



Minimal Dependency Translation: A Framework for Computer-Assisted Translation for Under-Resourced Languages

Michael Gasser^(✉)

Indiana University, Bloomington, IN, USA
gasser@indiana.edu
<http://homes.soic.indiana.edu/gasser/>

Abstract. This paper introduces Minimal Dependency Translation (MDT), an ongoing project to develop a rule-based framework for the creation of rudimentary bilingual lexicon-grammars for machine translation and computer-assisted translation into and out of under-resourced languages as well as initial steps towards an implementation of MDT for English-to-Amharic translation. The basic units in MDT, called **groups**, are headed multi-item sequences. In addition to wordforms, groups may contain lexemes, syntactic-semantic categories, and grammatical features. Each group is associated with one or more translations, each of which is a group in a target language. During translation, constraint satisfaction is used to select a set of source-language groups for the input sentence and to sequence the words in the associated target-language groups.

1 Introduction

For the majority of the world's languages we lack adequate resources to make use of the machine learning techniques that have become the standard for modern computational linguistics. For machine translation (MT) and computer-assisted translation (CAT), the lack is even more serious because what is required for machine learning is sentence-aligned translations, which are even less common than monolingual corpora. However, linguistic descriptions and sizable communities of native speakers do exist for many under-resourced languages, including Asian languages such as Telugu and Burmese, African languages such as Amharic and Hausa, and indigenous American languages such as Quechua and Guarani. There is thus a need for frameworks that facilitate the rapid creation of computational grammars and lexica by people and their automatic extension through the limited corpora that are available.

The lack of computational linguistic resources for a language usually correlates with a lack of written material in the language, an even more serious disadvantage for the community of speakers. The gap in available material is easily seen on Wikipedia, where the Amharic edition currently has 13,767 articles

and the Hausa edition 1,504 articles. Compare these numbers with the editions for more privileged languages with comparable numbers of speakers: 1,909,454 articles for Dutch and 1,237,519 for Polish [14]. One way to alleviate this gap between the more privileged and less privileged languages is to accelerate the translation of documents into the under-resourced languages.

For this reason, this project focuses on MT, and especially CAT, into languages like Amharic. The long-term goal is a system that allows users with little or no linguistic experience to write bilingual lexicon-grammars for low-resource languages that can also be updated on the basis of corpora, when these are available, and that can be easily integrated into a CAT system, where they are also updated on the basis of feedback from users.

This paper describes Minimal Dependency Translation (MDT), a lexical-grammatical framework for MT and CAT. The core of MDT is a lexicon of phrasal units called groups. A group's entry specifies translations to groups in one or more other languages. Our focus to date has been on the language pairs Spanish-Guarani and English-Amharic. Examples and implementation details discussed in this paper are from the English-Amharic system, called MIT'MIT'A.

2 Lexica and Grammars

2.1 Phrasal Lexica

The idea of treating phrases rather than individual words as the basic units of a language goes back at least to the proposal of a Phrasal Lexicon by Becker [3]. In recent years, the idea has gained currency within the related frameworks of Construction Grammar [13] and Frame Semantics [6] as well as in phrase-based statistical machine translation (PBSMT). Arguments in favor of phrasal units are often framed in terms of the ubiquity of idiomaticity, that is, departure from strict compositionality. Seen another way, phrasal units address the ubiquity of lexical ambiguity. If a verb's interpretation depends on its object or subject, then it may make more sense to treat the combination of the verb and particular objects or subjects as units in their own right.

Arguments based on idiomaticity and ambiguity are semantic, but they extend naturally to translation. If the meaning of a source-language phrase fails to be the strict combination of the meanings of the words in the phrase, then it is unlikely that the translation of the phrase will be the combination of the translations of the source-language words. Adding lexical context to an ambiguous word may permit an MT system to select the appropriate translation.

2.2 A Simple Phrasal Lexicon

The basic lexical entries of MDT are multi-word units called **groups**. Each group represents a catena [11]. Catenae go beyond constituents, including all combinations of elements that are continuous in the vertical dimension within a dependency tree. For example, in the sentence *I gave her a piece of my mind*, $\{I,gave\}$ and $\{gave,her,piece\}$ are catenae but not constituents of the sentence.

A catena has a head, and each MDT group must also have a head, which indexes the group within the lexicon. The other elements in the group are dependents of the head, but the group has no further structure; it is thus a *minimal* dependency structure. A group’s entry also specifies translations to groups in one or more other languages. For each translation, the group’s entry gives an **alignment**, representing inter-group correspondences between elements, as in the phrase tables of PBSMT. Entry 1 shows a simple group entry of this sort. The English group <one way or the other> with head way¹ has as its Amharic translation the group <በዘሀም ሆነ በዚያ> (*bezzihim hone bezziya*) which has its own entry in the Amharic lexicon. In the alignment, three of the words in the English group are associated with positions in the Amharic group; the others (indicated with “0”) correspond to no word in the Amharic group.

Entry 1 Group entry for *one way or the other* and its Amharic translation

<one [way] or the other>

→amh በዘሀም ሆነ በዚያ
align: [0,1,2,0,3]

2.3 The Lexicon-Grammar Tradeoff

A rudimentary lexicon with entries of this sort is simple in two senses: given an appropriate interface, a user with no formal knowledge of linguistics can add entries in a straightforward manner, and the resulting entries are easily understood. Such a lexicon permits the translation of sentences that are combinations of the wordforms in the group entries, as long as group order is preserved across the languages and there are no constraints between groups that would affect the form of the target-language words. However, such a lexicon permits no *generalization* to combinations of wordforms that are not explicit in the lexicon. It would require a group entry for every reasonably possible combination of wordforms.

At the other extreme from this simple lexicon is a full-blown grammar that is driven by the traditional linguistic concern with parsimony: every possible generalization must be “captured”. Although such a grammar has the advantage of compactness and of reflecting general principles of linguistic structure, it is difficult to write, to debug, and to understand, requiring significant knowledge of linguistics.

In the MDT project, the goal is a range of possibilities along the continuum from purely lexical (and phrasal) to syntactic/grammatical, with the emphasis on ease of entry creation and interpretation.

2.4 Lexemes

We can achieve significant generalization over simple groups consisting of wordforms by permitting lexemes in groups. As an example, consider the English

¹ In the figures, heads are enclosed in brackets.

group <lose_v hope>, where lose_v is the verb lexeme *lose*. In order to make such a group usable, the system requires knowledge of verb morphology, either in the form of a morphological analyzer or a list of wordforms associated with each lexeme in the lexicon. For example, the system needs to be able to recognize that *loses* is the third person singular present tense form of the lexeme *lose_v*. MDT assumes such a resource for the source language and in addition a part-of-speech tagger to reduce the syntactic and morphological ambiguity that can result when words are analyzed in isolation.

Because a source-language lexeme will normally be translated as a lexeme rather than a wordform, the system also requires knowledge of target-language morphology, specifically either a morphological generator or a list of wordform associated with each combination of lexeme and grammatical features. For example, the system needs to know that for the Amharic verb ጠፋ (*t'effa*)², one translation of *lose*, the forms corresponding to *loses* are ይጠፋል (*yit'efal*) and ትጠፋለች (*tit'efalečč*), for masculine and feminine respectively. Entry 2 shows the entry for the expression *lose hope* and its Amharic translation ተስፋ ቆረጠ (*tesfak'orret'e*), literally 'cut hope'.

Entry 2 Group entry for *lose hope* and its Amharic translation

<[lose_v] hope>

→amh ተስፋ ቆረጠ

align: [2,1] ; agr: [([2,1] , (tam:tam, sb:sb))]

Because this entry accommodates multiple sequences of English wordforms, we need to map these onto appropriate target-language sequences. This is accomplished through pairs of agreement features for the lexeme, constraining the corresponding target language form to agree with the source form on those features. In the example, the head lose_v and its translation in the Amharic group agree on the tense-aspect-modality (tam) and subject (sb) features. For example, if this group is selected in the translation of the sentence *John loses hope*, the head of the corresponding Amharic group will be constrained to be third person singular present tense: ጆን ተስፋ ይቆረጣል (*jon tesfa yik'ort'al*)³ For more about where the features of source-language words come from, see the section on morphosyntactic transformations below.

2.5 Lexical/Grammatical Categories

Another straightforward way to generalize across groups is to introduce syntactic or semantic categories. Consider the English expression *make fun of somebody*.

² For simplification, Amharic verb lexemes are given in their usual citation form, the third person singular masculine perfect form. In fact they are represented internally in terms of their abstract roots, the sequence of consonants that characterize words in Semitic languages.

³ In fact the system would also generate the (incorrect) feminine of the verb in this case since the group does not include the subject itself: ትቆረጠለች.

We can generalize across specific word sequences such as *made fun of her* and *made fun of the mayor* by replacing the specific wordforms in position 4 in the group with a category that includes the wordforms that can fill that position. This requires a dictionary of category labels for wordforms. Entry 3 shows how this appears in the lexicon. Category names are preceded by \$.

Entry 3 Group entry for *make fun of somebody* and its Amharic translation

<[make_v] fun of \$sbd>

→ amh \$sbd[+acc] አሻፈ.

align: [2,0,0,1] ; agr: [[(2,1), (tns:tns, sb:sb)], ([4,1), (num:num)]]

Because group positions that are filled by categories do not specify a surface form, during translation they must be merged with other groups that match the category and do specify a form. For example, in the translation of the sequence *made fun of the mayor*, position 4 in the group <make_v fun of \$sbd> may be filled by the head of the group <the mayor>. This **node merging** process is illustrated in Fig. 1.




Fig. 1. Merging of two groups in *make fun of the mayor*

Finally, the Amharic group that is the translation of *make fun of \$sbd* includes the constraint that whatever word fills the role of \$sbd in the Amharic sentence must be accusative since it is the direct object of the verb አሻፈ.

2.6 Morphosyntactic Transformations

For languages pairs, such as English and Amharic, that differ greatly in their syntax and morphology, a further elaboration of the framework permits many generalizations that save on the number of groups required. Amharic verb morphology is extremely complex, including tense-aspect-modality, subject and sometimes object agreement, as well as morphemes indicating whether the verb is negative and the main verb of a relative clause. In English almost all of these features are indicated syntactically rather than morphologically. Consider the sentence *they do not know her* and its Amharic translation አዳውቁትሃን (ayawk^w’atim). With a different subject, different object, or an affirmative rather than a negative verb, the single Amharic verb translating this sentence would differ. Clearly without including the subject, the object, or the word *not* in groups, there is no way the

Amharic morphology can end up correct. However, doing this entails a significant combinatorial explosion: for each English verb there would need to be hundreds of groups to cover all the different combinations of subject, object, polarity, and tense-aspect-modality.

We can deal with some aspects of this problem by incorporating a pre-processing phase that modifies the English input to make it more like Amharic. Following morphological analysis of the input sentence, it is matched against a set of morphosyntactic transformation rules, each of which may delete words and/or modify the features of words. One of these is shown in Entry 4. This rule matches sequences in input sentences consisting of the words *they do not* followed by a verb. It modifies the features of the verb to make its subject third person plural, its aspect imperfective (corresponding roughly to English present and future), and its polarity negative, and it deletes the first three words of the sequence.

Entry 4 Morphosyntactic transformation rule for *they do not \$v*

they do not \$v[*sb=3p,tam=impf,+neg*] ; del 1, 2, 3

Getting all of this morphology right is a challenge, and the MDT approach only handles some of the cases. Since the morphosyntactic transformation phase does not actually involve a syntactic parse of the input sentence, gaps present a problem. For example, adverbs can intervene between *not* and the main verb in English, in which case a rule such as that in Entry 4 would not apply.

3 Constraint Satisfaction and Translation

The steps in MDT translation are illustrated in Fig. 2 for the input sentence *she made fun of the mayor*.

Following tokenization of the input sentence, the wordforms in the sentence are tagged for part-of-speech and analyzed morphologically (2). Next, the sequence of analyzed words is matched against the morphosyntactic transformation rules (3). In the example sentence two rules match, one for *she* followed by a past tense verb, one for *the* followed by a noun. The first rule assigns perfective aspect to the verb and deletes *she*. The second assigns definiteness to the noun and deletes *the* (definiteness is expressed by a suffix in Amharic). Next the words or lexemes resulting from this first pass are used to look up candidate groups in the groups dictionary (4).

To complete sentence analysis, the system assigns a set of groups to the input sentence. A successful group assignment associates as many words in the sentence as possible with a group, and no word to more than one group, unless that word represents a node merging (see below). Longer groups over sequences of shorter groups. Node merging takes place during this phase; in the example, the \$sbd node in the instantiation of the group <make_v fun of \$sbd> is merged with an instantiation of the group <mayor_n> (5).

Group selection is implemented in the form of constraint satisfaction, making use of insights from the Extensive Dependency Grammar framework (XDG) [5].

aiming to develop systems that are unrestricted with respect to target language. In fact we are more interested in MT systems with low-resource languages as target languages because of the lack of documents in such languages.

Although MDT is not intended as a linguistic theory, it is worth mentioning which theories it has the most in common with. Like Construction Grammar [13] and Frame Semantics [6], it treats linguistic knowledge as essentially phrasal. Like synchronous context-free grammar (SCFG) [4], it associates multiword units in two languages, aligning the elements of the units and representing word order within each. MDT differs from SCFG in having nothing like rewrite rules or non-terminals. MDT belongs to the family of dependency grammar (DG) theories because the heads of its phrasal units are words or lexemes rather than non-terminals. However, it remains an extremely primitive form of DG, permitting only flat structures with unlabeled arcs and no relations between groups other than through the merge operation described in Sect. 2.5. This means that complex grammatical phenomena such as long-distance dependencies and word-order variability can only be captured through specific groups.

5 English-Amharic Implementation

We are in the process of creating an English–Amharic implementation of MDT, called MIT’MIT’A (ሚጥሚጥ). In doing so, we have relied on the tokenizer and POS tagger from spaCy [12], the Amharic-English dictionary of Amsalu Aklilu [2], the Amharic morphological generator within the HornMorpho system for morphological processing [8], the extensive grammatical descriptions of Amharic, and the author’s own knowledge of the grammars of English and Amharic. Although far from finalized, the implementation already contains approximately 7000 groups and 500 morphosyntactic transformation rules. The MDT framework offers a range of possibilities with respect to how many grammatical generalizations are captured through the use of morphosyntactic transformations and category nodes in groups, and MIT’MIT’A falls on the heavily grammatical end of the spectrum. Thus there are transformation rules accommodating combinations of pronoun subjects with all English tenses and modal verbs in both affirmative and negative forms. The result is that the system often, though by no means always, gets Amharic morphology right.

Evaluation of an MDT implementation should be of two types: for accuracy of the translations and for usability of the system with CAT. Since MIT’MIT’A is still under development, we have not undertaken a systematic evaluation of either of these sorts. However, we have begun to informally compare the system’s accuracy with that of existing statistical MT systems.

There are several commercial English–Amharic machine translation systems, including Google Translate [9] and Abyssinica Translator [1], developed by EthioCloud. We can highlight the strengths of MIT’MIT’A by examining the grammaticality of the output of Google Translate for the grammatical patterns that MIT’MIT’A is designed to capture. MIT’MIT’A knows how to translate all patterns consisting of a pronoun subject followed by a negative or affirmative verb in any

of eight possible English tenses/aspects and combinations of modal and main verbs, as well as combinations of transitive verbs with personal pronoun objects. Given the roughly 2,700 English verbs that the system has Amharic translations for, the result is hundreds of thousands of translatable patterns such as *she is about to break it*. Randomly selecting from the possible verb patterns, pronouns, and six common verbs, we get an idea of how well Google Translate performs on such combinations. Of the 54 resulting sentences, Google Translate outputs only one grammatically correct verb. MIT'MIT'A, on the other hand, makes only one minor mistake on three sentences, treating *him* in *write him* as a direct rather than an indirect object. MIT'MIT'A also has the advantage of returning multiple translations when there is ambiguity, for example, translating English *you* in three ways (feminine singular, masculine singular, plural). Needless to say, this is not really a fair comparison since it is based on examples that MIT'MIT'A is designed to handle, but it does give an idea of what sorts of advantages this rule-based system can have over a statistical system for a morphologically complex language in the context of limited training data.

6 Status of Project, Ongoing and Future Work

MDT code, including implementations for Spanish–Guarani and English–Amharic, is available at <https://github.com/hltldi/mainumby> and <https://github.com/hltldi/mitmita> under the GPL license.

In order to develop more complete lexicon-grammars for English–Amharic and Spanish–Guarani, we are working on methods for automatically extracting groups from the limited bilingual corpora that are available. We are also implementing an interface for the use of MDT implementations for CAT; working versions for Spanish–Guarani and English–Amharic can be found at <https://plogs.soic.indiana.edu/mainumby/> and <https://plogs.soic.indiana.edu/mitmita/>. Here it will be important to evaluate to what extent translators find their task simplified through the use of the system. Finally, since the interface records the user's translations whether or not they make use of the suggestions from MDT, there is the opportunity to update the system's lexicon-grammar on the basis of those translations.

7 Conclusions

Relatively sophisticated computational grammars, parsers, and/or generators exist for perhaps a dozen languages, and usable MT systems exist for at most dozens of pairs of languages. This leaves the great majority of languages and the communities who speak them even more disadvantaged than they were before the digital revolution. What is called for are methods that can be quickly and easily deployed to begin to record the grammars and lexica of these languages and to use these tools for the benefit of the linguistic communities. The MDT project is designed with these needs in mind. Though far from achieving our ultimate goals, we have developed a simple, flexible, and robust framework for

bilingual lexicon-grammars and MT/CAT that we hope will be a starting point for a large number of under-resourced languages.

References

1. Abyssinica Translator. <http://translator.abysinnica.com/>
2. Aklilu, A.: Amharic-English Dictionary. Addis Ababa, Kuraz (1979)
3. Becker, J.: The phrasal lexicon. In: Schank, R., Nash-Webber, B. (eds.) *Theoretical Issues in Natural Language Processing*, pp. 38–41. Association for Computational Linguistics (1975)
4. Chiang, D.: Hierarchical phrase-based translation. *Comput. Linguist.* **33**, 201–228 (2007)
5. Debusmann, R.: Extensible dependency grammar: a modular grammar formalism based on multigraph description. Ph.D. thesis, Universität des Saarlandes (2007)
6. Fillmore, C.J., Baker, C.F.: Frame semantics for text understanding. In: *Proceedings of WordNet and Other Lexical Resources Workshop, NAACL* (2001)
7. Forcada, M.L., Ginestí-Rosell, M., Nordfalk, J., O’Regan, J., Ortiz-Rojas, S., Pérez-Ortiz, J.A., Sánchez-Martínez, F., Ramírez-Sánchez, G., Tyers, F.M.: *Aperitium*: a free/open-source platform for rule-based machine translation. *Mach. Transl.* **25**, 127–144 (2011)
8. Gasser, M.: *HornMorpho*: a system for morphological processing of Amharic, Oromo, and Tigrinya. In: *Proceedings of Conference on Human Language Technology for Development, Alexandria, Egypt* (2011)
9. Google Translate: English-Amharic. <https://translate.google.com/#en/am/>
10. McShane, M., Nirenburg, S., Cowie, J., Zacharski, R.: Embedding knowledge elicitation and MT systems within a single architecture. *Mach. Transl.* **17**, 271–305 (2002)
11. Osborne, T., Putnam, M., Gross, T.: *Catena*: introducing a novel unit of syntactic analysis. *Syntax* **15**, 354–396 (2012)
12. Honnibal, M.: *spaCy*: industrial-strength natural language processing in python (2016). <https://spacy.io/>
13. Steels, L. (ed.): *Design Patterns in Fluid Construction Grammar*. John Benjamins, Amsterdam (2011)
14. Wikipedia (English): List of Wikipedias (2017). https://en.wikipedia.org/w/index.php?title=List_of_Wikipedias&oldid=796394102. Accessed 29 Aug 2017