Computational Methods in Semantics

Gábor Recski

Ph.D. Dissertation

Supervisor:
András Kornai D.Sc.

Ph.D. School of Linguistics
Vilmos Bárdosi D.Sc.

Theoretical Linguistics Ph.D Program
Zoltán Bánréti C.Sc.

Department of Theoretical Linguistics
Eötvös Loránd University, Budapest

Budapest, 2016

Acknowledgements

First and foremost I would like to thank my advisor András Kornai for his ongoing support of my studies and research over the past seven years, during which he supervised all work related to both my masters theses and virtually all projects and assignments related to my studies at ELTE and my work at the Institute for Computer Science (SZTAKI). Not only was András always available and willing to spend his time and energy on lengthy discussions of any topic, it was his efforts that secured a truly inspiring work environment for my studies and research, sometimes against all odds.

For acquainting me with natural language processing, sharing their tremendous experience, teaching me to code, and for being two of the smartest and coolest people I know, I'd like to thank former MOKK colleagues Péter Halácsy and Dániel Varga. I am forever grateful for their patience and helpfulness.

I am thankful to the professors at the Department of Theoretical Linguistics and I am particularly indebted to László Kálmán, who over the last 10 years has been my professor of logic, semantics, syntax, computer science, computational linguistics, and many other language-related topics. I thank him for the countless thought-provoking discussions during his Tuesday seminars, after classes, on nyelveszforum, and at semantics camp.

I would also like to say thanks to all my current and former colleagues at the Research Group for Mathematical Linguistics: Judit Ács, Gábor Borbély, Márton Makrai, Dávid Nemeskey, Katalin Pajkossy, and Attila Zséder. I feel truly honored to be working with a team of such amazing people.

Contents

Chapter 1

Introduction

TODO: a kind of-ot is leszedni u1gy, mint a somethingot. A Zen-es pe1lda is fasza1bb lesz

TODO: STS ML re1szletek, Kata instrukcio1i alapja1n, ro2viden

TODO: 2016-os submissionu2nk

TODO: Evalua1cio1 to2ke2letes parszokon, angolra e1s magyarra

This thesis presents computational methods for creating semantic representations of natural language utterances and some early applications of such representations in various computational semantics tasks. All software presented in this thesis is downloadable from the `4lang` repository at `http://github.com/kornai/4lang` and may be freely distributed under an MIT license. The state of the `4lang` codebase at the time of submission of this thesis is preserved in the branch `recski_thesis`.

The thesis is structured as follows: Chapter 2 gives a short review of existing theories of word meaning, with a special focus on their applicability to natural language processing. Chapter 3 will provide an overview of the `4lang` formalism for modeling meaning, but will not attempt a full discussion, since the `4lang` formalism is the product of joint work by half a dozen researchers (Kornai et al., 2015), rather than being a contribution of this thesis. Chapter 4 presents the `dep_to_4lang` pipeline, which creates `4lang`-style meaning representations from running text, Chapter 5 describes its application to monolingual dictionary definitions, `dict_to_4lang`, used to create large concept lexica automatically. Chapter 6 presents applications of the `text_to_4lang` module to various tasks in Computational Semantics, including a competitive system for measuring semantic textual similarity (STS) (Recski & Ács, 2015) and an experimental framework for natural language understanding (Nemeskey et al., 2013). Chapter 7 presents the architecture of the ca. 3000-line `4lang` codebase, serving both as an overview of how the main tools presented in this thesis are implemented and as comprehensive software documentation. Finally, Chapter 8 discusses our plans for future applications of the `4lang` system including question answering and recognizing textual entailment.

1.1 Main contributions

The main contributions of the present work are the pipelines for building semantic representations from raw text and dictionary definitions, and their application to common tasks in computational semantics, presented in Chapters 4, 5, and 6, respectively. Papers presenting parts of this thesis will be cited along the way. The `4lang` principles outlined in Chapter 3 are the result of collaboration with current and former members of the Research Group for Mathematical Linguistics at the Hungarian Academy of Sciences: Judit Ács, András Kornai, Márton Makrai, Dávid Nemeskey, Katalin Pajkossy, and Attila Zséder. The systems presented in Chapters 4 and 5 constitute the author's work with only minor exceptions: the functions performing graph expansion (Section 5.3) are a result of joint work with Gábor Borbély (Budapest University of Technology), and a parser for the Collins Dictionary was contributed by Attila Bolevácz (Eötvös Loránd University). The SemEval systems presented in Section 6.1 were built in collaboration with Judit Ács and Katalin Pajkossy, the experimental systems described in Section 6.2 were implemented together with Dávid Nemeskey and Attila Zséder.

Chapter 2

Theories of word meaning

This chapter gives a survey of approaches to modeling the semantics of natural language, focusing on key ideas in representing word meaning. Our overview is neither complete, nor does it provide a full introduction to any theory in particular, it is merely an overview of major contributions to word meaning representation. We begin with a short overview of the historically central Katz and Fodor's *Structure of a Semantic Theory* (Section 2.1), followed by reviews of graph-based models of word meaning in Section 2.2, in particular the Semantic Memory Model of Quillian, the KL-ONE family of formalisms, and the more recent Abstract Meaning Representation framework. An overview of Montagovian approaches to word meaning is given in Section 2.3. Finally, in Section 2.4, we discuss continuous vector space semantics, the approach to representing word meaning that is currently most widely used in natural language processing.

2.1 Katz and Fodor's semantics

In their paper *The Structure of a Semantic Theory*, Katz and Fodor (1963) set a lower bound on what a theory of semantics must include. Their examples show three skills of a competent speaker to be independent of their knowledge of grammar: (i) handling ambiguity (the bill is large, but need not be paid), (ii) detecting anomaly (the paint is silent) and (iii) paraphrasing (What does the note say? Does it say X?).

In setting an upper bound on the domain of semantics, they disown the issue of disambiguating between various readings of the same sentence (in isolation) based on context, since that would require modeling all extralinguistic knowledge:

> "...if a theory of setting selection is to choose the correct reading for the sentence *Our store sells alligator shoes*, it must represent the fact that, to date, alligators

> do not wear shoes, although shoes for people are sometimes made from alligator skin". (Katz & Fodor, 1963, p.178)

Katz and Fodor conclude that the upper bound on a semantic theory should be that of semantic interpretation - a function that maps each sentence to a set of semantic representations, one corresponding to each possible reading of the sentence. They make clear that they impose this limit merely for practicality, because they "cannot in principle distinguish between the speaker's knowledge of his language and his knowledge of the world, because (...) part of the characterization of a LINGUISTIC ability is a representation of virtually all knowledge about the world that speakers share." (Katz & Fodor, 1963, p.179) In Section 3.3 we shall also argue that any apparatus capable of representing the meaning of natural language utterances must be capable of representing all of (naive, non-technical) world knowledge.)

In describing the components of a semantic theory, Katz and Fodor define the lexicon to contain separate entries for multiple *senses* of each word, and at the same time they state that the grammar and the lexicon together are still insufficient for a deterministic semantic interpretation, because of the multiple senses associated with most word forms. A *projection rule* that selects the appropriate sense of each word form in a sentence is postulated. This rule requires the senses of each word to be structured in the lexicon as exemplified in Figure 2.1. In Chapter 3 we shall describe the `4lang` representation of word meaning that is radically *monosemic*, i.e. makes as little use of *word senses* as possible and would map a word such as *bachelor* to a single representation that is compatible with all uses of the word.

Note that the representation of lexical items in Figure 2.1 also includes a theory of semantic primitives (*human, male, animal,* etc., Katz and Fodor refer to these as *semantic markers*), much in the spirit of Prague-style phonological theory (Trubetzkoy, 1958). A significant problem with this approach is that they have little to say about where the set of all semantic markers available might come from, i.e. what the primitives of their representation should be. All remaining lexical information about a word sense that is not contained in the semantic markers, i.e. the parts in square brackets in Figure 2.1 are called *distinguishers*. This distinction between the layers of markers and distinguishers is not unlike that between Aristotle's *genus* and *differentia* (Smith, 2015). Katz and Fodor also claim that distinguishers are out of reach for a theory of semantics:

> "The distinction between markers and distinguishers is meant to coincide with the distinction between that part of the meaning of a lexical item which is systematic for the language and that part which is not. In order to describe the

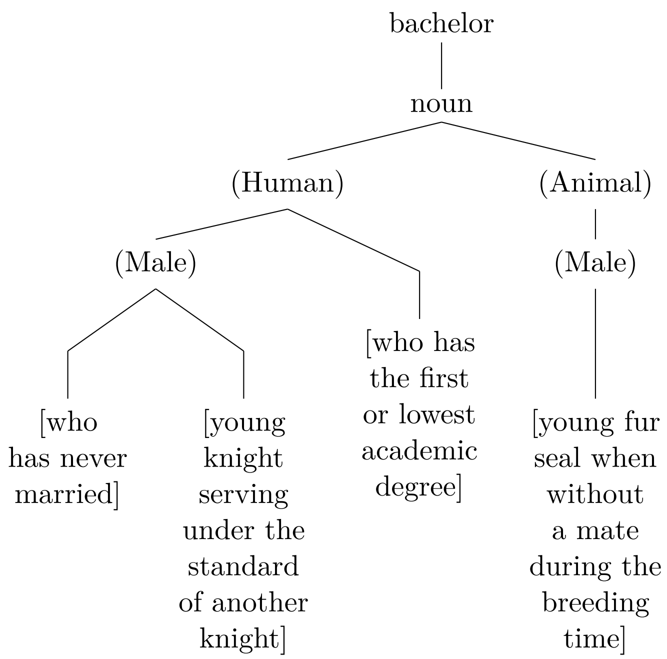

Figure 2.1: Decomposition of lexical items (Katz & Fodor, 1963, p.186)

> systematicity in the meaning of a lexical item, it is necessary to have theoretical constructs whose formal interrelations compactly represent this systematicity. The semantic markers are such constructs. The distinguishers, on the other hand, do not enter into theoretical relations within a semantic theory. The part of the meaning of a lexical item that a dictionary represents by a distinguisher is the part of which a semantic theory offers no general account." (Katz & Fodor, 1963, p.178)

What this last statement amounts to is that the (finite) set of semantic markers is a set of universal primitives that is sufficient for representing the language-independent component of word meaning. Then, if some non-English word is a hypernym of `bachelor`$_1$ - *man who has never married*, then its set of semantic markers must be a subset of the markers in the entry for `bachelor`$_1$. On the other hand, if we find a word in some language that is the hyponym of `bachelor`$_1$, e.g. a word `w` that means *a man who has never married and lives with his parents*, we must conclude that our original representation for `bachelor`$_1$ was inadequate, since the components of its meaning beyond *male* and *human*, whatever they may be, are shared with the entry `w` and should therefore be encoded by semantic markers, not distinguishers. Since the potential absence of such a word `w` from all human languages can only be accidental, we have to conclude that the distinction between meaning encoded by markers and by distinguishers is also arbitrary. Bolinger (1965, p.560) makes a similar argument, demonstrating that for virtually any component of any distinguisher in Figure 2.1 it is possible to construct an example that justifies 'promoting' that particular component to marker status, and concluding that "it is possible to do away with the dualism by converting the distinguisher into a string of markers". We shall return to his examples in Section 3.5 when we argue for a theory of meaning representation that encodes word meaning using language-independent primitives – and nothing else!

Finally, Katz and Fodor claim that word meaning representations may contain limitations on the semantic content of elements with which the given word can combine. In their example, an excerpt from *The shorter Oxford English dictionary*, the entry *honest* contains the definition '... of women: chaste, virtuous'; such requirements they would represent by adding constraints such as (Human) and (Female) on certain *paths* of the representation (paths in the sense of Figure 2.1). Section 3.2 will discuss how such constraints may be enforced by a `4lang`-based system that lacks a notion of *paths* or *senses*.

2.2 Graph-based models of semantics

This section reviews popular systems for representing meaning using graphs – networks of nodes and edges connecting them. We shall summarize the basic principles of Quillian's 1960s *Memory Model* in Section 2.2.1, the KL-ONE family of Knowledge Representation systems, widely used between the late 1970s and early 1990s, in Section 2.2.2, and finally in Section 2.2.3 the most recent formalism of *Abstract Meaning Representations* which has been gaining popularity in the past 4 years. All these systems share some common principles of representations with each other and with `4lang`, e.g. that each map lexical items to nodes in some graph and use directed edges to represent asymmetric relationships between them. Where they differ significantly is their elements of representation or their notions of a syntax-semantics interface.

2.2.1 Quillian's Semantic Memory Model

Memory model Quillian's theory of *word concepts* (1968) is of particular interest to us. Not only does he propose to represent word meaning by means of directed graphs of concepts (much like the `4lang` theory that serves as the basis of this thesis and will be introduced in Chapter 3), it also defines graph configurations that are in many ways similar to those in `4lang`. Quillian also suggests that definitions of concepts should be learned automatically, which is exactly what our module `dict_to_4lang` does (see Chapter 5).

Quillian proposes to encode meaning as a graph of nodes representing concepts, and *associative links* between nodes, which may encode a variety of semantic relationships between these concepts. Figure 2.2 reproduces Quillian's original presentation of associative link types. Types 1 and 2, which stand for hypernymy and attribution respectively – encode relationships that `4lang` will treat as a single relation (along with predication, see Section 3.1). Also, his links of type 5 and 6 are not unlike the binary configuration in `4lang` graphs.

Quillian proposes two types of nodes: *type nodes* are unique for each concept and serve to define them as networks of other concepts. *Token nodes* occur multiple times for each concept when they themselves are used in definitions. In Section 6.2, when we review early attempts at inferencing on `4lang` representations, we shall see that this distinction is not unlike that of *active* and *static* nodes made by (Nemeskey et al., 2013). Quillian organizes nodes into *planes*, one for each type node and its definition graph, and emphasizes the need to perform an exhaustive search of an arbitrary number of such planes for a complete definition of any concept:

> *"a word's full concept is defined in the model memory to be all the nodes that can be reached by an exhaustive tracing process, originating at its initial, patriarchal type node, together with the total sum of relationships among these nodes specified by within-plane, token-to-token links* (...) This information will start off with the more "compelling" facts about machines, such as that they are usually man-made, involve moving parts, and so on, and will proceed "down" to less and less inclusive facts, such as that typewriters are machines, and then eventually will get to much more remote information about machines, such as the fact that a typewriter has a stop that prevents its carriage from flying off everytime it is returned." (Quillian, 1968, p.413, emphasis in the original)

Quillian concludes that the bulk of information associated with a concept such as *machine* must be an unstructured list of all concepts that refer to types of machines and as such have edges directed towards tokens of *machine*. Thus a distinction is made, then between the *definition* of some concept, i.e. the tokens accessible (in the digraph sense) from its type node, and the network of all nodes connected to any token of the concept, all potentially carrying information about the concept – in Section 3.2 we shall argue that it is the latter that must be accessible to any language understanding mechanism.

Unlike Katz and Fodor, Quillian suggests not to represent the complex meaning of a word by means of a hierarchical structure of word senses. Instead he suggests that the unified network of all concepts that linked to either the type node or to some token node

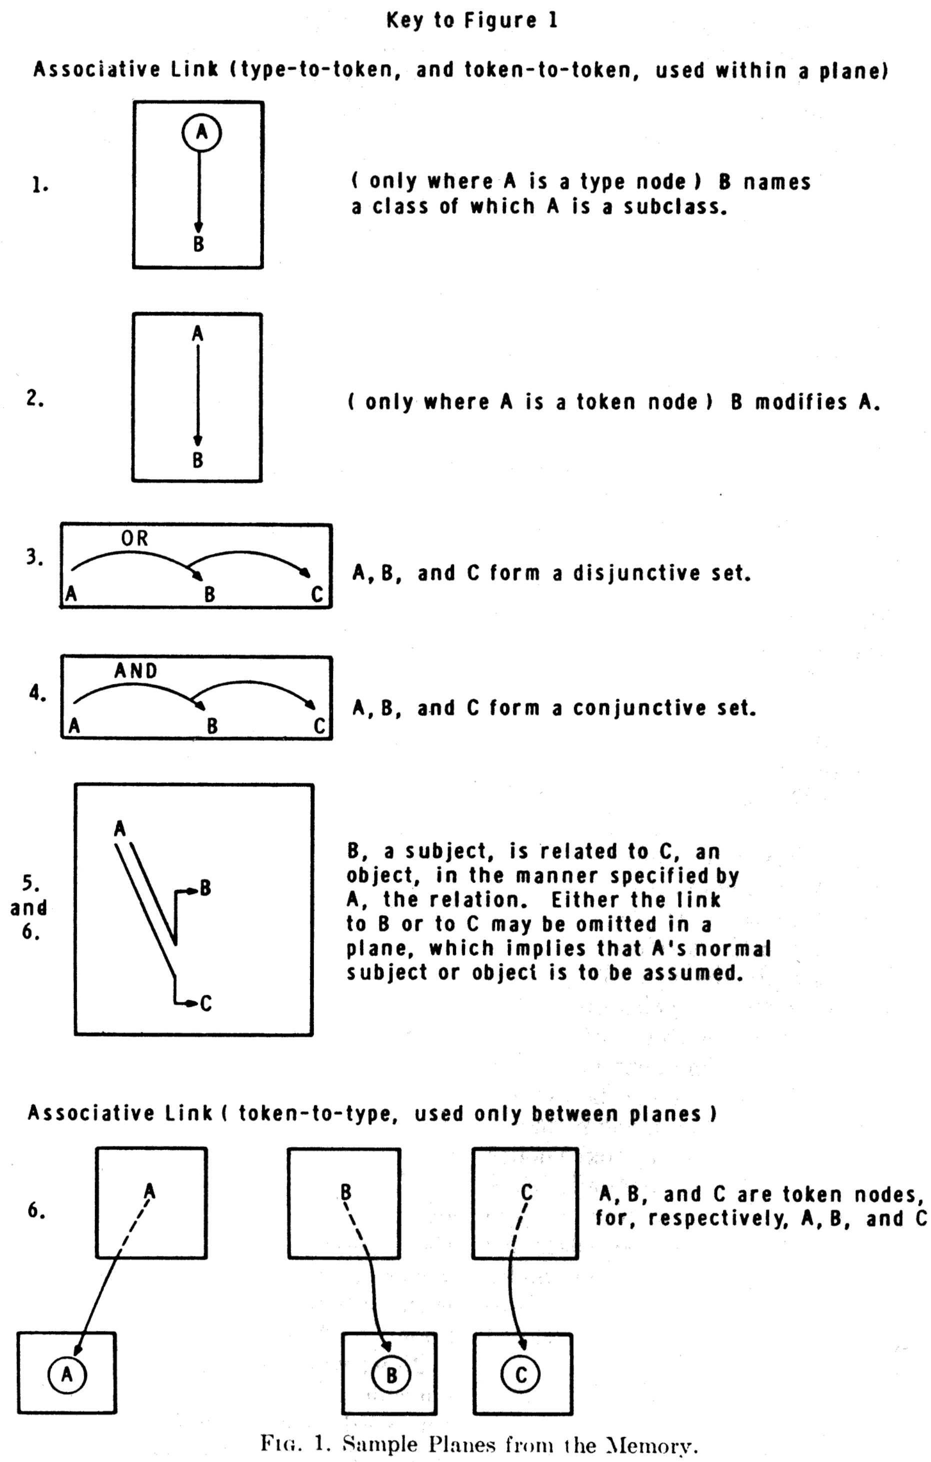

Figure 2.2: Associative links (Quillian, 1968, p.412)

of the concept being defined should by itself serve as a store of all knowledge associated with some word. He criticizes hierarchical structures of word senses commonly found in explanatory dictionaries by pointing out that "the common elements within and between various meanings of a word are many, and any outline designed to get some of these together under common headings must at the same time necessarily separate other common elements, equally valid from some other point of view" (Quillian, 1968, p.419). Nevertheless, the memory model still makes use of *word senses* and the proposed mechanism for building semantic representations from any given sentence still requires to select for each word exactly one of several encoded senses. In Section 3.2 we shall propose a *radically monosemic* approach to representing word meaning which abolishes the concept of multiple word senses (with the exception of true homonyms such as the *trunk* of a car and the *trunk* of an elephant).

Quillian also suggests that most concept definitions could be acquired algorithmically given a small set of predefined primitives and definitions written in natural language:

> "if one could manage to get a small set of basic word meanings adequately encoded and stored in computer memory, and a workable set of combination rules formalized as a computer program, he could then bootstrap his store of encoded word meanings by having the computer itself "understand" sentences that he had written to constitute the definitions of other single words" (Quillian, 1968, p.416)

It is precisely this bootstrapping process that the `dict_to_4lang` module of the `4lang` library, described in detail in Chapter 5, performs using definitions from explanatory dictionaries of English and Hungarian as well as a set of some 2,200 manually predefined concepts.

Language understanding The above model of semantic memory serves as the basis of a full-fledged language understanding system introduced in (Quillian, 1969). The process the *Teachable Language Comprehender* (TLC) applies to language understanding involves retrieving for each entity in the input text a list of concepts and entities in its memory that the text may be mentioning. For these newly created copies of concepts, the TLC also initializes pointers for each valency of the given concept: e.g. given a mention of *client*, defined as seen in Figure 2.3, pointers to employer and employee are created as such that should eventually be filled in the process of comprehending the full text. TLC then conducts for each pointer a search for compatible properties present in its current representation of the input, thus generating a list of candidates for the pointer. E.g. given the phrase *lawyer's client*, `lawyer` will eventually be found as compatible with the

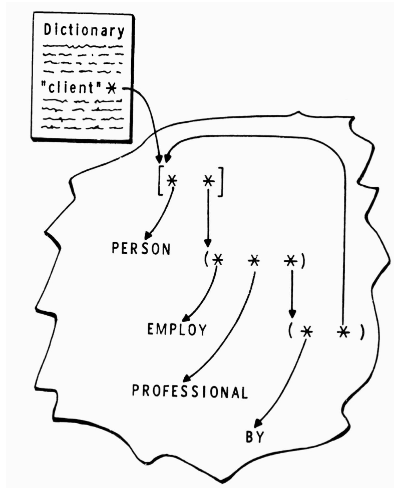

Figure 2.3: Quillian's definition of *client* (Quillian, 1969, p.462)

property `employer` of `client`, since both are linked to the property `professional`. This iterative search process also incorporates anaphora resolution: pointers may be filled with referents already present in the model of the current input. The next step involves trying to justify connections from syntax: TLC's memory also contains a set of *form tests*, each of which encode some particular configuration that is typical of a semantic relation (e.g. in this case "X's Y" or "Y of X") An example of a sample TLC session is reproduced from (Quillian, 1969) in Figure 2.4.

Note that Quillian's model is that of a *teachable* language comprehender; his account also involves feedback given by human supervisors of the process, teaching the system e.g. new form tests for each link of each concept as they occur. Such a system could be trained through human labor to make highly reliable judgments as to whether some entities in a text refer to a client and her employer. Human supervision would be necessary for practically all concepts with arguments. The framework we propose in this thesis is intended to be more robust by using more generic concept representations. The `4lang` representation of `client` may be as simple as `work` $\xleftarrow{1}$ `FOR` $\xrightarrow{2}$, but this is with the intention of leaving open as many interpretations as possible (see Section 3.2 for more discussion).

2.2.2 The KL-ONE family

The KL-ONE system (R. Brachman & Levesque, 1985) and its successors (Moser, 1983; R. J. Brachman et al., 1983) are systems for Knowledge Representation (KR) rather than models of linguistic semantics. They are of great historical significance in the field of Artificial Intelligence and their formalisms are in many ways similar to both `4lang` and phrase, *any* other property having EMPLOY as an attribute is investigated, the newly added form test will again be available, with no intervention by the monitor required. For instance, if the memory contains properties stating that agents are employed by actors and that bookkeepers are employed by companies, the form test just added will provide the syntactic capability TLC needs to comprehend input phrases such as "agent for Marlon Brando" or "accountant for Bolt Beranek and Newman"

KEY TO FIGURE 6. Numbers represent the example number. When the program is run in a more closely monitored mode, as in example 12, it prints out two lines of information each time it uses a property to help comprehend the input. This output always names what it will print out, followed by a colon, followed by the information named. The meaning of the names used are as follows:

- USING: The attribute and value of the data property it is currently using.
- ATR*: A word in the input which it has identified with the attribute of the data property.
- VAL*: A word in the input which it has identified with the value of the data property.
- SOURCE: The word of the input whose meaning provided the data property.
- PER: The form test used. Form tests always are named T1, T2, . . ., Tn. Any words *preceding* the form test name describe *how* it was used: ATRIB means it was used because the property's attribute was intersected; CKBACK means the intersection occurred during a "check back"; NESTED means the property used is a subproperty: PENDING means the property has been held pending before use.
- HEAD: The word chosen as the syntactic head of the words currently used.
- NOW-CAN-USE: This is used in place of USING if a property's use has been dependent on the use of one of its subproperties.

```
1.  READ(YOUNG CLIENT)

 ((CLIENT (AGE (YOUNG))))

NOW WE ARE TALKING ABOUT A YOUNG CLIENT.

2.  READ(THE LAWYER 'S YOUNG CLIENT)

 ((CLIENT (AGE (YOUNG))
    (EMPLOY (LAWYER)
      (BY (*THIS* . CLIENT)))))

HERE WE ARE CONCERNED WITH A YOUNG CLIENT; HE IS A CLIENT
WHO EMPLOYS A LAWYER.

3.  READ(CLIENT 'S LAWYER)

 ((LAWYER ((AOR REPRESENT ADVISE)
       (CLIENT)
       (BY (*THIS* . LAWYER))
       (IN (MATTER (TYPE LEGAL))))))

AT THIS POINT WE ARE DISCUSSING A LAWYER WHO REPRESENTS
OR ADVISES A CLIENT IN A LEGAL MATTER.

4.  READ(MAN 'S LAWYER)

 ((LAWYER ((AOR REPRESENT ADVISE)
       (MAN)
       (BY (*THIS* . LAWYER))
       (IN (MATTER (TYPE LEGAL))))))

NOW WE ARE TALKING ABOUT A LAWYER WHO REPRESENTS OR ADVISES
A MAN IN A LEGAL MATTER.

5.  READ(DOCTOR 'S LAWYER)

 ((LAWYER ((AOR REPRESENT ADVISE)
       (DOCTOR)
       (BY (*THIS* . LAWYER))
       (IN (MATTER (TYPE LEGAL))))))

HERE WE ARE CONCERNED WITH A LAWYER WHO REPRESENTS OR
ADVISES A DOCTOR IN A LEGAL MATTER.

6.  READ(LAWYER 'S DOCTOR)

 ((DOCTOR (CURE (LAWYER)
       (BY (*THIS* . DOCTOR)))))

HERE WE ARE CONCERNED WITH A DOCTOR WHO CURES A LAWYER

7.  READ(LAWYER OF THE CLIENT)

 ((LAWYER ((AOR REPRESENT ADVISE)
       (CLIENT)
       (BY (*THIS* . LAWYER))
       (IN (MATTER (TYPE LEGAL))))))

AT THIS POINT WE ARE DISCUSSING A LAWYER WHO REPRESENTS
OR ADVISES A CLIENT IN A LEGAL MATTER.

8.  READ(LAWYER 'S REPRESENT ATION OF THE CLIENT)

 ((REPRESENT ((*THIS* . REPRESENT)
       (CLIENT)
       (BY (LAWYER))
       (IN (MATTER (TYPE LEGAL))))))

NOW WE ARE TALKING ABOUT THE REPRESENTING OF A CLIENT
BY A LAWYER IN A LEGAL MATTER.

9.  READ(THE CLIENT ADVISE ED BY THE LAWYER)

 ((CLIENT ((ADVISE)
       (*THIS* . CLIENT)
       (BY (LAWYER))
       (IN (MATTER (TYPE LEGAL))))))

HERE WE ARE CONCERNED WITH A CLIENT WHO IS ADVISED BY
A LAWYER IN A LEGAL MATTER.

10.  READ(CLIENT EMPLOY S A LAWYER)

 ((TO ((*THIS* . EMPLOY)
       (LAWYER)
       (BY (CLIENT)))))

AT THIS POINT WE ARE DISCUSSING THE EMPLOYING OF A LAWYER
BY A CLIENT.

11.  READ(THE CLIENT CURE ED BY THE DOCTOR)

 (((AND CLIENT PATIENT)
     ((CURE)
       (*THIS* . CLIENT)
       (BY (DOCTOR)))))

NOW WE ARE TALKING ABOUT A CLIENT, WHO IS A PATIENT, WHO
IS CURED BY A DOCTOR.

12.  READ (THE CLIENT HEAL ED BY THE DOCTOR EMPLOY S THE
LAWYER)

USING: CURE PATIENT. ATR*: HEAL. VAL*: CLIENT
SOURCE: DOCTOR. PER: ATRIB T29. HEAD: CLIENT

USING: BY DOCTOR. VAL*: DOCTOR
SOURCE: DOCTOR. PER: NESTED T21. HEAD: CLIENT

USING: EMPLOY PROFESSIONAL. ATR*: EMPLOY. VAL*: LAWYER
SOURCE: CLIENT. PER: ATRIB CKBACK T17. HEAD: EMPLOY

USING: BY CLIENT. VAL*: CLIENT
SOURCE: CLIENT. PER: NESTED CKBACK T18. HEAD: EMPLOY

OUTPUT1:
  (EMPLOY ((*THIS* . EMPLOY)
      (LAWYER)
      (BY ((AND CLIENT PATIENT)
          ((HEAL)
            (*THIS* . CLIENT)
            (BY (DOCTOR)))))))

OUTPUT2:

 AT THIS POINT WE ARE DISCUSSING THE EMPLOYING OF A LAWYER
 BY A CLIENT, WHO IS A PATIENT, WHO IS HEALED BY A DOCTOR

13.  READ (LAWYER FOR THE CLIENT)

USING: BY LAWYER. VAL*: LAWYER
SOURCE: LAWYER. PER: NESTED T32. HEAD: LAWYER

NOW-CAN-USE: (AOR REPRESENT ADVISE) CLIENT. VAL*: CLIENT
SOURCE: LAWYER. PER: NESTED T31. HEAD: LAWYER

OUTPUT1:
 (LAWYER ((AOR REPRESENT ADVISE)
      (CLIENT)
      (BY (*THIS* . LAWYER))
      (IN (MATTER (TYPE LEGAL)))))

OUTPUT2:

 NOW WE ARE TALKING ABOUT A LAWYER WHO REPRESENTS OR ADVISES
 A CLIENT IN A LEGAL MATTER.
```

Figure 2.4: Sample session of the Teachable Language Comprehender (Quillian, 1969, p.470)

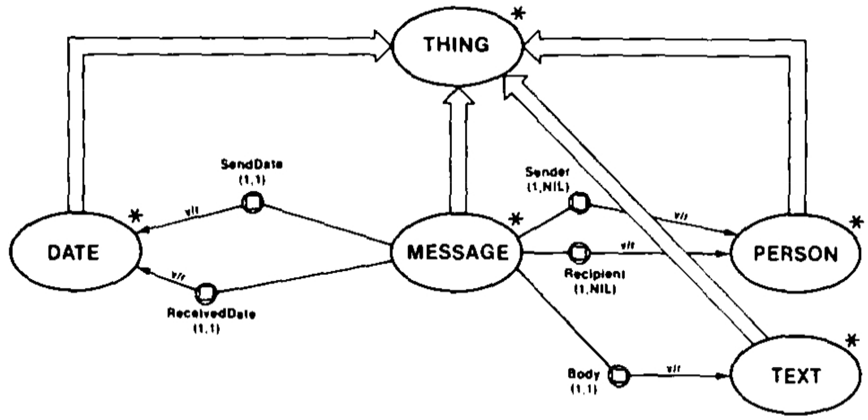

"A MESSAGE is, among other things, a THING with at least one Sender, all of which are PERSONs, at lease one Recipient, all of which are PERSONs, a Body, which is a TEXT, a SendDate, which is a DATE, and a ReceivedDate, which is a DATE."

Figure 2.5: A primitive concept in KL-ONE and its specification in JARGON (R. J. Brachman & Schmolze, 1985, p.183)

the other graph-based models mentioned in this section.

Representation Like many other approaches, KL-ONE adopts the tradition of representing information as a network of nodes and links between them. Nodes in KL-ONE networks represent *Concepts*, which are defined by three components: a list of *super-concepts*, whose properties they inherit, a list of *Roles*, describing the relationships between the concept and other concepts, and *structural descriptions*, which describe the relationship between Roles. *RoleSets* specify attributes that hold for all fillers occupying some Role, e.g. that in case of the concept `message`, the `sender` must be a `person`; such conditions are known as *Value Restrictions*. Structural Descriptions (SDs) of KL-ONE concepts serve to characterize the relationship between Roles of a Concept, e.g. that an `important message` is such that the sender is the supervisor of the recipient. A sample KL-ONE concept is depicted in Figure 2.5, along with its equivalent in JARGON, an English-like, human-readable specification language for KL-ONE.

KL-ONE explicitly forbids any violations of Value Restrictions, a clear symptom that it is a formalism for the representation of (formalized) knowledge rather than a tool for modeling language meaning directly. To account for exceptions, it is the inheritance of properties between concepts that may be defined in a way that allows for potential violations; e.g. `elephants` are defined as `four-legged-mammals`, "unless you have information to the contrary" (R. J. Brachman & Schmolze, 1985, p.190). The relationship between a concept and its super-concepts is known in KL-ONE as *subsumption*. RoleSets may enter in to a similar relationship called *restriction*, which results in the RoleSet of some concept inheriting the properties of a RoleSet of some super-concept – similar to how classes inherit functions from their superclasses in programming languages.

Semantic parsing The outline of a system mapping natural language input to KL-ONE representation is also presented in (R. J. Brachman & Schmolze, 1985). We briefly review its capabilities, since the main contribution of our thesis is also a system for mapping raw text to its meaning representation. Similar to the `text_to_4lang` system, which we describe in Chapter 4, the natural language understanding system described by Brachman and Schmolze relies on a syntactic parser (Bobrow, 1979a), the output of which is then used to build semantic representations. For the latter step, the PSI-KLONE tool is used (Bobrow, 1979b), the output of which can then serve as the input to a component responsible for handling pragmatics, bookkeeping of knowledge acquired in various contexts, etc.

The main idea behind the PSI-KLONE system is that the syntactic representation serving as its input is already encoded in a KL-ONE network, with Concepts such as `NP`, RoleSets such as `PP-modifier`, etc. The system processes a sentence by fragments received from the parser, providing feedback to it if the semantic interpretation fails and the parsing hypothesis cannot be maintained. The interpretation process itself relies on maps from words to lemmas and from lemmas to Concepts, e.g. *teaches* is mapped to the `TEACH-VERB` concept via *teach*, *professor* is mapped to `TEACHER-NOUN`, etc. Concepts retrieved this way are combined with the *syntaxonomy*, the KL-ONE network describing the relationships between syntactic units, e.g. that `VERB` is a sub-concept of `CLAUSE` which is a sub-concept of `PHRASE`. An example representation is shown in Figure 2.6.

A somewhat more recent account of PSI-KLONE (Sondheimer et al., 1984) sheds light on the next steps of semantic interpretation. *Frames* are KL-ONE concepts that describe a 'semantically distinguishable type of phrase'; e.g. the frame associated with the sending of messages is represented by the `SEND-CLAUSE` concept, whose Roles encode the selection restrictions that apply to such an event and map syntactic functions to semantic relations. For example, a `SEND-CLAUSE` must contain a `TRANSMISSION-VERB` and `MESSAGE-NOUN`, among others, and semantic restrictions on each are imposed in the form of Value Restrictions. The process of mapping a syntactic parse to a KL-ONE network is therefore directly responsible for producing semantically felicitous representations, unlike the `text_to_4lang` pipeline described in this thesis, which will produce `4lang` graphs describing any states-of-affairs based on its input. Slots of KL-ONE frames are tied to concepts via rules of the form `Paraphrase-as X`. The frame depicted in Figure 2.7 provides two example rules, stating that the indirect and direct object of a `SEND-CLAUSE` are to

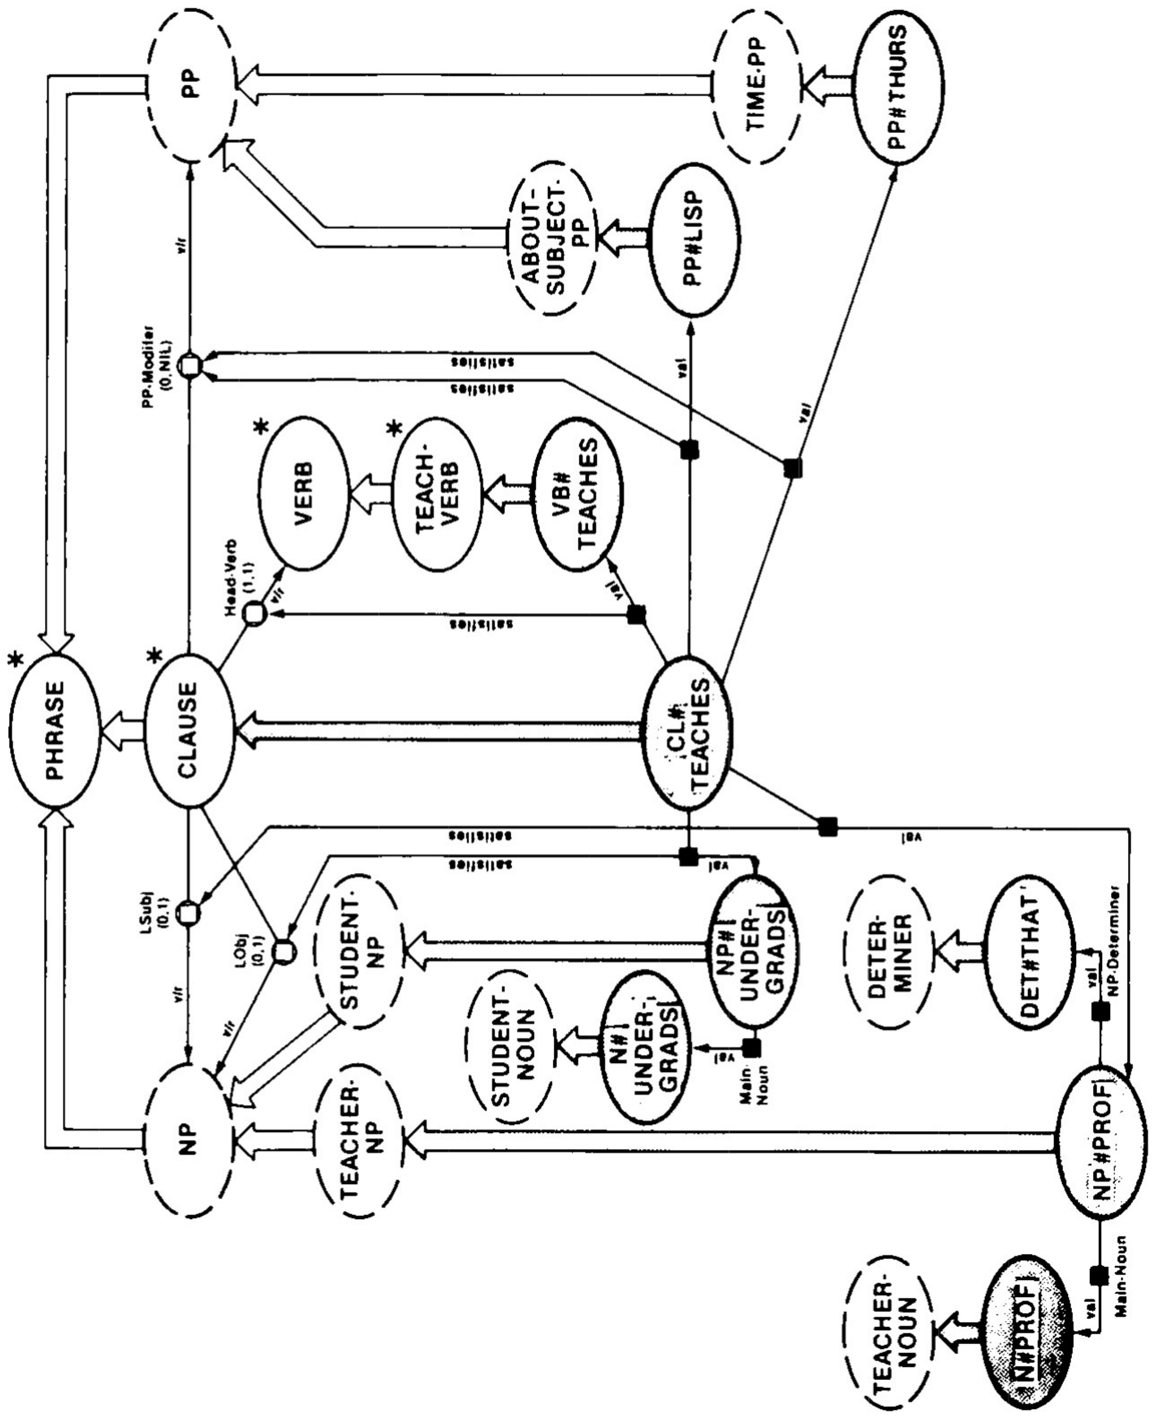

Figure 2.6: KL-ONE representation of *That professor teaches undergraduates about Lisp on Thursday* produced by `PSI-KLONE` (R. J. Brachman & Schmolze, 1985, p.214)

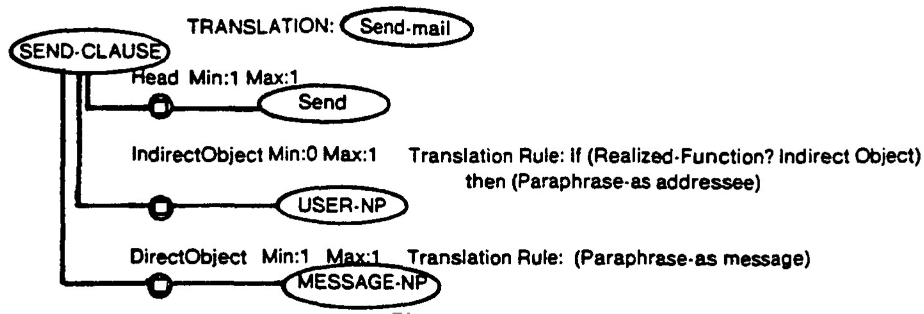

Figure 2.7: Example of a KL-ONE *frame* (Sondheimer et al., 1984, p.104)

be paraphrased as `ADDRESSEE` and `MESSAGE`, respectively. Semantic generalizations over groups of frames can be captured via common super-concepts, known as *abstract case frames*, e.g. all Concepts describing completion of an activity, such as *come*, *reach*, *finish* or *arrive*, can be grouped under an abstract frame from which they inherit the potential to accept time-modifiers.

2.2.3 Abstract Meaning Representations

Abstract Meaning Representation, or AMR (Banarescu et al., 2013), is a more recent formalism for representing the meaning of linguistic structures as directed graphs. The last few years have seen a rise in AMR-related work, including a corpus of AMR-annotated text (Banarescu et al., 2013), several approaches to generating AMRs from running text (Vanderwende et al., 2015; Peng et al., 2015; Pust et al., 2015), and various applications to computational semantics (Pan et al., 2015; Liu et al., 2015).

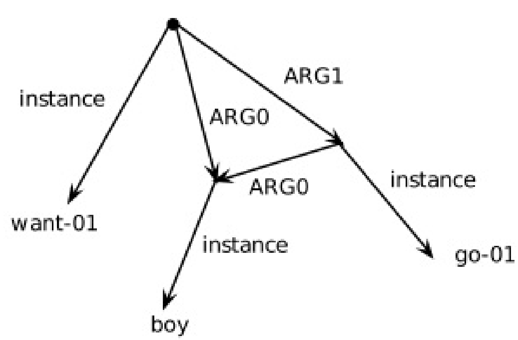

Figure 2.8: `AMR` representation of *The boy wants to go* (Banarescu et al., 2013, p.179)

Nodes of AMR graphs represent concepts of two basic types: they are either English words, or *framesets* from PropBank (Palmer et al., 2005), used to abstract away from English syntax. PropBank framesets are essentially English verbs (or verb-particle con-

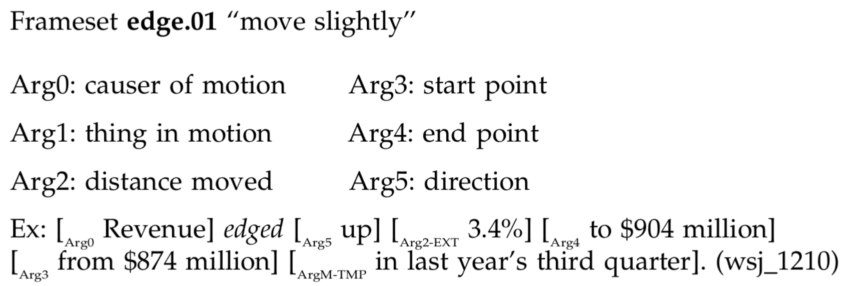

Frameset **edge.01** "move slightly"

Arg0: causer of motion	Arg3: start point
Arg1: thing in motion	Arg4: end point
Arg2: distance moved	Arg5: direction

Ex: [$_{\text{Arg0}}$ Revenue] *edged* [$_{\text{Arg5}}$ up] [$_{\text{Arg2-EXT}}$ 3.4%] [$_{\text{Arg4}}$ to $904 million] [$_{\text{Arg3}}$ from $874 million] [$_{\text{ArgM-TMP}}$ in last year's third quarter]. (wsj_1210)

Figure 2.9: A PropBank frameset (Palmer et al., 2005, p.76)

structions) with a list of possible arguments along with their semantic roles; an example frameset can be seen in Figure 2.9. Unlike the `4lang` representation used in this thesis, AMR also makes a distinction similar to Quillian's type and token nodes by separating nodes that represent some entity, event, property, etc. from nodes that are arguments of some frameset, linking the latter with an *instance* relation to the former. The AMR representation of the sentence *The boy wants to go* would hence be that in Figure 2.8 as opposed to the `4lang` representation in Figure 2.10. AMRs also handles a wide range of phenomena that `4lang` currently doesn't: the formalism provides relations to encode negation, modals, copulars, and questions. It also includes special relations to encode named entities – in the broader sense, i.e. including not only proper names but also e.g. dates, quantities, etc. The formalism accommodates a wide range of phenomena typical of English, AMR creators admit that "AMR is heavily biased towards English. It is not an Interlingua." (Banarescu et al., 2013, p.179).

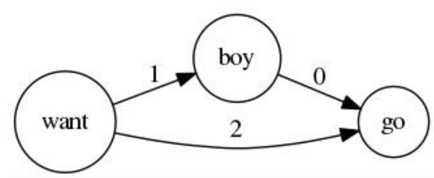

Figure 2.10: `4lang` representation of *The boy wants to go*

2.3 Montague-style theories

A considerable amount of the literature on the semantics of natural language has in the past few decades focused on *Montagovian* representations of meaning (Montague, 1970a, 1970b, 1973; Kamp, 1981; Groenendijk & Stokhof, 1991). The shared agenda of these approaches is to provide a mapping from linguistic structures to logical formulae; the bulk of actual work is concerned with handling particular portions of syntax. Nearly all such accounts take Montague's original treatment of word meaning for a given. It has been shown that at least 84 percent of the information content of an average utterance is encoded by word meaning (Kornai, 2012), yet most proposed interpretations of sentences such as *Every man loves a woman such that she loves him* rarely have anything to say about the concepts `man, woman`, or `love`. There are some generic principles of how word meaning should be represented in logical formulae: nouns like *man* are typically thought of as functions that decide for all objects of the world whether they are men or not, verbs like *love* are thought of as describing an event such that for any event in the world one can decide whether an act of loving has taken place. Such principles have little practical value, however, when linking particular utterances to states-of-affairs. To our knowledge, no lexicon with a substantive list of meaning postulates has ever benn built. In Chapter 5 we shall construct `4lang`-style meaning representations for all headwords of monolingual dictionaries of English.

If common nouns like *giraffe* and adjectives like *blue* are both seen as selecting a subset of all objects in the world, then an NP such as *blue giraffe* might map to the intersection of these subsets. The same mechanism fails for *enormous fleas*: the representation of *enormous* must be updated to accommodate the fact that you cannot tell if some size is enormous unless you know whose size it is (e.g. half an inch is enormous for a flea but tiny for a giraffe). Clearly there does not exist a function that selects a universal set of enormous fleas – what constitutes large may depend e.g. on the speaker's previous experience. Yet if we are to account for the fact that people can use this phrase successfully in conversations, we must map *enormous* to some function that might take as its parameter not only an entry encoding shared beliefs of speakers about defining properties of fleas, but also some information regarding their beliefs of the size of fleas. It is tempting to handle such a phenomenon by simply defining the interfaces with *extra-linguistic* knowledge, after which the meaning of *small blue giraffe* can be a formula with parameters for speakers' knowledge of what size range counts as small for a giraffe, what shades of color counts as blue, perhaps even what set of characteristics would make something/somebody a giraffe. Travis (1997) describes this approach in *A Companion to the Philosophy of Language*:

> What some words say, or contribute to what is said in using them, varies across speakings of them. Where this is so, the meaning of the words does two things. First, it determines on just what facts about a speaking the semantic contribution of the words so spoken depends. Second, it determines just how their semantics on a speaking depends on these facts. Specifically, it determines a specifiable function from values of those factors to the semantics the words would have, if spoken where those values obtain. (Travis, 1997, p.92)

Proponents of Montagovian theories of semantics may claim that the subject of their study (*meaning* in a narrow sense) is the component of the effect an utterance has on the information state of speakers that is unchanged across "speakings". Nevertheless, such a representation of e.g. *small blue giraffe* must contain information about the meaning of each of the individual concepts `small, blue,` and `giraffe`. It is one thing to disown the issue of inter-speaker variation on which colors are blue, what sizes of giraffes are small, etc., but surely what makes the phrase more informative than e.g. *small blue animal* is that the variation among all giraffes is considerably smaller than the variation among all animals. That MG accounts of semantics do not decompose the meaning of content words is problematic because we have seen that to construct the meaning of even the simplest kinds of phrases, one needs to account for how their meanings interact. Any mechanism with a chance to interpret *small giraffe* or *young giraffe* will have to make reference to a particular set of components of the meaning of `giraffe`, otherwise we cannot make predictions about the size or age of the giraffe as we would about some unknown X given the phrases *small X* and *young X*. The necessity of decomposing word meaning has already been argued for by (Katz & Fodor, 1963), but the actual use of meaning postulates in MG remains restricted to the resolution of technical problems caused by handling intensionality; for a survey, see (Zimmermann, 1999). In Chapter 3 we shall present a theory of meaning representation that encodes word meaning as a network of concepts, making them accessible to mechanisms responsible for constructing the meaning of larger structures.

2.4 CVS representations

The most widely used models of word meaning today are continuous vector spaces (CVS). State-of-the-art systems in most standard NLP tasks rely on *word embeddings*, mappings from words of a language to real-valued vectors, trained on datasets containing 10^6-10^{10} words. In this section we review key aspects of CVS semantics, which set the current standard for representing word meaning (cf. Section 2.4.1). Remarkably, they do so using elements of representations that – unlike `4lang` representations – do not lend themselves to compositionality in any obvious way (cf. Section 2.4.2).

2.4.1 Vectors as word representations

Methods used to obtain mappings from words to vectors are based on the *distributional hypothesis* (Harris, 1954), which states that words are similar if they appear in similar contexts. When training word embeddings on large bodies of unannotated text, the most commonly used algorithms (Mikolov, Chen, et al., 2013; Pennington et al., 2014) will take into account all contexts the word has occurred in (typically some fixed-size sequence of surrounding words) and attempt to find vectors for each word that minimizes the difference between the predicted and observed probability of the word appearing in those contexts. Embeddings trained this way can be evaluated by using them as the initial layers of neural network models trained for a variety of NLP tasks such as named entity recognition, chunking, POS-tagging, etc. (Collobert & Weston, 2008; Turian et al., 2010). Word vectors are also often measured by their direct applicability to particular tasks such as answering word analogy questions (Mikolov, Yih, & Geoffrey, 2013) or finding missing words in text (Zweig et al., 2012). Analogical questions such as "man is to woman as king is to X" can be answered successfully by taking the vectors associated with each word ($\vec{m}$, $\vec{w}$, $\vec{k}$ for man, woman, and king, respectively) and finding the word whose vector has the greatest cosine similarity to $\vec{k} + \vec{w} - \vec{m}$. The fact that this strategy is relatively successful indicates that the relational hypothesis holds to some extent: word representations trained based on distribution are at least implicitly related to word meaning, making them candidates for use in computational semantics systems. Indeed, word embeddings have been used successfully in state of the art systems for e.g. Semantic Role Labeling (Foland Jr & Martin, 2015), Knowledge Base Construction (Nickel et al., 2015), and Semantic Textual Similarity (Han et al., 2015). Vector representations are also practical for establishing a connection between linguistic and non-linguistic data, a striking indication is the work presented in (Karpathy et al., 2014), mapping text fragments to pictures for information retrieval (image search).

2.4.2 Vectors beyond the word level

In this section we mention only a few examples that are relevant to our thesis. For a generic overview of compositionality in CVS semantics, the reader is referred to Section 2 of (Grefenstette & Sadrzadeh, 2015). An example of training vectors that represent linguistic units larger than a single word is the Compositional Vector Grammar (CVG) parser introduced in (Socher, Bauer, et al., 2013), which outperforms by a significant margin the state of the art in syntactic parsing by combining the standard PCFG approach with recursive neural networks (RNNs) trained on each layer of a parse tree, assigning vectors not only to words but all nonterminals of the grammar. The `text_to_4lang` system introduced in Chapter 4 relies on CVGs for syntactic parsing, therefore we now provide a very brief overview of them as presented in (Socher, Bauer, et al., 2013).

PCFG parsers such as that implemented by the Stanford Parser will return for some

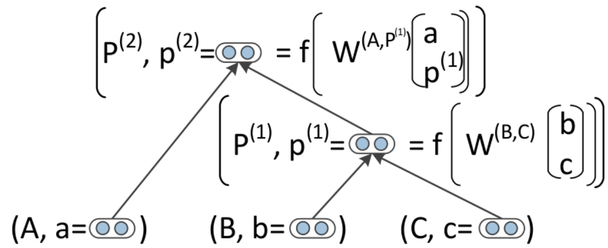

Figure 2.11: Example of a syntactically untied RNN (Socher, Bauer, et al., 2013, p.459)

input sentence a ranked list of candidate parses. If a grammar is able to generate the correct parse tree for nearly all sentences, i.e. the correct parse can be expected to be among the candidates returned for some sentence, then increasing parsing accuracy amounts to improving the component responsible for ranking candidates based on their likelihood. CVGs combine the power of PCFGs and RNNs by devising a method to rerank parse trees in the output of a standard PCFG parser using neural networks trained on a treebank. The core idea is that in calculating the score of a given syntactic derivation (parse tree) for a sentence, the likelihood of each derivation step should be assigned based on not only the observed frequency of the given structure, but rather its likeliness to cover the particular sequence of words, and that this calculation should factor in word forms via a distributional model, approximating the properties of rare or unseen words using more frequent ones that appear in similar contexts. *Syntactically untied* networks (SU-RNNs) learn separate parameters for each rewrite rule. The parameters for a rule of the form $A \rightarrow BC$ are encoded by the *syntactic triplet* $((A,a),(B,b),(C,c))$, where b and c are vectors of R^n assigned to the non-terminals B and C, respectively, and A is computed as $f(W^{(B,C)}([b,c]))$, where $[b,c]$ is a vector in R^{2n} obtained by concatenating b and c, and $W^{B,C}$ is a matrix in $R^{n\times 2n}$ which is learnt during the training process. f is the element-wise nonlinearity function tanh. The process is summarized in Figure 2.11.

Compositionality of word vectors has also been explored in the context of Sentiment Analysis (Socher, Perelygin, et al., 2013; Zhu et al., 2015) and Semantic Textual Similarity (Sultan et al., 2015). The latter work assigns vectors to sentences by calculating the componentwise average of all word vectors, Socher, Perelygin, et al. (2013) use *Recursive Neural Tensor Networks* (RNTNs) to obtain vectors for each node in the parse tree of a sentence.

Chapter 3

The 4lang system

This chapter describes the `4lang` system for representing meaning using directed graphs of concepts. Since the underlying theory is not the main contribution of this thesis, but rather the work of half a dozen researchers over the course of 6 years, we shall not attempt a full presentation of the `4lang` principles. Instead we shall introduce the formalism in Section 3.1, then continue to discuss some specific aspects relevant to this thesis. `4lang`'s approach to multiple word senses is summarized in Section 3.2, Section 3.3 is concerned with reasoning based on `4lang` graphs. Treatment of extra-linguistic knowledge is discussed in Section 3.4. Finally, Section 3.5 considers the primitives of the `4lang` representation and contrasts them with some earlier approaches mentioned in Chapter 2.

For a complete presentation of the theory of lexical semantics underlying `4lang` the reader is referred to (Kornai, 2010) and (Kornai, 2012). (Kornai et al., 2015) compares `4lang` to contemporary theories of word meaning. `4lang` is also the name of a manually built dictionary[1] mapping 2,200 English words to concept graphs (as well as their translations in Hungarian, Polish, and Latin, hence its name). The dictionary is described in (Kornai & Makrai, 2013). For work on extending `4lang` to include the top 40 languages (by Wikipedia size), see (Ács et al., 2013).

3.1 The formalism

4lang represents the meaning of words, phrases and utterances as directed graphs whose nodes correspond to language-independent concepts and whose edges may have one of three labels, based on which they'll be referred to as 0-edges, 1-edges, and 2-edges. (The `4lang` theory represents concepts as Eilenberg-machines (Eilenberg, 1974) with three *partitions*, each of which may contain zero or more pointers to other machines and therefore also

[1] https://github.com/kornai/4lang/blob/master/4lang

represent a directed graph with three types of edges. The additional capabilities offered by Eilenberg-machines have not so far been applied by the author, some of them have not even been implemented yet, therefore it makes more sense to consider the representations under discussion as plain directed graphs.) First we shall discuss the nature of `4lang` *concepts* - represented by the nodes of the graph, then we'll introduce the types of relationships encoded by each of the three edge types.

3.1.1 Nodes

Nodes of `4lang` graphs correspond to *concepts*. `4lang` concepts are not words, nor do they have any grammatical attributes such as part-of-speech (category), number, tense, mood, voice, etc. For example, `4lang` representations make no difference between the meaning of *freeze* (N), *freeze* (V), *freezing*, or *frozen*. Therefore, the mapping between words of some language and the language-independent set of `4lang` concepts is a many-to-one relation. In particular, many concepts will be defined by a single link to another concept that is its hypernym or synonym, e.g. $\texttt{above} \xrightarrow{0} \texttt{up}$ or $\texttt{grasp} \xrightarrow{0} \texttt{catch}$. Encyclopedic information is omitted, e.g. `Canada`, `Denmark`, and `Egypt` are all defined as `country` (their definitions also containing a pointer to an external resource, typically to Wikipedia). In general, definitions are limited to what can be considered the shared knowledge of competent speakers – e.g. the definition of `water` contains the information that it is a colorless, tasteless, odorless liquid, but not that it is made up of hydrogen and oxygen. We shall now go through the types of links used in `4lang` graphs.

3.1.2 The 0-edge

The most common relation between concepts in `4lang` graphs is the 0-edge, which represents attribution ($\texttt{dog} \xrightarrow{0} \texttt{friendly}$); the `IS_A` relation (hypernymy) ($\texttt{dog} \xrightarrow{0} \texttt{animal}$); and unary predication ($\texttt{dog} \xrightarrow{0} \texttt{bark}$). Since concepts do not have grammatical categories, this uniform treatment means that the same graph can be used to encode the meaning of phrases like *water freezes* and *frozen water*, both of which would be represented as $\texttt{water} \xrightarrow{0} \texttt{freeze}$.

3.1.3 1- and 2-edges

Edge types 1 and 2 connect binary predicates to their arguments, e.g. $\texttt{cat} \xleftarrow{1} \texttt{catch} \xrightarrow{2} \texttt{mouse}$). The formalism used in the `4lang` dictionary explicitly marks binary (transitive) elements – by using UPPERCASE printnames. The pipeline that we'll introduce in Chapter 4 will

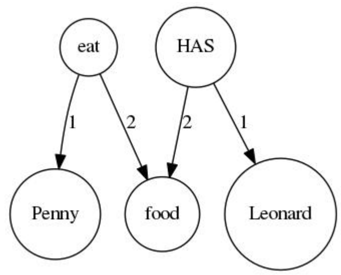

Figure 3.1: `4lang` graph with two types of binaries.

HAS	`shirt` $\xleftarrow{1}$ `HAS` $\xrightarrow{2}$ `collar`
IN	`letter` $\xleftarrow{1}$ `IN` $\xrightarrow{2}$ `envelope`
AT	`bridge` $\xleftarrow{1}$ `AT` $\xrightarrow{2}$ `river`
CAUSE	`humor` $\xleftarrow{1}$ `CAUSE` $\xrightarrow{2}$ `laugh`
INSTRUMENT	`sew` $\xleftarrow{1}$ `INSTRUMENT` $\xrightarrow{2}$ `needle`
PART_OF	`leaf` $\xleftarrow{1}$ `PART_OF` $\xrightarrow{2}$ `plant`
ON	`smile` $\xleftarrow{1}$ `ON` $\xrightarrow{2}$ `face`
ER	`slow` $\xleftarrow{1}$ `ER` $\xrightarrow{2}$ `speed`
FOLLOW	`Friday` $\xleftarrow{1}$ `FOLLOW` $\xrightarrow{2}$ `Thursday`
MAKE	`bee` $\xleftarrow{1}$ `MAKE` $\xrightarrow{2}$ `honey`

Table 3.1: Most common binaries in the `4lang` dictionary

not make use of this distinction, any concept can have outgoing 1- and 2-edges. Binaries marked with uppercase are nevertheless clearly set apart from other concepts by the fact that they are *necessarily* binary, i.e. they must always have exactly two outgoing edges. We retain the uppercase marking for those binary elements that do not correspond to any word in a given phrase or sentence, e.g. the meaning of the sentence *Penny ate Leonard's food* will be represented by the graph in Figure 3.1[2]. The top ten most common binaries used in `4lang` are listed in Table 3.1 and examples are shown for each.

Given two concepts c_1 and c_2 such that c_2 is a predicate that holds for c_1, `4lang` will allow for one of two possible connections between them: $c_1 \xrightarrow{0} c_2$ if c_2 is a one-place predicate and $c_2 \xrightarrow{1} c_1$ if c_2 is a two-place predicate. The mutual exclusiveness of these two configurations is both counter-intuitive and unpractical for the `4lang`-based

[2]Evidence for different patterns of linking predicates and their arguments could be obtained from ergative languages (?, ?), these shall not be discussed here.

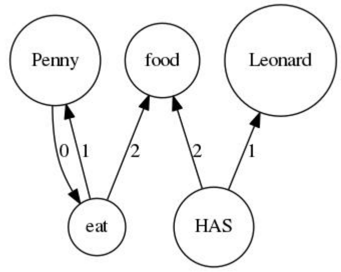

Figure 3.2: Revised `4lang` graph with two types of binaries.

systems presented in this thesis. Two-place predicates often appear with a single argument (e.g. *John is eating*), and representing such a statement as $\texttt{John} \xrightarrow{0} \texttt{eat}$ while the sentence *John is eating a muffin* warrants $\texttt{John} \xleftarrow{1} \texttt{eat} \xrightarrow{2} \texttt{muffin}$ would mean that we consider the relationship between `John` and `eat` dependent on whether we have established the object of his eating. Therefore we choose to adopt a modified version of the `4lang` representation where the 0-connection holds between a subject and predicate regardless of whether the predicate has another argument. The example graph in Figure 3.1 can then be revised to obtain that in Figure 3.2[3].

The meaning of each `4lang` concept is represented as a `4lang` graph over other concepts – a typical definition in the `4lang` dictionary can be seen in Figure 3.3; this graph captures the facts that birds are vertebrates, that they lay eggs, and that they have feathers and wings. The generic applicability of the `4lang` relations introduced in Section 3.1 have the consequence that to create, understand, and manipulate `4lang` representations one need not make the traditional distinction between entities, properties, and events. The relationships $\texttt{dog} \xrightarrow{0} \texttt{bark}$ and $\texttt{dog} \xrightarrow{0} \texttt{inferior}$ (Kornai, in preparation) can be treated in a uniform fashion, when making inferences based on the definitions of each concept, e.g. that $\texttt{dog} \xleftarrow{1} \texttt{MAKE} \xrightarrow{2} \texttt{sound}$ or that calling another person a *dog* is insulting.

3.2 Ambiguity and compositionality

`4lang` does not allow for multiple senses when representing word meaning, all occurrences of the same word form – with the exception of true homonyms like *trunk* 'the very long

[3] Since the `text_to_4lang` pipeline presented in Chapter 4 assigns `4lang` graphs to raw text based on the output of dependency parsers that treat uniformly the relationship between a subject and verb irrespective of whether the verb is transitive or not, the `4lang` graphs we build will include a 1-edge between all verbs and their subjects. We do not consider this a shortcoming: for the purposes of semantic analysis we do not see the practicality of a distinction between transitive and intransitive verbs – we only recognize the difference between the likelihood (based on data) of some verb taking a certain number of arguments.

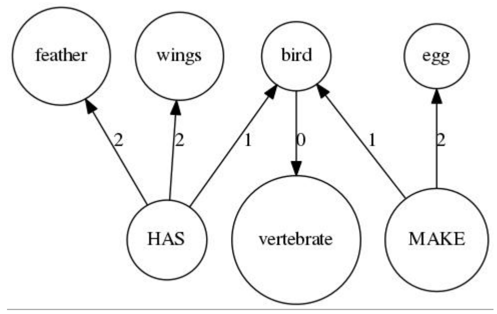

Figure 3.3: 4lang definition of `bird`.

nose of an elephant' and *trunk* 'the part at the back of a car where you can put bags, tools etc'[4] – must be mapped to the same concept, whose definition in turn must be generic enough to allow for all possible uses of the word. As Jakobson famously noted, such a monosemic approach might define the word *bachelor* as 'unfulfilled in typical male role' (Fillmore, 1977). Such definitions place a great burden on the process responsible for combining the meaning of words to create representations of phrases and utterances (see Chapter 4), but it has the potential to model the flexibility and creativity of language use:

> "we note here a significant advantage of the monosemic approach, namely that it makes interesting predictions about novel usage, while the predictions of the polysemic approach border on the trivial. To stay with the example, it is possible to envision novel usage of *bachelor* to denote a contestant in a game who wins by default (because no opponent could be found in the same weight class or the opponent was a no-show). The polysemic theory would predict that not just seals but maybe also penguins without a mate may be termed bachelor – true but not very revealing."(Kornai, 2010, p.182)

One typical consequence of this approach is that `4lang` definitions will not distinguish between `bachelor` and some concept `w` that means 'unfulfilled male' – both could be defined in `4lang` as e.g. `male, LACK`. This is not a shortcoming of the representation, rather it is in accordance with the principles underlying it; the concepts `unfulfilled` and `male` cannot be combined (e.g. to create a representation describing an *unfulfilled male*) without making reference to some nodes of the graph representing the meaning of `male`; if something is a 'typical male role', this should be indicated in the definition graph of `male` (if only by inbound pointers), and without any such information, *unfulfilled male* cannot be interpreted at all.

[4]All example definitions, unless otherwise indicated, are taken from the Longman Dictionary of Contemporary English (Bullon, 2003)

This does not mean that `male` cannot be defined without listing all stereotypes associated with the concept. If the piece of information that 'being with a mate at breeding time' is a typical male role – which is necessary to account for the interpretation of *bachelor* as 'young fur seal when without a mate at breeding time' – is to be accessed by some inference mechanism, then it must be present in the form of some subgraph containing the nodes `seal`, `mate`, `male`, and possibly others. Then, a `4lang`-based natural language understanding system that is presented with the word *bachelor* in the context of mating seals for the first time may explore the neighborhood of these nodes until it finds this piece of information as the only one that 'makes sense' of this novel use of *bachelor*. Note that this is a model of novel language use in general. Humans produce and understand without much difficulty novel phrases that most theories would label 'semantically anomalous'. In particular, all language use that is commonly labeled *metaphoric* involves accessing a lexical element for the purpose of activating some of its meaning components, while ignoring others completely. It is this use of language that `4lang` wishes to model, as it is most typical of everyday communication (Richards, 1937; Wilks, 1978; Hobbs, 1990).

Another `4lang` principle that ensures metaphoric interpretation is that any link in a `4lang` definition can be overridden. In fact, the only type of negation used in `4lang` definitions, `LACK`, carries the potential to override elements that might otherwise be activated when definitions are expanded: e.g. the definition of `penguin`, which undoubtedly contains $\xrightarrow{0}$ `bird`, may also contain $\xleftarrow{1}$ `LACK` $\xrightarrow{2}$ `fly` to block inference based on `bird` $\xrightarrow{0}$ `fly`. That any element can freely be overridden ensures that novel language use does not necessarily cause contradiction: "[T]o handle 'the ship plowed through the sea', one lifts the restriction on 'plow' that the medium be earth and keeps the property that the motion is in a substantially straight line through some medium" (Hobbs, 1990, p.55). Since a `4lang` definition of `plow` must contain some version of $\xrightarrow{2}$ `earth`, there must be a mechanism allowing to override it and not make inferences such as `sea` $\xrightarrow{0}$ `earth`[5].

3.3 Reasoning

The `4lang` principles summarized so far place a considerable burden on the inferencing mechanism. Given the possibility of defining all concepts using only a small set of primitives, and a formalism that strictly limits the variety of connections between concepts, we claim to have laid the groundwork for a semantic engine with the chance of understanding

[5]Note that such an inference must access some form of world knowledge in addition to the definition of each concept: the definition of `ship` will contain $\xleftarrow{1}$ `ON` $\xrightarrow{2}$ `water` (or similar), but to infer that this makes it incompatible with the `earth` in the definition of `plow` one must also be aware that water and earth cancel each other out in the context of where a vehicle runs

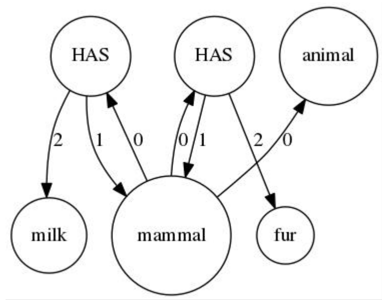

Figure 3.4: 4lang definition of `mammal`.

creative language use. Generic reasoning has not yet been implemented in `4lang`, we only present early attempts in Section 5.3 and some specific applications in Chapter 6. Here we shall simply outline what we believe could be the main mechanisms of such a system.

The simplest kind of lexical inference in `4lang` graphs is performed by following paths of 0-edges from some concept to determine the relationships in which it takes part. The concept `mammal` is defined in `4lang` as an animal that has fur and milk (see Figure 3.4), from which one can conclude that the relations $\xleftarrow{1}$ `HAS` $\xrightarrow{2}$ `milk` and $\xleftarrow{1}$ `HAS` $\xrightarrow{2}$ `fur` also hold for all concepts whose definition includes $\xrightarrow{0}$ `mammal` (we shall assume that this simple inference can be made when we construct `4lang` definitions from dictionary definitions in Chapter 5). Similar inferences can be made after *expanding* definitions, i.e. replacing concept nodes with their definition graphs (see Section 5.3 for details). If the definition of `giraffe` contains $\xrightarrow{0}$ `mammal`, to which we add edges $\xleftarrow{1}$ `HAS` $\xrightarrow{2}$ `fur` and $\xleftarrow{1}$ `HAS` $\xrightarrow{2}$ `milk`, this expanded graph will allow us to infer the relations `giraffe` $\xleftarrow{1}$ `HAS` $\xrightarrow{2}$ `fur` and `giraffe` $\xleftarrow{1}$ `HAS` $\xrightarrow{2}$ `milk`. As mentioned in the previous section, this process requires that relations present explicitly in a definition override those obtained by inference: penguins are birds and yet they cannot fly, humans are mammals without fur, etc.

A more complicated procedure is necessary to detect connections between nodes of an expanded definition and nodes connected to the original concept. Recall Quillian's example in Section 2.2.1: given the phrase *lawyer's client* his iterative search process will eventually find `lawyer` to be compatible with the `employer` property of `client`, since both are `professionals`. A similar process can be implemented for `4lang` graphs; consider the definition graphs for `lawyer` and `client` in Figures 3.5 and 3.6, built automatically from definitions in the Longman dictionary, as described in Chapter 5, then pruned manually. (These graphs, being the output of the `dict_to_4lang` system and not manual annotation, have numerous issues: the word *people* in the Longman dictionary definition of `lawyer` was not mapped to `person`, nor have the words *advice* and *advise* been mapped to the same concept. After correcting these errors manually, nodes with identical names in the

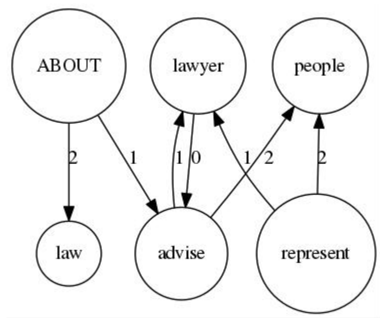

Figure 3.5: Definition graph for `lawyer`

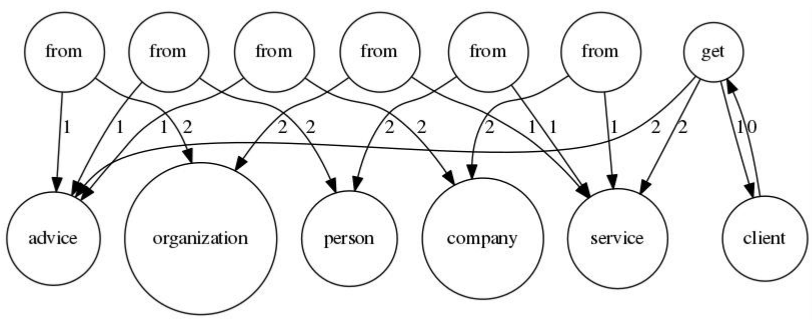

Figure 3.6: Definition graph for `client`

graph for *lawyer's client* (Figure 3.7) can form the starting point of the inference process. Let us now go over the various steps of inference necessary to reduce this graph to the most informative representation of *lawyer's client*. Note that we do not wish to impose any logical order on these steps; they should rather be the 'winners' of a process that considers many transformations in parallel and ends up keeping only some of them. A simple example of such a system will be described in Section 6.2.

We should be able to realize that the `person` who is `advice`d (and is `represent`ed by) the `lawyer` can be the same as the `client` who `get`s `advice` from the lawyer. To this end we must be able to make the inference that $\texttt{X} \xleftarrow{1} \texttt{get} \xrightarrow{2} \texttt{advice}$ and $\texttt{advice} \xrightarrow{2} \texttt{X}$ are synonymous. We believe a `4lang`-based system should be able to make such an inference in at least one of two independent ways. First, we expect our inference mechanism to compute, based on the definitions of `get` and `advice`, that $\texttt{X} \xleftarrow{1} \texttt{get} \xrightarrow{2} \texttt{advice}$ entails $\texttt{advice} \xrightarrow{2} \texttt{X}$ (and vice versa). Secondly, we'd like to be able to accommodate *constructions* in the `4lang` system (see also Section 8.4) that may explicitly pair the above two configurations for some concepts but not for others (e.g. $\texttt{X} \xleftarrow{1} \texttt{get} \xrightarrow{2} \texttt{drink}$ should not trigger $\texttt{drink} \xrightarrow{2} \texttt{X}$).

We should also consider unifying the `person` node in $\texttt{person} \xleftarrow{1} \texttt{from} \xrightarrow{2} \texttt{advice}$ with `lawyer` in $\texttt{advice} \xrightarrow{1} \texttt{lawyer}$, which would once again require either some construction

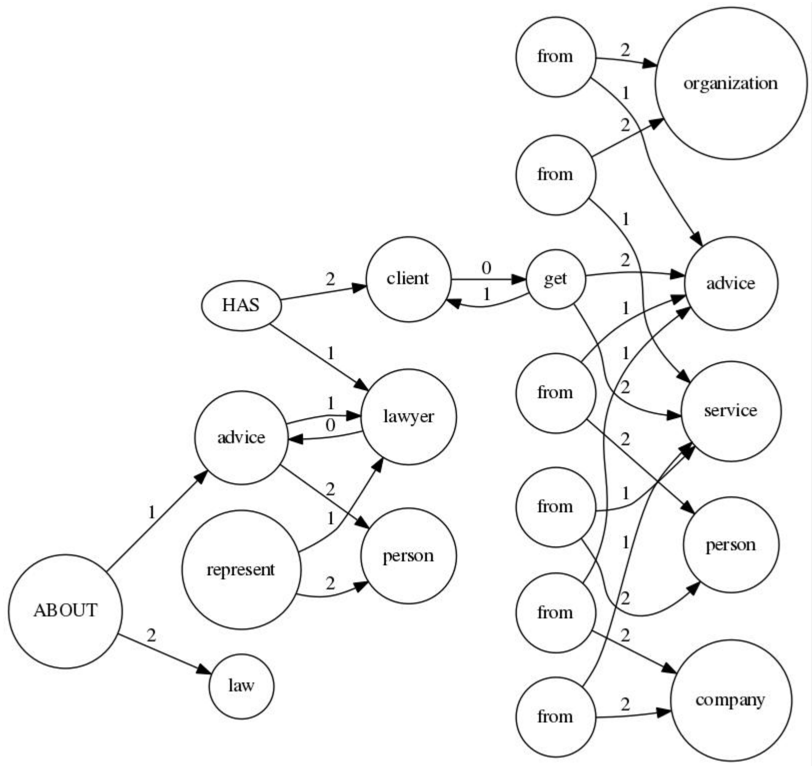

Figure 3.7: Corrected graph for *lawyer's client*

that states that when someone *advises*, then the *advice* is *from* her, or a generic rule that can guess the same connection. Given these inferences, the two `advice` can also be merged as likely referring to the same action, resulting in the final graph in Figure 3.8. The nodes `organization`, `company`, and `service` have been omitted from the figure to improve readability.

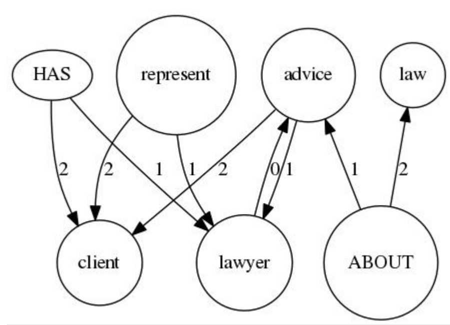

Figure 3.8: Inferred graph for *lawyer's client*

3.4 Extra-linguistic knowledge

Chapter 3 of (Kornai, in preparation) argues that knowledge representation for the purposes of natural language understanding requires a distinction between analytic and synthetic knowledge, and that the `4lang` theory is adequate to represent all analytic knowledge. When we discuss inference in terms of `4lang` representations, we only make reference to knowledge that is clearly within the boundaries of the naive theories described by Kornai. We emphasize that we do not even need to establish any particular piece of knowledge as essential to our inferencing capabilities, just as in mathematics, where we do not need to establish the truth of the axioms. Returning to one of the simplest examples above, where $\texttt{bird} \xrightarrow{0} \texttt{fly}$ is overridden to accommodate both $\texttt{penguin} \xleftarrow{1} \texttt{LACK} \xrightarrow{2} \texttt{fly}$ and $\texttt{penguin} \xrightarrow{0} \texttt{bird}$, we need not decide whether the particular piece of information that penguins cannot fly is part of the meaning of `penguin`. Clearly it is possible for one to learn of the existence of penguins and that they are a type of bird without realizing that they cannot fly, and this person could easily make the incorrect *inference* that they can. Some components of word meaning, on the other hand, appear to be essential to the understanding of a particular concept, e.g. if a learner of English believes that *nephew* refers to the child of one's sibling, male or female (perhaps because in her native language a single word stands for both nephews and nieces, and because she has heard no contradicting examples), we say that she does not know the meaning of the word; $\texttt{nephew} \xrightarrow{0} \texttt{male}$ is internal to the concept `nephew` in a way that $\texttt{penguin} \xleftarrow{1} \texttt{LACK} \xrightarrow{2} \texttt{fly}$ is to `penguin`. This distinction is commonly made in semantics under the heading analytic vs. synthetic knowledge, but imperfections in acquiring analytic knowledge are common and a normal part of the language acquisition process. Carrying a conversation successfully only requires that the participants' representations of word meaning does not contradict each other in a way *relevant to the conversation at hand*[6]. Static lexical resources such as LDOCE or the `4lang` concept dictionary must make decisions about which pieces of information to include, and may do so based on some notion of how 'technical' or 'commonplace' they are. A person's ignorance of the fact that somebody's nephew is necessarily male is probably itself the result of one or several conversations about nephews that somehow remained consistent despite his incomplete knowledge about how the word is typically used.

[6]This is also reflected in The Urban Dictionary's definition of *semantics: The study of discussing the meaning/interpretation of words or groups of words within a certain context; usually in order to win some form of argument* (http://www.urbandictionary.com)

3.5 Primitives of representation

In the following two chapters this thesis will present methods for 1) building `4lang` representations from raw text and 2) building `4lang` definition graphs for virtually all words based on monolingual dictionaries. Given these two applications, any text can be mapped to `4lang` graphs and nodes of any graph can be expanded to include their `4lang` definitions. Performing this expansion iteratively, all representations can be traced back to a small set of concepts; in case the Longman Dictionary is used to build definition graphs, the concepts listed in the `4lang` dictionary will suffice to cover all of them, since it contains all words of the Longman Defining Vocabulary (LDV), the set of all words used in definitions of the Longman Dictionary (Boguraev & Briscoe, 1989). The set of concepts necessary to define all others can be further reduced: we show in (Kornai et al., 2015) that as few as 129 `4lang` concepts are enough to define all others in the `4lang` dictionary, and thus, via monolingual dictionaries, practically all words in the English language.

In response to Katz and Fodor's markers and distinguishers (see Section 2.1), Bolinger (1965) argues that any component of word meaning that Katz and Fodor may consider to belong to the domain of distinguishers, and as such out of grasp for a semantic theory, can be further decomposed into markers. He demonstrates his point by providing example uses of the word *bachelor* that allow a competent speaker to disambiguate between the senses listed by Katz and Fodor, but only based on properties of senses that are below the last marker in K&F's decomposition (cf. Figure 2.1). Since each of these examples is self-contained argument for the existence of some semantic category, we shall use some of them to demonstrate `4lang`'s ability to decompose meaning. In Figure 3.9 we present Bolinger's first five examples along with his original explanation of how each necessitates the introduction of some semantic marker:

Our account of these examples will be incomplete given the current limitations of our implemented systems, e.g. its current lack of treatment for modality, negation, and temporal relations. These already concern the first example: what is implemented of `4lang` so far does not have a sophisticated system for representing temporal relations. The concepts `after` and `before` are used in `4lang` definitions to encode event structure, e.g. the definition of `discover` contains `know` $\xrightarrow{0}$ `after` and `effort` $\xrightarrow{0}$ `before`. Whether the inference indicated by Bolinger can be made depends on how the definition of `marry` (Figure 3.10) is negated – given proper treatment, a man who *has never married* will be established as one for whom (`before` $\xrightarrow{0}$) `marriage` $\xleftarrow{2}$ `IN` $\xrightarrow{0}$ `NOT` holds, and `become` should entail that for some predicate `before` $\xrightarrow{0}$ is false, rendering it incompatible with the *unmarried man* interpretation of `bachelor`.

1. *He became a bachelor.* This rules out the 'man who has never married' – it is impossible to become one who has never done something. We can extract the -ever part of never from the distinguisher and set up a marker (Nonbecoming).
2. *The seven-year-old bachelor sat on the rock.* The definition 'male who has never married' was deficient. It should have been something like 'adult male who has never married,' and from that expanded distinguisher we now extract the marker (Adult).
3. *Lancelot was the unhappiest of all the bachelors after his wife died.* This seems to justify raising (Unmarried) to marker status and wipes out the distinguisher on one of the branches: *bachelor*-noun-(Human)–(Male)-(Adult)-(Non-becoming)-(Unmarried).
4. *That peasant is a happy bachelor.* Being a peasant is not compatible with being a knight. There must be a marker of status lying around somewhere. A knight has to be of gentle birth. Let us extract (Noble) from the distinguisher (leaving the degree of nobility for the moment undisturbed as still part of the knight's distinguisher).
5. *George is one bachelor who is his own boss.* This eliminates the knight, and turns 'serving under' into another status marker that might be called (Dependent).

Figure 3.9: Examples and arguments for new markers (Bolinger, 1965, p.558-560)

Example 2 requires us to derive the incompatibility of `adult` with `7-year-old`. Since the definitions of *adult* in both Longman and `en.wiktionary` contain the term *fully grown*, this inference requires us to make reference to knowledge about the average age at which humans stop growing. The third example can be handled in `4lang` similarly to the first: the *unmarried (adult) male* reading of `bachelor` must entail that at no time in the past could $\xleftarrow{1}$ `IN` $\xrightarrow{2}$ `marriage` have been true. Example 4 requires a contradiction to be detected between `knight` and `peasant` – this can be straightforward given the right definition, but given our method of building definitions from dictionary definitions, we cannot expect our definition graphs to be as comprehensive as to include `noble` and `LACK` $\xrightarrow{0}$ `noble` in the respective graphs for `knight` and `peasant`. Instead we should be able to infer these relations from the definitions we do encounter: the Longman definition of *knight*: *'a man with a high rank in the past who was trained to fight while riding a horse'* should result in the subgraph `knight` $\xleftarrow{1}$ `HAS` $\xrightarrow{2}$ `rank` $\xrightarrow{0}$ `high`[7], the definition of `peasant`: *a poor farmer*

[7] Incidentally, to construct this graph we would also need to overcome a parsing error: the Stanford Parser analyses this noun phrase as describing a man whose rank was trained and the rank is in the past. Parser errors such as this one will be discussed in Sections 4.4.1 and 8.4

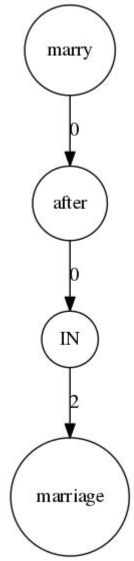

Figure 3.10: `4lang` definition of `marry`.

who owns or rents a small amount of land, either in past times or in poor countries will yield `peasant` $\xrightarrow{0}$ `poor`. These relations are not strictly incompatible, the original example also depends upon the assumption that being a peasant entails being of low rank – we have much better chances given a definition that makes this assumption itself, such as the one in the English Wiktionary: *A member of the lowly social class which toils on the land (...)*. In the latter case, all that remains is making the connection between `rank` and `class` ($\xrightarrow{0}$ `social`), but the former should also allow us, given a probabilistic system, to establish that a peasant is not likely to be a knight.

Finally, in Example 5, it is the incompatibility of 'being one's own boss' and the 'serving under' component of the *young knight serving under the standard of another knight* that must be established. The Longman definition of `boss`: *the person who employs you or who is in charge of you at work* will allow us to map *George is his own boss* to `George` $\underset{2}{\overset{1}{\Leftarrow}}$ `employ`, contradicting `George` $\xrightarrow{0}$ `serve` $\xleftarrow{1}$ `under` $\xrightarrow{2}$ `X` if the identity of `X` and `George` cannot be established, in this case explicitly excluded by the phrase *another knight*. We refrain from discussing the remaining 10 examples in (Bolinger, 1965). Details of the processes presented here are yet to be worked out, but we have shown that each inference is possible given our current set of semantic primitives.

3.6 Theoretical significance

This chapter provided a brief summary of the main principles behind the `4lang` system for representing the meaning of linguistic structures. Before we proceed to present a set of tools for building and manipulating `4lang` representations, as well as their applications to some tasks in computational semantics, let us point out some of the most important characteristics of `4lang` representations that make it our formalism of choice in the remainder of this thesis.

No categories `4lang` does not differentiate between concepts denoting actions, entities, attributes, etc., there are no categories of concepts equivalent to part-of-speech categories of words. This ensures, among other things, that words with a shared root are typically mapped to the same concept, and that ultimately utterances with the same information content can be mapped to inferentially identical `4lang` representations.

No polysemy `4lang` will only accommodate multiple senses of a word as a last resort. Distant but related uses of the same word must be interpreted via the same generic concept. This virtually eliminates the difficulty of word sense disambiguation.

Requires powerful inference The above principles require a mechanism for deriving all uses of a word from minimalistic definitions. Such a mechanism may stand a real chance at handling creative language use typical of everyday human communication (and responsible for polysemy in the first place).

No failure of interpretation No combinations of concepts and connections between them are forbidden by the formalism itself. Inference may judge certain states-of-affairs impossible, but the formalism will not fail the interpretation process.

Chapter 4

Phrases

In this chapter we present our work on combining word representations like those described in Chapter 3 to create graphs that encode the meaning of phrases. We relegate the task of syntactic parsing to the state of the art Stanford Parser (DeMarneffe et al., 2006; Socher, Bauer, et al., 2013). The pipeline presented in this chapter processes sets of dependency triplets emitted by the Stanford Parser to create `4lang`-style graphs of concepts (our future plans to implement syntactic parsing in `4lang` are outlined in Section 8.4). This chapter is structured as follows: dependency parsing is briefly introduced in Section 4.1, the central `dep_to_4lang` module which maps dependencies to `4lang` graphs is presented in Sections 4.2 and 4.3. Major issues are discussed in Section 4.4, some solutions are presented in Section 4.5, manual evaluation of the `text_to_4lang` system is provided in Section 4.6. Finally, Section 4.7 presents the adaptation of `text_to_4lang` to Hungarian. The module presented in this chapter is accessible via the `text_to_4lang`[1] module of the `4lang` repository. Besides the ability to map chunks of running text to semantic representations, `text_to_4lang` will see another application that is crucial to the system described in this thesis: we process definitions of monolingual dictionaries to acquire word representations for lexical items that are not covered by `4lang`. The resulting module `dict_to_4lang` will be presented in Chapter 5. The modules `dep_to_4lang` and `dict_to_4lang` are also presented in (Recski & Borbély, 2016), the adaptation to Hungarian is published in (Recski et al., 2016).

[1]https://github.com/kornai/4lang/blob/master/src/text_to_4lang.py

4.1 Dependency parsing

We use a robust, state of the art tool, the Stanford Parser[2] to obtain dependency relations that hold between pairs of words in an English sentence. Unlike dependency parsers that have been trained on manually annotated dependency treebanks, the Stanford Parser discovers relations by matching templates against its parse of a sentence's constituent structure (DeMarneffe et al., 2006). This approach is more robust, since phrase structure parsers, and in particular the PCFG parser in the Stanford toolkit (Klein & Manning, 2003), are trained on much larger datasets than what is available to standard dependency parsers.

The Stanford Dependency Parser is also capable of returning *collapsed* dependencies, which explicitly encode relations between two words that are encoded in the sentence by a function word such as a preposition or conjunction. E.g. in case of the sentence *I saw the man who loves you*, standard dependency parse would contain the relation `nsubj(loves, who)` but not `nsubj(loves, man)`, even though *man* is clearly the subject of *loves*. Collapsed dependency parses contain these implicitly present dependencies and are therefore more useful for extracting the semantic relationships between words in the sentence. Furthermore, the Stanford Parser can postprocess *conjunct dependencies*: in the sentence *Bills on ports and immigration were submitted by Senator Brownback, Republican of Kansas*, the NP *Bills on ports and immigration* will at first be parsed into the relations `prep_on(Bills, ports)` and `cc_and(ports, immigration)`, then matched against a rule that adds the relation `prep_on(Bills, immigration)`. For our purposes we enable both types of postprocessing and use the resulting set of relations (or *triplets*) as input to the `dep_to_4lang` module, which uses them to build `4lang` graphs and will be introduced in Section 4.2.

The list of dependency relations extracted from a sentence (for a detailed description of each dependency relation see (De Marneffe & Manning, 2008)) is clearly not intended as a representation of meaning; it will nevertheless suffice for constructing good quality semantic representations because of the nature of `4lang` relations: for sentences and phrases such as *Mary loves John* or *queen of France*, `4lang` representations are as simple as $\texttt{Mary} \xleftarrow{1} \texttt{love} \xrightarrow{2} \texttt{John}$ and $\texttt{France} \xleftarrow{1} \texttt{HAS} \xrightarrow{2} \texttt{queen}$ which can be straightforwardly constructed from the dependency relations `nsubj(love, Mary)`, `dobj(love, John)`, and `prep_of(queen, France)`. Any further details that one may demand of a semantic representation, e.g. that John is an experiencer or that France does not physically possess the queen, will be inferred from the `4lang` definitions of the concepts `love` and `queen`, in

[2] http://nlp.stanford.edu/software/lex-parser.shtml

the latter case also accessing the definitions of `rule` or `country`.

4.2 From dependencies to graphs

To construct `4lang` graphs using dependency relations in the parser's output, we created manually a mapping from relations to `4lang` subgraphs, assigning to each dependency one of nine possible configurations (see Table 4.1). Additionally, all remaining relations of the form `prep_*` and `prepc_*` are mapped to binary subgraphs containing a node corresponding to the given preposition. To map words to `4lang` concepts, we first lemmatize them using the `hunmorph` morphological analyzer and the `morphdb.en` database (Trón et al., 2005). Graph edges for each dependency are added between the nodes corresponding to the lemmas returned by `hunmorph`. The full mapping from dependencies to `4lang`-subgraphs is presented in Table 4.1. Figure 4.1 provides an example of how `4lang` subgraphs correspond to dependency triplets.

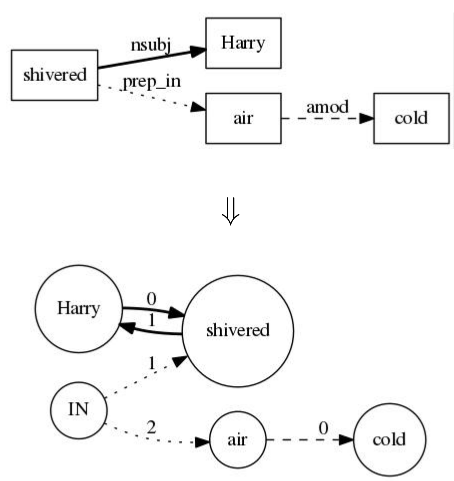

Figure 4.1: Constructing the graph for *Harry shivered in the cold night air*

4.3 Utterances

Dependency relations obtained from multiple sentences can be used to update graphs over a single set of nodes, therefore the `text_to_4lang` pipeline presented in this chapter can be applied to documents of arbitrary size. Some of our preliminary experiments showed coreference resolution to be a significant challenge posed by processing several sentences into a single concept graph; we have therefore extended the `text_to_4lang` module to

Dependency	Edge
amod advmod npadvmod acomp dep num prt	$w_1 \xrightarrow{0} w_2$
nsubj csubj xsubj agent	$w_1 \underset{0}{\overset{1}{\rightleftharpoons}} w_2$
dobj pobj nsubjpass csubjpass pcomp xcomp	$w_1 \xrightarrow{2} w_2$
poss prep_of	$w_2 \xleftarrow{1} \texttt{HAS} \xrightarrow{2} w_1$
tmod	$w_1 \xleftarrow{1} \texttt{AT} \xrightarrow{2} w_2$
prep_with	$w_1 \xleftarrow{1} \texttt{INSTRUMENT} \xrightarrow{2} w_2$
prep_without	$w_1 \xleftarrow{1} \texttt{LACK} \xrightarrow{2} w_2$
prep_P	$w_1 \xleftarrow{1} \texttt{P} \xrightarrow{2} w_2$

Table 4.1: Mapping from dependency relations to `4lang` subgraphs

run the Stanford Coreference Resolution system (Lee et al., 2011) and use its output to unify nodes in the concept graphs. An example is shown in Figure 4.2.

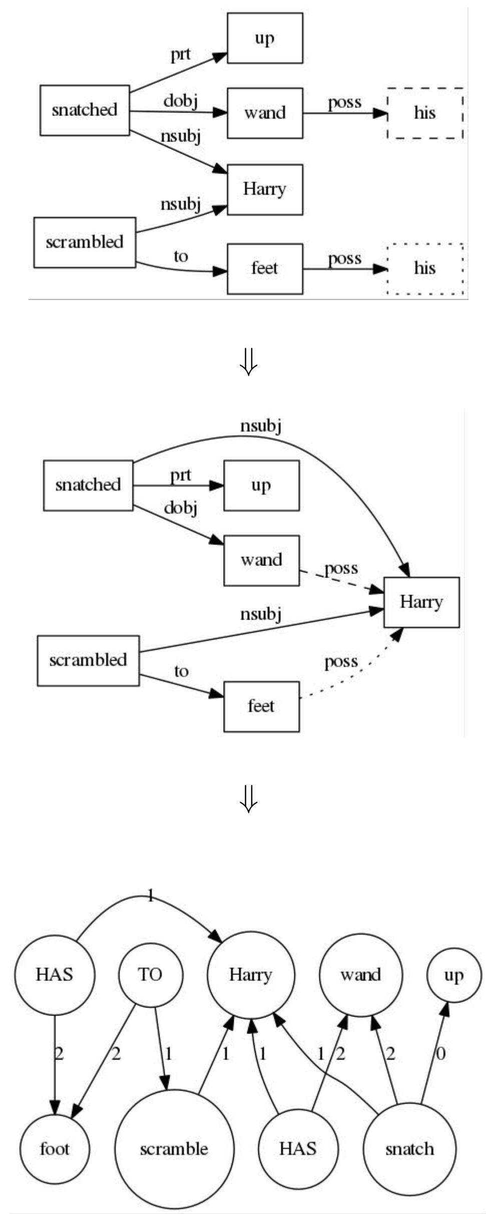

Figure 4.2: `text_to_4lang` processing of *Harry snatched up his wand and scrambled to his feet* with coreference resolution

A proper treatment of complete utterances would require us to encode the logical relationships that hold among clauses, signaled by linking constructions such as *because*, *unless*, *although*. The Stanford parser detects the pairs of clauses between which these words constitute dependencies, but there is no component in `4lang` to handle them. Some of these relationships could be handled straightforwardly, e.g. sentences expressing causal relationships, such as *if ... then ...* constructions, could be mapped to $\texttt{X} \xleftarrow{1} \texttt{CAUSE} \xrightarrow{2} \texttt{Y}$, where `X` and `Y` are root nodes of the `4lang` representations corresponding to each clause in the sentence.

4.4 Issues

4.4.1 Parsing errors

Using the Stanford Parser for dependency parsing yields high-quality output, it is however limited by the quality of the phrase structure grammar parser. Parsing errors constitute a major source of errors in our pipeline, occasionally resulting in dubious semantic representations that could be discarded by a system that integrates semantic analysis into the parsing process. Our long-term plans include implementing such a process within the `4lang` framework using *constructions* (see Section 8.4), currently we rely on independent efforts to improve the accuracy of phrase structure grammar parsers using semantic information.

Results of a pioneering effort in this direction are already included in the latest versions of the Stanford Parser (including the one used in the `4lang` system) and was introduced in Section 2.4.2: (Socher, Bauer, et al., 2013) improves the accuracy of the Stanford Parser by using *Compositional Vector Grammars*, RNN-based models that learn for each terminal rule $R^n \rightarrow R^{2n}$ linear transformations that can be applied to pairs of word vectors of length n to obtain an $n \times n$ matrix representing the nonterminal that is the result of applying the given rule. The purpose of this model is to account for the semantic relationships between words in the text that is to be parsed and words that have occurred in the training data. E.g. the sentence *He ate spaghetti with a spoon* can be structurally distinguished from *He ate spaghetti with meatballs* even if in the training phase the model has only had access to *[eat [spaghetti] [with a fork]]*, by grasping the similarity between the words *spoon* and *fork*.

This phenomenon of incorrect PP-attachment is the single most frequent source of anomalities in our output. For example, syntactic ambiguity in the Longman definition of `basement`: *a room or area in a building that is under the level of the ground*, which has the constituent structure in Figure 4.3 is incorrectly assigned the structure in Figure 4.4, resulting in the incorrect semantic representation in Figure 4.5 (instead of the expected graph in Figure 4.6). Most such ambiguities are easily resolved by humans based on lexical facts (in this case e.g. that buildings with some underground rooms are more common than buildings that are entirely under the ground, if the latter can be called buildings at all) but it seems that such inferencing is beyond the capabilities even for parsers using word embeddings. As already discussed in Section 3.3, such deductions can be made based on `4lang` representations.

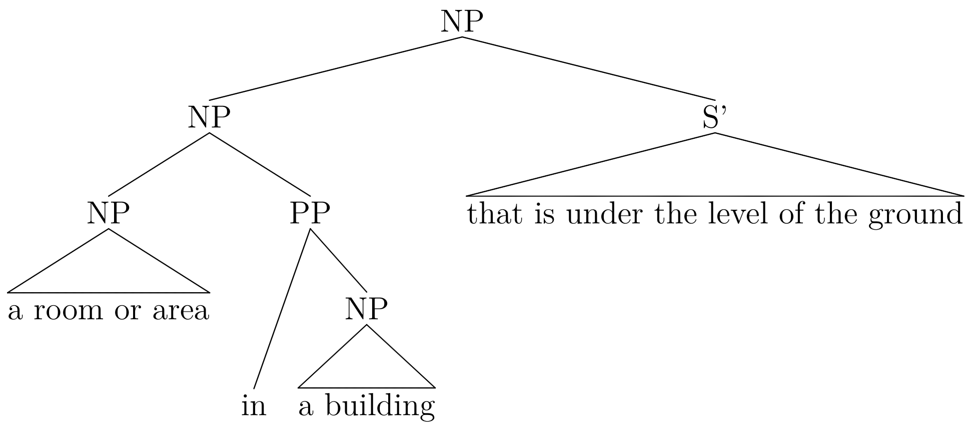

Figure 4.3: Constituent structure of *a room or area in a building that is under the level of the ground*

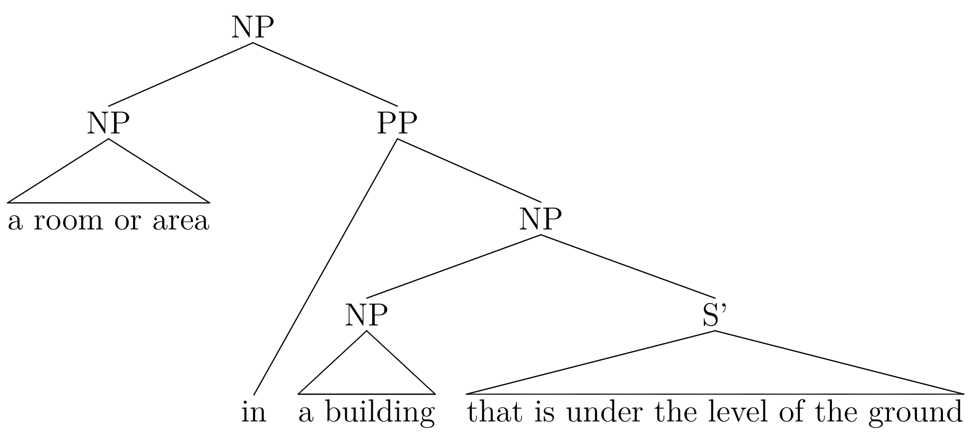

Figure 4.4: Incorrect parse tree for *a room or area in a building that is under the level of the ground*

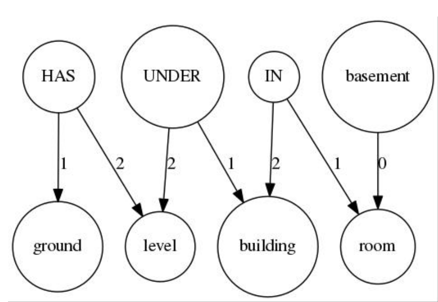

Figure 4.5: Incorrect definition graph for `basement`.

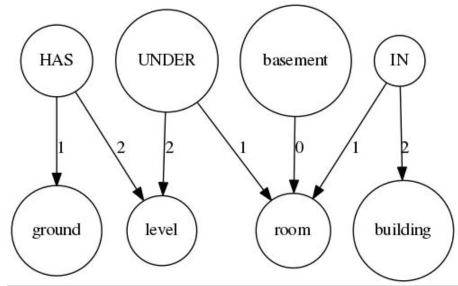

Figure 4.6: Expected definition graph for `basement`.

4.4.2 Relationships among clauses

The `text_to_4lang` system does not currently detect relationships between multiple clauses of a sentence expressed by conjunctions such as *because*, *unless*, *although*, etc., since they do not appear as syntactic dependency relations in the output of dependency parsers (unlike e.g. clausal modifiers of noun phrases, which are processed by the Stanford Parser to obtain e.g. `nsubj(appears, liquid)` $\xrightarrow{f}$ rom the definition of `perspiration`: *liquid that appears on your skin when you are hot or nervous*. Such conjunctions should be treated on a case-by-case basis by constructions enforcing simple rules. Such a construction might state that for some sentence *X, because Y*, the `4lang` graphs corresponding to X and Y should be joined by $\xleftarrow{1}$ `CAUSE` $\xrightarrow{2}$. The definition of *perspiration* could then map to the graph in Figure 4.7.

4.5 Postprocessing dependencies

Some of the typical issues of the graphs constructed by the process described in Section 4.2 can be resolved by postprocessing the dependency triplets in the parser's output before passing them to `dep_to_4lang`. Currently the `dependency_processor` module handles two configurations: coordination (Section 4.5.1) and copular sentences (Section 4.5.2)

4.5.1 Coordination

One frequent class of parser errors related to PP-attachment (cf. Section 4.4.1) involve constituents modifying a coordinated phrase which are analyzed as modifying only one of the coordinated elements. E.g. in the Longman entry **casualty** - *someone who is hurt or killed in an accident or war*, the parser fails to detect that the PP *in an accident or war* modifies the constituent *hurt or killed*, not just *killed*. Determining which of two possible parse trees is the correct one is of course difficult - once again, `casualty` may as well mean

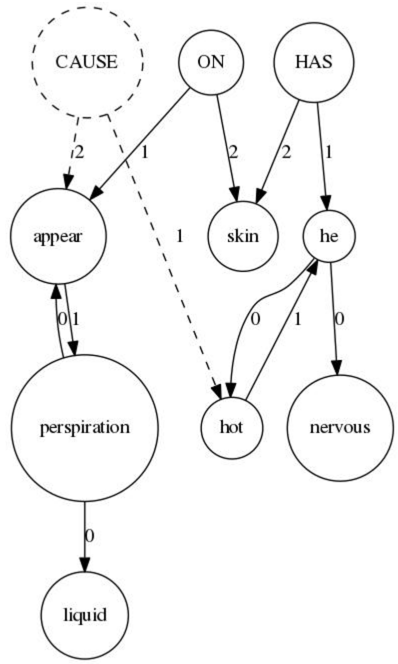

Figure 4.7: Definition graph built from **perspiration**: *liquid that appears on your skin because you are hot or nervous*

someone who is killed in an accident or war or someone who is hurt (in any way) and that such a misunderstanding is unlikely in real life is a result of inference mechanisms well beyond what we are able to model.

Our simple attempt to improve the quality of graphs built is to process all pairs of words between which a coordinating dependency holds (e.g. `conj_and, conj_or`, etc.) and copy all edges from each node to the other. This could hardly be called a solution, as it may introduce dependencies incorrectly, but in practice it has proved an improvement. In our current example this step enables us to obtain missing dependencies and thus build the correct `4lang` graph (see Figure 4.8).

4.5.2 Copulars and prepositions

Two further postprocessing steps involve copular constructions containing prepositional phrases. In simple sentences such as *The wombat is under the table*, the parser returns the pair of dependencies `nsubj(is, wombat)` and `prep_under(is, table)`, which we use to generate `prep_under(wombat, table)`. Similarly, when PPs are used to modify a noun, such as in the Longman definition of `influenza`: *an infectious disease that is like a very bad cold*, for which the dependency parser returns, among others, the triplets `rcmod(disease, is)` and `prep_like(is, cold)`, we let a simple rule add the triplet

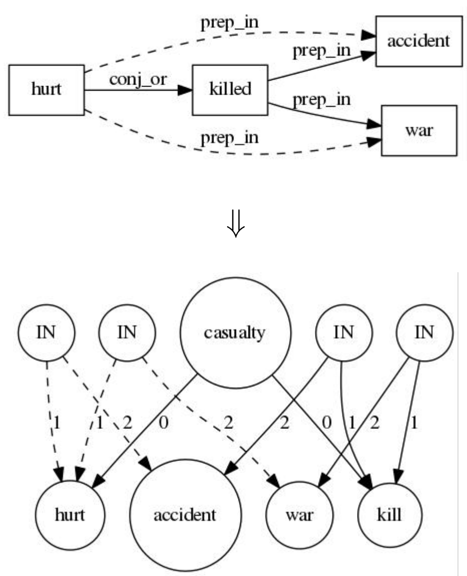

Figure 4.8: Definition graph built from: **casualty** - *someone who is hurt or killed in an accident or war*, with extra dependencies added by the postprocessor

`prep_like(disease, cold)` (see Figure 4.9). In both cases we finish by removing the copular verb in order to simplify our final representation.

4.6 Evaluation

We performed manual evaluation of the `text_to_4lang` module on a sample from the *UMBC Webbase* corpus (Han et al., 2013), a set of 3 billion English words based on a 2007 webcrawl performed as part of the *Stanford Webbase*[3] project. We used the GNU utility `shuf` to extract a random sample of 20 sentences, which we processed with `text_to_4lang` and examined manually both the final output and the dependencies output by the Stanford Parser in order to gain a full understanding of each anomaly in the graphs created. The sentences in this corpus are quite long (25.5 words/sentence on average, compared to 18.8 of the Wall Street Journal section of the Penn Treebank), therefore most graphs are affected by multiple issues; we shall now take stock of those that affected more than one sentence in our sample.

Parser errors remain the single most frequent source of error in our final `4lang` graphs: 7 sentences in our sample of 20 were assigned dependencies erroneously. 3 of these cases are related to PP-attachment (see Section 4.4.1). Parser errors are also virtually the only issue that cause incorrect edges to be added to the final graph – nearly all remaining issues will result in missing connections only. The second largest source of errors in this

[3] http://dbpubs.stanford.edu:8091/~testbed/doc2/WebBase/

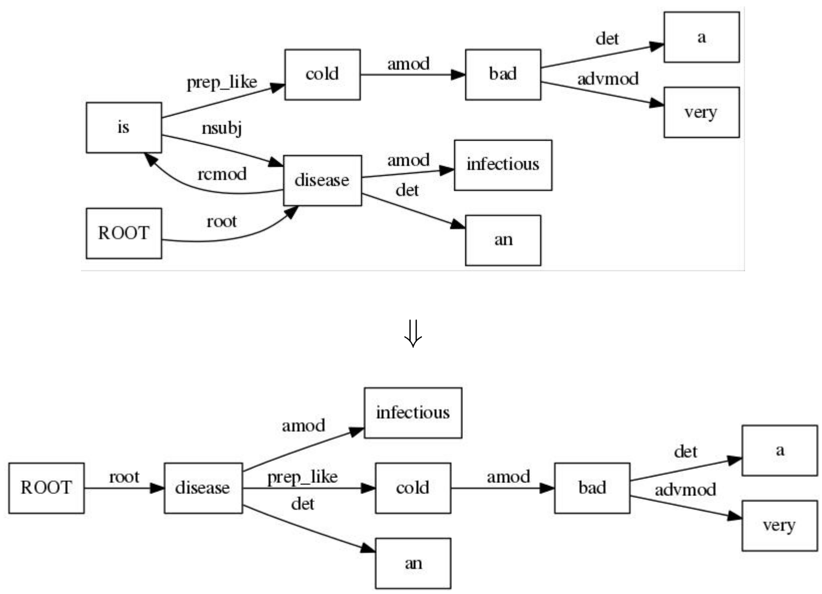

Figure 4.9: Postprocessing the definition *an infectious disease that is like a very bad cold*

dataset are related to connectives between clauses that our pipeline does not currently process (see Section 4.3). Our sample contains 6 such examples, including pairs of clauses connected by *that* (3x), *unless*, *regardless*, etc. The output of our pipeline for these sentences typically consists of two graphs that are near-perfect representations of the two clauses, but are not connected to each other in any way – an example is shown in Figure 4.10, we shall briefly return to this issue in Section 8.1.

There are three more error classes to be mentioned, each of which affects three sentences in our sample. The first are recall errors made by the Stanford Coreference Resolution system: in these cases connections of a single concept in the final graph are split among two or more nodes, since our pipeline failed to identify two words as referring to the same entity (Figure 4.11 shows an example). The second group of errors is caused by sentences that are assigned the `vmod` dependency. This relation holds between a noun and a *reduced non-final verbal modifier*, which "is a participial or infinitive form of a verb heading a phrase (which may have some arguments, roughly like a VP). These are used to modify the meaning of an NP or another verb."(DeMarneffe et al., 2006, p.10). This dependency is not processed by `dep_to_4lang`, since it may encode the relation between a verb and either its subject or object; e.g. the example sentences in the Stanford Dependency Manual, *Truffles picked during the spring are tasty* and *Bill tried to shoot, demonstrating his incompetence* will result in the triplets `vmod(truffles, picked)` and `vmod(shoot, demonstrating`, but should be represented in `4lang` by the edges `pick` $\xrightarrow{2}$ `truffles` and `shoot` $\xrightarrow{0}$ `demonstrate`, respectively. When we extend our tools to handle Hungarian input (see Section 4.7, we add

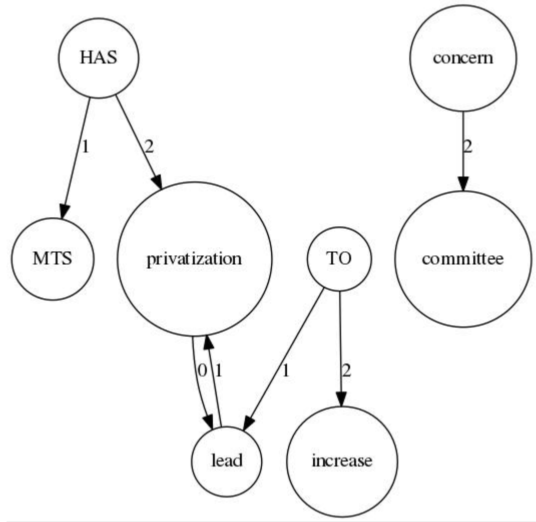

Figure 4.10: `4lang` graph built from the sentence *The Manitoba Action Committee is concerned that the privatization of MTS will lead to rate increases.*. The dependency `ccomp(concerned, lead)` was not processed.

to `dep_to_4lang` the capability of differentiating between words based on morphological analysis. English POS-tags are not currently processed, but this feature would make it straightforward to handle the `vmod` dependency using two rules, one for gerunds and one for participle forms. While most of the representations evaluated suffer from multiple errors, 5 out of 20 sentences were assigned perfect or near-perfect `4lang` representations.

4.7 Hungarian

We have created an experimental version of our pipeline for Hungarian, using the NLP library `magyarlanc` for dependency parsing and a mapping to `4lang` graphs that is sensitive to the output of morphological analysis, to account for the rich morphology of Hungarian encoding many relations that a dependency parse cannot capture. We describe the output of `magyarlanc` and the straightforward components of our mapping in Section 4.7.1. In Section 4.7.2 we discuss the use of morphological analysis in our pipeline and in Section 4.7.3 we present some arbitrary postprocessing steps similar to those described in Sections 4.5.1 and 4.5.2. Finally, in Section 4.7.4 we discuss the performance and main issues of the Hungarian subsystem.

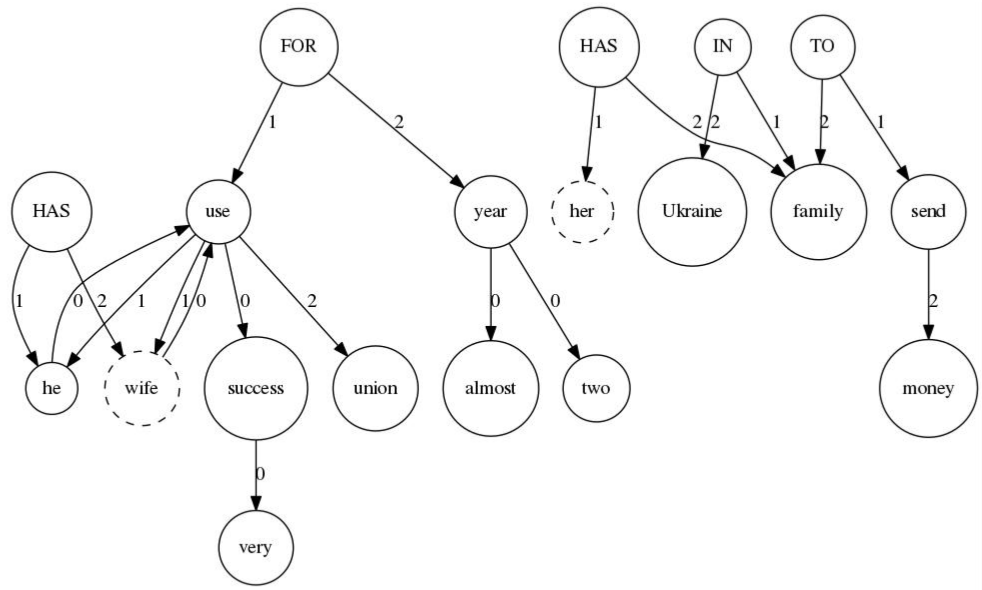

Figure 4.11: `4lang` graph built from the sentence *My wife and I have used Western Union very successfully for almost two years to send money to her family in Ukraine.*. Nodes with dashed edges should have been unified based on coreference resolution.

4.7.1 Dependencies

The `magyarlanc` library[4] (Zsibrita et al., 2013) contains a suite of standard NLP tools for Hungarian, which allows us, just like in the case of the Stanford Parser, to process raw text without building our own tools for tokenization, POS-tagging, etc. The dependency parser component of `magyarlanc` is a modified version of the Bohnet parser (Bohnet, 2010) trained on the Szeged Dependency Treebank (Vincze et al., 2010). The output of `magyarlanc` contains a much smaller variety of dependencies than that of the Stanford Parser. Parses of the ca. 4700 entries of the NSzT dataset (to be introduced in Section 5.1) contain nearly 60,000 individual dependencies, 97% of which are covered by the 10 most frequent dependency types (cf. Table 4.2). We shall first discuss dependencies that can be handled straightforwardly in the `dep_to_4lang` framework introduced in Section 4.2.

The dependencies `att, mode,` and `pred`, all of which express some form of unary predication, can be mapped to the 0-edge. `subj` and `obj` are treated in the same fashion as the Stanford dependencies `nsubj` and `dobj`. The dependencies `from, tfrom, locy, tlocy, to,` and `tto` encode the relationship of a predicate and an adverb or postpositional phrase answering the question 'from where?', 'from when?', 'where?', 'when?', 'where to?', and 'until when?', respectively.

[4] http://www.inf.u-szeged.hu/rgai/magyarlanc

att	26.0%
punct	16.1%
coord	15.0%
obl	9.6%
root	7.8%
conj	6.6%
mode	5.0%
det	4.7%
obj	3.7%
subj	2.6%

Table 4.2: Most common dependencies in `magyarlanc` output

4.7.2 Morphology

Hungarian is a language with rich morphology, and in particular the relationship between a verb and its NP argument is often encoded by marking the noun phrase for one of 17 distinct cases – in English, these relations would typically be expressed by prepositional phrases. The Stanford Parser maps prepositions to dependencies and the sentence *John climbed under the table* yields the dependency `prep_under(table, climb)`. The Hungarian parser does not transfer the morphological information to the dependencies, all arguments other than subjects and direct objects will be in the `OBL` relation with the verb. Therefore we updated the `dep_to_4lang` architecture to allow our mappings from dependencies to `4lang` subgraphs to be sensitive to the morphological analysis of the two words between which the dependency holds. The resulting system maps the phrase *a késemért jöttem* the knife-POSS-PERS1-CAU come-PAST-PERS1 'I came for my knife' to `FOR(come, knife)` based on the morphological analysis of *késem*, performed by `magyarlanc` based on the `morphdb.hu` database (Trón et al., 2005).

This method yields many useful subgraphs, but it also often leaves uncovered the true semantic relationship between verb and argument, since nominal cases can have various interpretations that are connected to their 'primary' function only remotely, or not at all. The semantics of Hungarian suffixes *-nak/-nek* (dative case) or *-ban/-ben* (inessive case) exhibit great variation – not unlike that of the English prepositions *for* and *in*, and the 'default' semantic relations `FOR` and `IN` are merely one of several factors that must be considered when interpreting a particular phrase. Nevertheless, our mapping from nominal cases to binary relations can serve as a strong baseline, just like interpreting English *for*

Dependency	Edge
att mode pred	$w_1 \xrightarrow{0} w_2$
subj	$w_1 \xrightarrow{1} w_2$
obj	$w_1 \xrightarrow{2} w_2$
from	$w_1 \xleftarrow{1}$ `FROM` $\xrightarrow{2} w_2$
tfrom	$w_1 \xleftarrow{1}$ `since` $\xrightarrow{2} w_2$
locy tlocy	$w_1 \xleftarrow{1}$ `AT` $\xrightarrow{2} w_2$
to	$w_1 \xleftarrow{1}$ `TO` $\xrightarrow{2} w_2$
tto	$w_1 \xleftarrow{1}$ `until` $\xrightarrow{2} w_2$

Table 4.3: Mapping from `magyarlanc` dependency relations to `4lang` subgraphs

and *in* as `FOR` and `IN` via the Stanford dependencies `prep_for` and `prep_in`. The mapping from `magyarlanc` dependencies to `4lang` graphs is shown in Table 4.3, nominal cases of `OBL` arguments are mapped to `4lang` binaries according to Table 4.4.

4.7.3 Postprocessing

Copulars

In the Szeged Dependency Treebank, and consequently, in the output of `magyarlanc`, copular sentences will contain the dependency relation `pred`. Hungarian only requires a copular verb in these constructions when a tense other than the present or a mood other than the indicative needs to be marked (cf. Table 4.5). While the sentence in (1) is analyzed as `subj(Ervin, álmos)`, all remaining sentences will be assigned the dependencies `subj(Ervin, volt)` and `pred(volt, álmos)`. The same copular structures allow the predicate to be a noun phrase (e.g. *Ervin tüzoltó* 'Ervin is a firefighter'). In each of these cases we'd like to eventually obtain the `4lang` edge `Ervin` $\xrightarrow{0}$ `sleepy` (`Ervin` $\xrightarrow{0}$ `firefighter`), which could be achieved in several ways: we might want to detect whether the nominal predicate is a noun or an adjective and add the `att` and `subj` dependencies accordingly. Both of these solutions would result in a considerable increase the complexity of the `dep_to_4lang` system and neither would simplify its input: the simplest examples (such as (1) in Table 4.5) would still undergo different treatment. With

Case	Suffix	Subgraph
sublative superessive	*-ra/-re* *-on/-en/-ön*	$w_1 \xleftarrow{1} \texttt{ON} \xrightarrow{2} w_2$
inessive illative	*-ban/-ben* *-ba/-be*	$w_1 \xleftarrow{1} \texttt{IN} \xrightarrow{2} w_2$
temporal adessive	*-kor* *-nál/nél*	$w_1 \xleftarrow{1} \texttt{AT} \xrightarrow{2} w_2$
elative ablative delative	*-ból/-ből* *-tól/-től* *-ról/-ről*	$w_1 \xleftarrow{1} \texttt{FROM} \xrightarrow{2} w_2$
allative terminative	*-hoz/-hez/-höz* *-ig*	$w_1 \xleftarrow{1} \texttt{TO} \xrightarrow{2} w_2$
causative	*-ért*	$w_1 \xleftarrow{1} \texttt{FOR} \xrightarrow{2} w_2$
instrumental	*-val/-vel*	$w_1 \xleftarrow{1} \texttt{INSTRUMENT} \xrightarrow{2} w_2$

Table 4.4: Mapping nominal cases of `OBL` dependants to `4lang` subgraphs

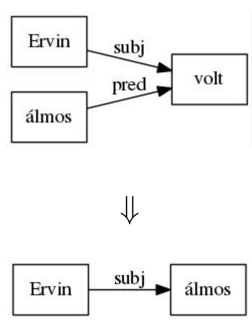

Figure 4.12: Postprocessing dependencies of a copular sentence

these considerations in mind we took the simpler approach of mapping all pairs of the form `nsubj(x, c)` and `pred(c, y)` (such that `c` is a copular verb) to the relation `subj(x, y)` (see Figure 4.12), which can then be processed by the same rule that handles the simplest copulars (as well as verbal predicates and their subjects.)

Coordination

Unlike the Stanford Parser, `magyarlanc` does not propagate dependencies across coordinated elements. Therefore we introduced a simple postprocessing step where we collect words of the sentence governing a `coord` dependency, then find for each the words accessible via `coord` or `conj` dependencies (the latter connects coordinating conjunctions such as *és* ‘and’ to the coordinated elements). Finally, we unify the dependency relations of all

(1)	*Ervin*	*álmos*			
	Ervin	sleepy			
	'Ervin is sleepy'				
(2)	*Ervin*	*nem*	*álmos*		
	Ervin	not	sleepy		
	'Ervin is not sleepy'				
(3)	*Ervin*	*álmos*	*volt*		
	Ervin	sleepy	was		
	'Ervin was sleepy'				
(4)	*Ervin*	*nem*	*volt*	*álmos*	
	Ervin	not	was	sleepy	
	'Ervin was not sleepy'				

Table 4.5: Hungarian copular sentences

coordinated elements – Figure 4.13 shows a simple example[5]

4.7.4 Evaluation and issues

As in the case of the English system, we have randomly chosen 20 sentences to manually evaluate `text_to_4lang` on Hungarian data. The source of our sample is the Hungarian Webcorpus (Halácsy et al., 2004). As before, we shall start by providing some rough numbers regarding the average quality of the 20 `4lang` graphs, then proceed to discuss some of the most typical issues, citing examples from the used sample. 10 of the 20 graphs were correct `4lang` representations, or had only minor errors. An example of a correct transformation can be seen in Figure 4.15. Of the remaining graphs, 4 were mostly correct but had major errors, e.g. 1-2 content words in the sentence had no corresponding node, or several erroneous edges were present in the graph. The remaining 6 graphs had many major issues and can be considered mostly useless.

When investigating the processes that created the more problematic graphs, nearly all errors seem to have been caused by sentences with multiple clauses. When a clause is introduced by a conjunction such as *hogy* 'that' or *ha* 'if', the dependency trees of each graph are connected via these conjunctions only, i.e. the parser does not assign dependencies that hold between words from different clauses. We are able to build good quality subgraphs from each clause, but further steps are required to establish the semantic relationship between them based on the type of conjunction involved – a process that

[5]This step introduces erroneous edges in a small fraction of cases: when a sentence contains two or more clauses that are not connected by any conjunction – i.e. no connection is indicated between them – a `coord` relation is added by `magyarlanc` to connect the two dependency trees at their root nodes.

Csengő,	*vidám,*	*kellemes*	*kacagás*	*hangzott*	*a*	*magasból*
Ringing	joyful	pleasant	giggle	sound-PST	the	height-ELA

‘Ringing, merry, pleasant laughter sounded from above’

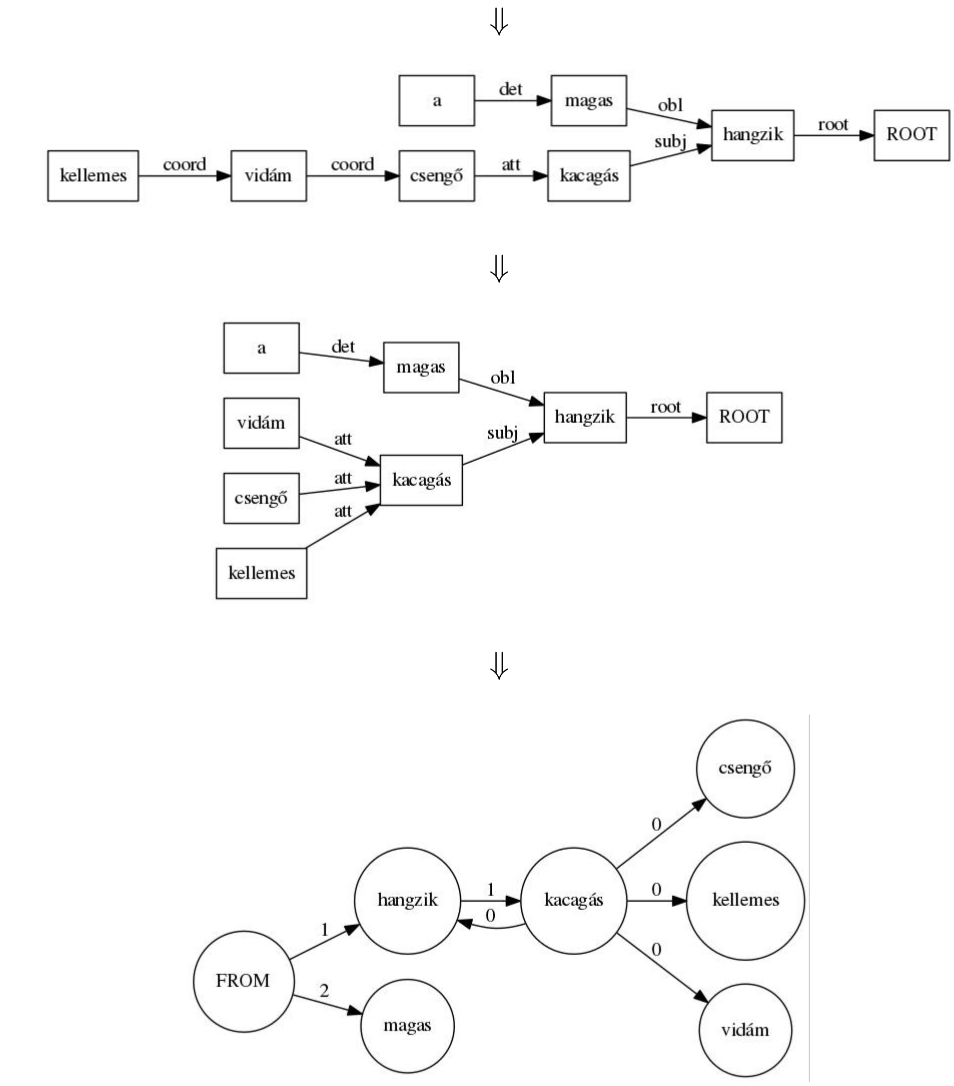

Figure 4.13: Processing a coordinated sentence

requires case-by-case treatment and would even then be non-trivial. An example from our sample is the sentence in Figure 4.14; here a conditional clause is introduced by a phrase

Örülnénk,	*ha*	*a*	*konzultációs*	*központok*
rejoice-COND-1PL	if	the	consultation-ATT	center-PL
közötti	*kilométerek*	*nem*	*jelentenének*	
between-ATT	kilometer-PL	not	mean-COND-3PL	
az	*emberek*	*közötti*	*távolságot.*	
the	person-PL	between-ATT	distance-ACC	
'We'd be glad if the kilometers between consultation centers did not mean distance between people'				

Figure 4.14: Subordinating conjunction

that roughly translates to 'We'd be glad if...'. Even if we disregard the fact that a full analysis of how this phrase affects the semantics of the sentence would require some model of the speaker's desires – we could still interpret the sentence literally by imposing some rule for conditional sentences, e.g. that given a structure of the form A if B, the `CAUSE` relation is to hold between the root nodes of B and A. Such rules could be introduced for several types of conjunctions in the future. A further, smaller issue is caused by the general lack of personal pronouns in sentences: Hungarian is a *pro-drop* language: if a verb is inflected for person, pronouns need not be present to indicate the subject of the verb, e.g. *Eszem.* 'eat-1SG' is the standard way of saying 'I'm eating' as opposed to *?Én eszem* 'I eat-1G' which is only used in special contexts where emphasis is necessary. Currently this means that `4lang` graphs built from these sentences will have no information about who is doing the `eating`, but in the future these cases can be handled by a mechanism that adds a pronoun subject to the graph based on the morphological analysis of the verb. Finally, the lowest quality graphs are caused by very long sentences containing several clauses and causing the parser to make multiple errors.

1995	*telén*	*vidrafelmérést*	*végeztünk*
1995	winter-POSS-SUP	otter-survey-ACC	conduct-PST-1PL
az	*országos*	*akció*	*keretében.*
the	country-ATT	action	frame-POSS-INE
'In the winter of 1995 we conducted an otter-survey as part of our national campaign'			

⇓

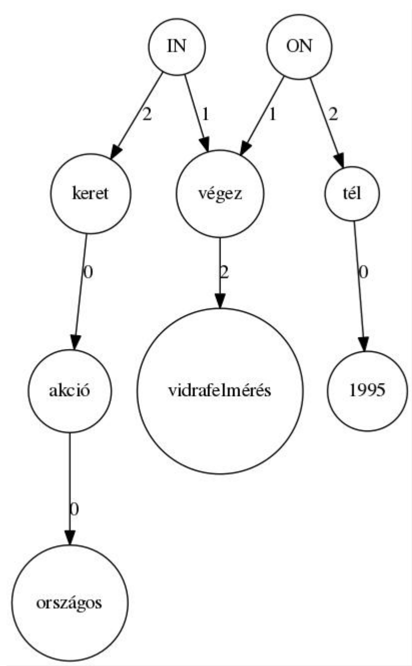

Figure 4.15: Example of perfect `dep_to_4lang` transformation

Chapter 5

Building definition graphs

One application of the text_to_lang module is of particular importance to us. By processing entries in monolingual dictionaries written for humans we can attempt to build definition graphs like those in 4lang for practically any word. This section presents the dict_to_4lang module, which extends the text_to_4lang pipeline with parsers for several major dictionaries (an overview of these is given in Section 5.1) as well as some preprocessing steps specific to the genre of dictionary definitions – these are presented in Section 5.2. Section 5.3 discusses *expansion* of 4lang representations, the process of copying links in definition graphs (both hand-written and built by dict_to_4lang) to 4lang representations created by text_to_4lang. Finally, Section 5.4 points out several remaining issues with definition graphs produced by the dict_to_4lang pipeline. Applications of dict_to_4lang, both existing and planned, shall be described in Chapter 6. The entire pipeline is available as part of the 4lang library, implemented by the dict_to_4lang module[1].

5.1 Data sources

We've built parsers for three large dictionaries of English and two of Hungarian. Custom parsers have been built for all five sources and are distributed as part of the 4lang module.

5.1.1 Longman Dictionary of Contemporary English

The Longman Dictionary of Contemporary English (Bullon, 2003) contains ca. 42,000 English headwords. Its definitions are constrained to a small vocabulary, the Longman

[1] https://github.com/kornai/4lang/blob/master/src/dict_to_4lang.py

Defining Vocabulary (LDV, (Boguraev & Briscoe, 1989)). The `longman_parser` tool processes the xml-formatted data and extracts for each headword a list of its senses, including for each the plain-text definition, the part-of-speech tag, and the full form of the word being defined, if present: e.g. definitions of acronyms will contain the phrase that is abbreviated by the headword. No component of `4lang` currently makes use of this last field, `AAA` will not be replaced by `American Automobile Association`, but this may change in the future.

5.1.2 Collins Cobuild Dictionary

The Collins-COBUILD dictionary (Sinclair, 1987) contains over 84 500 headwords. Its definitions use a vocabulary that is considerably larger than LDOCE, including a large technical vocabulary (e.g. **adularia:** *a white or colourless glassy variety of orthoclase in the form of prismatic crystals.*, rare words (**affricare:** *to rub against*), and multiple orthographic forms (**adsuki bean:** *variant spelling of adzuki bean*). Since many definitions are simply pointers to other headwords, the average entry in Collins is much shorter than in LDOCE. Given the technical nature of many entries, the vocabulary used by definitions exhibits a much larger variety: Longman definitions, for the greatest part limited to the LDV, contain less than 9000 English lemmas, not including named entities, numbers, etc., Collins definitions use over 38 000 (these and subsequent figures on vocabulary size are approximated using the `hunmorph` analyzer and the morphological databases `morphdb.en` and `morphdb.hu`).

5.1.3 English Wiktionary

Our third source of English definitions, the English Wiktionary at http://en.wiktionary.org is the most comprehensive database, containing over 128 000 headwords and available via public data dumps that are updated weekly. Since Wiktionaries are available for many languages using similar – although not standardized – data formats, it has long been a resource for various NLP tasks, among them an effort to extend the `4lang` dictionary to 40 languages (Ács et al., 2013). While for most languages datasets such as Longman and Collins may not be publicly available (e.g. at the time of writing this thesis, both Hungarian dictionaries were only available to the author based on personal requests), wiktionaries currently contain over 100 000 entries for nearly 40 languages, and over 10 000 for a total of 76.

5.1.4 Dictionaries of Hungarian

We've also run the `dict_to_4lang` pipeline on two explanatory dictionaries of Hungarian: volumes 3 and 4 of the *Magyar Nyelv Nagyszótára* (NSzt), containing nearly 5000 headwords starting with the letter *b* (Ittzés, 2011)[2], and over 120 000 entries of the complete *Magyar Értelmező Kéziszótár* (Pusztai, 2003), which has previously been used for NLP research (Miháltz, 2010). Basic figures for all five datasets are presented in Table 5.1.

Dict	headwords	av. def. length	approx. vocab. size
LDOCE	30 126	11.6	9 000
Collins	82 026	13.9	31 000
en.wikt	128 003	8.4	38 000
EKsz	67 515	5.0	33 700
NSzt (b)	4 683	10.7	9 900

Table 5.1: Basic figures for each dataset

5.2 Parsing definitions

5.2.1 Preprocessing

Before passing dictionary entries to the parser, we match them against some simple patterns that are then deleted or changed to simplify the phrase or sentence without loss of information. A structure typical of dictionary definitions are noun phrases with very generic meanings, e.g. *something*, *one*, *a person*, etc. For example, LDOCE defines `buffer` as *someone or something that protects one thing or person from being harmed by another*. The frequency of such structures makes it worthwhile to perform a simple preprocessing step: phrases such as *someone*, *someone who*, *someone*, etc. are removed from definitions in order to simplify them, thus reducing the chance of error in later steps. The above definition of `buffer`, for example, can be reduced to *protects from being harmed*, which can then be parsed to construct the definition graph $\texttt{protect} \xleftarrow{1} \texttt{FROM} \xrightarrow{2} \texttt{harm}$. A similar step replaced all occurences of the strings *a type of* and *a kind of* with *a*, once again simplifying both the input of the syntactic parser and the final representation without loss of information in definitions such as **lizard**: *a type of reptile that has four legs and a long tail*.

[2]The author gratefully acknowledges editor-in-chief Nóra Ittzés for making an electronic copy available.

5.2.2 Constraining the parser

Since virtually all dictionary definitions of nouns are single noun phrases, we constrain the parser to only allow such analyses for the definitions of all noun headwords. The command-line interface of the Stanford Parser does not support adding constraints on parse trees, but the Java API does; we implemented a small wrapper in `jython` that allowed us to access the classes and functions necessary to enforce this constraint (see Section 7.4.3 for more details). This fixes many incorrect parses, e.g. when a defining noun phrase with the structure in Figure 5.1 could also be parsed as a complete sentence, as in Figure 5.2.

```
(S
  (NP
    (NP (DT the) (NN size))
    (PP (IN of)
      (NP (DT a) (NN radio) (NN wave)))
    (VP (VBN used)
      (S
        (VP (TO to)
          (VP (VB broadcast)
            ( ... )))))))
```

Figure 5.1: Expected parse tree for the definition of **wavelength**: *the size of a radio wave used to broadcast a radio signal*

```
(S
  (NP
    (NP (DT the) (NN size))
    (PP (IN of)
      (NP (DT a) (NN radio) (NN wave))))
  (VP (VBD used)
    (S
      (VP (TO to)
        (VP (VB broadcast)
          ( ... ))))))
```

Figure 5.2: Incorrect parse tree from the Stanford Parser for the definition of **wavelength**: *the size of a radio wave used to broadcast a radio signal*

5.2.3 Building definition graphs

The output of the – possibly constrained – parsing process is passed to the `dep_to_4lang` module introduced in Chapter 4. The `ROOT` dependency in each parse, which was ignored in the general case, is now used to identify the head of the definition, which is a hypernym of the word being defined. This allows us to connect, via a 0-edge, the node of the concept being defined to the graph built form its definition. We can perform this step safely because the vast majority of definitions contain a hypernym of the headword as their root element – exceptions will be discussed in Section 5.4.2.

Dict	# graphs	av. nodes
LDOCE	24 799	6.1
Collins	45 311	4.9
en.wikt	120 670	5.4
EKsz	67 397	3.5
NSzt	4676	6.4

Table 5.2: Graphs built from each dataset

5.3 Expanding definition graphs

The `4lang` dictionary contains by design all words of the Longman Defining Vocabulary (LDV, (Boguraev & Briscoe, 1989)). This way, if we use `dict_to_4lang` to define each headword in LDOCE as a graph over nodes corresponding to words in its dictionary definition, these graphs will only contain concepts that are defined in the hand-written `4lang` dictionary. To take advantage of this, we implement an *expansion* step in `4lang`, which adds the definition of each concept to a `4lang` graph by simply adjoining each definition graph to G at the node corresponding to the concept being defined. This can be stated formally as follows:

Definition 1. *Given the set of all concepts C, a `4lang` graph G with concept nodes $V(G) = c_1, c_2, \ldots, c_i \in C$, a set of definition graphs D, and a lexicon function $L : C \to D$ such that $\forall c \in C : c \in V(L(c))$, we define the expansion of G as*

$$G^* = G \cup \bigcup_{c_i \in L} L(G)$$

Hand-written definitions in the `4lang` dictionary may also contain pointers to arguments of the definiendum, e.g. `stand` is defined as `upright` $\xleftarrow{0}$ `=AGT` $\xleftarrow{1}$ `ON` $\xrightarrow{1}$ `feet`, indicating that it is the agent of `stand` that is $\xrightarrow{0}$ `upright`, etc. Detecting the thematic role of a verb's arguments can be difficult, yet we handle the majority of cases correctly using a simple step after expansion: all edges containing `=AGT` (`=PAT`) nodes are moved to the machine(s) with a 1-edge (2-edge) pointing to it from the concept being defined. This allows us to create the graph in Figure 5.3 based on the above definition of `stand`. Expansion will affect all nodes of graphs built from LDOCE; when processing generic English text using `text_to_4lang` we may choose to limit expansion to manually built `4lang` definitions, or we can turn to dictionaries built using `dict_to_4lang`, allowing ourselves to add definitions to nearly all nodes. `4lang` modules can be configured to select the approach most suitable for any given application.

5.4 Issues and evaluation

In this section we will describe sources of errors in our pipeline besides those caused by incorrect parser output (see Section 4.4.1). We shall also present the results of manual error analysis conducted on a small sample of graphs in an effort to determine both the average accuracy of our output graphs as well as to identify the key error sources.

5.4.1 Error analysis

To perform manual evaluation of the `dict_to_4lang` pipeline we randomly selected 20 headwords from the Longman Dictionary[3]. In one round of evaluation we grouped the 20 definition graphs by quality, disregarding the process that created them. We found that 11 graphs were perfect or near-perfect definitions (see e.g. Figure 5.4) and a further 4 were mostly accurate, with only minor details missing or an incorrect relation present in addition to the correct ones. Of the remaining 6 graphs, 2 still encoded several true relationships, the last 4 were essentially useless. Our sample is too small to conclude that 75% of the graphs we build are of acceptable quality, but these results are nevertheless promising. Our second round of manual inspection was directed at the entire process of building the 20 graphs and aimed to identify the source of errors. Out of the 9 graphs that had errors at all, 6 were clearly caused by parser errors (discussed in Section 4.4.1), while the other 3 were connected to non-standard definitions (see Section 5.4.2).

[3] The 20 words in our sample, selected randomly using GNU `shuf` were the following: *aircraft, characteristic, clothesline, contrived, cypress, dandy, efface, frustrate, incandescent, khaki, kohl, lizard, nightie, preceding, residency, rock-solid, scant, transference, whatsit, Zen*

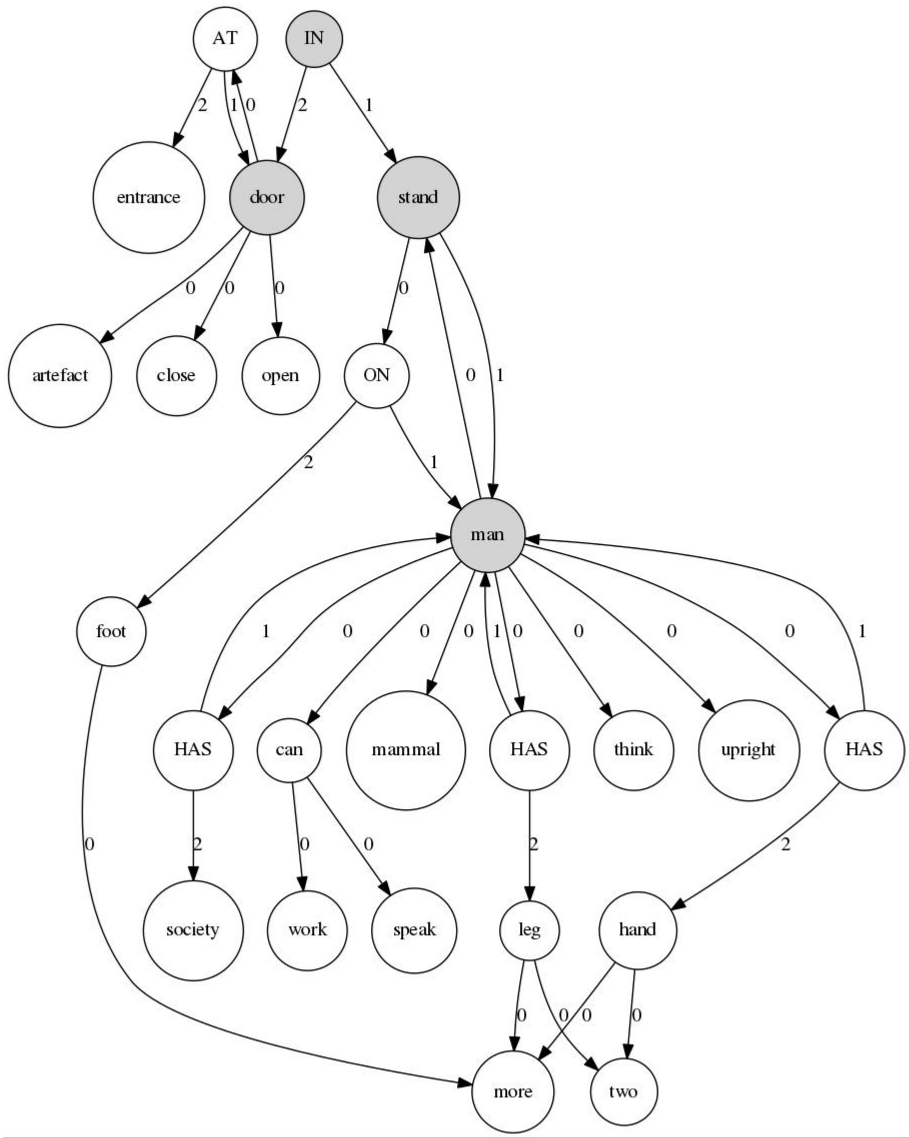

Figure 5.3: Expanded graph for *A man stands in the door*. Nodes of the unexpanded graph are shown in gray

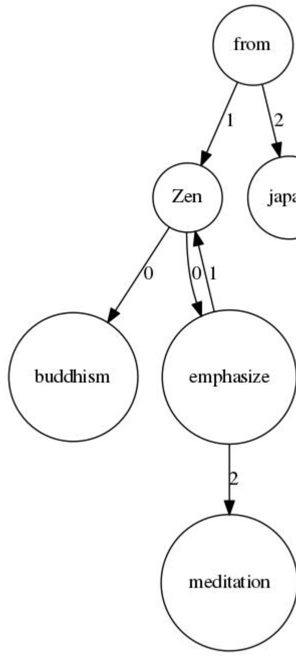

Figure 5.4: Graph constructed from the definition of **Zen**: *a kind of Buddhism from Japan that emphasizes meditation*

5.4.2 Non-standard definitions

Our method for building `4lang` definitions can be successful in the great majority of cases because most dictionary definitions – or at least their first sentences, which is all we make use of – are rarely complex sentences; in most cases they are single phrases describing the concept denoted by the headword – a typical example would be the definition of `koala`: *an Australian animal like a small grey bear with no tail that climbs trees and eats leaves.* It is these kinds of simple definitions that are prevalent in the dictionaries we process and that are handled quite accurately by both the Stanford Parser and our mapping from dependencies to `4lang` relations.

In some cases, definitions use full sentences to explain the meaning of a word in a more straightforward and comprehensible way, e.g.:

- **playback** - *the playback of a tape that you have recorded is when you play it on a machine in order to watch or listen to it*
- **indigenous** - *indigenous people or things have always been in the place where they are, rather than being brought there from somewhere else*
- **ramshackle** - *a ramshackle building or vehicle is in bad condition and in need of repair*

These sentences will result in a higher number of dependency relations, and consequently a denser definition graph; often with erroneous edges. In the special case when the Stanford Parser's output does not contain the `ROOT` relation, i.e. the parser failed to identify any of the words as the root of the sentence, we skip the entry entirely – this affects 0.76% of LDOCE entries, 0.90% of entries in `en.wiktionary`. That such definitions are problematic is also reflected in the fact that earlier editions of the Longman dictionary did not allow them, using the headword in the definition text was forbidden.

5.4.3 Word senses

As discussed in Section 3.2, the `4lang` theory assigns only one definition to each word form, i.e. it does not permit multiple word senses. All usage of a word must be derived from a single concept graph. Explanatory dictionaries like the ones listed in Section 5.1 provide several definitions for each word, of which we always process the first one. This decision is somewhat arbitrary, but produces good results in practice; the first definition typically describes the most common sense of the word, as in the case of `tooth`:

1. one of the hard white objects in your mouth that you use to bite and eat food
2. one of the sharp or pointed parts that sticks out from the edge of a comb or saw

We cannot expect to construct from this entry a generic definition such as `sharp, one_of_many`. Instead, to capture at a later stage that objects other than those in your mouth could be instances of `tooth`, we must turn to the principle that any link in a `4lang` definition can be overridden (see Section 3.2). Not only are we unable to predict the particular subset of links in the definition of `tooth` that will be shared across various uses of the word *tooth*, we *shouldn't* make any such predictions: it is no more than an accident that teeth turned out to be metaphors for small, sharp objects lined up next to one another and not for e.g. small, white, cube-shaped objects.

While in most cases the various sense defined for a word are metaphoric uses of the first, there remain words whose first definition is not generic enough to accommodate all others even if we assume powerful inferencing capabilities. Consider e.g. the definitions of **shower** from LDOCE below:

1. a piece of equipment that you stand under to wash your whole body
2. an act of washing your body while standing under a shower
3. a short period of rain or snow

4. a lot of small, light things falling or going through the air together
5. a party at which presents are given to a woman who is going to get married or have a baby
6. a group of stupid or lazy people
7. to wash your whole body while standing under a shower
8. to give someone a lot of things
9. to scatter a lot of things onto a person or place, or to be scattered in this way

A `4lang` definition generic enough so that one could derive at least the majority of these cases would be most similar to definition #4: showers are occurrences of many things falling, typically through the air. Understanding the word *shower* in the context of e.g. baby showers (#5) would remain a difficult task, including among others that of understanding that `fall` may refer to an object changing place not only physically but also in terms of ownership. In the above LDOCE entry, since we use the first definition to build the `4lang` graph, we lose any chance of recovering any of the meanings #3-6 and #8-9.

5.4.4 Hungarian

We also conducted manual error analysis on our Hungarian output, in this case choosing 20 random words from the EKsz dictionary. The 20 words, selected once again using `shuf`, are the following: *állomásparancsnok, beköt, biplán, bugás, egyidejűleg, font, főmufti, hajkötő, indikál, lejön, munkásőr, nagyanyó, nemtelen, összehajtogat, piff-puff, szét, tipográfus, túlkiabálás, vakolat, zajszint.* The graphs built by `dict_to_4lang` were of very good quality (see Figure 5.5 for an example), with only 3 out of 20 containing major errors. This is partly due to the fact that `NSzt` contains many very simple definitions, e.g. 4 of the 20 headwords in our random sample contained a (more common) synonym as its definition.

All 3 significant errors are caused by the same pattern: the analysis of possessive constructions by `magyarlanc` involve assigning the `att` dependency to hold between the possessor and the possessed, e.g. the definition of `piff-puff` (see Figure 5.6) will receive the dependencies `att(hang, kifejezés)` and `att(lövöldözés, hang)`, resulting in the incorrect `4lang` graph in Figure 5.7 instead of the expected one in Figure 5.8. `kifejezés` $\xrightarrow{0}$ `hang` $\xrightarrow{0}$`lövöldözés` instead of `kifejezés` $\xleftarrow{2}$ `HAS` $\xrightarrow{1}$ `hang` $\xleftarrow{2}$ `HAS` $\xrightarrow{1}$ `lövöldözés`.

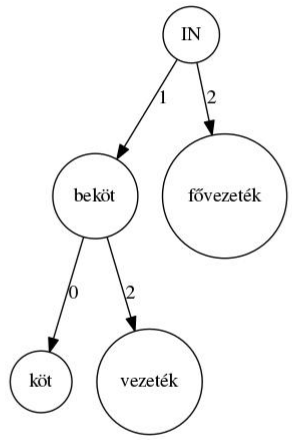

Figure 5.5: `4lang` graph built from the definition of **beköt**: *Vezetéket a fővezetékbe köt* 'cable-ACC the main-cable-ILL tie'

These constructions cannot be handled even by taking morphological analysis into account, since possessors are not usually marked (although in some structures they receive the dative suffix *-nak/-nek*, e.g. in embedded possessives like our current example (*hangjának* 'sound-POSS-DAT' is marked by the dative suffix as the possessor of *kifejezésére*). Unless possessive constructions can be identified by `magyarlanc`, we shall require an independent parsing mechanism in the future. The structure of Hungarian noun phrases can be efficiently parsed using the system described in (Recski, 2014), the grammar used there may in the future be incorporated into a `4lang`-internal parser (see Section 8.4).

Lövöldözés	*vagy*	*ütlegelés*	*hangjának*	*kifejezésére*
Shooting	or	thrashing	sound-POSS-DAT	expression-POSS-SUB

'Used to express the sound of shooting or thrashing'

⇓

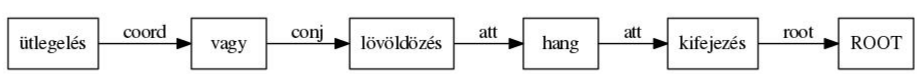

Figure 5.6: Dependency parse of the `EKsz` definition of the (onomatopoeic) term `piff-puff`

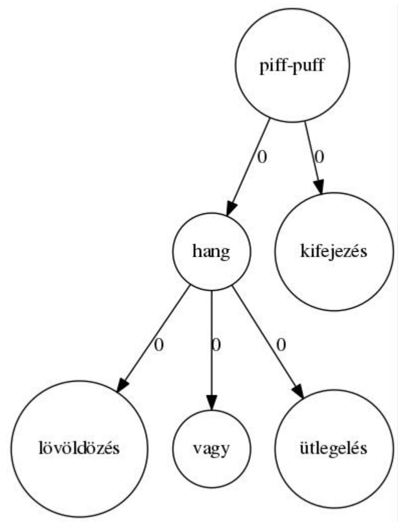

Figure 5.7: Incorrect graph for `piff-puff`

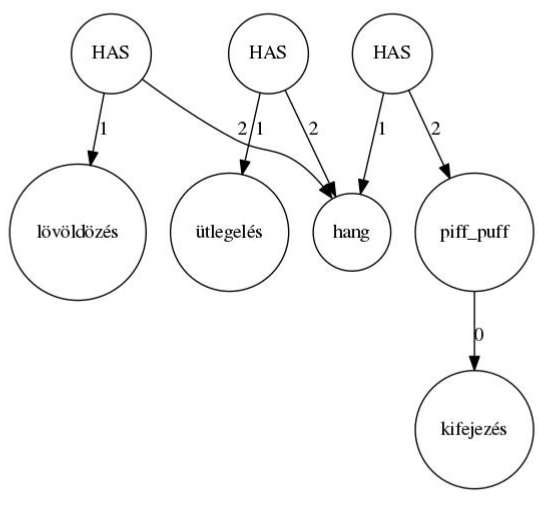

Figure 5.8: Expected graph for `piff-puff`

Chapter 6

Applications

This chapter presents applications of the `4lang` system. Section 6.1 presents our approaches to measuring semantic similarity between words and sentences using `4lang` graphs, and resulting systems submitted to the *Semantic Textual Similarity* tasks of two SemEval conferences[1]. Section 6.2 presents two early attempts at natural language understanding systems that use *spreading activation* over `4lang` graphs. The 2014 and 2015 SemEval systems described in Sections 6.1.5 and 6.1.6 are results of joint work with Judit Ács and Katalin Pajkossy, respectively. The 2012 systems presented in Section 6.2 were built in cooperation with Dávid Nemeskey and Attila Zséder.

6.1 Semantic similarity

To demonstrate the use for concept graphs built using `dict_to_4lang`, we participated in SemEval tasks concerned with measuring semantic similarity. The methods used in state of the art systems to measure sentence similarity rely heavily on word similarity, typically derived from word embeddings (see Section 2.4). We demonstrate that a simple measure of similarity between `4lang` graphs is a competitive measure of semantic word similarity. In the systems we submitted to Semeval competitions in 2015 and 2016 we combined `4lang` similarity with features derived from various word embeddings, lexical resources like WordNet, and surface forms of words. Our participation in the SemEval competitions of 2015 and 2016 are described in detail in (Recski & Ács, 2015) and (Recski & Pajkossy, 2016), respectively.

[1] `http://alt.qcri.org/semeval2016/`

Compare Two Similar Sentences

Score how similar two sentences are to each other according to the following scale.

The sentences are:

(5) Completely equivalent, as they *mean the same thing*.
(4) Mostly equivalent, but some *unimportant details differ*.
(3) Roughly equivalent, but some *important information differs/missing*.
(2) Not equivalent, but *share some details*.
(1) Not equivalent, but are *on the same topic*.
(0) On different topics.

Select a similarity rating for each sentence pair below:

Figure 6.1: Instructions for annotators of the STS datasets (Agirre et al., 2012, p.3)

6.1.1 The STS task

The SemEval conferences, which organize shared tasks in various applications of computational semantics, have featured tracks on Semantic Textual Similarity (STS) every year since 2012. While the datasets used have changed annually, the task has remained unchanged in all evaluations: participating systems are expected to measure the degree of semantic similarity between pairs of sentences. Datasets used in recent years were taken from a variety of sources (news headlines, image captions, answers to questions posted in online forums, answers given by students in classroom tests, etc.). Gold annotation was obtained by crowdsourcing (using Amazon Mechanical Turk), annotators were required to grade sentence pairs on a scale from 0 to 5; Figure 6.1 shows the instructions they were given. Inter-annotator agreement was calculated to ensure the high quality of annotations.

6.1.2 Architecture of the `MathLingBudapest` systems

Our framework for measuring semantic similarity of sentence pairs is based on the system of (Han et al., 2013), who were among the top scorers in all STS tasks since 2013 (Kashyap et al., 2014; Han et al., 2015). Their architecture, *Align and Penalize*, involves computing an alignment score between two sentences based on some measure of word similarity. We have chosen to reimplement this system because it allowed us to experiment with various measures of word similarity, including those based on `4lang` graphs built by `dict_to_4lang`, which we shall present in Section 6.1.4. We reimplemented virtually all rules and components described by (Han et al., 2013) for experimentation but will now describe only those that ended up in at least one of the 8 configurations submitted to

SemEval in 2015 and 2016 (the particular setups are described in Sections 6.1.5 and 6.1.6).

The core idea behind the *Align and Penalize* architecture is, given two sentences S_1 and S_2 and some measure of word similarity, to align each word of one sentence with some word of the other sentence so that the total similarity of word pairs is maximized. The mapping need not be one-to-one and is calculated independently for words of S_1 (aligning them with words from S_2) and words of S_2 (aligning them with words from S_1). The score of an alignment is the sum of the similarities of each word pair, normalized by sentence length, the final score assigned to a pair of sentences is the average of the alignment scores for each sentence.

In our top-scoring 2015 system, as well as in all configurations in 2016, we used supervised learning to establish the weights with which each source of word similarity contributes to the similarity score assigned to a pair of words. For out-of-vocabulary (OOV) words, i.e. those that are not covered by the component used for measuring word similarity, we rely on string similarity: we measure the Dice- and Jaccard-similarities (Dice, 1945; Jaccard, 1912) over the sets of character n-grams in each word for $n = 1, 2, 3, 4$. Additionally, we use simple rules to detect acronyms and compounds: if a word of one sentence that is a sequence of 2-5 characters (e.g. *ABC*) has a matching sequence of words in the other sentence (e.g. *American Broadcasting Company*), all words of the phrase are aligned with this word and receive an alignment score of 1. If a sentence contains a sequence of two words (e.g. *long term* or *can not*) that appear in the other sentence without a space and with or without a hyphen (e.g. *long-term* or *cannot*), these are also aligned with a score of 1.

The word similarity component can also be influenced by a boost feature based on WordNet (Miller, 1995). Scores are assigned if one word is a hypernym of the other, if one appears frequently in glosses of the other, or if they are derivationally related. For the exact cases covered and a description of how the boost is calculated, the reader is referred to (Han et al., 2013).

The similarity score may be reduced by a variety of penalties, which we only enabled in our submissions for Task 1 of the 2015 SemEval (Semantic Similarity in Twitter) – they haven't improved our results on any other dataset. Of the penalties described in (Han et al., 2013) we only used the one which decreases alignment scores if the word similarity score for some word pair is very small (< 0.05). For the Twitter tasks in 2015 we also introduced two new types of penalties based on our observations of error types in Twitter data: if one sentence starts with a question word and the other one does not or if one sentence contains a past-tense verb and the other does not, we reduce the overall score by $1/(L(S_1) + L(S_2))$, where $L(S_1)$ and $L(S_2)$ are the numbers of words in each sentence.

6.1.3 Machine learning

In our 2015 submissions our hybrid systems were trained using plain least squares regression on training data available from earlier years. By 2016 we developed a more sophisticated framework, allowing us to test various ML methods and perform feature selection.

6.1.4 Word similarity in `4lang`

The `4lang`-similarity of two words is the similarity between the `4lang` graphs defining them. We developed a measure of graph similarity by testing simple versions directly in our STS systems described in Section 6.1.2. To define the similarity of two `4lang` graphs, we start by the intuition that similar concepts will overlap in the elementary configurations they take part in: they might share a 0-neighbor, e.g. `train` $\xrightarrow{0}$ `vehicle` $\xleftarrow{0}$ `car`, or they might be on the same path of 1- and 2-edges, e.g. `park` $\xleftarrow{1}$ `IN` $\xrightarrow{2}$ `town` and `street` $\xleftarrow{1}$ `IN` $\xrightarrow{2}$ `town`.

For ease of notation we define the *predicates* of a node as the set of elementary configurations it takes part in. For example, based on the definition graph in Figure 3.3, we say that the predicates of the concept `bird` (P(bird)) are {`vertebrate; (HAS, feather); (HAS, wing); (MAKE, egg)`}. Our initial version of graph similarity is the Jaccard similarity of the sets of predicates of each concept, i.e.

$$S(w_1, w_2) = J(P(w_1), P(w_2)) = \frac{|P(w_1) \cap P(w_2)|}{|P(w_1) \cup P(w_2)|}$$

Early experiments lead us to extend the definition of predicates by allowing them to be inherited via paths of 0-edges, e.g. `(HAS, wing)` is considered a predicate of all concepts for which $\xrightarrow{0}$ `bird` holds. We have also experimented with similarity measures that take into account the sets of all nodes accessible from each concept in their respective definition graph ($N(w)$). This proved useful in establishing that two concepts which would otherwise be treated as entirely dissimilar are in fact somewhat related. For example, given the definitions of the concepts `casualty` and `army` in Figure 6.2, the node `war` will allow us to assign nonzero similarity to the pair `(army, casualty)`. We found it most effective to use the maximum of these two types of similarity.

Testing several versions of graph similarity on past years' STS data, we found that if two words w_1 and w_2 are connected by a path of 0-edges, it is best to assign to them a similarity of 1. This proved very efficient for determining semantic similarity of the most common types of sentence pairs in the SemEval datasets. Two descriptions of the

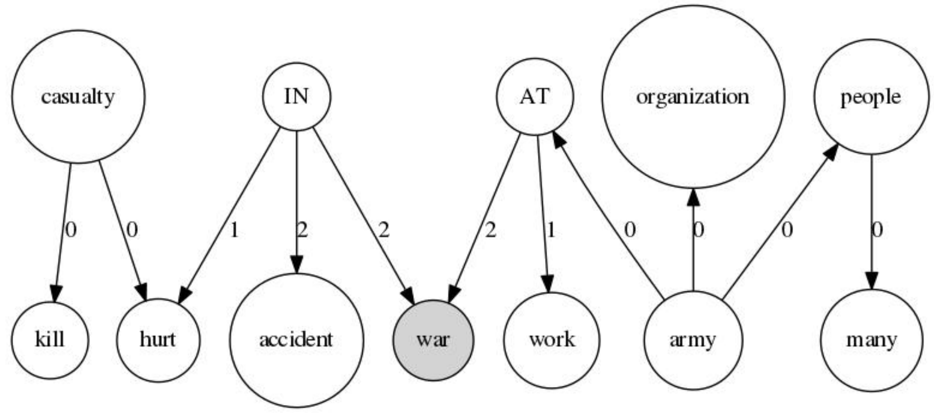

Figure 6.2: Overlap in the definitions of `casualty` (built from LDOCE) and `army` (defined in `4lang`)

same event (common in the *headlines* dataset) or the same picture (in *images*) will often only differ in their choice of words or choice of concreteness. In a dataset from 2014, for example, two descriptions, likely of the same picture, are *A bird holding on to a metal gate* and *A multi-colored bird clings to a wire fence*. Similarly, a pair of news headlines are *Piers Morgan questioned by police* and *Piers Morgan Interviewed by Police*. *wire* is by no means a synonym for *metal*, nor does being *questioned* mean exactly the same as being *interviewed*, but treating them as perfect synonyms proved to be an efficient strategy for the purpose of assigning similarity scores that correlate highly with human annotators' judgments.

The single word similarity score defined above was used as a component in our 2015 STS systems (see Section 6.1.5 for details). In our 2016 architecture (described in Section 6.1.3) the machine learning component has direct access to each component of the `4lang` similarity score; for each word pair, `4lang` returns a set of scores, each of which serve as individual features. Three of these features are binary: e.g. 1 if the two words are on the same path of 0-edges and 0 otherwise, three others can take any value between 0 and 1, e.g. the Jaccard-similarity of the sets of predicates for each word; all six features are listed in Table 6.1. `4lang` similarity was calculated separately using each of the three dictionaries described in Chapter 5, and all three configurations were run with and without expansion of the definition graphs (see Section 5.3). Altogether the `4lang` similarity component returns $3 \times 2 \times 6 = 36$ features.

feature	definition
links_jaccard	$J(P(w_1), P(w_2))$
nodes_jaccard	$J(N(w_1), N(w_2))$
links_contain	1 if $w_1 \in P(w_2)$ or $w_2 \in P(w_1)$, 0 otherwise
nodes_contain	1 if $w_1 \in N(w_2)$ or $w_2 \in N(w_1)$, 0 otherwise
0_connected	1 if w_1 and w_2 are on the same directed path of 0-edges, 0 otherwise

Table 6.1: 4lang similarity features (2016)

6.1.5 STS 2015

In 2015 we participated in two SemEval Tasks: Task 1 - *Paraphrase and Semantic Similarity in Twitter* (Xu et al., 2015) involved detecting paraphrases among tweets (Task 1a) and measuring the semantic similarity between them (Task 1b). Task 2 - *Semantic Textual Similarity* (Agirre et al., 2015) involved measuring the similarity between sentence pairs from a variety of sources.

Datasets

In 2015, STS systems were evaluated on a mixed dataset compiled from 5 sources: the `headlines` data contained titles of news articles gathered from several sources. The `images` dataset contained descriptions of images taken sampled from a set of 1000 images with 10 descriptions each. Half of sentence pairs were descriptions of the same image, the other half described different ones. The `answers-student` dataset contains answers given by pupils to an automated tutoring system during a session on basic electronics. Pairs of one-sentence answers were selected based on string similarity. The `answers-forums` dataset contains pairs of responses from the StackExchange Q&A website; some pairs are responses to the same question, others were written in reply to different ones. Finally, the `belief` data contains pairs of user comments on online discussion forums. Pairs were sampled based on string similarity, then annotated and filtered based on inter-annotator agreement. For details on the origins of each dataset, see (Agirre et al., 2015).

Submissions

For Task 1 we submitted two systems: `twitter-embed` uses a single source of word similarity, a word embedding built from a corpus of word 6-grams from the Rovereto Twitter N-Gram Corpus[2] using the `gensim`[3] package's implementation of the method presented in

[2] http://clic.cimec.unitn.it/amac/twitter_ngram/
[3] http://radimrehurek.com/gensim

	embedding	hybrid
Task 1a: *Paraphrase Identification*		
Precision	0.454	0.364
Recall	0.594	0.880
F-score	**0.515**	**0.515**
Task 1b: *Semantic Similarity*		
Pearson	0.229	**0.511**

Table 6.2: Performance of submitted systems on Task 1.

	embedding	machine	hybrid
Task 2a: *Semantic Similarity*			
answers-forums	0.704	0.698	0.723
answers-students	0.700	0.746	0.751
belief	0.733	0.736	0.747
headlines	0.769	0.805	0.804
images	0.804	0.841	0.844
mean Pearson	0.748	0.777	**0.784**

Table 6.3: Performance of submitted systems on Task 2.

(Mikolov, Chen, et al., 2013). Our second submission, `twitter-mash` combines similarities based on character ngrams, two word embeddings (built from 5-grams and 6-grams of the Rovereto corpus, respectively) and the `4lang`-based word similarity described in Section 6.1.4. For Task 2 (Semantic Textual Similarity) we were allowed three submissions. The `embedding` system uses a word embedding built from the first 1 billion words of the English Wikipedia using the `word2vec`[4] tool (Mikolov, Chen, et al., 2013). The `machine` system uses the word similarity measure described in Section 6.1.4 (both systems use the character ngram baseline as a fallback for OOVs). Finally, for the `hybrid` submission we combined these two systems and the character-similarity.

Evaluation

Our results on each task are presented in Tables 6.2 and 6.3. In case of Task 1a (Paraphrase Identification) our two systems performed equally in terms of F-score and ranked 30th among 38 systems. On Task 1b the hybrid system performed considerably better than the purely vector-based run, placing 11th out of 28 runs. On Task 2 our hybrid system ranked 11th among 78 systems, the systems using the word embedding and the `4lang`-based similarity alone (with string similarity as a fallback for OOVs in each case) ranked 22nd and 15th, respectively.

[4] `https://code.google.com/p/word2vec/`

6.1.6 STS 2016

Datasets

Submissions

Evaluation

6.1.7 Difficulties

We have obtained from `4lang` graphs a measure of word similarity that we successfully combine with vector-based metrics to create state of the art STS system, but we cannot expect our metric to outperform distributional similarity on its own. Here we discuss some of the more typical issues that we encountered.

Lack of inferencing

Without performing some inference on the concept graphs built from dictionary definitions, the near-synonyms **wizard** - *a man who is supposed to have magic powers* and **magician** - *a man in stories who can use magic* will be assigned a score of only 0.182 by our system; a higher score is not warranted by the knowledge that both concepts refer to men and that both have some connection to magic. In this example the task is as difficult as realizing that the subgraphs $\texttt{X} \xleftarrow{1} \texttt{HAS} \xrightarrow{2} \texttt{power} \xrightarrow{0} \texttt{magic}$ and $\texttt{X} \xleftarrow{1} \texttt{CAN} \xrightarrow{2} \texttt{use} \xrightarrow{2} \texttt{magic}$ refer to roughly the same state-of-affairs. This kind of inference is beyond the system as currently implemented, but well within the capabilities of `4lang`, see (Kornai, 2016) for a discussion.

OOVs

Another significant source of errors were out-of-vocabulary words (OOVs). Given the sources of input data, named entities (e.g. in `headlines`) and non-standard orthography (e.g. `forums`) are often unknown for both word embeddings and `4lang`. Character similarity can mitigate these effects significantly, but in the future we must reduce OOV-rates of all components, e.g. by training embeddings on larger datasets, building `4lang` definitions from additional resources (e.g. the Urban Dictionary) and by improving the quality of lemmatization.

6.2 Natural language understanding

We now summarize our earliest application of the `4lang` representation, a dialogue system using spreading activation over `4lang`-machines, presented in detail in (Nemeskey et al., 2013). Two systems mimicking the actions of a ticket clerk at a Hungarian railway station (one selling tickets and another responding to timetable inquiries) use Eilenberg-machines – the formal objects behind `4lang` graphs that are viewed as directed graphs of concepts throughout this thesis – to represent user input at all levels of analysis. Words and chunks detected in user input are represented by machines, as are entire utterances after processing. User input is first processed by standard tools: a morphological analyzer (Trón et al., 2005) and an NP chunker (Recski & Varga, 2009). *Constructions* over machines take over in the next step, pairing surface structures with arbitrary actions, in this case filling slots of Attribute Value Matrices (AVMs) with domain-specific fields such as DESTINATION[5]. For example, when encountering *Gödre* 'to Göd'), a noun phrase in sublative case that also contains the name of a Hungarian town, the DESTINATION field can be populated.

Simple rules such as this one are responsible for storing domain-specific knowledge extracted from user input, but a domain-independent activation of machines corresponding to `4lang` concepts governs the actions taken by the system. For each concept found in the input, machines are added to the set of *active* machines and expanded, using either their `4lang` definitions (e.g. in the case of `ticket`) or an external dictionary storing domain-specific information, e.g. that `student` and `pensioner` can be synonyms for `half-price` in the context of train tickets. At every iteration of the activation process, concepts are also activated if all concepts in their definitions are active at the end of the previous iteration. Other interfaces of the system can activate machines and fill AVMs, e.g. the location of the user can activate the concepts `ticket` and `schedule`), and populate the ticket-AVM field SOURCE with the name of the station (which may later be overridden based on user input).

The system was built to respond perfectly to ca. 40 real-life dialogues – transcribed by the author over a 30-minute period at a Budapest railway station and informally referred to as the MÁV-corpus (MÁV is the largest railroad company in Hungary). Our system was never formally evaluated with human users, but was presented to the public, spawning considerable interest (Szedlák, 2012; nyest.hu, 2012). All code is available under an MIT license from http://www.github.com/kornai/pymachine/, but while most components of the software are still used by the `4lang` module, a working system for serving railroad-related requests is no longer actively maintained.

[5] `Construction` objects in the `pymachine` module – a dependency of `4lang`– are not introduced in this thesis, but Section 8.4 will briefly mention some more applications. AVM filling is performed by subtypes of the `Operator` class, also not documented here.

Chapter 7

System architecture

This chapter describes the main building blocks of the `4lang` system. The most up-to-date version of this document is available under https://github.com/kornai/4lang/tree/master/doc. Besides introducing the main modules `dep_to_4lang` (Section 7.3) and `dict_to_4lang` (Section 7.4), which were introduced in Chapters 4 and 5 repsectively, this chapter also describes auxiliary components such as the `Lemmatizer` and `Lexicon` classes (Sections 7.6 and 7.5) as well as some modules of the `pymachine` library used by `4lang` (Section 7.7). Section 7.2 lists the external dependencies of the `4lang` module along with brief instructions on how to obtain and install them. The purpose of the first section (7.1) is to make this chapter accessible on its own, those who have read Chapters 3 through 6 of this thesis may safely skip it. Finally, Section 7.8 gives detailed instructions on how to customize each `4lang` tool using configuration files.

7.1 Overview

The `4lang` library provides tools to build and manipulate directed graphs of concepts that represent the meaning of words, phrases and sentences. `4lang` can be used to

- build concept graphs from plain text (`text_to_4lang`)
- build concept graphs from dictionary definitions (`dict_to_4lang`)
- measure semantic similarity of concept graphs
- (experimental) measure entailment between concept graphs

Both `text_to_4lang` and `dict_to_4lang` rely on the Stanford CoreNLP (English) and the `magyarlanc` (Hungarian) toolchains for generating dependency relations from text, which are in turn processed by the `dep_to_4lang` module.

The top-level file 4lang contains a manually built concept dictionary, mapping ca. 3000 words to 4lang-style definition graphs. Graphs are specified using a simple human-readable format, partially documented in (Kornai et al., 2015) (a more complete description is forthcoming). Definitions in the 4lang dictionary can be processed using the definition_parser module of the pymachine library (see Section 7.7).

The text_to_4lang module takes as its input raw text, passes it to the Stanford CoreNLP package for dependency parsing and coreference resolution, than calls the dep_to_4lang module to convert the output into interconnected Machine instances. The dict_to_4lang tool builds graphs from dictionary definitions by extending the pipeline with parsers for several machine-readable monolingual dictionaries and some genre-specific preprocessing steps.

7.2 Requirements

7.2.1 pymachine

the pymachine library is responsible for implementing machines, graphs of machines, and some more miscellaneous tools for manipulating machines. The library is documented in Section 7.7. The library can be downloaded from http://www.github.com/kornai/pymachine and installed by running python setup.py install from the pymachine directory.

7.2.2 hunmorph and hundisambig

The lemmatizer class in 4lang, documented in Section 7.6 uses a combination of tools, two of which are the hunmorph open-source library for morphological analysis and the hundisambig tool for morphological disambiguation. The source code for both can be downloaded from http://mokk.bme.hu/en/resources/hunmorph/, the pre-built models for English and Hungarian, morphdb.en and morphdb.hu, are also made available. Alternatively, pre-compiled binaries for both hunmorph and hundisambig are available at http://people.mokk.bme.hu/~recski/4lang/huntools_binaries.tgz, they can be expected to work on most UNIX-based systems. The archive should be extracted in the 4lang working directory, which will create the huntools_binaries directory. If binaries need to be recompiled, they should also be copied to this directory, or the value of the parameter hunmorph_path must be changed in default.cfg to point to an alternative directory.

7.2.3 Stanford Parser and CoreNLP

4lang runs the Stanford Parser in two separate ways. When parsing dictionary definitions, the stanford_wrapper module launches the Jython-based module stanford_parser.py, which can communicate directly with the Stanford Parser API to enforce constraints on the parse trees (see Section 5.2.2 for details). These modules require the presence of the Stanford Dependency Parser, which can be obtained from http://nlp.stanford.edu/software/lex-parser.shtml#Download and the Jython tool, available from http://www.jython.org/downloads.html. After downloading and installing these tools, all you need to do is edit the 'stanford' and 'corenlp' sections of the default configuration file 'conf/default.cfg' so that the relevant fields point to your installations of each tool and your copy of the englishRNN.ser.gz model (details on the config file will be given in Section 7.8).

The text_to_4lang tool, on the other hand, runs parsing as well as coreference resolution using the Stanford CoreNLP package. To save the overhead of loading multiple models each time text_to_4lang is run, CoreNLP is run using the corenlp-server tool, which takes care of downloading CoreNLP, then launching it and keeping it running in the background, allowing text_to_4lang to pass requests to it continuously. The corenlp-server tool can be downloaded from https://github.com/kowey/corenlp-server, then instructions in its README should be followed to launch the server.

7.3 dep_to_4lang

The core module for building 4lang graphs from text is the dep_to_4lang module which processes the output of dependency parsers. The text_to_4lang module only contains glue code for feeding raw text to Stanford CoreNLP and passing the output to dep_to_4lang. The dict_to_4lang module, which parses and preprocesses dictionary definitions before passing them to CoreNLP, will be described in the next section.

The dep_to_4lang module processes for each sentence the output of a dependency parser, i.e. a list of relations (or *triplets*) of the form $R(w_1, w_2)$, and optionally a list of coreferences, i.e. indications that a group of words in the sentence all refer to the same entity (this is currently available for English, using the Stanford Coreference Resolution system from the CoreNLP library). The configuration passed to the DepTo4lang class upon initialization must point to a file containing a map from dependencies to 4lang edges and/or binary relations. For English the default map is the dep_to_4lang.txt file in the project's root directory.

The core method of the dep_to_4lang module is
DepTo4lang.get_machines_from_deps_and_corefs, which expects as its parameter not

just a list of dependencies but also the output of coreference resolution, which is called by `text_to_4lang` but not by `dict_to_4lang`. This function will ultimately return a map from surface word forms to `Machine` instances. To create machines, the function requires the dependencies to also contain each word's lemma - for Hungarian data these are extracted from the output of `magyarlanc` by `magyarlanc_wrapper`, for English data the `Lemmatizer` module is called (see Section 7.6). Dependency triplets are iterated over, Machines are instantiated for each lemma, and the `apply_dep` function is called for each triple of `(relation, machine1, machine2)`.

The `apply_dep` function matches such triplets against `Dependency` instances that have been created by parsing the `dep_to_4lang.txt` file containing the mapping from dependency relations to `4lang` configurations. In order to handle morphological features in Hungarian data, these patterns may make reference to the `MSD` labels of words which have also been extracted from the `magyarlanc` output. In case of a match, `Operators` associated with the dependency are run on the machines to enforce the specific configurations[1].

7.4 `dict_to_4lang`

The `dict_to_4lang` module implements the pipeline that builds `4lang` graphs from dictionary entries by connecting a variety of dictionary parsers, a module for preprocessing dictionary entries (`EntryPreprocessor`), and a custom wrapper for the Stanford Parser (`stanford_parser.py`) written in `Jython` that allows adding custom constraints to the parsing process. The output from dependency parsers is passed by `dict_to_4lang` to `dep_to_4lang`, the resulting graph of `4lang` concepts is used to construct the definition graph for each headword in the dictionary, which are then saved using the `Lexicon` class (see Section 7.5).

7.4.1 Parsing dictionaries

`dict_to_4lang` supports 5 input data formats:

- an XML version of the Longman Dictionary of Contemporary English
- a typographer's tape version of the Collins COBUILD Dictionary from the ACL/DCI dataset (https://catalog.ldc.upenn.edu/LDC93T1)

[1]We do not document the `Operator` class, which is used to define complex actions over `Machines` that may be sensitive to some input data. In its current state the codebase makes no more use of them as it does of `Machines`: they are elaborate structures performing one or two very simple tasks; in this case, adding edges between machines. They do however play a significant role in the experimental system presented in Section 6.2 and will likely play a crucial part in `4lang`-based parsing (see Section 8.4).

- XML dumps of the English Wiktionary (https://dumps.wikimedia.org/enwiktionary/)
- an XML version of the *Magyar Nyelv Nagyszótára* (Hungarian)
- a preprocessed XML format of the *Magyar Értelmező Kéziszótár*. (Hungarian)

These datasets are processed by the modules `longman_parser`, `collins_parser`, `wiktionary_parser`, `nszt_parser`, and `eksz_parser`, respectively. All except `collins_parser` are subclasses of the `xml_parser` module. Each parser extracts a dictionary containing a list of definitions for each headword, each with part-of-speech tag (where available), and possibly other data which is not currently used by `dict_to_4lang`. Parsers also perform format-specific preprocessing if necessary (e.g. replacing abbreviated forms of frequent words with their full form in Hungarian definitions). If run as standalone applications, all five parsers will print their output in human-readable format, useful for testing.

7.4.2 Preprocessing entries

The output from parsing dictionary data is passed to the `EntryPreprocessor` module, which performs various steps that clean and simplify data before it is passed to external syntactic parsers. This module defines a list of regex patterns to be removed or replaced in definitions, and each pattern can be associated with one or more flags that are added to the entry if a replacement took place. It is therefore straightforward to define, given a new datasource, rules that will e.g. remove the string *of person* from a definition and simultaneously add the flag `person` to the entry being processed. The preprocessor also performs sentence tokenization (via `nltk.punkt`) and by default keeps only the first sentence of the first definition for each headword (but see Section 7.8 on how to change this).

7.4.3 Parsing definitions

Definitions returned by `EntryPreprocessor` are passed to one of two external tools for dependency parsing: the Stanford Parser for English definitions and the `magyarlanc` tool for Hungarian, both accessed via the python wrappers `stanford_wrapper.py` and `magyarlanc_wrapper.py`. Both wrappers use the `Subprocess` module to launch external tools; `magyarlanc` is launched directly and the Stanford Parser is used via a Jython wrapper.

The Jython wrapper

Since the `dict_to_4lang` module requires access to the Stanford Parser's API (see below for details), a wrapper (`stanford_parser.py` was written in Jython, a Java implementation of the Python interpreter that allows direct access to Java classes from Python code. The Jython module `stanford_parser.py` is not to be confused with the python module `stanford_wrapper.py`: the latter can be imported by any Python application and will launch a Jython session running the former.

Access to the Stanford Parser API is necessary to pass custom *constraints* to the parser before processing sentences, limiting the types of possible parse trees. Currently this feature is used to enforce that dictionary definitions of nouns get parsed as noun phrases (NPs). When using the `parse_definitions` function for parsing, part-of-speech tags for each entry are passed to the `get_constraints` function, which returns a list of `ParserConstraint` instances – currently a list of length 0 or 1 (more ParserConstraints can be created from regex `Patterns`).

7.5 The `Lexicon` class

The `Lexicon` class stores `4lang` definitions for words, separating the manually written ones in the `4lang` dictionary from those built by the `dict_to_4lang` module. When invoked from the command line, `Lexicon.py` processes the `4lang` dictionary (using the `definition_parser` module of the `pymachine` library) and saves the resulting `Lexicon` instance in pickle format. `dict_to_4lang` loads the lexicon built from `4lang`, adds definitions built from dictionaries, and saves the output. All other applications can load any of the pickle files to use the corresponding `Lexicon` instance. Applications typically use the `get_machine` function to obtain the `4lang` definition graph for some word. By default, `get_machine` first searches for definitions of a word in `4lang`, then among words for which graphs have been built automatically, and finally falls back to creating a new `Machine` instance with no definition (i.e. no connections to other `Machines`). The `expand` function implements expansion of definitions (see Section 5.3), adding links to all nodes in a definition taken from their own definitions. Stopwords are omitted by default, the user can specify other words that are to be skipped. Expansion does not affect definition graphs stored in the lexicon.

7.6 The Lemmatizer class

The Lemmatizer combines various external tools in trying to map words to 4lang concepts. For each word processed, the lemmatize function invokes the hunmorph morphological analyzer (using wrappers around ocamorph and hundisambig from the hunmisc library), as well as the Porter stemmer. lemmatize caches the results of each analysis step, storing for each word form it encounters the stem (according to the Porter stemmer), the list of possible morphological analyses (according to ocamorph) and the analysis chosen by hundisambig. In using all these to select the lemma to be returned, the lemmatize function supports several strategies for different applications.

If no flags are passed, lemmatize returns the output of hundisambig. The option defined can be used to pass the list of all lemmas from which lemmatize should try to return one (e.g. the list of all concepts defined) – if specified, lemmatize will return the word itself if it is defined, then try the lemma from *hundisambig*, and then go through all other lemmas proposed by ocamorph. If no match is found, the stemmed form is tried as a last resort before returning None. If the flag stemmed_first is set to True, lemmatize will run the above process on the stem first and only return to the original word form if no defined lemma is found. If defined is left unspecified and stem_first is set to true at the same time, lemmatize will act as a plain Porter stemmer, and a warning is issued. By default, Lemmatizer loads on startup a cache file of previously analyzed words. To save a new cache file (or overwrite an old one), the program using Lemmatizer must call its write_cache function.

7.7 The pymachine library

Concept graphs built by 4lang are encoded using the external library pymachine (http://www.github.com/kornai/pymachine, which implements Eilenberg machines via the Machine class. Currently 4lang uses these objects simply as graph nodes, not as Eilenberg machines. pymachine.utils provides, among others, the MachineGraph class for building, manipulating, (de)serializing and visualizing graphs of Machines. This class relies on the open-source library networkx as its backend for encoding directed graphs. The pymachine.definition_parser module provides a parser for the format used by the 4lang dictionary, generation is currently not supported, i.e. graphs created with 4lang cannot be saved in this format. pymachine also contains several modules that form the codebase of the system described in (Nemeskey et al., 2013), these are not used by the 4lang library.

7.8 Configuration

All `4lang` modules can be configured using standard Python configuration files, command line parameters have been avoided nearly everywhere. All parameters left unspecified in the cfg file passed to a module will be set to the values specified in `default.cfg`. If no configuration file is passed, defaults are used everywhere, running simple tests for most modules on data in the `test/input` directory. Options are documented in `default.cfg`, see Appendix A.

Chapter 8

Outlook

This chapter outlines our future plans for using `4lang` to solve some of the most challenging tasks in computational semantics. In Section 8.1 we mention some outstanding issues in the `4lang` library which we plan to address in the near future. We shall then proceed to briefly discuss the tasks of measuring *sentence similarity* and *entailment* (Section 8.2), *question answering* (Section 8.3), and semantics-based *parsing* (Section 8.4), arguing that each of these should be approached via the single generic task of determining the *likelihood* of some `4lang` representation based on models of context trained on other `4lang` graphs relevant to the task at hand (the context). Our plans for such a generic component are outlined in Section 8.5.

8.1 Outstanding issues

8.1.1 True homonyms

At present we do not treat multiple entries for the same word, e.g.

- club_1: an organization for people who share a particular interest or enjoy similar activities, or a group of people who meet together to do something they are interested in
- club_2: a long thin metal stick used in golf to hit the ball
- club_3: one of the four suits in a set of playing cards, which has the design of three round black leaves in a group together

In the future these will have to be accommodated by three separate `4lang` concepts, at which point we will have to implement some form of word sense disambiguation process in our applications.

8.1.2 Alternate word forms, synonyms

When processing dictionaries with `dict_to_4lang`, we do not currently handle definitions that consist of a single synonym of the headword. Resulting graphs such as `purchase` $\xrightarrow{0}$ `buy` are adequate representations of meaning, since the 0-edge warrants inheritence of all links, but explicitly replacing such words with their synonyms may have its practical advantages. The Collins Dictionary also lists alternate forms of many headwords, these could also be added to the concept dictionary, e.g. `realise` could point to the graph built from the definition of *realize*. Sometimes dictionaries give identical definitions for (perfect) synonyms, e.g. Longman defines both *vomit* and *upchuck* as *to bring food or drink up from your stomach and out through your mouth because you are ill or drunk*. Such duplicates can be detected to add the edges `vomit` $\underset{0}{\overset{0}{\rightleftharpoons}}$ `upchuck`.

8.2 Sentence similarity and entailment

In Section 6.1 we have introduced a measure of semantic similarity between words based on their `4lang` definitions which helped achieve state of the art performance on the task of measuring sentence similarity. Most top STS systems find a way to reduce this task to that of word similarity, where lexical resources such as `WordNet` and surface features such as character-based similarity can play an important role. Our current systems are no exception. We believe that the task of directly quantifying the similarity of two meaning representations amounts to detecting entailment between parts of such representations. The nature of the similarity scale (e.g. what it means for two sentences to be 70% similar) is unclear, but it can be assumed that (i) if two sentences S_1 and S_2 are perfectly similar (i.e. mean exactly the same thing), then each of them must entail the other, and (ii) if S_1 and S_2 are similar *to some extent* then there must be exist some substructures of the meanings of S_1 and S_2 such that these substructures are perfectly similar, i.e. entail each other.

The nature of these substructures is less obvious. A straightforward approach is to consider subgraphs, and assume that similarity of two representations is connected to the intersection of graphs (i.e. the intersection of the sets of edges over the intersection of the sets of nodes). For example, the sentences *John walks* and *John runs*, when interpreted in `4lang` and properly expanded, will map to graphs that share the subgraph `John` $\underset{1}{\overset{0}{\rightleftharpoons}}$ `move` $\xleftarrow{1}$ `INSTRUMENT` $\xrightarrow{2}$ `foot`. Other common configurations between graphs can also warrant similarity, e.g. *John walks with a stick* and *John fights with a stick* both map to `John` $\underset{1}{\overset{0}{\rightleftharpoons}}$ `X` $\xleftarrow{1}$ `INSTRUMENT` $\xrightarrow{2}$ `stick` for some X. If our notion of similarity could refer to shared subgraphs only, no connection could be made between `John` and `stick` and these sentences could not be judged more similar to each other than to virtually any sentence about John or about a stick being an instrument. We are therefore inclined to include such common templates in determining the similarity of two `4lang` graphs – templates are essentially graphs with some unspecified nodes. The number of such templates matching a given graph grows exponentially with the number of nodes, but we can expect the relevant templates to be of limited size and a search for common templates in two graphs seems feasible[1].

If similarity can be defined in terms of common substructures of `4lang` graphs, a definition of entailment can follow that takes into account the substructures in one graph that are also present in the other. Simply put, *John walks* entails *John moves* because the representation of the latter, $\texttt{John} \overset{0}{\underset{1}{\rightleftharpoons}} \texttt{move}$, is contained in that of the former, but entailment does not hold the other way round, because many edges for *John walks* are left uncovered by *John moves*, e.g. those in $\texttt{move} \overset{1}{\leftarrow} \texttt{INSTRUMENT} \overset{2}{\rightarrow} \texttt{foot}$. Since this asymmetric relationship between graphs – the ratio of templates in one that are present in the other – is also of a gradual nature, it is more intuitive to think of it as the extent to which some utterance *supports* the other – the term *entailment* is typically used as a strictly binary concept. *John moves* may not entail *John walks*, it nevertheless *supports* it to a greater extent than e.g. *John sings*.

How similarity and support between `4lang` graphs should be measured exactly cannot be worked out without considerable experimenting (we are trying to approximate human judgment, as in the case of the STS task in Section 6.1.1), what we argued for here is that `4lang` representations are powerful and expressive enough that the semantic relatedness of utterances can be measured through them effectively.

8.3 Question Answering

In the previous section we discussed the task of measuring the extent to which one utterance *supports* another – a relationship that differs from entailment in being gradual. A workable measure of support can take part in question answering: it can be used to rank candidates in order to determine answers that are more supported by a given context. There remains the task of finding candidates that are relevant answers to the question

[1] The `4lang` theory of representing meaning using networks of Eilenberg machines – of which our graphs are simplifications – will have the machines `walk` and `fight` inherit all properties of all machines to which they have pointers on their 0th partition; in other words they will end up with all properties of concepts that are accessible through a path of IS_A relationships, and will probably share at least some very generic properties such as `voluntary action`. The machine-equivalent of templates could then be networks of machines whose sets of properties do not necessarily contain the properties of any concept.

asked. The `text_to_4lang` pipeline offers no special treatment for questions. A wh-question such as *Who won the 2014 World Cup* are handled by all components in the same way as indicatives, creating e.g. the edges $\texttt{who} \xleftarrow{1} \texttt{win} \xrightarrow{2} \texttt{cup}$. Yes-no questions are simply not detected as such, *Did Germany win the 2014 World Cup* and *Germany won the 2014 World Cup* will map to the same `4lang` graph. In the future we plan to experiment with simple methods for finding candidates: e.g. searching for wh-questions allows us to identify the template $\texttt{X} \xleftarrow{1} \texttt{win} \xrightarrow{2} \texttt{cup(...)}$ and match it against graphs already in the context; we shall discuss how such a context might be modeled in Section 8.5.

8.4 Parsing in 4lang

For the purposes of the `4lang` modules and applications presented in this thesis, we relegate syntactic analysis to dependency parsers. In Section 4.4.1 we have seen examples of errors introduced by the parsing component, and in sections on evaluation we observed that they are in fact the single greatest source of errors in most of our applications. Our long-term plans for the `4lang` library include an integrated module for semantics-assisted parsing. Since most of our plans are unimplemented (with the exception of some early experiments documented in (Nemeskey et al., 2013)), here we shall only provide a summary of our basic ideas.

Since generic parsing remains a challenging task in natural language processing, many NLP applications rely on the output of chunkers for high-accuracy syntactic information about a sentence. Chunkers typically identify the boundaries of phrases at the lowest level of the constituent structure, e.g. in the sentence *A 61-year old furniture salesman was pushed down the shaft of a freight elevator* they would identify the noun phrases *[A 61-year old furniture salesman], [the shaft],* and *[freight elevator].* Since chunking can be performed with high accuracy across languages ((Kudo & Matsumoto, 2001; Recski & Varga, 2009)), and some of our past experiments suggest that the internal syntactic structure of chunks can also be detected with high accuracy (Recski, 2014), our first goal for `4lang` is to detect phrase-internal semantic relations directly.

The aim of parsing with `4lang` is to make the process sensitive to (lexical) semantics. Currently the phrase *blue giraffe* would be mapped to the graph $\texttt{giraffe} \xrightarrow{0} \texttt{blue}$ on the basis of the dependency relation `amod(giraffe, blue)`, warranted by a particular fragment of the parse-tree, something along the lines of $[_{NP}\ [_{A}\ \textit{blue}]\ [_{N}\ \textit{giraffe}\ ]\ ]$, which is again constructed with little or no regard to the semantics of `blue` or `giraffe`. The architecture we propose would still make use of the constituent structure of phrases, but it would create a connection between *blue giraffe* and $\texttt{giraffe} \xrightarrow{0} \texttt{blue}$ by means of a *construction* that pairs the rewrite rule `NP` $\rightarrow$ `A N` with the operation that adds the 0-edge between the concepts corresponding to the words *blue* and *giraffe*[2].

Since many dependency parsers, among them the Stanford Parser used by `dict_to_4lang`, derive their analyses from parse trees using template matching, it seems reasonable to assume that a direct mapping between syntactic patterns and `4lang` configurations can also be implemented straightforwardly. The task of ranking competing parse trees can then be supplemented by some module that ranks `4lang` representations by likelihood; what likelihood means and how such a module could be designed is discussed in Section 8.5. Thus, the problem of resolving ambiguities such as the issue of PP-attachment discussed in Section 4.4.1, e.g. to parse the sentence *He ate spaghetti with meatballs*, becomes no more difficult then predicting that $\texttt{eat} \xrightarrow{2} \texttt{meatball}$ is significantly more likely than $\texttt{eat} \xleftarrow{1} \texttt{INSTRUMENT} \xrightarrow{2} \texttt{meatballs}$. Since we plan to make such predictions based on statistics over `4lang` representations seen previously, our approach can be seen as the semantic counterpart of *data-oriented parsing* (Bod, 2008), a theory that estimates the likelihood of syntactic parses based on the likelihood of its substructures, learned from structures in some training data.

8.5 Likelihood of `4lang` representations

We have proposed the notion of support, the extent to which parts of one utterance entail parts of another, in Section 8.2, and we have also indicated in Section 8.3 that we require a model of context that allows us to measure the extent to which the context supports some utterance. Finally, in Section 8.4, we argued that a method for ranking `4lang` (sub)graphs by the extent to which the context supports them could be used to improve the quality of syntactic parsing and thereby reduce errors in the entire `text_to_4lang` pipeline. We shall refer to this measure as the *likelihood* of some `4lang` graph (given some context); we conclude this chapter by presenting our ideas for the design of a future `4lang` module that models context and measures likelihood. Given a system capable of comparing the likelihoods of competing semantic representations, we will have a chance of successfully addressing more complex tasks in artificial intelligence, such as the Winograd-schema Challenge (Levesque et al., 2011).

In Section 8.2 we introduced `4lang` *templates* – sets of concepts and paths of edges between them – as the structures shared by `4lang` graphs that are semantically related.

[2]As mentioned in Section 3.1, the directed graphs used throughout this thesis are simplifications of our formalism; the constructions in `4lang` actually map surface patterns to operations over Eilenberg-machines, in this case one that places a pointer to a `blue` machine on the 0th partition of a `giraffe` machine

Templates are more general structures than subgraphs, two graphs may share many templates over a set of nodes in spite of having only few shared edges; a previous example was the pair of sentences *John walks with a stick* and *John fights with a stick*, sharing the template $\texttt{John} \underset{1}{\overset{0}{\rightleftharpoons}} \texttt{X} \xleftarrow{1} \texttt{INSTRUMENT} \xrightarrow{2} \texttt{stick}$. Our initial approach is to think of the likelihood of some graph as some product of the likelihood of matching templates, given a model of the context. We believe that both the likelihood of templates in some context and the way they can be combined to obtain the likelihood of an utterance should be learned from the set of `4lang` graphs associated with the context. E.g. if we are to establish the likelihood of the utterance *Germany won the 2014 World Cup* and the context is a set of `4lang` graphs obtained by processing a set of newspaper articles on sports using `text_to_4lang`, our answer should be based on (i) the frequency of templates in the target `4lang` graph, as observed in the set of context graphs and (ii) our knowledge of how important each template is, e.g. based on their overall frequency in the context or among all occurrences over their sets of nodes[3].

In theory there is an enormous number of templates to consider over some graph (doubly exponential in the number of nodes), but the search space can be effectively reduced in a fashion similar to the way standard language modeling reduces the space of all possible word sequences to that of trigrams. If e.g. we consider templates of no more than 4 nodes, and we use expansion to reduce all graphs to some form of 'plain English' with a vocabulary no greater than 10^5 (in (Kornai et al., 2015) we have shown that an even greater reduction is possible, by iterative expansion `4lang` representations can be reduced to 131 primitives, possibly fewer), then the number of node sets will remain in the 10^{15} range, and while the total number of theoretically possible `4lang` graphs over 4 nodes is as high as $2^{6\binom{4}{2}} \approx 10^{12}$, we cannot expect to observe more than a fraction of them: the present `4lang` architecture in itself determines a much smaller variety.

Note that templates likely to occur in data are also mostly meaningful: e.g. templates over the graph for *Germany won the 2014 World Cup* are representations for states-of-affairs such as 'Germany won a 2014 something' ($\texttt{Germany} \xleftarrow{1} \texttt{win} \xrightarrow{2} \texttt{X} \xrightarrow{0} \texttt{2014}$), 'somebody won a world cup' ($\texttt{X} \xleftarrow{1} \texttt{win} \xrightarrow{2} \texttt{cup} \xrightarrow{0} \texttt{world}$), or 'Germany did something to a world something' ($\texttt{Germany} \xleftarrow{1} \texttt{X} \xrightarrow{2} \texttt{Y} \xrightarrow{0} \texttt{world}$) – our proposed parameters are the likelihoods of each of these parameters based on what we've learned from previous experience.

What we outlined here are merely directions for further investigation – the exact ar-

[3] At this point we must note that likelihood is not (directly related to) truth; in fact none of our previous discussions leading up to this notion makes reference to truth. Neither do we suggest that calculating likelihood can take the place of *inference* – a context may entail or contradict an utterance regardless of how *likely* the latter is; our notion is rather motivated by the various applications discussed in this chapter.

chitecture, the method of learning (including reduction of the parameter space) need to be determined by experiments, as does the question of how far such an approach can scale across many domains, genres, and large amounts of data. Our purpose was once again to argue for the expressiveness of `4lang` representations, and to indicate our plans for future research in computational semantics.

Appendices

Appendix A

Configuration file of the 4lang module

```
#When loading some cfg file in a 4lang module, unspecified parameters are
#assigned default values from this file
#Wherever possible, these values correspond to the most typical settings and
#test datasets distributed with 4lang

#Stanford Parser
[stanford]
#may in the future support using remote servers for parsing, leave it False for now
remote = False

#full path of Stanford Parser directory
dir = /home/recski/projects/stanford_dp/stanford-parser-full-2015-01-30/

#name of parser JAR file
parser = stanford-parser.jar

#name of model to load
model = englishRNN.ser.gz

#full path of jython executable
jython = /home/recski/projects/jython/jython

#Stanford CoreNLP
[corenlp]
#name of Java class to load
class_name = edu.stanford.nlp.pipeline.StanfordCoreNLP

#full path of Stanford CoreNLP directory
#CAUTION: when you change this path to point to your download, make sure it
#still ends with /*
```

```
classpath = /home/recski/projects/stanford_coreNLP/stanford-corenlp-full-2015-04-20/*

[magyarlanc]
path = magyarlanc/magyarlanc-2.0.jar

#miscellaneous data
[data]
#directory to save output of dependency parsing
deps_dir = test/deps
#directory for temporary files
tmp_dir = test/tmp

#dictionary data
[dict]
#input format
#possible values are: longman, collins, wiktionary, eksz, nszt
input_type = longman

#path to input file
input_file = test/input/longman_test.xml

#path to JSON file containing parsed dictionary entries
output_file = test/dict/longman_test.json

#text_to_4lang options
[text]
#path to input data
input_sens = test/input/mrhug_story.sens

#set to True to perform expansion on graphs built from text
expand = False

#set True to print dot files for each sentence
print_graphs = True

#path to save dot files
graph_dir = test/graphs/text

#if True, only dependency parsing will run and its output saved, but 4lang
#graphs won't be built. Useful when working with large datasets.
parse_only = False

#path to save output of parsers
deps_dir = test/deps/text
```

```
#options to control which definitions are included by dict_to_4lang
[filter]

#include multiword expressions
keep_multiword = False

#include words with apostrophes
keep_apostrophes = False

#discard all but the first definition of each headword
first_only = True

[lemmatizer]
#full path of hunmorph binaries and models
hunmorph_path = /home/recski/sandbox/huntools_binaries

#path of cache (loaded but not updated by default, see docs)
cache_file = data/hunmorph_cache.txt

#options related to 4lang graphs
[machine]
#file containing 4lang dictionary
definitions = 4lang

#extra data for 4lang, currently not in use
plurals = 4lang.plural
primitives = 4lang.primitive

#pickle file to load 4lang graphs from
definitions_binary = data/machines/4lang.pickle

#pickle file to save 4lang graphs
definitions_binary_out = test/machines/wikt_test.pickle

#pickle file to save expanded 4lang graphs
expanded_definitions = test/machines/wikt_test_expanded.pickle

#path of directory for printing dot graphs
graph_dir = test/graphs/wikt_test

[deps]
#path to the map from dependencies to 4lang edges
dep_map = dep_to_4lang.txt
#language of the mapping (en or hu)
lang = en
```

```
#options for testing the word similarity module
[word_sim]
4langpath = /home/recski/sandbox/4lang
definitions_binary = %(4langpath)s/data/machines/longman_firsts.pickle
dep_map = %(4langpath)s/dep_to_4lang.txt
graph_dir = %(4langpath)s/data/graphs/sts
batch = true

#options for experimental sentence similarity system
[sim]
similarity_type = word_test
word_test_data = ws_data/wordsim_similarity_goldstandard.txt
graph_dir = test/graphs/sts_test
deps_dir = test/deps/sts_test

#options for experimental question answering system
[qa]
input_file = test/input/clef_qa_sample.xml
output_file = test/qa/clef_qa_sample.answers
graph_dir = test/graphs/qa_test
deps_dir = test/deps/qa_test
```

References

Ács, J., Pajkossy, K., & Kornai, A. (2013). Building basic vocabulary across 40 languages. In *Proceedings of the sixth workshop on building and using comparable corpora* (pp. 52–58). Sofia, Bulgaria: ACL.

Agirre, E., Banea, C., Cardie, C., Cer, D., Diab, M., Gonzalez-Agirre, A., ... Wiebe, J. (2015). SemEval-2015 Task 2: Semantic Textual Similarity, English, Spanish and Pilot on Interpretability. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015).* Denver, CO, U.S.A..

Agirre, E., Cer, D., Diab, M., & Gonzalez-Agirre, A. (2012). SemEval-2012 Task 6: A Pilot on Semantic Textual Similarity. In *First Joint Conference on Lexical and Computational Semantics (*SEM)* (pp. 385–393). Montréal, Canada: Association for Computational Linguistics.

Banarescu, L., Bonial, C., Cai, S., Georgescu, M., Griffitt, K., Hermjakob, U., ... Schneider, N. (2013). Abstract meaning representation for sembanking. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse* (pp. 178–186). Sofia, Bulgaria: Association for Computational Linguistics. Retrieved from http://www.aclweb.org/anthology/W13-2322

Bobrow, R. J. (1979a). The RUS natural language parsing framework. *Research in Natural Language Understanding, annual report*.

Bobrow, R. J. (1979b). Semantic interpretation in PSI-KLONE. *Research in Natural Language Understanding, annual report*.

Bod, R. (2008). *The data-oriented parsing approach: Theory and application.* Springer.

Boguraev, B. K., & Briscoe, E. J. (1989). *Computational Lexicography for Natural Language Processing.* Longman.

Bohnet, B. (2010). Top accuracy and fast dependency parsing is not a contradiction. In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)* (pp. 89–97). Beijing, China: Coling 2010 Organizing Committee.

Bolinger, D. (1965). The atomization of meaning. *Language*, 555–573.

Brachman, R., & Levesque, H. (1985). *Readings in knowledge representation.* Morgan

Kaufman Publishers Inc., Los Altos, CA.

Brachman, R. J., Fikes, R. E., & Levesque, H. J. (1983). KRYPTON: A functional approach to knowledge representation. *IEEE Computer*, *10*, 67–73.

Brachman, R. J., & Schmolze, J. G. (1985). An overview of the KL-ONE knowledge representation system. *Cognitive science*, *9*(2), 171–216.

Bullon, S. (2003). *Longman dictionary of contemporary English 4*. Longman.

Collobert, R., & Weston, J. (2008). A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th International Conference on Machine Learning* (pp. 160–167). New York, NY, USA: ACM. Retrieved from http://doi.acm.org/10.1145/1390156.1390177 doi: 10.1145/1390156.1390177

DeMarneffe, M.-C., MacCartney, W., & Manning, C. (2006). Generating typed dependency parses from phrase structure parses. In *Proceedings of the 5th International Conference on Language Resources and Evaluation (LREC)* (Vol. 6, pp. 449–454). Genoa, Italy.

De Marneffe, M.-C., & Manning, C. D. (2008). Stanford typed dependencies manual [Computer software manual]. Retrieved from http://nlp.stanford.edu/software/dependencies_manual.pdf (Revised for the Stanford Parser v. 3.5.1 in February 2015)

Dice, L. R. (1945). Measures of the amount of ecologic association between species. *Ecology*, *26*(3), 297–302.

Eilenberg, S. (1974). *Automata, languages, and machines* (Vol. A). Academic Press.

Fillmore, C. J. (1977). Scenes-and-frames semantics. In A. Zampolli (Ed.), *Linguistic structures processing* (pp. 55–88). North Holland.

Foland Jr, W. R., & Martin, J. H. (2015). Dependency-based semantic role labeling using convolutional neural networks. In *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics* (pp. 279–288).

Grefenstette, E., & Sadrzadeh, M. (2015). Concrete models and empirical evaluations for the categorical compositional distributional model of meaning. *Computational Linguistics*, *41*(1), 71–118.

Groenendijk, J., & Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and philosophy*, *14*(1), 39–100.

Halácsy, P., Kornai, A., Németh, L., Rung, A., Szakadát, I., & Trón, V. (2004). Creating open language resources for Hungarian. In *Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC2004)* (pp. 203–210).

Han, L., Kashyap, A., Finin, T., Mayfield, J., & Weese, J. (2013). UMBC_EBIQUITY-

CORE: Semantic textual similarity systems. In *Proceedings of the 2nd Joint Conference on Lexical and Computational Semantics* (pp. 44–52).

Han, L., Martineau, J., Cheng, D., & Thomas, C. (2015). Samsung: Align-and-Differentiate Approach to Semantic Textual Similarity. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)* (pp. 172–177). Denver, Colorado: Association for Computational Linguistics.

Harris, Z. S. (1954). Distributional structure. *Word*, *10*, 146–162.

Hobbs, J. R. (1990). *Literature and cognition* (No. 21). Center for the Study of Language (CSLI).

Ittzés, N. (Ed.). (2011). *A magyar nyelv nagyszótára III-IV.* Akadémiai Kiadó.

Jaccard, P. (1912). The distribution of the flora in the alpine zone. *New phytologist*, *11*(2), 37–50.

Kamp, H. (1981). A theory of truth and semantic representation. In J. Groenendijk, T. Jansen, & M. Stokhof (Eds.), *Formal methods in the study of language* (pp. 277–322). Amsterdam: Mathematisch Centrum.

Karpathy, A., Joulin, A., & Li, F. F. F. (2014). Deep fragment embeddings for bidirectional image sentence mapping. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, & K. Weinberger (Eds.), *Advances in neural information processing systems 27* (pp. 1889–1897). Curran Associates, Inc.

Kashyap, A., Han, L., Yus, R., Sleeman, J., Satyapanich, T., Gandhi, S., & Finin, T. (2014). Meerkat Mafia: Multilingual and Cross-Level Semantic Textual Similarity Systems. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)* (pp. 416–423). Dublin, Ireland: Association for Computational Linguistics and Dublin City University.

Katz, J., & Fodor, J. A. (1963). The structure of a semantic theory. *Language*, *39*, 170–210.

Klein, D., & Manning, C. D. (2003). Accurate unlexicalized parsing. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics* (pp. 423–430).

Kornai, A. (2010). The algebra of lexical semantics. In C. Ebert, G. Jäger, & J. Michaelis (Eds.), *Proceedings of the 11th Mathematics of Language Workshop* (pp. 174–199). Springer.

Kornai, A. (2012). Eliminating ditransitives. In P. de Groote & M.-J. Nederhof (Eds.), *Revised and Selected Papers from the 15th and 16th Formal Grammar Conferences* (pp. 243–261). Springer.

Kornai, A. (2016). Truth or dare. In C. Condoravdi (Ed.), *Lkf.*

Kornai, A. (in preparation). *Semantics*. Retrieved from http://kornai.com/Drafts/sem.pdf

Kornai, A., Ács, J., Makrai, M., Nemeskey, D. M., Pajkossy, K., & Recski, G. (2015). Competence in lexical semantics. In *Proceedings of the fourth joint conference on lexical and computational semantics* (pp. 165–175). Denver, Colorado: Association for Computational Linguistics. Retrieved from http://www.aclweb.org/anthology/S15-1019

Kornai, A., & Makrai, M. (2013). A 4lang fogalmi szótár. In A. Tanács & V. Vincze (Eds.), *IX. Magyar Számitógépes Nyelvészeti Konferencia* (pp. 62–70).

Kudo, T., & Matsumoto, Y. (2001). Chunking with support vector machines. In *Proceedings of the second meeting of the North American Chapter of the Association for Computational Linguistics (NAACL 2001)* (pp. 1–8).

Lee, H., Peirsman, Y., Chang, A., Chambers, N., Surdeanu, M., & Jurafsky, D. (2011). Stanford's multi-pass sieve coreference resolution system at the CoNLL-2011 shared task. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning: Shared Task* (pp. 28–34).

Levesque, H. J., Davis, E., & Morgenstern, L. (2011). The winograd schema challenge. In *Aaai spring symposium: Logical formalizations of commonsense reasoning.*

Liu, F., Flanigan, J., Thomson, S., Sadeh, N., & Smith, N. A. (2015). Toward abstractive summarization using semantic representations. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 1077–1086). Denver, Colorado: Association for Computational Linguistics. Retrieved from http://www.aclweb.org/anthology/N15-1114

Miháltz, M. (2010). *Semantic resources and their applications in Hungarian natural language processing* (Doctoral dissertation, Pázmány Péter Catholic University). Retrieved from https://itk.ppke.hu/uploads/articles/163/file/Mihaltz_diss.pdf

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. In Y. Bengio & Y. LeCun (Eds.), *Proceedings of the ICLR 2013.*

Mikolov, T., Yih, W.-t., & Geoffrey, Z. (2013). Linguistic regularities in continuous space word representations. In *Proceedings of NAACL-HLT 2013* (pp. 746–751).

Miller, G. A. (1995). Wordnet: a lexical database for English. *Communications of the ACM*, *38*(11), 39–41.

Montague, R. (1970a). English as a formal language. In R. Thomason (Ed.), *Formal* *philosophy* (Vol. 1974, pp. 188–221). Yale University Press.

Montague, R. (1970b). Universal grammar. *Theoria*, *36*, 373–398.

Montague, R. (1973). The proper treatment of quantification in ordinary English. In R. Thomason (Ed.), *Formal philosophy* (pp. 247–270). Yale University Press.

Moser, M. (1983). An overview of NIKL, the new implementation of KL-ONE. *Research in Knowledge Representation and Natural Language Understanding*, 7–26.

Nemeskey, D., Recski, G., Makrai, M., Zséder, A., & Kornai, A. (2013). Spreading activation in language understanding. In *Proc. CSIT 2013* (pp. 140–143). Yerevan, Armenia: Springer.

Nickel, M., Murphy, K., Tresp, V., & Gabrilovich, E. (2015). A review of relational machine learning for knowledge graphs: From multi-relational link prediction to automated knowledge graph construction. *arXiv preprint arXiv:1503.00759*.

nyest.hu. (2012). Miből lesz a robot-MÁV-pénztáros. *Nyelv és Tudomány*. Retrieved from http://www.nyest.hu/hirek/mobol-lesz-a-robbot-mav-penztaros

Palmer, M., Gildea, D., & Kingsbury, P. (2005). The Proposition Bank: An annotated corpus of semantic roles. *Computational linguistics*, *31*(1), 71–106.

Pan, X., Cassidy, T., Hermjakob, U., Ji, H., & Knight, K. (2015, May–June). Unsupervised entity linking with abstract meaning representation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 1130–1139). Denver, Colorado: Association for Computational Linguistics. Retrieved from http://www.aclweb.org/anthology/N15-1119

Peng, X., Song, L., & Gildea, D. (2015). A synchronous hyperedge replacement grammar based approach for AMR parsing. In *Proceedings of CoNLL 2015* (p. 32).

Pennington, J., Socher, R., & Manning, C. (2014). Glove: Global vectors for word representation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*.

Pust, M., Hermjakob, U., Knight, K., Marcu, D., & May, J. (2015). Parsing English into Abstract Meaning Representation using syntax-based machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP 2015)* (p. 1143-1154).

Pusztai, F. (Ed.). (2003). *Magyar értelmező kéziszótár*. Akadémiai Kiadó.

Quillian, M. R. (1968). Word concepts: A theory and simulation of some basic semantic capabilities. *Behavioral Science*, *12*, 410–430.

Quillian, M. R. (1969). The teachable language comprehender. *Communications of the ACM*, *12*, 459-476.

Recski, G. (2014). Hungarian noun phrase extraction using rule-based and hybrid methods. *Acta Cybernetica*, *21*, 461–479.

Recski, G., & Ács, J. (2015). MathLingBudapest: Concept networks for semantic similarity. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)* (pp. 543–547). Denver, Colorado: Association for Computational Linguistics.

Recski, G., & Borbély, G. (2016). *Building concept graphs from monolingual dictionary entries.* (to appear at LREC2016)

Recski, G., Borbély, G., & Bolevácz, A. (2016). Building definition graphs using monolingual dictionaries of Hungarian. In A. Tanács & V. Vincze (Eds.), *XI. Magyar Számitógépes Nyelvészeti Konferencia [11th Hungarian Conference on Computational Linguistics].*

Recski, G., & Pajkossy, K. (2016). *MathLingBudapest: Measuring semantic similarity using 4lang definitions.* (to be submitted to Semeval-2016)

Recski, G., & Varga, D. (2009). A Hungarian NP Chunker. *The Odd Yearbook. ELTE SEAS Undergraduate Papers in Linguistics*, 87–93.

Richards, I. (1937). *The philosophy of rhetoric.* Oxford University Press.

Sinclair, J. M. (1987). *Looking up: an account of the COBUILD project in lexical computing.* Collins ELT.

Smith, R. (2015). Aristotle's logic. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Summer 2015 ed.). http://plato.stanford.edu/archives/sum2015/entries/aristotle-logic/.

Socher, R., Bauer, J., Manning, C. D., & Ng, A. Y. (2013). Parsing with compositional vector grammars. In *The 51st Annual Meeting of the Association for Computational Linguistics (ACL 2013).*

Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., & Potts, C. (2013). Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing* (pp. 1631–1642). Seattle, Washington, USA: Association for Computational Linguistics.

Sondheimer, N. K., Weischedel, R. M., & Bobrow, R. J. (1984). Semantic interpretation using KL-ONE. In *Proceedings of the 10th International Conference on Computational Linguistics and 22nd annual meeting of the Association for Computational Linguistics* (pp. 101–107).

Sultan, M. A., Bethard, S., & Sumner, T. (2015). DLS@CU: Sentence similarity from word alignment and semantic vector composition. In *Proceedings of the 9th International*

Workshop on Semantic Evaluation (SemEval 2015) (pp. 148–153). Denver, Colorado: Association for Computational Linguistics.

Szedlák, A. (2012). Felsögödig kérek egy ilyen nyugdijas. *Origo Techbázis*. Retrieved from http://www.origo.hu/techbazis/20120928-felsogodig-kerek-egy-ilyen-nyugdijas-robot-mavpenztarost-epit-a-sztaki.html

Travis, C. (1997). Pragmatics. In B. Hale & C. Wright (Eds.), *A companion to the philosophy of language.* Oxford: Blackwell.

Trón, V., Gyepesi, G., Halácsy, P., Kornai, A., Németh, L., & Varga, D. (2005). Hunmorph: open source word analysis. In M. Jansche (Ed.), *Proceedings of the ACL 2005 software workshop* (pp. 77–85). Ann Arbor: ACL.

Trubetzkoy, N. S. (1958). *Grundzüge der Phonologie.* Vandenhoeck & Ruprecht.

Turian, J., Ratinov, L., & Bengio, Y. (2010). Word representations: a simple and general method for semi-supervised learning. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics* (pp. 384–394).

Vanderwende, L., Menezes, A., & Quirk, C. (2015). An AMR parser for English, French, German, Spanish and Japanese and a new AMR-annotated corpus. In *Proceedings of NAACL-HLT* (pp. 26–30).

Vincze, V., Szauter, D., Almási, A., Móra, G., Alexin, Z., & Csirik, J. (2010). Hungarian dependency treebank. In *Proceedings of the Seventh Conference on International Language Resources and Evaluation (LREC'10).*

Wilks, Y. A. (1978). Making preferences more active. *Artificial Intelligence*, *11*, 197–223.

Xu, W., Callison-Burch, C., & Dolan, W. B. (2015). SemEval-2015 Task 1: Paraphrase and semantic similarity in Twitter (PIT). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval).*

Zhu, X., Guo, H., & Sobhani, P. (2015). Neural networks for integrating compositional and non-compositional sentiment in sentiment composition. In *Proceedings of the Fourth Joint Conference on Lexical and Computational Semantics* (pp. 1–9). Denver, Colorado: Association for Computational Linguistics.

Zimmermann, T. E. (1999). Meaning postulates and the model-theoretic approach to natural language semantics. *Linguistics and Philosophy*, *22*, 529–561.

Zsibrita, J., Vincze, V., & Farkas, R. (2013). magyarlanc: A toolkit for morphological and dependency parsing of Hungarian. In *Proceedings of RANLP* (pp. 763–771).

Zweig, G., Platt, J. C., Meek, C., Burges, C. J., Yessenalina, A., & Liu, Q. (2012). Computational approaches to sentence completion. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics* (pp. 601–610). Jeju Island, Korea: Association for Computational Linguistics.