
carry	$\rightarrow$	out	2.186
arguing	$\rightarrow$	that	3.641
awareness	$\rightarrow$	of	2.001
has	$\leftarrow$	been	3.493
has	$\rightarrow$	been	2.833
opposed	$\rightarrow$	by	1.245
approved	$\rightarrow$	by	2.17
	$\vdots$		

Directed influence

fewer than five times, in a half-million word sample of Reuters' reports. Such frequency statistics also show that handling rare words and implementing a learning lexicon would go hand in hand.

The word-pairs in the table have been selected randomly from the test-set with the criterion that their directed influence score (DIS) was "significant" (*ie*, > 1.2). The DIS in this case was defined as follows:

$$\begin{aligned} \text{MIS(xy)} &= \log(\tfrac{P(xy)}{P(x)P(y)}) \\ \text{OSIS(x, xy)} &= \log(\tfrac{P(xy)}{P(y)}) \\ \text{DIS(x, xy)} &= \text{OSIS(x, xy)} + \text{MIS(xy)} \end{aligned}$$

The arrow in each row indicates the direction of influence (or flow of information). The preponderance of word-pairs that do not influence each other to similar degree shows that no symmetric score could have captured the correlations in all of these phrases.

Although preliminary attempts at using DIS have not succeeded dramatically, there is a lot of ground to be explored yet. It is noteworthy that the DIS successfully captures the notions of lexical and lexico-semantic associations in one formula. Thus we have 'interest' and 'rates' influencing each other almost equally, while the asymmetric influence of 'approved' on 'by' is also detected by the DIS.

To continue this line of thought, what we really need is a generalised information score, a measure that takes into account several factors, such as:

- directionality in correlation
- multiple words participating in a lexical relationship
- different (morphological) forms of words *and*
- spatial flexibility in the components of a collocation

The generalised information score would capture all the variations that are introduced by the above factors, and allow for the variants to reflect a "normalised" measure of contextual influence. The concept of phrase edit distance suggested by Rohini is an example of a generalised information measure that allows for local variations. If we use phrase edit distance (rather than simple string distance) for comparing phrases, the phrase 'savings and lawns' would match up well with the (presumed) original, namely, 'savings and loans'.

4 My contributions, hopefully

4.1 A context window in the input stream

Although there have been several attempts in recent years to exploit context in improving recognition results, to my knowledge, no one has applied the concept of a sliding window of context. The activation levels of word candidates for the different word positions within the window grow and decay as a function of all the active candidates visible in the context window. Apart from having psychological validity in terms of activation ideas, this also makes for a practical, realtime use of context, and lays the foundation for broader application of the ideas developed herein. The collocational analysis system already implements the context window idea.

4.2 A context hierarchy

I have identified a hierarchy of contexts, most of which are likely to be of use in our system. The ordering in the context (or association) hierarchy is based on the form-content (or structure-function) continuum, and exhibits an almost parallel gradation in the rigidity demanded of the spatial relationship between the primer and primed. Although the form-content distinction is well-understood[‖] in classical philosophy, it will be useful to remember that i refer to the syntactic or structural aspects of the input when i use the term form, while the term content can be loosely paraphrased as "meaning" or "purport". Connine [90] (2) (2) describes content as "connections within the lexicon" (p. 283).

Position in the continuum	Type of association	Spatial rigidity	Example
content	lexical/associative	very inflexible (usually adjacent)	'parking lot'
⇑	lexico-semantic	word order rigid	'sneak off'
	theme-based	very flexible. almost complete absence of spatial constraints	'plumber' & 'money'
⋮			
	case frames	order isn't inflexible. spatial proximality needed.	'put' & locative
⋮			
	semantic	proximality rqd (order flexible)	'class' & <semester>
	part of speech	no concept of "word order". POS transitions only.	determiner & adjective
⇓			
form	constituent	"hypertag" transitions	NP & relative clause

The context hierarchy

In the case of lexical/associative context, the primer and the primed are objects of thought that are themselves constituents or components of a single object of thought. Hence, associative context is characterised by rigid word order, and usually implies that the primer and the primed together act as one lexical unit. *eg*: 'managing director' acts as a noun. you can pluralise it ('managing directors'), but you cannot change its "internal structure" (*'managed director' *and* *'managing direction').

The best examples of lexico-semantic association are phrasal verbs. This type of association is called so because of the presence of fixed words (lexico–) but with idiomatic meaning (–semantic). A lexico-semantic association is usually a verb-preposition combination. Very often, for a given verb, the number of prepositions that can form a phrasal verb with it are very limited (*eg*, 'tear up', 'tear down', or 'tear off', but not 'tear to', 'tear for', or 'tear by'). However, phrasal verbs are characterised by morphological complexity in the verb part and spatial flexibility in the phrase as a whole.

In topic/theme-based association, the primer and the primed can be loosely thought of as being members of a single sub-unit withing the semantic network of the mind, *eg*: 'lawyer' and 'court'. This is very close to the idea of priming or activation as it is used in the cognitive psychology literature.

[‖] sort of

Semantic association is similar to theme-based association, but is a much more local phenomenon that gives rise to the setting up of expectations for the "semantic category" of a given word-position based on the words encountered in its neighbourhood. Thus, a word like 'meeting' can raise our expectations of immediate mention of a "location-phrase" ('Pizza Piazza', or 'Burger Queen').

Thematic role and case-frame information can also serve to provide an important source of disambiguation data. These structures clearly provide a level of information that can allow us to decide between alternatives based on our prediction of the roles (agent, goal, *etc.*) that must necessarily be present to complete a sentence or even a clause. It can be seen that thematic roles and case frames straddle the form-content "non-divide". Manaster-Ramer's idea of lexical grammars seems to me to be a good illustration of this fact. He talks about ideas of partial templates for verbs, such as <*thing* **eat** *thing*> for 'eat' and <*thing* **put** *thing locative*> for 'put'.

Part-of-speech (POS) information can also provide a lot of information that helps disambiguate between candidates. [Rohini&Charlotte] have already explored this idea to much success. Syntactic categories, or POS's, (especially, content categories) can be very effectively interpreted as high level conceptual categories. For example, nouns correspond to the intuitive notion of "thing"s, while verbs correspond to "action"s. Thus, the syntax level can itself be seen to be more than "mere" form.

Constituent context (which has also been investigated by [Rohini&Charlotte]) is not likely to provide as much disambiguation information as any of the other levels in the hierarchy, mainly because we are pushing the idea of syntactic** context beyond applicable limits. This difficulty seems to stem from the fact that English tolerates a very high degree of flexibility in the constituent order, even though it is very intolerant of word-order violations within a constituent. This is further demonstrated by the distituent parsing techniques developed by Magerman and Marcus. However, there may yet be much promise in combining distituent parsing techniques with the use of context in disambiguation.

4.3 A unified framework for the above

I hope to demonstrate that a corpus-based approach (an approach that essentially implements unsupervised learning of transition statistics at different levels of the context hierarchy) will be very effective in improving recognition results. I also hope to demonstrate that this approach is not just of practical significance, but also provides a good theoretical model of human perceptual processing, albeit to differing degrees of accuracy at the different levels of the hierarchy. Others have already looked at the issue of integrating multiple models [Isotani 94, Rosenfeld 94], and it will be helpful to build on their ideas. But these attempts have been restricted to combining information from two or three sources only, and it is not clear (*prima facie*) that these techniques can be extended to handle all levels of the context hierarchy together.

In developing the unified framework, we do have to keep in mind the issues involved in combining (*ie*, integrating) modules *vs.* cascading them. Seidenberg *et al* [74] (3) have demonstrated that both types of contextual effects are present in certain tasks. Following Connine [90] (2), we can conclude that the human language-processing system has both modular and interactive subprocesses active within it. Hence, the ideal approach is probably a combination of the two methods. Our system will be

**Note that topic/theme-based association is not syntactic context, and is hence likely to be useful despite the fact that it demands even less spatial rigidity than constituent context — indeed, the very idea of imposing spatial constraints on topic strikes one as inherently contradictory.

modular at the topmost level, but the feedback loop from the language module to the recognition module provides the systems with most of the capabilities of interactive systems.

In this context, there has been much recent interest [Erickson 88, Kapur 92] in the application of entropy optimisation principles to the effective integration of information from multiple knowledge sources. Jaynes' MaxEnt method [57] has been described as "a universal method of finding a 'best' positive distribution constrained by incomplete data" [Skilling 88, p. 173]. The multiple knowledge sources act as constraints on the "random variable", and MaxEnt provides a method of minimising the bias in the distribution.

We need this unified framework so that contextual information from the various levels of the hierarchy can be integrated to make meaningful decisions when updating the trellis. Developing such a formalism will require qualitative and quantitative analysis of the kinds of information flow (through MIS and other influence measures). This effort will constitute a bulk of the theoretical work in this research.

5 Results so far

At this preliminary stage of implementing our ideas, we* have looked only at re-rankings and deletions, among the several different kinds of changes that the language module can make to the word-candidate lists that the recogniser outputs for the different word-positions.

Further, although we plan to use context at different levels of the context hierarchy in the overall scheme, we have so far confined ourselves to context analysis for collocations of fixed structure — in other words, we have not yet considered collocations that involve arbitrary separation among component words, and those that involve phrasal verbs (*eg*, 'give up'). Clearly, phrasal verbs present special challenges because of the need for morphological analysis, and variable (and usually unpredictable) inter-component distance (since the preposition is really a non-bound particle in such constructions). Examples like 'Portia gave up smoking.' and 'Desdemona considered giving a part of the project up.' illustrate these phenomena. [Aside: the latter example is a good illustration of the point that a unified framework is needed to perform disambiguation well and efficiently. There is nothing to prevent Desdemona from giving the extra-credit section of the project up (or the second half of the extra-credit section from the mid-term project, for that matter). Without a combined methodology, it may be very difficult to handle such examples of arbitrary separation.]

The corpus that we have used for these partial results is the Wall Street Journal (WSJ) corpus. However, in the long run, we plan to use an email corpus that is being built up here in CEDAR, for our training. The motivation behind an email corpus is that the style of English in an email† is a good approximation of the casual, unedited nature of the English that is used in speech and in online writing. Most of the existing corpora, including the WSJ corpus, represent formal English that has been edited and altered to conform to "correct" English. On the other hand, "e-language" exhibits many features that edited English rarely does. Such features include sentence fragments, typos, ungrammaticality, abbreviations, flexibility in punctuation styles, *etc.* Further, e-language provides a mixture of discourse and narrative structure of the kind found in speech and other spontaneous natural language sources.

*Here, and in the rest of the writeup, 'we' defaults to 'Rohini and me'.

†OK, piece of email

The section of the WSJ corpus that was used for training contains more than 102K sentences, totalling slightly over two million words. By using simple concordancing techniques, we identified some 30K word-pairs of significant correlation. These had occurred approximately 260K times in the corpus. This gives us roughly 2.6 "collocations" per sentence, or alternately, one "collocation-word" in approximately every four words.

The tables below provide a quick summary of preliminary results, and demonstrate very effectively that the unorthodox definition of MIS, which i dub "MIS^{++}", is much better at utilising context properly. The statistics are shown for a 1205-word test set that was chosen from a portion of the WSJ that has been set aside for this purpose (and not trained on, obviously).

The first table shows the actual performance of the collocational analysis in terms of the number of hits in the top n choices, for $n \in \{1,3,5,10\}$. The '||word-list||' column gives the average size of the word-lists for the word-positions, and thus measures the performance of the deletion subroutine. This is also the average height of the trellis, and so a measure of the "area" of the trellis, and of the amount of computation required by any post-processing modules.

	n = 1	= 3	= 5	= 10	‖word-list‖	topchoice%
From HWR	678	764	806	847	11.57	56.2
MIS	905	993	1012	1035	10.61	75.1
MIS^{++}	945	1004	1017	1039	10.69	78.4

The effect of context

The following tables look at MIS and MIS^{++} in more detail, and compare their performances. $f(xy)$ and $P(xy)$ respectively refer to the frequency (in the corpus) of the pair xy and the estimated probability of this co-occurrence.

MIS(xy) $= \log(\frac{P(xy)}{P(x)P(y)})$

	success	partial	error	total
reranking	424	131	157	712
deletion	17	207	4	228

MIS^{++}(xy) $= \log(\frac{f^2(xy)}{f(x)f(y)})$

	success	partial	error	total
reranking	538	168	112	818
deletion	17	205	4	226

Comparing the MIS'ses

A success in re-ranking is an operation where the truth moved to the top of the trellis, as a result of the re-ranking, while an error is a re-ranking operation that demotes the truth. A partial success in re-ranking is an operation which promoted the truth, but did not manage to make it the top choice for the word-position.

A success in deleting is a deletion where all word-candidates except the truth were deleted successfully, while an error is a deletion that deletes the truth itself. A partial success in deleting is an operation which deleted some word-candidates while preserving the truth, but did not succeeding in eliminating all competitors.

The above statistics demonstrate that non-standard measures of influence bear much promise. Measures like MIS^{++} which give more weightage to the collocation frequency need to be examined more thoroughly.

6 What needs to be done

It should be clear from the above that there is enormous potential for these ideas, and that the proper use of the techniques can yield significant benefits in terms of recogniser performance. As we have seen before, contextual information exists at multiple levels of the context hierarchy. But it is not very clear as to what is the best formalisation of this information. What is needed is a theoretically sound definition of contextual influence. The criteria for a generalised information score that we identified in §3 may prove useful here.

The thrust of my work will be towards developing a unified framework for these various levels, and in building a sound information-theoretic base from which to utilise the contextual information in the context hierarchy. This will involve the application of entropy optimisation principles to integrate the contextual information from the various sources. At the implementation level, however, we plan to concentrate on the lexical and lexico-semantic levels of the hierarchy. Both of these levels come under the collocational umbrella, in that the primer and the primed are fixed entries in the lexicon, and are not themselves syntactic (or other) categories.

To tackle both types of collocations uniformly within the language module, we need to incorporate more flexibility into the definition of what a collocation is. In other words, we should perform some For example, simple pre-processing of the input sequence, can yield the "root form" of the input sequence, by stripping simple suffixes such as number and tense markers. Similarly, we will also need to handle the different degrees of spatial rigidity that the different levels of the hierarchy demand. This would "normalise" the input, and facilitate the use of template-matching techniques in the context window.

Proposing words that did not exist in the word choices of the recogniser is a tougher problem. When the input is handwriting, we need to look at "visual neighbours" of words in the choices, and perform collocational analysis on the word-pairs thus formed. A naïve application of this idea would result in an explosion of word-pair possibilities, since there are multiple word-choices for any given word-position. The principal difficulty is in curtailing this explosive growth of possibilities, by choosing a sufficiently small neighbourhood that nevertheless contains the truth with high probability. In other words, efficient implementation of word proposing and feedback to the recogniser will involve sophisticated techniques of lexicon organisation. Rohini has suggested the concept of a visually similar neighbourhood (VSN), which would constitute a "probable cohort set" for a given word. It is well-known that short words cause maximum confusion in the recognition phase. By Zipf's "law", short words are also the most frequent words in natural language. Therefore, pre-compiling the VSN for short, frequent words is likely to facilitate the lexicon search process.

7 Loose ends

In this section we will briefly skim through some issues that range in importance from "implementation details" to "hopes on the horizon".

Different levels, different techniques: At different levels of the hierarchy, we will need to resort to different tricks and techniques to improve the performance of the system. As a simple example, the confidence with which we apply word-transition statistics at the "fixed collocation" level will differ from our use of tag-transition statistics. Similarly, handling phrasal verbs will require morhpological pre-processing before template-matching can be done.

Effects of domains, lexicons *etc*: Although we will assume that the input does not have any artificial "domain" imposed on it for computational purposes, it may be useful to exploit domain-specific features in tuning techniques. It would also be useful to identify domain-independent techniques that have general applicability.

Unlimited lexicons: I have made several oblique references like "any word in the English language" in the preceding discussion. The issue that is being swept onto the carpet here is that of managing a very comprehensive lexicon (say, a lexicon of a million or more words!), and the immediate issue of handling words not present in the lexicon. Ehrlich [92] has been working on understanding words from context, and it is likely that some of that work may find application in augmenting the unlimited lexicon. Implementing the learning lexicon will also require the proper handling of "rare events" (*cf.* [Dunning 93]). It would also be necessary for the system to perform morphological analysis on "new" words, so as to be able to handle previously unencountered forms of words that it already knows.

Numbers and symbols: Some researchers [Oostdijk 91, Towell 92, Briscoe 93] have evinced interest in the possibility of combining the advantages that symbolic systems and connectionist/statistical techniques offer. The need for this seems more and more pressing as we move down in the hierarchy. I would therefore like to develop a conceptual framework for emergent symbols within the essentially connectionist setup of our system, in line with the ideas proposed by Smolensky [88] and others. I propose that the frequencies and the probabilities that the system learns from the corpus are essentially the sub-symbols that Smolensky talks about.

Improper word segmentation: Providing the word recogniser with feedback about improper word segmentation, is very tough, since it requires "understanding" of the input at some level. Part of the difficulty here is the issue of deciding whether an improper segmentation took place or whether it is a new (*ie*, unknown) word that we are looking at. It is likely that POS-transition statistics can help us make such decisions.

8 Conclusion

Here we have seen a viable formalism for the proper utilisation of linguistic context in natural language sources. This contextual information, if used meaningfully, can contribute to significant improvement in the performance of an automatic recognition system. Several theoeretical issues still need to be resolved, however. Several types of recogniser improvement have been identified, and preliminary results indicate that all of these can be implemented within a unified framework of "lexical influence".

References

[1] Gerry TM Altmann. *Cognitive models of speech processing: psycholinguistic and computational perspectives*. MIT Press, 1990.

[2] Cynthia Connine. Effects of sentence context and lexical knowledge in speech processing. *in* Altmann 1990 (1), pages 281–294, 1990.

[3] MS Seidenberg, GS Waters, M Sanders, and P Langer. Pre- and post-lexical loci of contextual effects on word perception. *Memory and Cognition v. 12*, pages 315–328, 1984.

moreferences:

1. Rohini (proposal, charlotte etc)
2. Bill's book (especially stuff about semantics and reference)
3. my stuff on smolensky, langacker and cantwellsmith
4. Aleksander, Igor & Morton, Helen (1993), *Neurons and symbols: The stuff that mind is made of*, Chapman&Hall.
5. Brown, Peter F; Della Pietra, Stephen A; Della Pietra, Vincent J; Lai, Jennifer C & Mercer, Robert L (1992a), "An estimate of an upper bound for the entropy of English", *Computational Linguistics*, **18:1**, pages?
6. Brown, Peter F; Della Pietra, Stephen A; Desouza, Peter V; Lai, Jennifer C & Mercer, Robert L (1992b), "Class-based n-gram models of natural language", *Computational Linguistics*, **18:4**, 467-479.
7. Cover, Thomas M & Thomas, Joy A (1991), *Elements of information theory*, Wiley.
8. Doyle, Arthur Conan (1892), "The Boscombe Valley Mystery", *The adventures of Sherlock Holmes*, George Newnes Ltd.
9. Dunning, Ted (1993), "Accurate methods for the statistics of surprise and coincidence", *Computational Linguistics*, **19:1**, 61-74.
10. Ehrlich, Karen & Rapaport, William J (1992), "Automatic acquisition of word meanings from natural-language contexts", *Technical Report 92-03*, Center for Cognitive Science, SUNY, Buffalo.
11. Erickson, JG & Smith CR (ed.) (1988), *Maximum Entropy and Bayesian Methods in Science and Engineering* (v. 1&2), Kluwer Academic Publishers.
12. Garside, Roger; Leech, Geoffrey & Sampson, Geoffrey (1987), *The computational analysis of English*, Longman.
13. Geschwind, Norman (1979), "Specializations of the human brain", *Scientific American* **241-3**, 158-168.

14. Hofstadter, Douglas R (1985), *Metamagical themas: questing for the essence of mind and pattern*, Basic Books.

15. Isotani, Ryosuke & Matsunaga, Shoichi 1994, "Speech recognition using a stochastic language model integrating local and global constraints", *Proceedings of the ARPA workshop on human language technology*, March 1994, 87-92.

16. Jaynes, ET (1957), "Information Theory and Statistical Mechanics", *Physical Reviews*, **106**:620-630.

17. Kapur, JN & Kesavan, HK (1992), *Entropy Optimization Principles with Applications*, Academic Press.

18. Kripasundar V (1993), "Impressions of The Correspondence Continuum", term-paper for CS 718 (Spring '93).

19. Marslen-Wilson (1990), "Activation, Competition, and Frequency in Lexical Access", speechbook: 148-172.

20. Samuel, Arthur G, "Using perceptual-restoration effects to explore the architecture of perception", speechbook: 295-314.

21. Rosenfeld, Ronald 1994, "A hybrid approach to adaptive statistical language modeling", *Proceedings of the ARPA workshop on human language technology*, March 1994, 76-81.

22. Sampson, Geoffrey (1992), "Probabilistic parsing", in *Directions in corpus linguistics*, Mouton de Gruyter.

23. Sapir, Edward (1921), *Language: an introduction to the study of speech*, Harcourt.

24. Skilling, J (1988), "The axioms of maximum entropy", in *Maximum Entropy and Bayesian Methods in Science and Engineering*, Erickson, JG & Smith CR (ed.), 173-187.

25. Smolensky, Paul (1988), "On the proper treatment of connectionism", *Behavioral and Brain Sciences*, **11**:1-74.

26. Towell, Geoffrey G & Shavlik, Jude W (1992), "Using symbolic learning to improve knowledge-based neural networks", AAAI-92, 177-182.

27. Warren, RM (1970), "Perceptual restoration of missing speech sounds", *Science* **167**, 392-393.

28. Woods, WA (1983), "Language processing for speech understanding", *Computer Speech Processing*; (reprinted in *Readings in Speech Recognition*, Morgan-Kaufmann Publishers, 519-533).

29. Zernik, Uri (ed.) (1991), *Lexical Acquisition: Exploiting On-line Resources to Build a Lexicon*, Lawrence Erlbaum.

30. Zwitserlood, "The locus of the effects of sentential-semantic context in spoken word processing", *Cognition* 32:25-64.