

# Analogical Modeling: An Update

**David Eddington** and **Deryle Lonsdale**  
Department of Linguistics & English Language  
Brigham Young University  
Provo, UT, USA 84602  
{[eddington](mailto:eddington@byu.edu), [lonz](mailto:lonz@byu.edu)}@byu.edu

## Abstract

Analogical modeling is a supervised exemplar-based approach that has been widely applied to predict linguistic behavior. The paradigm has been well documented in the linguistics and cognition literature, but is less well known to the machine learning community. This paper sets out some of the basics of the approach, including a simplified example of the fundamental algorithm's operation. It then surveys some of the recent analogical modeling language applications, and sketches how the computational system has been enhanced lately to offer users increased flexibility and processing power. Some comparisons and contrasts are drawn between analogical modeling and other language modeling and machine learning approaches. The paper concludes with a discussion of ongoing issues that still confront developers and users of the analogical modeling framework.

## 1 Introduction

Analogical modeling (AM) is an exemplar-based modeling approach designed to predict linguistic behavior on the basis of stored memory tokens (Skousen, 1989; Skousen, 1992; Skousen, 1995; Skousen, 1998). It is founded on the premise that previous linguistic experience is stored in the mental lexicon. When the need arises to determine some linguistic behavior (pronunciation, morphological relationship, word, etc.), the lexicon itself is

accessed. A search is conducted for the stored exemplars that are most similar to the one whose behavior is being predicted. The behavior of highly similar stored entities generally predicts the behavior of the one in question, although less similar ones have a small chance of applying as well.

AM has been implemented in a succession of computer implementations that allow users to specify an exemplar base and then to test unseen instances of language behavior against this accumulated store to predict (and quantify) the relevant and possible outcome(s).

A comprehensive examination of the system is beyond the scope of this paper<sup>1</sup>. Instead, an attempt will be made to situate AM within the space of exemplar-based modeling systems and, to a certain extent, within the larger context of machine learning systems and cognitive modeling systems in general.

This paper begins by giving an informal overview of the algorithm used to perform AM. To accomplish this it traces in schematic form an early but typical application of the AM approach: Spanish gender guessing. The subsequent section discusses AM with respect to other types of modeling systems, and gives an overview of the types of linguistic phenomena modeled by AM. The paper then discusses several items of recent and ongoing work with the AM program. The conclusion mentions possible areas for future improvement and investigation.

---

<sup>1</sup>For further information including a bibliography and software downloads see the project webpage at: <http://humanities.byu.edu/am>

Subcontexts	Members of the Subcontext	Pointers	# of Disagreements
pan	none	none	0
$\bar{p}an$	<i>plan</i> M	<i>plan</i> M > <i>plan</i> M	0
$p\bar{a}n$	none	none	0
$pa\bar{n}$	<i>paz</i> F, <i>par</i> M	<i>paz</i> F > <i>paz</i> F <i>par</i> M > <i>par</i> M * <i>paz</i> F > <i>par</i> M * <i>par</i> M > <i>paz</i> F	2
$\bar{p}a\bar{n}$	<i>cal</i> F	<i>cal</i> F > <i>cal</i> F	0
$p\bar{a}\bar{n}$	<i>tren</i> M, <i>ron</i> M	<i>tren</i> M > <i>tren</i> M <i>ron</i> M > <i>ron</i> M <i>tren</i> M > <i>ron</i> M <i>ron</i> M > <i>tren</i> M	0
$p\bar{a}\bar{n}$	none	none	0
$\bar{p}a\bar{n}$	<i>rey</i> M	<i>rey</i> M > <i>rey</i> M	0

Table 1: Subcontexts and their disagreements.

## 2 The AM algorithm

Perhaps the best way to understand the AM algorithm is with a concrete illustration. Predictions are always made in terms of specific exemplars, therefore, for the purposes of the example, the following seven monosyllabic Spanish nouns and their corresponding gender will be considered<sup>2</sup>:

*paz* F  
*plan* M  
*tren* M  
*cal* F  
*ron* M  
*par* M  
*rey* M

The task will be to predict the gender of *pan* 'bread' on the basis of these seven items as the set of exemplars. The three variables used are the phoneme or phoneme cluster of the onset, nucleus, and rhyme. First, all possible exemplar items are assigned to a series of subcontexts which are defined in terms of the given context *pan*. This yields eight subcontexts (Table 1). A bar over a variable indicates that any value of the variable except the barred value is permitted.

By assigning members to subcontexts it is possible to determine disagreements. Disagreements

<sup>2</sup>This example is a simplification; see (Eddington, 2002b) for a complete AM treatment of this phenomenon.

are marked with asterisks in Table 1. A disagreement occurs when words that are equally similar to the given context, exhibit different behaviors, in this case, different genders. The number of disagreements is determined by pairing all members of a subcontext with every other member, including itself, by means of unidirectional pointers, and counting the number of times the members of the pair have different behaviors. In this example, the only subcontext containing any disagreement is  $pa\bar{n}$ .

Subcontexts are then arranged into more comprehensive groups called supracontexts as shown in Table 2, where a hyphen indicates a wildcard.

In the subcontextual analysis a tally of the all of the subcontextual disagreements is made. The supracontextual analysis consists of analyzing all of the words that appear in a given supracontext, and again tallying disagreements (Table 3).

In the supracontextual analysis, words that have more than one variable in common with *pan* appear in more than one supracontext. This is AM's way of allowing the gender of words which are more similar to *pan* to influence the gender assignment of *pan* to a greater extent.

The purpose of AM's algorithm is to determine which members of the exemplar base are most likely to affect the gender assignment of *pan*, and also to calculate the extent of analogical influence exerted.

Supracontext	Subcontexts in Supracontext	# Subcontextual Disagreements
p a n	pan	0
p a –	pan, *pañ	2
p – n	pan, p̄an	0
– a n	pan, p̄an	0
p – –	pan, p̄an, *pañ, p̄añ	2
– a –	pan, p̄an, *pañ, p̄añ	2
– – n	pan, p̄an, p̄an, p̄an	0
– – –	pan, p̄an, p̄an, *pañ, p̄añ, p̄an, p̄añ, p̄añ	2

Table 2: Subcontextual analysis.

This is accomplished by calculating heterogeneity. Heterogeneity is determined by comparing the number of disagreements in the supracontextual and subcontextual analyses. If there are more disagreements in the supracontextual analysis, the supracontext is heterogenous, and its members are eliminated from consideration as possible analogs. If the number of disagreements does not increase, the supracontext is homogenous.

Words belonging to homogenous supracontexts comprise the analogical set (Table 4). In the example under consideration, disagreements increase in the supracontexts – – –, and – a –. Therefore, their members are eliminated from consideration. *Cal*, and *rey* appear exclusively in these heterogenous supracontexts. As a result, they do not form part of the analogical set. The word *plan* is also a member of both – – –, and – a –, however, it is also a member of the homogenous supracontexts – a n, and – – n, so it will still be available to influence *pan*.

It should not be surprising that *rey* would be eliminated through heterogeneity; it has no phonemes in common with *pan*. However, consider the words *ron* and *cal*. Both share only one feature with *pan*, yet heterogeneity eliminates only *cal*, and not *ron*. This is due to the fact that *ron* appears in the supracontext – – n, and all of the members of that supracontext are masculine, therefore, there is no disagreement. *Cal* F, on the other hand, competes with masculine words in all of the supracontexts in which it appears.

The analogical set contains all of the exemplar items that can possibly influence the gender assignment of *pan*. There are two methods for calculating

the influence of the analogical set on the behavior of the given context. One is to assign the most frequently occurring behavior to the given context (selection by plurality). Of the 18 pointers in the set, 14 point to masculine. Therefore, selection by plurality would assign masculine gender to *pan*. The other method (random selection) involves randomly selecting a pointer, and assigning the behavior of the word indicated by the pointer to the given context. In this case, the probability of masculine gender assignment would be 77.78% (14/18).

### 3 Situating AM

The AM theoretical construct has been implemented computationally through several generations of application programs. Pascal code for the first version was listed as an appendix in (Skousen, 1989). A subsequent version was implemented in C. In recent years the system was reimplemented as a Perl module; extensive details on how to use the current version are available in (Skousen et al., 2002).

The system assumes *a priori* labeled input instances, and thus only works in supervised learning mode. Each exemplar must be represented as a fixed-length vector comprised of features that may or may not eventually be relevant to the phenomenon being modeled. The features themselves may be of fixed or variable length, as long as they are appropriately delimited. An outcome is specified for each exemplar instance vector. Vector features are purely symbolic and nominal, not admitting continuous values. When the use of continuous-valued features is needed (e.g. integers or real numbers),

Supracontexts	Words in Supracontext	Pointers	# of Subcontextual Disagreements
p a n	none	none	0
p a -	<i>par M, paz F</i>	<i>paz F &gt; paz F</i> <i>par M &gt; par M</i> <i>*paz F &gt; par M</i> <i>*par M &gt; paz F</i>	2
p - n	none	none	0
- a n	<i>plan M</i>	<i>plan M &gt; plan M</i>	0
p - -	<i>par M, paz F</i>	<i>paz F &gt; paz F</i> <i>par M &gt; par M</i> <i>*paz F &gt; par M</i> <i>*par M &gt; paz F</i>	2
- a -	<i>plan M, cal F, par M, paz F</i>	<i>plan M &gt; plan M</i> <i>*plan M &gt; cal F</i> <i>plan M &gt; par M</i> <i>*plan M &gt; paz F</i> <i>cal F &gt; cal F</i> <i>*cal F &gt; plan M</i> <i>*cal F &gt; par M</i> <i>cal F &gt; paz F</i> <i>*paz F &gt; plan M</i> <i>*paz F &gt; par M</i> (not all shown; all the rest involve agreement)	6
- - n	<i>plan M, tren M, ron M</i>	(not shown)	0
- - -	<i>plan M, tren M, ron M, cal F, paz F, par M, rey M</i>	(not shown)	12

Table 3: Supracontextual analysis.

the features must be symbolically quantized within pre-specified ranges in order to be meaningful to the system. More details on how to choose, quantize, and encode exemplar vectors is given in (Skousen et al., 2002).

As the system loads up the feature vectors, it builds the contextual lattice as sketched in the previous section. Once the lattice has been constructed, the system can process incoming test items to determine their outcome(s). Each test item consists of a feature vector containing the same number of features as the exemplar vectors do. Each test item also has the appropriate outcome specified, if known, or else some placeholder symbol. Each test item is matched against the lattice, its analogical set com-

puted, and outcomes reported. AM, unlike some other machine learning approaches, can report all possible outcomes, and associates each with a probability rating or confidence rating reported as a percentage.

### 3.1 Sample implementations

AM has been applied to a number of different linguistic phenomena in several different languages. Although a comprehensive enumeration is not possible here, a representative sampling of applications gives an appreciation for the breadth and depth of coverage possible.

The earliest documented AM results with language data—and indeed the preponderance of work since—focus primarily on phonology, morphology,

Homogenous Supracontext	Words in Supracontext	Pointers	# of Pointers to Masculine	# of Pointers to Feminine
p a –	<i>par</i> M, <i>paz</i> F	<i>par</i> M > <i>par</i> M <i>par</i> M > <i>paz</i> F <i>paz</i> F > <i>par</i> M <i>paz</i> F > <i>paz</i> F	2	2
– a n	<i>plan</i> M	<i>plan</i> M > <i>plan</i> M	1	0
p – –	<i>par</i> M, <i>paz</i> F	<i>par</i> M > <i>par</i> M <i>par</i> M > <i>paz</i> F <i>paz</i> F > <i>par</i> M <i>paz</i> F > <i>paz</i> F	2	2
– – n	<i>plan</i> M, <i>tren</i> M, <i>ron</i> M	(not shown)	9	0

Table 4: Analogical set non-empty homogenous supracontexts.

and lexical selection. Early applications (Skousen, 1989) included predicting Finnish past tense forms, describing allomorphic variation (e.g. the *a/an* distinction in English), and predicting lexical selection based on sociolinguistic variables (e.g. Arabic forms for “friend”). Subsequent efforts (Skousen et al., 2002) addressed such phenomena as German pluralization, Dutch compound linkers, Spanish gender on nouns, stress in Dutch, and Turkish consonantal alternation.

Most of these implementations involve feature vectors that specify relatively low-level linguistic features such as sound segments, phonetic or orthographic environments, syllable structure, and word boundaries. Often the information for constructing these features comes from widely-used annotated language resources such as lexicons or corpora. For example, several English, German, and Dutch applications have used features partially derived the CELEX lexical database (Baayen et al., 1993); others have used WordNet (for English) (Fellbaum, 1998), TELL (for Turkish)<sup>3</sup>, and LEX-ESP (for Spanish) (Gallés et al., 2000).

With the advent of large-scale corpora, both annotated and unannotated, increasing amounts of information are available to language modelers. Recent AM efforts have benefited from speech and text corpora. The popular TIMIT speech database provided exemplars for recent work on flapping in English (Eddington, 2007). In probably the first AM

application to diachronic language description, a recent paper (Chapman and Skousen, 2005) used the Helsinki corpus<sup>4</sup> for modeling the historical development of English negative prefixes.

Increasingly, modelers are using their own experimentally obtained results to compare against more standard resources. The Arabic lexical selection example involved only features encoding novel researcher-collected situational observations. A recent paper on the comparative form of English adjectives (Elzinga, 2006) matches human judgments against occurrences from the internet collected via the Google search engine.

The choice of which features to include in the exemplar and test instances is an important issue for the AM user, as it is for all exemplar models. Since features have to be enumerated and encoded *a priori*, at least some of the features chosen should reflect salient properties of the phenomenon in question. This does not presuppose a prescient omniscience; it is not the case that all and only the relevant features must appear in the vectors. Any features that end up not contributing to the outcome simply constitute irrelevant overhead. Of course, in an exponential system such overhead should be avoided wherever possible.

### 3.2 Comparison to other approaches

The algorithm sketched in Section 2 is what sets the AM paradigm apart from other machine learn-

<sup>3</sup>See <http://socrates.berkeley.edu/TELLhome.html>.

<sup>4</sup>See <http://icame.uib.no>.

ing and language modeling approaches. The system does not employ sub-symbolic representations or continuous-valued features and thus differs in fundamental ways from connectionist systems. The computation of analogical (dis)similarity resembles in many ways nearest-neighbor approaches, though in some cases a “gang” of similar items, none of which is the nearest neighbor, may conspire to overcome the nearest neighbor and impose its own outcome. These gang effects are difficult to identify, though some have been discussed in the literature, and represent a unique aspect of AM results. When gang effects are not present, AM processing generally obtains results comparable to memory-based learning (e.g. in TiMBL). The latter, though, allows for continuous-valued features.

Another aspect where AM and some approaches differ is that AM considers each feature in any given vector of equal weight; only across the exemplar base do features emerge as more or less relevant. In this respect AM does not perform feature weighting in a separate processing stage across the exemplar base before testing new items. Similarly, no dependencies can be specified explicitly across features.

Extensive comparisons have been made between AM and other machine learning approaches including memory-based learning (Daelemans et al., 2004), version spaces, neural networks, and Optimality Theory; see (Skousen et al., 2002) for complete details.

The concept of an analogical set with quantifiable support is also unique to AM, and its relevance is being explored. For example, recent work (Chandler, 2005) has shown how AM results can map onto response times in certain types of psycholinguistic experiments. Eddington (2007 forthcoming) demonstrates how AM is able to mirror the results of psycholinguistic production studies that are presented elsewhere as evidence against analogy in processing regular morphology.

In fact, another notable feature of AM is its ability to model human language errors and variants that also result from analogy. Skousen, in early AM work, discusses how “leakage” (i.e. erroneous AM guesses) often tends to follow observable usage errors (Skousen, 1989).

Eddington’s (2002b) simulations of gender assignment in Spanish yielded similar errors rates to

those of children of different ages, and in a simulation of Spanish diminutives (Eddington, 2002a), 52% of the incorrect predictions made by AM were actually found to be attested alternative diminutive forms.

A comparison of corpus data for Russian verbs of motion—and subsequent modeling in AM combining lexical, morphological, and semantic variables—showed a tendency for the system’s errors to coincide with those of language learners (Dodge and Lonsdale, 2006).

Certainly further comparison of the behavior observed in AM modeling tasks versus human performance constitutes a promising future direction to pursue.

When considering the place of AM in the machine learning realm, one area deserves considerable more effort. The natural language processing (NLP) field has profited greatly from machine learning, in part due to the increased availability of large-scale language resources. Though early work in AM illustrated the promise of the approach to typical NLP tasks (Jones, 1996), the paradigm has been underrepresented in recent NLP shared tasks and competitive evaluations addressing such problems as word-sense disambiguation, part-of-speech tagging, shallow parsing, semantic role labeling, and phoneme classification. Thanks to recent developments in the implementation as discussed in the next section, we expect this situation to improve in the future.

## 4 Recent progress in AM

The AM system has undergone substantial development since the last documented release of the system (Skousen et al., 2002), and this has permitted applications in areas that would have been previously impossible. In this section we survey some of the recent AM developments and consequent research they have enabled.

### 4.1 Algorithmic refinement

As explained above, the AM system’s basic data structure is a fully connected lattice representing (dis)agreements among data items. The original algorithm for handling the lattice involved a straightforward but exhaustive traversal and was hence very costly. The most current version of the system in-

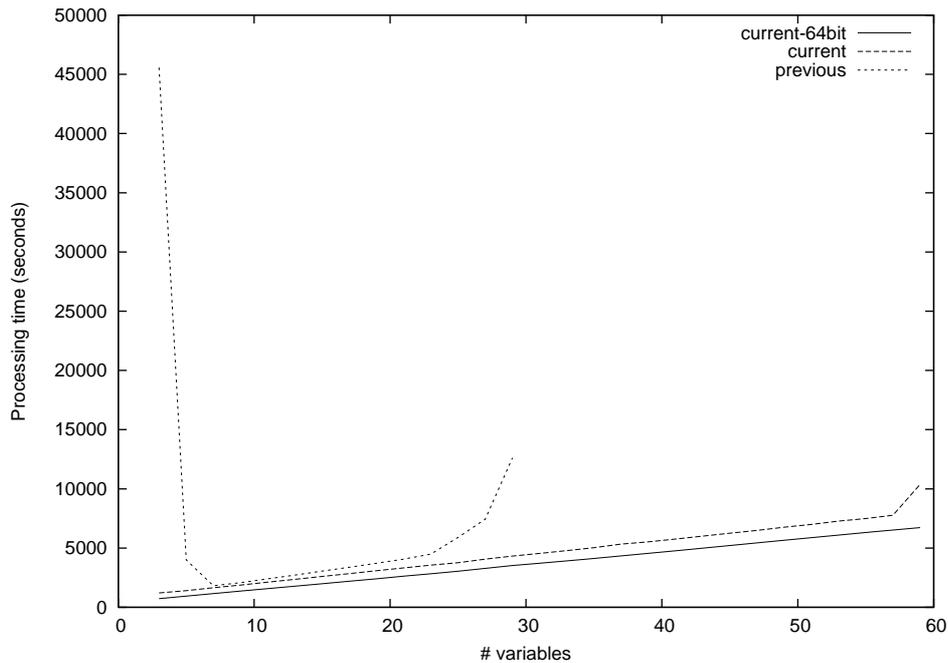


Figure 1: Performance statistics for previous and current versions of AM

volves a completely new, but equivalent, method of extracting information from the lattice.

The crucial insight is to focus on heterogeneity instead of homogeneity. For each supracontext in the lattice, heterogeneity can be locally determined by looking for plurality of outcomes and plurality of subcontexts. When heterogeneity is detected, all more general supracontexts can be automatically discarded from that point on.

Implementing this insight has resulted in a substantial speed-up in processing time, and the ability to handle a greater number of features, exemplars, and test items. Figure 1 illustrates the run time for different versions of the AM program on a typical task (English part-of-speech tagging) involving 50,000 exemplars, 1,000 test items, and 58 outcomes. It shows that the previous version of the system grew exponentially as the number of variables grew (though the maximum number of variables possible was fixed at 29); the algorithm also performed particularly poorly for a vector size of 3 features. This contrasts with the current version of the system, which runs largely linearly (on a dual-processor Sun Sparc Blade 25000) until the current maximum vector size of 59 features, at which point the exponential curve begins. A new 64-bit version

of the system continues to behave linearly even at 59 features.

The rewrite of the algorithm also incorporated more decomposition in the handling of the lattice, in theory laying the groundwork for a truly parallel implementation. To date, though, no concurrent parallel implementation has been developed.

Even with this improvement, exponentiality still poses a problem, if only theoretically. Two independent efforts have been pursued recently to address the processing bottleneck issue. One involves using Monte Carlo sampling to estimate the results that traditional AM processing provides (Johnsen and Johansson, 2005). The other method is to compute the dissimilarities via a quantum algorithm (Skousen, 2005). To discuss either would be beyond the scope of this paper.

## 4.2 Modularization

The newest version of the AM implementation has been repackaged as a Perl module. This was done primarily to improve flexibility in using the system. Its deployment within Perl supports better cross-architecture functionality. The provision of a number of processing hooks allows for various user-programmable add-on procedures for customizing

the input, matching, and output stages of processing.

One example of the use of customized output involves the integration of AM with a cognitive modeling system for natural language called NL-Soar. In this work, which involves parsing, the cognitive agent must occasionally resolve prepositional-phrase attachment ambiguities (Lonsdale and Manookin, 2004; Manookin and Lonsdale, 2004); it does so by invoking AM to resolve the ambiguity using a PP-attachment exemplar set previously mined by others<sup>5</sup>. This integration would have been quite difficult with prior versions of the AM system.

Some modeling problems are sequential in nature, meaning that a particular test's outcome might be dependent on that of a preceding test's outcome: dynamically chaining results from one test to another is necessary in such cases and now possible with AM. Similarly, collecting results of individual tests, integrating them, and reporting them in another format is also now possible in the current AM system.

For example, recent work on Romanizing Farsi names (Lonsdale, 2006) profited from the improved AM system input/output hooks. Input to the system consisted of an exemplar base of several thousand Farsi names in Arabic script form<sup>6</sup>, and their Romanized counterparts; often there is more than one way to Romanize such names. The system processes each name, for which all possible Romanizations are generated from successive analogical sets. In addition, a client-specified metric measuring the confidence of each possible outcome can also be generated via the output hook and further post-processing. Figure 2 illustrates sample results. The first input, *bhnAm*, represents the Arabic-script name; two possible Romanized forms (*Bahnaam* and *Behnaam*) are posited by the system, with respective confidence scores of 436.044 and 402.424. Note that the third input name has ellipsis dots (“...”) representing the Divine Name, which is pronounced but not written; even in this case the system can compute the proper Romanized substitutions via analogy.

```
input=bhnAm
436.044 Bahnaam
402.424 Behnaam

input=HfiZi
450.000 Hafiz
450.000 Hafeez

input=lit ...bhbhAni
91.11 Ayat Allah Bahbahani
91.11 Ayat Allah Bahbahaani
88.89 Ayat Allah Bahbahane
88.89 Ayat Allah Bahbahaanee
88.89 Aayat Allah Bahbahaani
88.89 Aayat Allah Bahbahaanee
88.89 Aayat Allah Bahbahaani
88.89 Ayat Allah Bahbahaanee
86.67 Aayat Allah Bahbahaanee
86.67 Aayat Allah Bahbahane
86.67 Aayat Allah Bahbahaanee

input=jmCidi
399.414 Jamsheed
396.716 Jamshid
394.940 Jamshaid
384.322 Jamasheed
```

Figure 2: Sample summary report from Farsi name Romanization process showing inputs and outputs with alternatives scored per custom algorithm.

### 4.3 Lattice visualization

For many users of the AM system, insight into the operation of the AM algorithm has proven difficult, with the system operating as a black box. A new visualization tool allows users to load up their exemplars, enter test items, and step through the lattice incrementally to view the effects of matches, homogeneity/heterogeneity, and gangs.

Figure 3 shows that testing the item *pan* against the exemplars in the supracontext *p ā n̄* results in non-determinism (i.e. the supracontext is heterogeneous), specified by *par M* and *paz F*.

### 4.4 Architectural improvements

Loading a large number of exemplar items is time-consuming. This is keenly felt in sequential implementations, where the outcome of one test might depend on the results of a previous one (cf. the Farsi example described above).

To alleviate this problem, a new implementation provides a client/server infrastructure. The exemplar data can be loaded on a server machine once, and the server stands by for test item queries (over a socket)

<sup>5</sup>See <ftp://ftp.cis.upenn.edu/pub/adwait/PPattachData>.

<sup>6</sup>For convenience here we do not use actual Arabic-script characters for the input but rather transliterations.

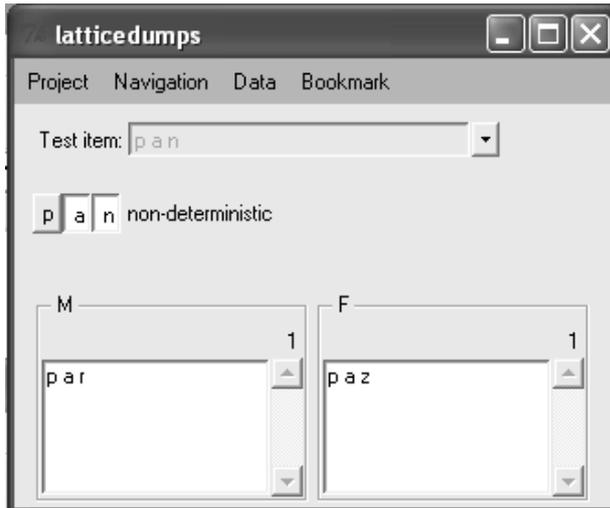


Figure 3: A screenshot of the AM lattice visualizer.

from client machines.

#### 4.5 User tools

Recent use of the system by users has resulted in two other tools worthy of note.

Development of exemplar and test item feature vectors can sometime result in ill-formed vectors which are difficult to locate. A new Perl script allows users to check their feature files and receive a report of which vectors do not comply to formatting specifications.

For users unable to configure the basic system parameters within the Perl code, use of the system can be awkward. A basic Perl/TK graphical user interface has been developed to provide users top-level control to the overall system, allowing them to select pertinent parameters and input files. In this way the users will not need to modify any Perl code.

### 5 Conclusions

The AM algorithm has been applied to a number of different language phenomena where small numbers of variables and long processing times were not an issue. For psycholinguists, the most important factors in a model are often its ability to mirror the responses and errors of language speakers, along with the psychological plausibility of the model. In NLP, on the other hand, processing speed and the ability to incorporate large numbers of variables are often more highly favored. Given the advances in the cur-

rent implementations of the AM algorithm we trust that AM will become a more attractive tool for a wider range of linguistic applications.

### Acknowledgements

We would like to thank Theron Stanford for his programming support work.

### References

- R. H. Baayen, R. Piepenbrock, and H. van Rijn. 1993. The CELEX Lexical Database [CD-ROM]. Linguistic Data Consortium, University of Pennsylvania; Philadelphia, PA.
- Steve Chandler. 2005. Computer simulations of RT in analogical modeling. Paper presented at the Thirty-second LACUS Forum at Dartmouth College, August 4, 2005.
- Don Chapman and Royal Skousen. 2005. The negative prefix in english and analogical modeling of language. *Journal of English Language and Linguistics*, 9:333–357.
- Walter Daelemans, Jakub Zavrle, Ko van der Sloot, and Antal van den Bosch. 2004. TiMBL: Tilburg memory-based learner, version 5.1 reference guide. Induction of linguistic knowledge technical report 04-02, ILK Research Group, Tilburg University.
- Inna Danielyan Dodge and Deryle Lonsdale. 2006. Modeling Russian verbs of motion: an analogical account. In *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, pages 1239–1244. Lawrence Erlbaum.
- David Eddington. 2002a. Spanish diminutive formation without rules or constraints. *Linguistics*, 40:395–419.
- David Eddington. 2002b. Spanish gender assignment in an analogical framework. *Journal of Quantitative Linguistics*, 9:49–75.
- David Eddington. 2007. Flapping and other variants of /t/ in american english: Allophonic distribution without constraints, rules, or abstractions. *Cognitive Linguistics*, 18:23–46.
- David Eddington. 2007, forthcoming. Spanish verbal inflection: A single- or dual-route system? *Linguistics*.
- Dirk Elzinga. 2006. English adjective comparison and analogy. *Lingua*, 116:757–770.
- Christiane Fellbaum, editor. 1998. *WordNet: An electronic lexical database*. MIT Press, Cambridge, MA.

- Núria Sebastián Gallés, Fernando Cuetos Vega, Manuel F. Carreiras Valina, and M. Antònia Martí Antonin. 2000. *Lexesp: Léxico informatizado del español*. Technical report, Barcelona: Biblioteca Universitària.
- Lars Johnsen and Christer Johansson. 2005. Efficient modeling of analogy. In *Proceedings of the CiCLing Conference*, volume 3406 of *Lecture Notes in Computer Science*, pages 694–703. Springer Berlin/Heidelberg.
- Daniel Jones. 1996. *Analogical Natural Language Processing*. Studies in Computational Linguistics. University College of London Press Limited, London, England.
- Deryle Lonsdale and Michael Manookin. 2004. Combining learning approaches for incremental on-line parsing. In *Proceedings of the Sixth International Conference on Cognitive Modeling*, pages 160–165. Lawrence Erlbaum.
- Deryle Lonsdale. 2006. Exploring syllables, Romanization, and analogy in names. In *Proceedings of the Sixth Annual Family History Technology Workshop*. BYU Computer Science Department.
- Michael Manookin and Deryle Lonsdale. 2004. Resolving automatic prepositional phrase attachments by non-statistical means. In *LACUS Forum XXX: Language, Thought, and Reality*, pages 291–300. The Linguistic Association of Canada and the United States.
- Royal Skousen, Deryle Lonsdale, and Dilworth B. Parkinson, editors. 2002. *Analogical Modeling: An exemplar-based approach to language*. John Benjamins, Amsterdam/Philadelphia.
- Royal Skousen. 1989. *Analogical Modeling of Language*. Kluwer Academic Publishers, Dordrecht, Netherlands.
- Royal Skousen. 1992. *Analogy and Structure*. Kluwer Academic Publishers, Dordrecht, Netherlands.
- Royal Skousen. 1995. Analogy: A non-rule alternative to neural networks. *Rivista di linguistica*, 7:213–231.
- Royal Skousen. 1998. Natural statistics in language modeling. *Journal of Quantitative Linguistics*, 5:246–255.
- Royal Skousen. 2005. Quantum analogical modeling: A general quantum computing algorithm for predicting language behavior. Posted at arxiv.org.