

An investigation into feature construction to assist word sense disambiguation

Lucia Specia · Ashwin Srinivasan · Sachindra Joshi ·
Ganesh Ramakrishnan · Maria das Graças Volpe Nunes

Received: 19 November 2008 / Revised: 2 April 2009 / Accepted: 17 April 2009 /
Published online: 12 June 2009
Springer Science+Business Media, LLC 2009

Abstract Identifying the correct sense of a word in context is crucial for many tasks in natural language processing (machine translation is an example). State-of-the art methods for Word Sense Disambiguation (WSD) build models using hand-crafted features that usually capturing shallow linguistic information. Complex background knowledge, such as semantic relationships, are typically either not used, or used in specialised manner, due to the limitations of the feature-based modelling techniques used. On the other hand, empirical results from the use of Inductive Logic Programming (ILP) systems have repeatedly shown that they can use diverse sources of background knowledge when constructing models. In this paper, we investigate whether this ability of ILP systems could be used to improve the predictive accuracy of models for WSD. Specifically, we examine the use of a general-purpose ILP system as a method to construct a set of features using semantic, syntactic and lexical information. This feature-set is then used by a common modelling technique in the field

Editors: Filip Zelezny and Nada Lavrac.

A.S. is also an Adjunct Professor at the Department of Computer Science and Engineering, University of New South Wales; and a Visiting Professor at the Computing Laboratory, University of Oxford.

L. Specia

Xerox Research Centre Europe, 6 Chemin de Maupertuis, Meylan 38240, France
e-mail: lucia.specia@xrce.xerox.com

A. Srinivasan · S. Joshi (✉)

IBM India Research Laboratory, 4-C, Institutional Area, Vasant Kunj, New Delhi 110 070, India
e-mail: jsachind@in.ibm.com

A. Srinivasan

e-mail: ashwin.srinivasan@in.ibm.com

G. Ramakrishnan

Department of Computer Science & Engineering, IIT Bombay, Poway, Mumbai, India
e-mail: ganesh@cse.iitb.ac.in

M. das Graças Volpe Nunes

ICMC—Universidade de São Paulo, Trabalhador São-Carlense, 400, São Carlos 13560-970, Brazil
e-mail: gracan@icmc.usp.br

(a support vector machine) to construct a classifier for predicting the sense of a word. In our investigation we examine one-shot and incremental approaches to feature-set construction applied to monolingual and bilingual WSD tasks. The monolingual tasks use 32 verbs and 85 verbs and nouns (in English) from the SENSEVAL-3 and SemEval-2007 benchmarks; while the bilingual WSD task consists of 7 highly ambiguous verbs in translating from English to Portuguese. The results are encouraging: the ILP-assisted models show substantial improvements over those that simply use shallow features. In addition, incremental feature-set construction appears to identify smaller and better sets of features. Taken together, the results suggest that the use of ILP with diverse sources of background knowledge provide a way for making substantial progress in the field of WSD.

Keywords ILP · Word sense disambiguation · Feature construction · Randomised search

1 Introduction

Word Sense Disambiguation (WSD) aims to identify the correct sense of an ambiguous word in a sentence. Usually described as an “intermediate task” (Wilks and Stevenson 1998), it is necessary in many natural language tasks like machine translation, information retrieval, question answering, and so on. That it is extremely difficult to completely solve WSD is a long-standing view (Bar-Hillel 1960) and accuracies with state-of-the-art methods are substantially lower than in other areas of text understanding. Part-of-speech tagging accuracies, for example, are now over 95%; in contrast, the best WSD results for fine-grained sense distinctions are still below 80%.

The principal approach adopted for the automatic construction of WSD models is a “shallow” one. In this, sample data consisting of sentences with the ambiguous words and their correct sense are represented using features capturing some limited context around the ambiguous words in each sentence. For example, features may denote words on either side of an ambiguous word and the part-of-speech tags of those words. Sample data represented in this manner are then used by a statistical model constructor to build a general predictive model for disambiguating words. Results from the literature on benchmark data like those provided under the various SENSEVAL competitions¹ suggest that support vector machines (SVMs) yield models with one of the highest accuracies. As this competition shows, some improvements have been achieved in the accuracy of predictions, despite of the use of limited, shallow information sources. On the other hand, it is generally thought that significant progress in automatic WSD would require a “deep” approach in which access to substantial body of linguistic and world knowledge could assist in resolving ambiguities. However, the incorporation of large amounts of domain knowledge has been hampered by the following: (a) access to such information in electronic form suitable for constructing models; and (b) modelling techniques capable of utilising diverse sources of domain knowledge. The first of these difficulties is now greatly alleviated by the availability in electronic form of very large semantic lexicons like WordNet (Miller et al. 1990), dictionaries, parsers, grammars and so on. In addition, there are now very large amounts of “shallow” data in the form of electronic text corpora from which statistical information can be readily extracted. Using these diverse sources of information is, however, beyond the capabilities of existing general-purpose statistical methods that have been used for WSD. Arguably, Inductive

¹See: <http://www.senseval.org>.

Logic Programming (ILP) systems provide a more general-purpose framework for dealing with such data: there are explicit provisions made for the inclusion of background knowledge of any form; the representation language is powerful enough to capture the contextual relationships that arise; and modelling is not restricted to being of a particular form (for example, classification only).

In this paper, we investigate the possibility that an ILP system equipped with deep and shallow knowledge sources can be used to improve the predictivity of models for WSD. Specifically, our interest is in ILP as a technique for constructing new features, and our hypothesis is the following:

ILP-features hypothesis. Feature-construction by ILP provides a general method of introducing background knowledge for improving the predictive accuracy of WSD models.

Empirical evidence is sought by investigating the use of ILP-based feature-construction on 124 datasets from three substantial WSD tasks (thus testing general applicability). When constructing features, an ILP system is provided with background knowledge drawn from 10 different sources (thus testing the ability to use a diverse range of lexical, syntactic and semantic information). In all cases, we find that when hand-crafted features are augmented by ILP-constructed features (we call these “ILP-assisted models”), the predictive power of models *does* improve significantly.

The rest of the paper is organised as follows. In Sect. 2 we present some related work on WSD. A specification for ILP implementations that construct features for use in ILP-assisted models is in Sect. 3.1. We then describe an implementation that meets these specifications in Sect. 3.2. The empirical evaluation comprising our investigation is described in Sect. 4. This includes materials (Sect. 4.1) and methods (Sect. 4.2). A summary of results and a related discussion are in Sect. 5. Section 6 concludes the paper. The paper is accompanied by an appendix that contains detailed tabulations of the experimental results.

2 Models for word sense disambiguation

The earliest computer-executable models for WSD were manually constructed, capturing specific aspects of human disambiguation expertise in symbolic structures like semantic networks (Quillian 1961) and semantic frames (Hirst 1987; McRoy 1992). Early reports also exist of sub-symbolic neural networks (Cottrell 1989). Most of these techniques appear to have suffered from the important difficulty in manual acquisition of expert knowledge, resulting in their application being limited to very small subsets of the languages.

The development of machine readable resources like lexical databases, dictionaries and thesauri has provided a turning point in WSD, enabling the development of techniques that used linguistic and extra-linguistic information extracted automatically from these resources (Agirre and Rigau 1996; Lesk 1986; Wilks and Stevenson 1997). While the resources provided ready access to large bodies of knowledge, the actual disambiguation models continued to be manually codified. This changed with the use of statistical and machine-learning techniques for constructing models. The characteristic of these methods is the use of a corpus of examples of disambiguation to construct automatically models for disambiguation (Pedersen 2002; Schutze 1998; Yarowsky 1995). The most common of these “corpus-based” techniques employ statistical methods that construct models based on features representing frequencies estimated from a corpus. For example, these may be the frequencies of some words on either side of the ambiguous word. While techniques using

such “shallow” features that refer to the local context of the ambiguous word have yielded better models than the previous ones, the accuracies obtained are still low, and significant improvements do not appear to be forthcoming.

More sophisticated corpus-based approaches such as (Stevenson and Wilks 2001) try to incorporate deeper knowledge using machine readable resources. These are special-purpose methods aimed at specific tasks and it is not clear how they could be scaled-up for use across a wide range of WSD tasks. ILP provides a general-purpose approach that can be tailored to a variety of NLP tasks by the incorporation of appropriate background knowledge. The use of ILP to build WSD models was first investigated by Specia (Specia 2006a; Specia et al. 2007a). Preliminary work on the use of ILP to construct features to be used by standard propositional machine learning algorithms like SVMs can be found in (Specia et al. 2007b, 2007c). The work here extends these substantially by exploring alternate ways of using ILP to build features and in terms of experimental results.

3 Feature construction with ILP

We motivate the task of feature-construction using the “trains” problem, originally proposed by Ryzhard Michalski. The task is to construct a model that can discriminate between east-bound and westbound trains, using properties of their carriages, and the loads carried (see Fig. 1).

We will assume that the trains can be adequately described by background predicates that will become evident shortly. Further, let us assume that the 10 trains shown in the figure are denoted $t1, t2, \dots, t10$ and that their classifications are encoded as a set of logical statements of the form: $Class(t1, eastbound), Class(t2, eastbound), \dots, Class(t10, westbound)$. As is quite normal in the use of ILP for feature-construction, we will assume features to be Boolean valued, and obtained from some clause identified by the ILP program. For example, Fig. 2 shows five such features, found by an ILP engine, and the corresponding tabular representation of the 10 examples in Fig. 1.

Suppose that these 6 features are the only features that can be constructed by the ILP engine, and further, that it is our task to find the best subset of these that can result in the best model. Clearly, if we simply evaluated models obtained with each of the 63 subsets of the




Fig. 1 The trains problem. Trains are classified either as “eastbound” or “westbound”. They have open or closed carriages of different shapes, lengths, and so on. The carriages contain loads of different shapes and numbers. The task is to construct a model that, given the description of a train, can predict whether it will be eastbound or westbound

	<i>f1</i>	<i>f2</i>	<i>f3</i>	<i>f4</i>	<i>f5</i>	<i>f6</i>	Class
$f1 = \begin{cases} 1 & \text{iff has short and closed car} \\ 0 & \text{otherwise} \end{cases}$	1	1	0	0	0	0	East
$f2 = \begin{cases} 1 & \text{iff has closed car and contains triangle load} \\ 0 & \text{otherwise} \end{cases}$	1	0	0	0	0	0	East
$f3 = \begin{cases} 1 & \text{iff has double-wall car and contains triangle load} \\ 0 & \text{otherwise} \end{cases}$	1	1	0	0	0	0	East
$f4 = \begin{cases} 1 & \text{iff has car with jagged roof} \\ 0 & \text{otherwise} \end{cases}$	1	0	1	0	0	0	East
$f5 = \begin{cases} 1 & \text{iff has long car and contains rectangular load} \\ 0 & \text{otherwise} \end{cases}$	0	0	0	1	0	1	West
$f6 = \begin{cases} 1 & \text{iff has long car and can with jagged roof} \\ 0 & \text{otherwise} \end{cases}$	0	0	0	1	1	1	West
	0	0	0	0	0	0	West

Fig. 2 Some Boolean features for the trains problem, and a corresponding tabular representation of the trains in Fig. 1




Fig. 3 Identifying the best subset of features for a model-construction algorithm *A*. The X-axis enumerates the different subsets of features that can be constructed by an ILP engine (\mathcal{F} denotes the set of all possible features that can be constructed by the engine). The Y-axis shows the probability that an instance drawn randomly using some pre-specified distribution will be correctly classified by a model constructed by *A*, given the corresponding subset on the X-axis. We wish to identify the subset k^* that yields the highest probability, without actually constructing all the features in \mathcal{F}

set $\{f1, f2, f3, f4, f5, f6\}$ and return the subset that returned the best model, we would be done. Now, let us consider a more practical situation. Suppose the features that can be constructed by an ILP system are not in the 10s, but in the 1000s or even 100s of 1000s. This would make it intractable to construct models with all possible subsets of features. Further, suppose that constructing each feature is not straightforward, computationally speaking, making it impractical to even use the ILP engine to construct all the possible features in the first place. Are we able to nevertheless determine the subset that would yield the best model (which we will now interpret to mean the model with the highest classification accuracy)? The conceptual problem to be addressed is shown in Fig. 3.

Readers will recognise this as somewhat similar to the problem addressed by a randomised procedure for distribution-estimation like Gibb’s sampling. There, if the \mathcal{F} features are given (or at least can be enumerated), then the sampling procedure attempts to converge

on the best-performing subset without examining the entire space of $2^{|\mathcal{F}|}$ elements. Clearly, if we are unable to generate all possible features in \mathcal{F} beforehand, we are not in a position to use these methods. We first propose a minimal specification to be satisfied by any ILP-based feature constructor.

3.1 Specification

Functionally, ILP can be largely characterised by two classes of programs. The first, predictive ILP, is concerned with constructing models (sets of rules; or first-order variants of classification or regression trees) for discriminating accurately amongst two sets of examples (“positive” and “negative”). The partial specifications provided by (Muggleton 1994) have formed the basis for deriving programs in this class, and are shown in Fig. 4 (we refer the reader to (Nienhuys–Cheng and de Wolf 1997) for definitions of the logical terms used).

The second category of ILP systems, descriptive ILP, is concerned with identifying relationships that hold amongst the background knowledge and examples, without a view of discrimination. The partial specifications for programs in this class are based on the description in (Muggleton and De Raedt 1994), and are shown in Fig. 5.

The idea of exploiting a feature-based model constructor that uses first-order features can be traced back at least to the LINUS program (Lavrač et al. 1990). More recently, the task of identifying good features using a first-order logic representation has been the province of programs developed under the umbrella of “propositionalization” (see Kramer et al. 2001 for a review). Programs in this class are not easily characterised as either predictive or descriptive ILP. Conceptually, solutions involve two steps: (1) a feature-construction step that identifies (within computational reason) all the features that are consistent with the constraints provided by the background knowledge. This is characteristic of a descriptive ILP program; and (2) a feature-selection step that retains some of the features based on their utility in classifying the examples. This is characteristic of a predictive ILP program. To this

- B is background knowledge consisting of a finite set of clauses $= \{C_1, C_2, \dots\}$
- E is a finite set of examples $= E^+ \cup E^-$ where:
 - *Positive Examples.* $E^+ = \{p_1, p_2, \dots\}$ is a non-empty set of definite clauses;
 - *Negative Examples.* $E^- = \{\bar{n}_1, \bar{n}_2, \dots\}$ is a set of Horn clauses (this may be empty)
- H , the output of the algorithm given B and E is acceptable if the following conditions are met:
 - *Prior Satisfiability.* $B \cup E^- \not\models \square$
 - *Posterior Satisfiability.* $B \cup H \cup E^- \not\models \square$;
 - *Prior Necessity.* $B \not\models E^+$
 - *Posterior Sufficiency.* $B \cup H \models e_1 \wedge e_2 \wedge \dots$

Fig. 4 A partial specification for a predictive ILP system from (Muggleton 1994)

- B is background knowledge consisting of a finite set of clauses $= \{C_1, C_2, \dots\}$
- E is a finite set of examples (this may be empty)
- H , the output of the algorithm given B and E is acceptable if the following condition is met:
 - *Posterior Sufficiency.* $B \cup H \cup E \not\models \square$

Fig. 5 A partial specification for a descriptive ILP system based on (Muggleton and De Raedt 1994)

extent, we would like specifications for feature construction to reflect a combination of these two dominant flavours of ILP programs. In particular, we will assume the following:

1. Examples are taken to be some subset of the binary relation $\mathcal{X} \times \mathcal{Y}$, where \mathcal{X} denotes the set of individuals and \mathcal{Y} some finite set of classes. We therefore do not make any special distinction between positive and negative examples.
2. Given examples of the form described, an ILP system identifies one or more definite clauses of the form $h_i : \text{Class}(x, c) \leftarrow Cp_i(x)$, where x is a variable and c is some class in the set of classes \mathcal{Y} . Here, adopting terminology from (Ratnaparkhi 1996), $Cp_i : \mathcal{X} \mapsto \{0, 1\}$ is a “context predicate” that corresponds to a conjunction of literals that evaluates to *TRUE* (1) or *FALSE* (0) for any particular individual. We will require that Cp_i contains at least one literal: in logical terms, we therefore require the h_i to be definite clauses with at least two literals (let us call them “mixed definite clauses”, to denote the requirement for exactly one positive literal and at least one negative literal). To ensure that the clauses found are not trivial, we require each h_i to entail at least one example in E .
3. Features are conjunctions of literals. Given a clause $h_i : \text{Class}(x, c) \leftarrow Cp_i(x)$, found by the ILP system, we construct a one-to-one mapping to a feature f_i as follows: $f_i(x) = 1$ iff $Cp_i(x) = 1$ (and 0 otherwise).² An example from a WSD task studied in this paper is shown in Fig. 7.
4. We would like features identified to be relevant. The notion of relevance of a feature in a set $F = \{f_1, \dots, f_k\}$ can be captured by a probabilistic statement on feature-based representation of the examples. Let individuals be elements of the set $F_1 \times \dots \times F_k$ (where F_i is the set of values of the i th feature: for us all the F_i are $\{0, 1\}$); and examples be elements of $F_1 \times \dots \times F_k \times \mathcal{Y}$. Then, as in (John et al. 1994), (with some abuse of notation) we can distinguish between the *strong relevance* of a feature f_i in F , if $Pr(Y|F) \neq Pr(Y|F - \{f_i\})$; and its *weak relevance* if there is some $F' \subset F$ such that $f_i \in F'$ and f_i is strongly relevant in F' . That is, $Pr(Y|F') \neq Pr(Y|F' - \{f_i\})$. As a minimal requirement, we would like an ILP-based feature-constructor to identify features that are at least weakly relevant.

With these assumptions, a partial specification for an ILP-based feature constructor is shown in Fig. 6.

3.2 Implementation

Here we describe an implementation that can be tailored to meet the specification just described. Recall that all the specification, with our subsequent modifications, requires is the following: (1) The implementation returns a set in which each feature is represented as a conjunction of literals; (2) Each feature is constructed from a mixed definite clause that entails at least one example, given the background knowledge; and (3) Each feature is weakly relevant, given the background knowledge and examples. Clearly a very large number of features can be constructed that satisfy these requirements. Consequently, we introduce the following general constraints for restricting the number of features constructed:

Support. For each clause $h_i : \text{Class}(x, c) \leftarrow Cp_i(x)$ we compute $P(h_i) = \{x : e \in E \wedge e = \text{Class}(x, c) \wedge B \cup \{h_i\} \cup \{-e\} \models \square\}$. $P(h_i)$ is the set of individuals correctly classified by h_i , and we denote their number by $tp(h_i) = |P(h_i)|$ (more correctly of course, we should

²In (Ratnaparkhi 1996) the f_i are denoted by $f_{i,c}$. The notation here is more in line with the machine learning literature.

- B is background knowledge consisting of a finite subset of Horn clauses = $\{C_1, C_2, \dots\}$ drawn from some set \mathcal{B} .
- E is a finite set of examples of the form $Class(i, c)$ where i is a member of a set of individuals \mathcal{X} and c is a member of a set of classes \mathcal{Y} .
- \mathcal{H} is the set of mixed definite clauses of the form $h_i : Class(x, c) \leftarrow Cp_i(x)$, constructible with predicates, functions and constants in $B \cup E$; \mathcal{F} the set of conjunctions of literals that can be constructed with predicates, functions and constants in $B \cup E$.
- $F = \{f_1, f_2, \dots, f_k\} \subseteq \mathcal{F}$, the output of the algorithm given B and E is a finite set of features constructed from a set of clauses $H = \{h_1, h_2, \dots\} \subseteq \mathcal{H}$ if the following condition is met:

Posterior Sufficiency.

- * For each $h_i \in H$, $B \cup \{h_i\} \models e_1 \vee e_2 \vee \dots$, where $\{e_1, e_2, \dots\} \subseteq E$
- * $f_i(x) = Cp_i(x)$
- * f_i is weakly relevant given B and E .

Fig. 6 A partial specification for an ILP system for feature construction

Clause:

$$h_1 : \quad Class(x, voltar) \leftarrow Has_expression(x, come_back, voltar) \wedge Has_pos(x, pcwr_4, nn)$$

Context Predicate:

$$Cp_1(x) : \quad Has_expression(x, come_back, voltar) \wedge Has_pos(x, pcwr_4, nn)$$

Feature:

$$f_1(x) = \begin{cases} 1 & Has_expression(x, come_back, voltar) \wedge Has_pos(x, pcwr_4, n) = 1 \\ 0 & \text{otherwise} \end{cases}$$

Fig. 7 Example of a boolean feature constructed from a clause for WSD. The clause shown here identifies the Portuguese sense of the English verb “to come”. The meanings of the predicate symbols *Has_expression* and *Has_pos* are explained in Sect. 4

include B and E in these functions). Adopting data-mining terminology, we call this the “support” of a clause. In our implementation, we reduce the number of acceptable clauses found by the ILP system (and hence the number of features) by placing some minimal requirement on the support of any clause.

Precision. For each class $c \in \mathcal{Y}$, we denote $N_c = \{\neg Class(x, c) : Class(x, c') \in E \wedge c \neq c'\}$. Then for each clause $h_i : Class(x, c) \leftarrow Cp_i(x)$ we compute $N(h_i) = \{x : e \in N_c \wedge e = \neg Class(x, c) \wedge B \cup \{h_i\} \cup \{e\} \models \square\}$. $N(h_i)$ is the set of individuals incorrectly classified by a h_i as class c , and we denote their number by $fp(h_i) = |N(h_i)|$. Once again, using terminology from the data-mining literature, we obtain the precision of a clause h_i as the value $tp(h_i)/(tp(h_i) + fp(h_i))$. In our implementation, we reduce the number acceptable clauses found by the ILP system (and hence the number of features) by requiring the precision of any clause to be greater than some pre-specified number.

We note that ensuring that the support of a clause is at least 1 meets the requirement that the clause entails at least one example. The constraint on precision can be used to ensure that features constructed are weakly relevant. This follows trivially, since each clause h for class c results in a feature f and the following contingency table for class values:

		Class		
		c	\bar{c}	
f	1	$tp(h)$	$fp(h)$	n_a
	0	$N_c - tp(h)$	$N_{\bar{c}} - fp(h)$	n_b
		N_c	$N_{\bar{c}}$	N

If the precision constraint ensures that the precision of a clause h for a class c is greater than the prior probability of class c (this probability is estimated from the examples E), then $precision(h) = tp(h)/n_a > N_c/N = Pr(Class = c)$. Now, since $Pr(Class = c|f = 1) = precision(h)$, it follows that $Pr(Class|f) \neq Pr(Class)$. That is, the corresponding feature f is weakly relevant.

We will call any clauses (and, with some overloading of terminology, the corresponding features) that satisfies the support and precision requirements as “acceptable”. We distinguish between two different techniques to obtain features:

1. One-shot feature-set construction, that identifies a set of acceptable features without any feedback from the model constructor (this is, in some sense, a “LINUS-inspired” approach, Lavrac et al. 1990). In keeping with terminology sometimes used in the ILP literature, we will call this approach “static feature-set construction”.
2. Incremental feature-set construction, that identifies a set of acceptable features after repeated iterations in which feature-construction is guided by feedback from the model constructor. In keeping with terminology sometimes used in the ILP literature, we will call this approach “dynamic feature-set construction”.

We present these techniques as instances of the implementation below.

3.3 Randomised feature-set construction

In (Zelezny et al. 2006) an implementation using a randomised local search was proposed to meet the specification of a predictive ILP system. It is possible to adopt the same approach in an implementation designed to meet the specifications for feature construction (see Fig. 8).

In Fig. 8 R and M bound the number of random restarts and local moves allowed. Existing techniques for feature construction can be re-cast as special cases of this procedure, with appropriate values assigned to R and M ; and definitions of a starting point (Step 3a) and local moves (Step 3(f)i). For example, one-shot methods that use an ILP engine to construct a large number of features independent of the model constructor can be seen as an instance of the randomised procedure with $R = 1$ and $M = 0$, with some technique for generating the “starting subset” in Step 3a. The feature-set constructor provided within the Aleph (Srinivasan 1999) system, shown in Fig. 9(a), is an example of a technique that generates features

1. bestfeatures:= {}
2. bestaccuracy:= 0.0
3. for i = 1 to R do begin
 - (a) currentfeatures:= randomly selected set of features
 - (b) currentmodel:= model constructed with currentfeatures
 - (c) accuracy:= estimated accuracy of currentmodel
 - (d) if accuracy > bestaccuracy then begin
 - i. bestfeatures:= currentfeatures
 - ii. bestaccuracy:= accuracy
 - (e) end
 - (f) for j = 1 to M do begin
 - i. nextfeatures:= best local move from currentfeatures
 - ii. nextmodel:= model constructed with nextfeatures
 - iii. accuracy:= estimated accuracy of nextmodel
 - iv. if accuracy > bestaccuracy then begin
 - A. bestfeatures:= nextfeatures
 - B. bestaccuracy:= accuracy
 - v. end
 - vi. currentfeatures:= nextfeatures
 - (g) end
4. end
5. return bestfeatures

Fig. 8 A basic randomised local search procedure, adapted to the task of feature construction

using a stratified sample of the examples (a simpler approach of generating k acceptable features using simple random sampling of the examples is in Fig. 9(b)).

We consider now the more general procedure shown. In this, R and M can take on any value from the set of natural numbers (including 0). In addition, the starting subset is assumed to be drawn using some distribution that need not be known; and a local move is one that either adds a single new feature to the existing subset of features, or drops a feature from the existing subset. We are now immediately confronted with two issues that make it impractical to use the procedure as shown. First, we have the same difficulty that prevented us from using an enumerative technique like a Gibb's sampler: generating the local neighbourhood requires us to obtain all possible single-feature additions. Second, for each local move, we need to construct a model, which can often be computationally expensive. We address each of these in turn.

Reducing local moves. We consider a modification of the search procedure in Fig. 8 that results in only examining a small sample of all the local moves possible before deciding on the next move in Step 3(f)i. Ideally, we are interested in obtaining a sample that, with high probability, contains the best local move possible. Assuming there are no ties, and that the number of possible local moves is very large, it would clearly be undesirable to select the sample using a uniform distribution over local moves. We propose instead a selection that uses the errors made by the model-constructor to obtain a sample of local moves. As a result, features in the local neighbourhood that are relevant to the errors are more likely to be selected. In some sense, this is somewhat reminiscent of boosting methods: here, instead of increasing the weights of examples incorrectly classified, the representation language is enriched in a way that is biased to classify these examples correctly on subsequent iterations (the possibility of using misclassified examples to guide a stochastic local search amongst an existing set of features was demonstrated in Paes et al. 2007).

Recall that at any point, a local move from a feature-subset F is obtained by either dropping an existing feature in F or adding a new feature to F . We are specifically concerned with the addition step, since in principle, all possible features that can be constructed by the

strat_features($B, E, \mathcal{Y}, k, s, p, \mathcal{L}$) : Given background knowledge B ; a set of examples E ; a set of class-labels \mathcal{Y} taken here to be the set of numbers $\{1, 2, \dots, c\}$; an upper-bound on the number of features allowed for each class k ($1 \leq k$); the minimal support for acceptable clauses s ($1 \leq s$); the minimal precision for acceptable clauses p ($0 < p \leq 1$); and a language \mathcal{L} specifying constraints on acceptable clauses that can be considered in the search for any one feature; return a set F of at most k acceptable features for each class.

1. for each $i \in \mathcal{Y}$ do
2. begin
 - (a) Let $H_i := \emptyset$
 - (b) Let E_i be all examples in E with class i
 - (c) while ($|H_i| \leq k$) and ($E_i \neq \emptyset$) do
 - (d) begin
 - i. Randomly select an example $e \in E_i$
 - ii. Let $j := k - |H_i|$
 - iii. Let H be a set of at most j mixed definite clauses found by an ILP engine such that $H \cap H_i = \emptyset$, and each clause $h \in H$ satisfies: (1) $h \in \mathcal{L}$; (2) $B \cup \{h\} \models e$; and (3) h is acceptable (that is, $\text{support}(h) \geq s$; and $\text{precision}(h) \geq p$).
 - iv. Let $H_i := H_i \cup H$
 - v. Let $E_i := E_i - \{e\}$
 - (e) end
3. end
4. Let F be the set of features obtained from the bodies of clauses in $H_1 \cup H_2 \dots H_c$
5. return F

(a) A procedure for constructing a stratified sample of acceptable features.

srs_features($B, E, k, s, p, \mathcal{L}$) : Given background knowledge B ; a set of examples E ; a set of class-labels \mathcal{Y} taken here to be the set of numbers $\{1, 2, \dots, c\}$; an upper-bound on the number of features allowed k ($1 \leq k$); the minimal support for acceptable clauses s ($1 \leq s$); the minimal precision for acceptable clauses p ($0 < p \leq 1$); and a language \mathcal{L} specifying constraints on acceptable clauses that can be considered in the search for any one feature; return a set F of at most k acceptable features.

1. Let $H = \emptyset$
2. Let E' be a random subset of k examples from E
3. while ($|H| \leq k$) and ($E' \neq \emptyset$) do
4. begin
 - (a) Randomly select an example $e \in E'$
 - (b) Let $j := k - |H|$
 - (c) Let H' be a set of at most j mixed definite clauses found by an ILP engine such that $H \cap H' = \emptyset$, and each clause $h \in H'$ satisfies: (1) $h \in \mathcal{L}$; (2) $B \cup \{h\} \models e$; and (3) h is acceptable (that is, $\text{support}(h) \geq s$; and $\text{precision}(h) \geq p$).
 - (d) Let $H := H \cup H'$
 - (e) Let $E' := E' - \{e\}$
5. end
6. Let F be the set of features obtained from the bodies of clauses in H
7. return F

(b) A procedure for constructing a simple random sample of acceptable features.

Fig. 9 Candidates for using an ILP engine to generate the “starting subset” in the randomised search for features. *support* and *precision* are as described in the text

ILP engine could be considered as candidates. We curtail this in the following ways. First, we restrict ourselves to samples of features that are related to examples misclassified by the model-constructor using the current set of features F (by “related” we mean those that are TRUE for at least one of the examples in error). Second, in our implementation, we restrict

ourselves to a single new feature for each such example (by “new”, we mean a feature not already in F).³

Reducing models constructed. The principal purpose of constructing and evaluating models in the local neighbourhood is to decide on the best next move to make. This will necessarily involve either an addition of a new feature to the existing set of features, or the deletion of an existing feature from the current set of features. That is, we are looking to find the best new feature to add, or the worst old feature to drop (given the other features in the set, of course). Correctly, we would form models with each old feature omitted in turn from the current set and each new feature added in turn to the current set. The best model would then determine the next move. Using a model-constructor that assigns weights to features allows us to adopt the following sub-optimal procedure instead. First, we find the feature with the lowest weight in the current model: this is taken to be the worst old feature. Next, we construct a single model with all features (old and new). Let us call this the “extended model”. The best new feature is taken to be the new feature with the highest weight in the extended model.

The procedure in Fig. 8 with these modifications is shown in Fig. 10. It is evident that the number of additional models constructed at any point in the search space is now reduced to just 3: the price we pay is that we are not guaranteed to obtain the same result as actually performing the individual additions and deletions of features.

It is the randomised local search procedure in Fig. 10, with one small difference, that we implement in this paper. The difference arises from the comparison of models: in the procedure shown, this is always done using estimated accuracies only. In our implementation, if estimated accuracies for a pair of models are identical, then the model using fewer features is preferred (that is, comparisons are done on the pair (A, F) where A is the estimated accuracy of the model and F is the number of features used in the model). One additional point worth clarifying concerns over-fitting the data. Some combination of the following mechanisms are clearly possible: (a) Limiting the amount of local search (using M); (b) Requiring acceptable features have some reasonable support; and (c) Using a model constructor that can perform some appropriate trade-off of fit-versus-complexity to avoid over-fitting.

We make the following observations about this implementation:

Termination. It is evident that, provided the ILP engine used construct features terminates, the procedure in Fig. 10 terminates, since both R and M are finite.

Correctness. All features constructed are obtained from acceptable clauses. That is, the support of every clause is at least s ($s > 1$), and, precision is at least p . If p is ensured to be greater than the prior probability of any class c in the examples E , then from the discussion earlier, the corresponding features are all weakly relevant.

Incompleteness. The procedure does not identify all weakly relevant features, since the local search procedure is incomplete.

Output. We assume that the starting subset in Step 3a contains features proportional to n (for example, with c classes, a starting subset using the Aleph feature-set constructor in Fig. 9 will produce a set containing no more than cn features). On any restart, a local move either adds a single feature to the existing set of features; or drops a feature from this set. Since this done no more than M times, the output contains more than $\alpha n \pm M$ features ($1 \leq \alpha < \infty$).

³In the implementation, we select this feature using its discriminatory power given the original set of examples.

randsearch_features($B, E, \mathcal{Y}, n, s, p\mathcal{L}, R, M$) : Given background knowledge B ; a set of examples E ; a set of class-labels \mathcal{Y} taken here to be the set of numbers $\{1, 2, \dots, c\}$; a sample size n ($1 \leq n$); the minimal support for acceptable clauses s ($1 \leq s$); the minimal precision for acceptable clauses p ($0 < p \leq 1$); and a language \mathcal{L} specifying constraints on acceptable clauses that can be considered in the search for any one feature; the number of random restarts R ($1 \leq R$); and the number of local moves M ($1 \leq M$), return a set of acceptable features F .

1. $F :=$
2. $\text{bestaccuracy} := 0.0$
3. for $i = 1$ to R do begin
 - (a) $\text{currentfeatures} :=$ sample of acceptable features given $B, E, \mathcal{Y}, n, s, p\mathcal{L}$
 - (b) $\text{currentmodel} :=$ model constructed with currentfeatures
 - (c) $\text{accuracy} :=$ estimated accuracy of currentmodel
 - (d) if $\text{accuracy} > \text{bestaccuracy}$ then begin
 - i. $F := \text{currentfeatures}$
 - ii. $\text{bestaccuracy} := \text{accuracy}$
 - (e) end
 - (f) for $j = 1$ to M do begin
 - i. $F_{\text{new}} :=$ sample of new acceptable features that are TRUE for errors made by currentmodel
 - ii. $\text{extendedmodel} :=$ model constructed using currentfeatures and F_{new}
 - iii. $f_{\text{worst}} :=$ feature in currentfeatures with lowest weight in currentmodel
 - iv. $f_{\text{best}} :=$ feature in F_{new} with highest weight in extendedmodel
 - v. $F^- :=$ set with feature subset obtained by dropping f_{worst} from currentfeatures
 - vi. $F^+ :=$ set with feature subset obtained by adding f_{best} to currentfeatures
 - vii. $\text{localmoves} := F^- \cup F^+$
 - viii. $\text{nextfeatures} :=$ best subset in localmoves
 - ix. $\text{nextmodel} :=$ model constructed with nextfeatures
 - x. $\text{accuracy} :=$ estimated accuracy of nextmodel
 - xi. if $\text{accuracy} > \text{bestaccuracy}$ then begin
 - A. $F := \text{nextfeatures}$
 - B. $\text{bestaccuracy} := \text{accuracy}$
 - xii. end
 - xiii. $\text{currentfeatures} := \text{nextfeatures}$
 - xiv. $\text{currentmodel} := \text{nextmodel}$
 - (g) end
4. end
5. return F

Fig. 10 The randomised local search procedure for feature construction, modified using theory-guided sampling of local moves and the use of feature-weights to reduce model construction. Acceptable features are obtained using the same technique as the static feature-set constructor: an example for which the feature is required to be *TRUE* is selected, and then an acceptable clause found

Based on the observations that: (a) the procedure in Fig. 10 terminates; and (b) with a value of p that can be determined from the examples, it returns a finite set of weakly relevant features, each constructed from clauses that entail at least 1 example, we can view it as an algorithm for feature construction.

4 Empirical evaluation

Our objective is to evaluate empirically if ILP-based features can assist in constructing models for WSD. Specifically, given different sources of background information, we intend to investigate the performance of the following kinds of models:

Baseline models. These are models for WSD that use features that are obtained manually from the different knowledge sources.

ILP-assisted models. These are models constructed by augmenting the hand-crafted features with those found by either the static or dynamic feature-set construction techniques, when provided with access to a richer representation of the same background information as the Baseline models. When static feature-set construction is used, we denote the result-

ing ILP-assisted model “ILP-S”; and when dynamic feature-set construction is used, we denote the corresponding model “ILP-D”.

We wish to investigate whether the ILP-assisted models have a significantly higher predictive accuracy than the baseline models. Statistical evidence in favour of the ILP-assisted models will, in turn, be taken as providing support for the “ILP-features hypothesis”, which we reiterate here:

ILP-features hypothesis. Feature-construction by ILP provides a general method of introducing background knowledge for improving the predictive accuracy of WSD models.

Statistically speaking, the null hypothesis for the investigation is there is no difference in the (average) predictive accuracies of Baseline and ILP-S or ILP-D models. Inability to reject this hypothesis will be taken as evidence against the usefulness of ILP as a feature-constructor for WSD.

4.1 Materials

4.1.1 Data

We experiment with datasets contained in three different tasks. Each collection is concerned with the disambiguation of different sets of ambiguous words, the so called “target words”, including nouns and verbs, and contain different numbers of classes (possible senses or translations) and class distributions. The three tasks are:

Monolingual-1. Data consist of the 32 verbs from the SENSEVAL-3 competition. SENSEVAL⁴ is a joint evaluation effort for WSD and related tasks. We use all the verbs of the English lexical sample task from the third edition of the competition. The number of examples for each verb varies from 40 to 398 (average of 186). The number of senses varies from 3 to 12 with an average of 7 senses. The average accuracy of the majority class is about 55%. We refer the reader to (Mihalcea et al. 2004) for more information about the SENSEVAL-3 data.

Monolingual-2. Data consist of 85 verbs and nouns of the English lexical sample task from the SemEval-2007 competition, the last edition of SENSEVAL. The number of examples varies from 24 to 3,061 (average of 272.12). The number of senses used in the training examples for a given word varies from 1 to 13 (average of 3.6). The average accuracy of the majority class is about 78%. More information about this data set can be found in (Pradhan et al. 2007).

Bilingual. Data consist of 7 highly frequent and ambiguous, mostly content-light, verbs: *come, get, give, go, look, make, take*. The sample corpus comprises around 200 English sentences for each verb extracted from several corpora, including the European Parliament, the Bible and fiction books, with the verb translation automatically annotated (Specia et al. 2005). The number of possible translations varies in the corpus from 5 to 17, with an average of 11 translations. The average accuracy of the majority class in the test data is about 54%.

Taken together, the tasks specify 124 independent datasets, with disambiguation required ranging from 3 to 17 different senses. The total number of examples is approximately 27,000.

⁴<http://www.senseval.org>.

4.1.2 Background knowledge

To achieve accurate disambiguation in both tasks is believed to require a variety of lexical, syntactic and semantic information. In what follows, we describe the background knowledge available for the tasks and illustrate it using the following sentence (assuming that we are attempting to determine the sense of ‘coming’ as target word):

“If there is such a thing as reincarnation, I would not mind coming back as a squirrel!”

Background knowledge is available in the following categories:

- B0.** *Baseline features.* Manually identified features using the information encoded in the predicates below, conveying the same information, but represented by means of attribute-value vectors.
- B1.** *Bag-of-words.* The 5 words to the right and left of the target word, extracted from the corpus and represented using definitions of the form *Has_bag(sentence, word)*. For example:

$$\begin{aligned} &Has_bag(snt1, mind). \\ &Has_bag(snt1, not). \dots \end{aligned}$$

- B2.** *Narrow context.* Lemmas of 5 content words to the right and left of the target word, extracted from the corpus, previously lemmatized by MINIPAR (Lin 1993). These are represented using definitions of the form *Has_narrow(sentence, wordposition, word)*. For example:

$$\begin{aligned} &Has_narrow(snt1, first_content_word_left, mind). \\ &Has_narrow(snt1, first_content_word_right, back). \dots \end{aligned}$$

These are not provided for the Bilingual task, since it was thought to be adequately covered by B5 below.

- B3.** *Part-of-speech tags.* Part-of-speech (POS) tags of 5 content words to the right and left of the target word, are obtained using MXPOST (Ratnaparkhi 1996) and represented using definitions of the form: *Has_pos(sentence, wordposition, pos)*. For example:

$$\begin{aligned} &Has_pos(snt1, first_content_word_left, nn). \\ &Has_pos(snt1, first_content_word_right, rb). \dots \end{aligned}$$

- B4.** *Subject-Object relations.* Subject and object syntactic relations with respect to the target word—in case it is a verb. If it is a noun, the representation includes the verb of which it is a subject or object, and the verb/noun it modifies. These were obtained from parsing sentences using MINIPAR and represented using definitions of the form *Has_rel(sentence, type, word)*. For example:

$$\begin{aligned} &Has_rel(snt1, subject, i). \\ &Has_rel(snt1, object, nil). \dots \end{aligned}$$

- B5.** *Word collocations.* 11 collocations with respect to the target word, extracted from the corpus: 1st preposition to the right, 1st and 2nd words to the left and right, 1st noun, 1st adjective, and 1st verb to the left and right. These are represented using definitions of the form *Has_collocation(sentence, collocation_type, collocation)*. For example:

$$\begin{aligned} &Has_collocation(snt1, first_word_right, back). \\ &Has_collocation(snt1, first_word_left, mind). \dots \end{aligned}$$

- B6.** *Verb restrictions.* Selectional restrictions of the verbs, when these are the target words, defined in terms of the semantic features of their arguments in the sentence, extracted using LDOCE (Procter 1978). If the restrictions imposed by the verb are not part of the description of its arguments, WordNet relations are used to check whether they can be satisfied by synonyms or hyperonyms of those arguments. A hierarchy of feature types is used to account for restrictions established by the verb that are more generic than the features describing its arguments in the sentence. These are represented by definitions of the form *Satisfy_restrictions(sentence, rest_subject, rest_object)*. For example:

Satisfy_restrictions(snt1, [human], nil).
Satisfy_restrictions(snt1, [animal, human], nil).

- B7.** *Dictionary definitions.* A relative count of the overlapping words in dictionary definitions of each of the possible senses of the target word (extracted from Parker and Stahel 1998, for the bilingual task, and from Procter 1978, for the monolingual tasks) and the words surrounding it in the sentence, to find the sense with the highest number of overlapping words. These are represented by facts of the form *Has_overlap(sentence, translation)*. For example:

Has_overlap(snt1, voltar).

- B8.** *Phrasal verbs.* Phrasal verbs involving the target verb, possibly occurring in a sentence, according to the list of phrasal verbs given by dictionaries and the context of the verb (5 surrounding words). These are represented by definitions of the form *Has_expression(sentence, verbal_expression)*. For example:

Has_expression(snt1, 'come back').

This is not provided for Monolingual task (1) since it does not consider senses of the verbs occurring in phrasal verbs.

- B9.** *Frequent bigrams.* Bigrams consisting of pairs of adjacent words in a sentence (without the target word) which occur more than 10 times in the corpus and are represented by definitions of the form *Has_bigram(sentence, word₁, word₂)*. For example:

Has_bigram(snt1, such, a).
Has_bigram(snt1, back, as). . . .

These are available for Monolingual task (2) only.

- B10.** *Frequently related words.* Related words consisting of pairs in the sentence that occur in the corpus more than 10 times related by verb-subject, verb-object, verb-modifier, subject-modifier, and object-modifier syntactic relations, without including the word to be disambiguated. There are represented by facts of the type *Has_related_pair(sentence, word₁, word₂)*. For example:

Has_related_pair(snt1, there, is). . . .

These are available for Monolingual task (2) only.

Of these definitions, B0 is intended for constructing the Baseline model. B1–B10 are intended for use by an ILP system. For each task we used a different subset of these knowledge sources, according to their availability. Further, although the ILP implementation we use is entirely capable of exploring intensional definitions of each of B1–B10, we represent definitions in an extensional form (that is, as a set of ground facts), since this representation is more efficient. This results in about 1.4 million facts.

Algorithms We distinguish here between three separate procedures: (1) The feature-construction algorithm in Fig. 10 for constructing a set of acceptable features; (2) The ILP engine concerned with identifying clauses required by the feature-set construction algorithms; and (3) The procedure for constructing models given a set of features and their values for sample instances, along with a class label associated with each instance. We have an implementation, in the Prolog language, of Fig. 10. For static-feature construction, the stratified sampling procedure provided within the Aleph system shown in Fig. 9(a) is used to provide the “starting subset” of features. In other cases, we commence with a simple random sample of features using the procedure in Fig. 9(b). For (2) we use a basic branch-and-bound search procedure provided within Aleph, that attempts to find all clauses that satisfy a given set of constraints and entail at least one example. The feature-based model constructor (3) is a linear SVM (the specific implementation used is the one provided in the WEKA toolbox called SMO.⁵). We will refer to (1) as “the feature constructor”; (2) as “the ILP learner” and (3) as “the feature-based learner.”

4.2 Method

Our method is straightforward:

For each verb in each task (that is, 32 words in Monolingual-1, 85 words in Monolingual-2 and 7 words in the Bilingual task):

1. Obtain the best model possible using the feature-based learner and the features in B0. Call this the Baseline model.
2. Construct a set of features using, in turn, the static and dynamic feature construction that uses the ILP learner, equipped with background knowledge definitions B1–B10. Call the corresponding sets of features “Static” and “Dynamic”.
3. Obtain the best model possible using the feature-based learner supplied with data for features $B0 \cup \text{Static}$. Call this model “ILP-S”. Similarly, obtain “ILP-D” using features $B0 \cup \text{Dynamic}$. For simplicity, we will refer to these models collectively as “ILP-assisted models”.
4. Compare the performance of the Baseline model against that of the ILP-assisted models.

The following details are relevant:

- (a) For the monolingual tasks, we use the training/test data as provided by SENSEVAL-3 and SemEval-2007 benchmarks. These specify different percentages for training and test, depending on the target word. For the bilingual task, we use 25% of the data to estimate the performance of disambiguation models. For all the data sets, performance is measured by the accuracy of prediction on the test set (that is, the percentage of test examples whose sense is predicted correctly).
- (b) The ILP learner constructs a set of acceptable clauses in line with the specifications described in Sect. 3.1. Positive examples for the ILP learner are provided by the correct sense (or translation in the bilingual case) of the target word in a sentence. Negative examples are generated automatically using all other senses (or translations). Clauses for a class are thus found using a “one-versus-the-rest” approach.

⁵<http://www.cs.waikato.ac.nz/~ml/weka/>.

- (c) For each target word and task, constructing the “best possible model” requires determining optimal values for some parameters of the procedures involved. We estimate these values using an instance of the method proposed in (Kohavi and John 1995) that proceeds as follows. First, we decide on the relevant parameters. Second, we obtain, using the training set only, unbiased estimates of the predictive accuracy of the models for each target word arising from systematic variation across some small number of values for these parameters. Values that yielded the best average predictive accuracy across all target words are taken to be optimal ones. This procedure is not perfect: correctly, optimal values may change from one target word to another; and even if they did not, the results obtained may be a local maximum (that is, better models may result from further informed variation of values).
- (d) Full-scale experimentation for optimal settings for parameters requires systematic joint variation of values of critical parameters for the feature constructor, ILP learner and the feature-based learner. For reasons of tractability, we restrict this here to an exploration of parameter values for the feature-based learner (here, a linear SVM) only. For the other two procedures, we select values that allow reasonably large numbers of features to be generated (see below). The principal parameters for the feature-based learner are taken to be: the C parameter used by the linear SVM and F , the number of features to be selected. The following C values were investigated: 0.1, 1.0, 10, 100, and 1000. Given a total of N features, F values considered were $N/64$, $N/32$, $N/16$, $N/8$, $N/4$, $N/2$, and N (we obtain the best subset of such features using software tools for feature-selection provided within WEKA). The predictive accuracy with each (C, F) setting is estimated and the values that yield the best results are used to construct the final model (the predictive accuracy estimate is obtained using an average over 5 repeats of predictions on 20% of the training data sampled to form a “validation” set).
- (e) For the record, parameter settings for the feature-constructor and ILP learner were as follows. For the former the value of k (the maximum number of features constructed for each class) is 5000; the minimal support s required is 2; and the minimal precision 0.6. We note that this value of precision does always guarantee weak relevance in all datasets (where the prior probability of a class is > 0.6). In such cases, not all the features may be weakly relevant. In addition, the search is bounded by requiring clauses to have no more than 10 literals in the context-predicate (that is, the body of clauses has at most 10 literals). Search for clauses is restricted to no more than 10,000 nodes. All these settings are admittedly *ad hoc*, with the intent being to place, within computational reason, as few constraints as possible on feature generation.
- (f) Comparison of performance is done using the Wilcoxon signed-rank test (Siegel 1956). This is a non-parametric test of the null hypothesis that there is no significant difference between the median performance of a pair of algorithms. The test works by ranking the absolute value of the differences observed in performance of the pair of algorithms. Ties are discarded and the ranks are then given signs depending on whether the performance of the first algorithm is higher or lower than that of the second. If the null hypothesis holds, the sum of the signed ranks should be approximately 0. The probabilities of observing the actual signed rank sum can be obtained by an exact calculation (if the number of entries is less than 10), or by using a normal approximation. We note that the comparing a pair of algorithms using the Wilcoxon test is equivalent to determining if the area under the ROC curves of the algorithms differ significantly (Hand 1997).
- (g) We note also that given that the ILP-assisted models are constructed with access to all features available for the Baseline model, we would expect that the predictive accuracies of the ILP-assisted models should, in principle, never be worse. That is, the alternate hypothesis for the Wilcoxon test is a uni-directional one that the median accuracy of the

ILP-assisted model (ILP-S or ILP-D) is higher than the Baseline model. However, in practise, ILP-assisted models could be worse either due to overfitting, or due to limitations of the feature-selection procedure employed (if any). Given this, we adopt the bi-directional alternate hypothesis, that the median accuracies of the ILP-assisted models are not equal to the Baseline model. This means obtaining two-tailed probabilities for the signed rank-differences.

5 Results and discussion

Tables 1, 2 and 3 tabulate the performance of the Baseline and ILP-assisted models on the three different WSD tasks. It is also standard practise to include the performance of a classifier that simply predicts the most frequent sense of the target word, which we denote in these tabulations as the “Majority class” model. The principal details in these tabulations are these: (1) The majority class models are better, on average, than the baseline models for Monolingual-2. However, its performance is significantly worse than the ILP-based models. We will leave them out of any further discussion related to the principal hypothesis being investigate here; (2) For all three tasks, average accuracies of the baseline models are lower than the ILP-assisted models.

We turn now to the question of whether the differences in accuracies observed between the Baseline and ILP-assisted models are in fact significant. The relevant probabilities calculated by using the Wilcoxon test are shown in Table 4.⁶ The tabulations show that, overall (see Table 4(a)), there is very little evidence in favour of the null hypothesis that the median accuracies of the ILP-assisted models are the same as the Baseline model. Closer examination of the individual tasks (Table 4(b)–(d)) suggests that we can be quite confident about this with ILP-D. ILP-S is also better, but some of the probabilities on Monolingual (1) and Bilingual tasks are slightly lower than what would usually be considered significant.

Based on these tabulations we conclude that there is indeed evidence in favour of the ILP-features conjecture, using ILP-S or ILP-D as an example of what is achievable with ILP assistance.

5.1 Going beyond the “ILP-features hypothesis”

We turn now to some issues that are not immediately relevant the principal hypothesis investigated in the paper, but are nevertheless of interest:

1. Statistical evidence in favour of the ILP-features hypothesis is of little direct interest to those working in automatic machine-translation. How do the ILP-S and ILP-D models perform compared to the state-of-the-art in WSD?
2. It is of some interest to ILP-practitioners to know which of ILP-S or ILP-D is better as a technique for feature-construction.
3. The results here merely constitute a kind of empirical existence proof that it is possible to use an ILP system to construct useful features. There is nothing to suggest here that we have obtained the best possible results from using ILP: can we do better using a different ILP system for constructing features?⁷

⁶These were obtained from the program kindly provided by Richard Lowry at <http://faculty.vassar.edu/lowry/wilcoxon.html>.

⁷This does not, of course, address the question of whether features capturing some or all of the information found by the ILP learner *could* be hand-crafted. This is beyond the scope of this paper.

Table 1 Estimates of accuracies of disambiguation models for datasets in Monolingual-1 task. “Micro” and “Macro” refer to averages weighted, and un-weighted by the number of examples

Word	Majority class	Baseline	ILP-S	ILP-D
activate	82.46 ± 3.56	81.58 ± 3.63	82.45 ± 3.56	92.98 ± 2.39
add	45.80 ± 4.35	81.68 ± 3.38	83.21 ± 3.27	85.50 ± 3.08
appear	44.70 ± 4.33	70.45 ± 3.97	71.21 ± 3.94	88.64 ± 2.76
ask	27.78 ± 3.99	53.97 ± 4.44	53.17 ± 4.45	60.32 ± 4.36
begin	59.74 ± 5.59	63.64 ± 5.48	72.73 ± 5.08	74.03 ± 5.00
climb	55.22 ± 6.08	71.64 ± 5.51	86.57 ± 4.17	85.07 ± 4.35
decide	67.74 ± 5.94	79.03 ± 5.17	80.64 ± 5.02	79.03 ± 5.17
eat	88.37 ± 3.46	87.21 ± 3.60	88.37 ± 3.46	87.21 ± 3.60
encounter	50.77 ± 6.20	73.85 ± 5.45	72.30 ± 5.55	73.85 ± 5.45
expect	74.36 ± 4.94	79.49 ± 4.57	92.31 ± 3.02	92.31 ± 3.02
express	69.09 ± 6.23	67.27 ± 6.33	67.27 ± 6.33	78.18 ± 5.57
hear	46.88 ± 8.82	62.50 ± 8.56	62.50 ± 8.56	65.63 ± 8.40
lose	52.78 ± 8.32	52.78 ± 8.32	52.70 ± 8.32	52.78 ± 8.32
mean	52.50 ± 7.90	72.50 ± 7.06	75.00 ± 6.85	75.00 ± 6.85
miss	33.33 ± 8.61	36.67 ± 8.80	36.66 ± 8.80	40.00 ± 8.94
note	38.81 ± 5.95	53.73 ± 6.09	88.06 ± 3.96	88.06 ± 3.96
operate	16.67 ± 8.78	66.67 ± 11.11	72.22 ± 10.56	72.22 ± 10.56
play	46.15 ± 6.91	53.85 ± 6.91	55.77 ± 6.89	55.77 ± 6.89
produce	52.13 ± 5.15	63.83 ± 4.96	65.96 ± 4.89	74.47 ± 4.50
provide	85.51 ± 4.24	89.86 ± 3.63	86.96 ± 4.05	89.86 ± 3.63
receive	88.89 ± 6.05	88.89 ± 6.05	88.89 ± 6.05	88.89 ± 6.05
remain	78.57 ± 4.90	84.29 ± 4.35	85.71 ± 4.18	87.14 ± 4.00
rule	50.00 ± 9.13	66.67 ± 8.61	83.33 ± 6.80	86.67 ± 6.21
smell	40.74 ± 6.69	79.63 ± 5.48	75.92 ± 5.82	74.07 ± 5.96
suspend	35.94 ± 6.00	56.25 ± 6.20	56.25 ± 6.20	57.81 ± 6.17
talk	72.60 ± 5.22	73.97 ± 5.14	73.97 ± 5.14	73.97 ± 5.14
treat	28.07 ± 5.95	47.37 ± 6.61	57.89 ± 6.54	50.88 ± 6.62
use	71.43 ± 12.07	92.86 ± 6.88	92.86 ± 6.88	92.86 ± 6.88
wash	67.65 ± 8.02	73.53 ± 7.57	61.76 ± 8.33	64.71 ± 8.20
watch	74.51 ± 6.10	76.47 ± 5.94	72.54 ± 6.25	74.51 ± 6.10
win	44.74 ± 8.07	47.37 ± 8.10	57.89 ± 8.01	60.53 ± 7.93
write	26.09 ± 9.16	47.83 ± 10.42	39.13 ± 10.18	52.17 ± 10.42
Micro	56.26	69.94	73.14	76.60
Macro	55.31	68.67	71.63	74.22

We consider each of these questions in turn. A comparison of the ILP models against the state-of-the-art models reported for the Monolingual-1, Monolingual-2, and Bilingual datasets is shown in Tables 5, 6 and 7. A simple comparison of the averages show the ILP-based models to be better than the state-of-the-art on Monolingual-1 and Bilingual tasks, but

Table 2 Estimates of accuracies of disambiguation models for datasets in Monolingual-2 task

Word	Majority class	Baseline	ILP-S	ILP-D
allow	97.14 ± 2.82	97.14 ± 0.48	97.14 ± 0.48	94.29 ± 0.66
announce	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
approve	91.67 ± 7.98	91.67 ± 2.30	91.67 ± 2.30	91.67 ± 2.30
area	70.27 ± 7.51	62.16 ± 1.31	72.97 ± 1.20	72.97 ± 1.20
ask	51.72 ± 6.56	50.00 ± 0.86	50.00 ± 0.86	50.00 ± 0.86
authority	23.81 ± 9.29	47.62 ± 2.38	85.71 ± 1.67	47.62 ± 2.38
base	10.00 ± 6.71	30.00 ± 2.29	55.00 ± 2.49	50.00 ± 2.5
begin	56.25 ± 7.16	45.83 ± 1.04	66.67 ± 0.98	68.75 ± 0.97
believe	78.18 ± 5.57	67.27 ± 0.85	63.64 ± 0.87	63.64 ± 0.87
bill	75.49 ± 4.26	66.67 ± 0.46	96.08 ± 0.19	96.08 ± 0.19
build	73.91 ± 6.47	60.87 ± 1.06	60.87 ± 1.06	65.22 ± 1.04
buy	76.09 ± 6.29	69.57 ± 1.00	65.22 ± 1.04	67.39 ± 1.02
capital	96.49 ± 2.44	96.49 ± 0.32	96.49 ± 0.32	94.74 ± 0.39
care	28.57 ± 17.07	42.86 ± 7.07	42.86 ± 7.07	42.86 ± 7.07
carrier	71.43 ± 9.86	71.43 ± 2.15	71.43 ± 2.15	71.43 ± 2.15
chance	40.00 ± 12.65	20.00 ± 2.67	60.00 ± 3.27	60.00 ± 3.27
claim	80.00 ± 10.33	80.00 ± 2.67	80.00 ± 2.67	80.00 ± 2.67
come	23.26 ± 6.44	20.93 ± 0.95	62.79 ± 1.12	65.12 ± 1.11
complain	85.71 ± 9.35	85.71 ± 2.50	85.71 ± 2.50	85.71 ± 2.50
complete	93.75 ± 6.05	93.75 ± 1.51	93.75 ± 1.51	93.75 ± 1.51
condition	76.47 ± 7.27	73.53 ± 1.30	73.53 ± 1.30	64.71 ± 1.41
contribute	50.00 ± 11.79	38.89 ± 2.71	72.22 ± 2.49	33.33 ± 2.62
defense	28.57 ± 9.86	14.29 ± 1.67	42.86 ± 2.36	33.33 ± 2.24
describe	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
development	62.07 ± 9.01	62.07 ± 1.67	72.41 ± 1.54	72.41 ± 1.54
do	90.16 ± 3.81	90.16 ± 0.49	90.16 ± 0.49	80.33 ± 0.65
drug	86.96 ± 4.97	76.09 ± 0.93	73.91 ± 0.95	76.09 ± 0.93
effect	76.67 ± 7.72	63.33 ± 1.61	93.33 ± 0.83	93.33 ± 0.83
end	52.38 ± 10.90	47.62 ± 2.38	90.48 ± 1.40	95.24 ± 1.01
enjoy	57.14 ± 13.23	50.00 ± 3.57	57.14 ± 3.53	50.00 ± 3.57
examine	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
exchange	73.77 ± 5.63	63.93 ± 0.79	86.89 ± 0.55	86.89 ± 0.55
exist	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
explain	88.89 ± 7.41	88.89 ± 1.75	94.44 ± 1.27	94.44 ± 1.27
feel	68.63 ± 6.50	68.63 ± 0.91	60.78 ± 0.96	60.78 ± 0.96
find	82.14 ± 7.24	78.57 ± 1.47	89.29 ± 1.10	85.71 ± 1.25
future	86.30 ± 2.85	77.40 ± 0.29	89.04 ± 0.21	83.56 ± 0.25
go	45.90 ± 6.38	32.79 ± 0.77	68.85 ± 0.76	63.93 ± 0.79
grant	80.00 ± 17.89	80.00 ± 8.00	100.00 ± 0.00	100.00 ± 0.00
hold	37.50 ± 9.88	25.00 ± 1.80	45.83 ± 2.08	25.00 ± 1.80
hour	89.58 ± 4.41	85.42 ± 0.74	89.58 ± 0.64	87.5 ± 0.69
job	82.05 ± 6.15	76.92 ± 1.08	74.36 ± 1.12	74.36 ± 1.12
join	38.89 ± 11.49	38.89 ± 2.71	33.33 ± 2.62	27.78 ± 2.49
keep	56.25 ± 5.55	46.25 ± 0.62	57.5 ± 0.62	46.25 ± 0.62

Table 2 (Continued)

Word	Majority class	Baseline	ILP-S	ILP-D
kill	87.50 ± 8.27	87.50 ± 2.07	87.5 ± 2.07	87.5 ± 2.07
lead	38.46 ± 7.79	30.77 ± 1.18	69.23 ± 1.18	33.33 ± 1.21
maintain	90.00 ± 9.49	80.00 ± 4.00	100.00 ± 0.00	100.00 ± 0.00
management	71.11 ± 6.76	62.22 ± 1.08	75.56 ± 0.95	82.22 ± 0.85
move	97.87 ± 2.10	97.87 ± 0.31	97.87 ± 0.31	97.87 ± 0.31
need	71.43 ± 6.04	57.14 ± 0.88	89.29 ± 0.55	87.50 ± 0.59
network	90.91 ± 3.88	76.36 ± 0.77	89.09 ± 0.57	87.27 ± 0.61
occur	86.36 ± 7.32	77.27 ± 1.90	95.45 ± 0.95	77.27 ± 1.90
order	91.23 ± 3.75	91.23 ± 0.50	85.96 ± 0.61	85.96 ± 0.61
part	66.20 ± 5.61	50.7 ± 0.70	53.52 ± 0.70	52.11 ± 0.70
people	90.43 ± 2.74	86.09 ± 0.30	86.09 ± 0.30	86.09 ± 0.30
plant	98.44 ± 1.55	98.44 ± 0.19	98.44 ± 0.19	100.00 ± 0.00
point	81.33 ± 3.18	70.00 ± 0.31	90.00 ± 0.20	78.67 ± 0.27
policy	97.44 ± 2.53	97.44 ± 0.40	94.87 ± 0.57	94.87 ± 0.57
position	46.67 ± 7.44	28.89 ± 1.01	31.11 ± 1.03	31.11 ± 1.03
power	27.66 ± 6.52	53.19 ± 1.06	82.98 ± 0.8	76.60 ± 0.9
prepare	77.78 ± 9.80	55.56 ± 2.76	83.33 ± 2.07	88.89 ± 1.75
president	72.88 ± 3.34	68.36 ± 0.26	97.18 ± 0.09	96.61 ± 0.10
produce	75.00 ± 6.53	72.73 ± 1.01	63.64 ± 1.09	63.64 ± 1.09
promise	75.00 ± 15.31	75.00 ± 5.41	87.5 ± 4.13	75.00 ± 5.41
propose	85.71 ± 9.35	85.71 ± 2.50	85.71 ± 2.50	85.71 ± 2.50
prove	68.18 ± 9.93	45.45 ± 2.26	68.18 ± 2.12	54.55 ± 2.26
raise	14.71 ± 6.07	8.82 ± 0.83	29.41 ± 1.34	29.41 ± 1.34
rate	86.21 ± 2.86	79.31 ± 0.28	79.31 ± 0.28	79.31 ± 0.28
recall	86.67 ± 8.78	86.67 ± 2.27	86.67 ± 2.27	86.67 ± 2.27
receive	95.83 ± 2.88	95.83 ± 0.42	93.75 ± 0.5	95.83 ± 0.42
regard	71.43 ± 12.07	78.57 ± 2.93	85.71 ± 2.50	92.86 ± 1.84
remember	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
replace	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
report	91.43 ± 4.73	91.43 ± 0.80	91.43 ± 0.80	91.43 ± 0.80
rush	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00	100.00 ± 0.00
say	98.71 ± 0.49	98.52 ± 0.02	97.78 ± 0.03	97.78 ± 0.03
see	44.44 ± 6.76	37.04 ± 0.89	55.56 ± 0.92	57.41 ± 0.92
set	28.57 ± 6.97	19.05 ± 0.93	61.9 ± 1.16	57.14 ± 1.18
share	97.14 ± 0.73	95.62 ± 0.04	96.57 ± 0.03	96.57 ± 0.03
source	37.14 ± 8.17	25.71 ± 1.25	82.86 ± 1.08	60.00 ± 1.40
space	78.57 ± 10.97	92.86 ± 1.84	92.86 ± 1.84	92.86 ± 1.84
system	48.57 ± 5.97	40.00 ± 0.70	37.14 ± 0.69	37.14 ± 0.69
turn	38.71 ± 6.19	22.58 ± 0.67	54.84 ± 0.80	53.23 ± 0.80
value	98.31 ± 1.68	94.92 ± 0.37	96.61 ± 0.31	96.61 ± 0.31
work	55.81 ± 7.57	46.51 ± 1.16	65.12 ± 1.11	65.12 ± 1.11
Micro	78.10	73.14	82.50	80.30
Macro	71.69	67.53	78.52	73.39

Table 3 Estimates of accuracies of disambiguation models for datasets in the bilingual task

Word	Majority class	Baseline	ILP-S	ILP-D
come	50.30 ± 7.45	60.00 ± 7.30	75.56 ± 6.41	75.56 ± 6.41
get	21.00 ± 6.52	30.77 ± 7.39	64.10 ± 7.68	76.92 ± 6.75
give	88.80 ± 4.70	95.56 ± 3.07	100.00 ± 0.00	100.00 ± 0.00
go	68.50 ± 6.64	73.47 ± 6.31	77.55 ± 5.96	75.51 ± 6.14
look	50.30 ± 7.29	78.72 ± 5.97	76.59 ± 6.18	80.85 ± 5.74
make	70.00 ± 7.07	78.57 ± 6.33	78.57 ± 6.33	90.48 ± 4.53
take	28.50 ± 7.98	43.75 ± 8.77	56.25 ± 8.77	53.13 ± 8.82
Micro	55.69	67.56	76.59	79.93
Macro	53.91	65.83	75.52	78.92

Table 4 Probabilities of observing the differences in accuracies for the monolingual and bilingual tasks, under the null hypothesis that median accuracies of the pair of algorithms being compared are equal. Each entry consists two-tailed probability estimates of the null hypothesis being true. Note that the Majority class is significantly better than Baseline on Monolingual-2 which results in a relatively high probability of observing the differences between the two overall (0.60)

	Majority	Base	ILP-S	ILP-D
<i>(a) Overall</i>				
Majority	–	–	–	–
Base	0.69	–	–	–
ILP-S	< 0.0001	< 0.0001	–	–
ILP-D	< 0.0001	< 0.0001	0.42	–
<i>(b) Monolingual-1</i>				
Majority	–	–	–	–
Base	< 0.0001	–	–	–
ILP-S	< 0.0001	0.06	–	–
ILP-D	< 0.0001	0.0006	0.002	–
<i>(c) Monolingual-2</i>				
Majority	–	–	–	–
Base	0.0001	–	–	–
ILP-S	0.0002	0.0001	–	–
ILP-D	0.053	< 0.0001	0.0013	–
<i>(d) Bilingual</i>				
Majority	–	–	–	–
Base	0.02	–	–	–
ILP-S	0.02	0.08	–	–
ILP-D	0.02	0.02	> 0.10	–

to be worse than the top 5 submissions made for Monlingual(2).⁸ Thus, the best we can say

⁸We note that fourth-best model on Monolingual-2 is, in fact our ILP-S procedure with differences arising from random sampling of features.

Table 5 A comparison of the ILP-based models against the best submissions made for the SENSEVAL-3 competition, that is, in the Monolingual-1 dataset. The best model, MC-WSD, is a multi-class averaged perceptron classifier with one component trained on the data provided by SENSEVAL and on WordNet glosses

Models	Accuracy
ILP-D	76.60
ILP-S	73.14
MC-WSD (Ciaramita and Johnson 2004)	72.50
Syntalex-3 (Mohammad and Pedersen 2004)	67.60
Syntalex-1 (Mohammad and Pedersen 2004)	67.00
CLaC1 (Lamjiri et al. 2004)	67.00
Syntalex-2 (Mohammad and Pedersen 2004)	66.50
CLaC2 (Lamjiri et al. 2004)	66.00
Syntalex-4 (Mohammad and Pedersen 2004)	65.30
Majority class	55.31

Table 6 A comparison of the ILP-based models against the best submissions made for the Semeval competition, that is, in the Monolingual-2 dataset. The best submission, NUS-ML is a SVM model. Note that our accuracies do not include words that have a single class on training and test dataset. If these words were taken into account, we achieve slightly higher accuracies (83.2 for ILP-S and 81.1 for ILP-D)

Models	Accuracy
NUS-ML (Cai et al. 2007)	88.70
UBC-ALM (Agirre and Lopez de Lacalle 2007)	86.90
I2R (Niu et al. 2007)	86.40
USP-IBM-2 (Specia et al. 2007d)	85.70
USP-IBM-1 (Specia et al. 2007d)	85.10
ILP-S*	82.57
ILP-D*	80.36
Majority class	78.10

at this point is that the ILP-based models are good, but perhaps not always the best possible. Of course, this form of adversarial comparison is not how we intend the ILP-based feature constructors to be used. In practise, we would expect other model construction techniques to incorporate the ILP features when constructing their models. Unfortunately, we are unable to assess how well such a collaborative method would perform, as we have been unable to get executable model-construction procedures for any of the non ILP-based submissions.

The difference between static and dynamic feature construction techniques embodied within ILP-S and ILP-D. On the evidence in the Appendix of this paper and Table 4, it appears that, for WSD tasks at least, there is little to chose between the two methods. Although accuracies are not significantly different, ILP-D models may be preferable on grounds of requiring substantially fewer features—this is after the feature selection step—than ILP-S (and sometimes even the baseline models: see Table 8), arising from the use of a partially correct model to direct the construction of only those features that may be necessary to correct the model.

Table 7 A comparison of the ILP-D and ILP-S with other models reported earlier in Specia (2006b) for the Bilingual dataset

Models	Accuracy
ILP-D	79.93
ILP-S	76.59
Aleph	76.00
SVM	68.00
C4.5	63.00
Majority	55.69

Table 8 Average numbers of features used by the different models

Model	Task		
	Monolingual-1	Monolingual-2	Bilingual
Baseline	136	282	102
ILP-S	2429	2813	2918
ILP-D	179	133	63

Do either the ILP-S and ILP-D models here contain the best features constructible with ILP? Clearly not: we already have evidence in Table 6 of an ILP-based model (USP-IBM-2) that uses the ILP-S procedure to achieve better results, using the same ILP system as the one used here (the differences being due to sampling variations). Better features may also result by using a different ILP system. Experiments conducted with features constructed by the ILP system RSD (Železný and Lavrač 2006) obtain results roughly comparable to the models here (average macro-accuracies of 70.06, 78.54 and 74.35 on Monolingual-1, Monolingual-2 and Bilingual data respectively).

6 Concluding remarks

Word sense disambiguation, a necessary component for a variety of natural language processing tasks, remains amongst the hardest to model adequately. It is of course possible that the vagaries of natural language may place a limit on the accuracy with which a model could identify correctly the sense of an ambiguous word, but it is not clear that this limit has been reached with the modelling techniques that constitute the current state-of-the-art. The performance of these techniques depends largely on the adequacy of the features used to represent the problem. As it stands, these features are usually hand-crafted and largely of a lexical nature. For substantial, scalable progress it is believed that knowledge that accounts for more elaborate syntactic and semantic information needs to be incorporated. In this paper, we have investigated the use of Inductive Logic Programming as a mechanism for incorporating multiple sources of syntactic and semantic information into the construction of models for WSD. The investigation has been in the form of empirical studies of using ILP to construct features that are then used to construct predictive models for monolingual and bilingual WSD tasks and the results provide evidence that this approach can yield better models.

We believe much of the gains observed with ILP stems from the use of substantial amounts of background knowledge. For the work here, this knowledge has been obtained

by translations of information in standard corpora or electronic lexical resources. This is promising, as it suggests that these translators, in conjunction with ILP, may provide a set of tools for the automatic incorporation of deep knowledge into the construction of general WSD models. Turning specifically to the findings of this paper, a combination of randomised search and incremental feature-set construction by an ILP engine appears to be a promising method to construct good features for WSD.

There are a number of ways in which the work here can be improved and extended. We list the main limitations here under three categories. On the conceptual front, it is evident that we have not provided any guarantees of optimality on the feature-subset constructed. While this is typical of randomised methods of the type proposed here, it would nevertheless be useful to obtain some performance bounds, however loose. Our specification of a feature-construction algorithm—to the best of our knowledge, the only one to be proposed in the ILP literature—can be refined to require the algorithm to return strongly relevant features (in the sense described by John et al. 1994).

On the implementation front, our implementation of dynamic feature-set construction is based on the simplest kind of randomised search (GSAT). Better methods exist and need to be investigated (for example, WalkSat). Further, we could consider other neighbourhood definitions for the local search such as adding or dropping upto k features. Of course, we are not restricted to use SVMs, or even the specific variant of SVM here, as our model constructor. Our experiments on the monolingual tasks here suggest that there is no significant difference between a “1-norm” SVM and the standard approach we have used here, but other model construction techniques may yield better results. We note that the feature-set constructor when used in conjunction with an SVM is an instance of SVILP, as proposed in (Muggleton et al. 2005). More generally, by interleaving feature and statistical model construction in the dynamic feature-set constructor, we are effectively performing a form of statistical relational learning. The applicability of this approach to this wider area needs to be investigated.

On the application front, the background knowledge used here is by no means exhaustive of the kind available for the WSD tasks studied here. For example, for the bilingual task, the “translation context” for a target word may help greatly. This refers to the translations into the target language of the words forming the context of the target word. Our preliminary experiments show that this does improve accuracies on the bilingual task. Further, the algorithms we have described here are unlikely to be the only ones to satisfy the specifications for feature-constructors. While the procedures in the paper establish a case for ILP-based feature construction in WSD, other feature-construction methods such as the ones in (Landwehr et al. 2006) and (Davis et al. 2007) may yield even better models, thus strengthening the case for ILP further.

On the flip side, it is evident that the feature-construction algorithm we have proposed here are not specific to WSD data. The use of the algorithm to model data concerned with tasks other than WSD, while outside the scope of this paper, has obvious wider interest.

Acknowledgements This work was begun when Lucia Specia visited Ashwin Srinivasan at the IBM Research Laboratory in New Delhi. Ganesh Ramakrishnan was at the IBM India Research Laboratory for most of the work done in this paper. Jimmy Foulds, at the University of Waikato, helped beyond the call of duty in getting an executable version of code that constructed models using a 1-norm SVM within WEKA.

References

Agirre, E., & Lopez de Lacalle, O. (2007). UBC-ALM: combining k-NN with SVD for WSD. In *Proceedings of the fourth international workshop on semantic evaluations* (pp. 342–345).

- Agirre, E., & Rigau, G. (1996). Word sense disambiguation using conceptual density. In *16th international conference on computational linguistics* (pp. 16–22), Copenhagen.
- Bar-Hillel, Y. (1960). Automatic translation of languages. In F. Alt, D. Booth, & R. E. Meagher (Eds.), *Advances in computers*. New York: Academic Press.
- Cai, J. F., Lee, W. S., & Teh, Y. W. (2007). NUS-ML: improving word sense disambiguation using topic features. In *Proceedings of the fourth international workshop on semantic evaluations* (pp. 249–252).
- Ciaramita, M., & Johnson, M. (2004). Multi-component word sense disambiguation. In *SENSEVAL-3: 3rd international workshop on the evaluation of systems for the semantic analysis of text* (pp. 97–100), Barcelona.
- Cottrell, G. W. (1989). *A connectionist approach to word sense disambiguation. Research notes in artificial intelligence*. San Mateo: Morgan Kaufmann.
- Davis, J., Ong, I., Struyf, J., Burnside, E., Page, D., & Costa, V. S. (2007). Change of representation for statistical relational learning. In *International joint conferences on artificial intelligence*.
- Hand, D. J. (1997). *Construction and assessment of classification rules*. Chichester: Wiley.
- Hirst, G. (1987). *Semantic interpretation and the resolution of ambiguity. Studies in natural language processing*. Cambridge: Cambridge University Press.
- John, G. H., Kohavi, R., & Pfleger, K. (1994). Irrelevant features and the subset selection problem. In *Proceedings of the eleventh international conference on machine learning* (pp. 121–129). San Mateo: Morgan Kaufmann.
- Kohavi, R., & John, G. H. (1995). Automatic parameter selection by minimizing estimated error. In *12th international conference on machine learning*. San Francisco: Morgan Kaufmann.
- Kramer, S., Lavrac, N., & Flach, P. (2001). Propositionalization approaches to relational data mining. In S. Dzeroski & N. Lavrac (Eds.), *Relational data mining* (pp. 262–291). Berlin: Springer.
- Lamjiri, A., Demerdash, O., & Kosseim, F. (2004). Simple features for statistical word sense disambiguation. In *SENSEVAL-3: 3rd international workshop on the evaluation of systems for the semantic analysis of text* (pp. 133–136), Barcelona.
- Landwehr, N., Passerini, A., De Raedt, L., & Frasconi, P. (2006). kFOIL: learning simple relational kernels. In Y. Gil & R. Mooney (Eds.), *Proceedings of the twenty-first national conference on artificial intelligence*.
- Lavrac, N., Dzeroski, S., & Grobelnik, M. (1990). *Learning nonrecursive definitions of relations with LINUS* (Technical report). Jozef Stefan Institute.
- Lesk, M. (1986). Automated sense disambiguation using machine-readable dictionaries: how to tell a pine cone from an ice cream cone. In *SIGDOC conference* (pp. 24–26), Toronto.
- Lin, D. (1993). Principle based parsing without overgeneration. In *31st annual meeting of the association for computational linguistics* (pp. 112–120), Columbus.
- McRoy, S. (1992). Using multiple knowledge sources for word sense discrimination. *Computational Linguistics*, 18(1), 1–30.
- Mihalcea, R., Chklovski, T., & Kilgariff, A. (2004). The SENSEVAL-3 English lexical sample task. In *SENSEVAL-3: 3rd international workshop on the evaluation of systems for semantic analysis of text* (pp. 25–28), Barcelona.
- Miller, G. A., Beckwith, R. T., Fellbaum, C. D., Gross, D., & Miller, K. (1990). Wordnet: an on-line lexical database. *International Journal of Lexicography*, 3(4), 235–244.
- Mohammad, S., & Pedersen, T. (2004). Complementarity of lexical and simple syntactic features: the syntalex approach to SENSEVAL-3. In *SENSEVAL-3: 3rd international workshop on the evaluation of systems for the semantic analysis of text* (pp. 159–162), Barcelona.
- Muggleton, S. (1994). Inductive logic programming: derivations, successes and shortcomings. *SIGART Bulletin*, 5(1), 5–11.
- Muggleton, S., & De Raedt, L. (1994). Inductive logic programming: theory and methods. *Journal of Logic Programming*, 19(20), 629–679.
- Muggleton, S., Lodhi, H., Amini, A., & Sternberg, M. J. E. (2005). Support vector inductive logic programming. In *8th international conference on discovery science* (pp. 163–175). Berlin: Springer.
- Niu, Z. Y., Ji, D. H., & Tan, C. L. (2007). I2R: three systems for word sense discrimination, Chinese word sense disambiguation, and English word sense disambiguation. In *Proceedings of the fourth international workshop on semantic evaluations* (pp. 177–182).
- Nienhuys-Cheng, S., & de Wolf, R. (1997). *Foundations of inductive logic programming*. Berlin: Springer.
- Paes, A., Zaverucha, G., Page, C. D. Jr., & Srinivasan, A. (2007). *LNCS: Vol. 4455 ILP through propositionalization and stochastic k-term DNF learning. Sense disambiguation using inductive logic programming. Selected papers from the 16th international conference on inductive logic programming*. Berlin: Springer, (pp. 379–393).
- Parker, J., & Stahel, M. (1998). *Password: English dictionary for speakers of Portuguese*. São Paulo: Martins Fontes.

- Pedersen, T. (2002). A baseline methodology for word sense disambiguation. In *3rd international conference on intelligent text processing and computational linguistics*, Mexico City.
- Pradhan, S., Loper, E., Dligach, D., & Palmer, M. (2007). SemEval-2007 Task-17: English lexical sample, SRL and all words. In *Fourth international workshop on semantic evaluations* (pp. 87–92), Prague.
- Procter, P. (Ed.). (1978). *Longman dictionary of contemporary English*. Essex: Longman Group.
- Quillian, M. R. (1961). *A design for an understanding machine. Colloquium of semantic problems in natural language*. Cambridge: Cambridge University Press.
- Ratnaparkhi, A. (1996). *A maximum entropy part-of-speech tagger. Empirical methods in NLP conference*. Philadelphia: University of Pennsylvania Press.
- Schutze, H. (1998). Automatic word sense discrimination. *Computational Linguistics*, 24(1), 97–124.
- Siegel, S. (1956). *Nonparametric statistics for the behavioural sciences*. New York: McGraw-Hill.
- Specia, L. (2006a). A hybrid relational approach for WSD—first results. In *Student research workshop at Coling-ACL* (pp. 55–60), Sydney.
- Specia, L. (2006b). A hybrid relational approach for WSD—first results. In *Proceedings of the COLING/ACL 2006 student research workshop* (pp. 55–60).
- Specia, L., Nunes, M. G. V., & Stevenson, M. (2005). Exploiting parallel texts to produce a multilingual sense-tagged corpus for word sense disambiguation. In *RANLP-05, Borovets* (pp. 525–531).
- Specia, L., Nunes, M. G. V., & Stevenson, M. (2007a). Learning expressive models for word sense disambiguation. In *45th annual meeting of the association for computational linguistics* (pp. 41–48), Prague.
- Specia, L., Nunes, M. G. V., Srinivasan, A., & Ramakrishnan, G. (2007b). Word sense disambiguation using inductive logic programming. In *LNCS: Vol. 4455 Selected papers from the 16th international conference on inductive logic programming* (pp. 409–423). Berlin: Springer.
- Specia, L., Nunes, M. G. V., Srinivasan, A., & Ramakrishnan, G. (2007c). USP-IBM-1 and USP-IBM-2: the ILP-based systems for lexical sample WSD in SemEval-2007. In *4th international workshop on semantic evaluations* (pp. 442–445), Prague.
- Specia, L., Das, G. M., Nunes, M. G. V., Srinivasan, A., & Ramakrishnan, G. (2007d). USP-IBM-1 and USP-IBM-2: the ILP-based systems for lexical sample WSD in SemEval-2007. In *Proceedings of the fourth international workshop on semantic evaluations* (pp. 442–445).
- Srinivasan, A. (1999). *The aleph manual*. Available at <http://www.comlab.ox.ac.uk/oucl/research/areas/machlearn/Aleph/>.
- Stevenson, M., & Wilks, Y. (2001). The interaction of knowledge sources for word sense disambiguation. *Computational Linguistics*, 27(3), 321–349.
- Wilks, Y., & Stevenson, M. (1997). Combining independent knowledge sources for word sense disambiguation. In *3rd conference on recent advances in natural language processing* (pp. 1–7), Tzigov Chark.
- Wilks, Y., & Stevenson, M. (1998). The grammar of sense: using part-of-speech tags as a first step in semantic disambiguation. *Natural Language Engineering*, 4(1), 1–9.
- Yarowsky, D. (1995). Unsupervised word sense disambiguation rivaling supervised methods. In *33rd annual meeting of the association for computational linguistics* (189–196), Cambridge.
- Zelezny, F., Srinivasan, A., & Page, C. D. Jr. (2006). Randomised restarted search in ILP. *Machine Learning*, 64(1–3), 183–208.
- Železný, F. & Lavrač, N. (2006). Propositionalization-based relational subgroup discovery with RSD. *Machine Learning*, 62(1–2), 33–63.