[*Progress in Machine Translation*, ed. Sergei Nirenburg. Amsterdam: IOS Press, 1993.]

KIELIKONE Machine Translation Workstation

• H. JÄPPINEN, K. HARTONEN, L. KULIKOV, A. NYKÄNEN AND A. YLÄ-ROTIALA

The great majority of Finns speak a language which differs radically from main Indo-European languages. Finnish is highly inflectional and words have potentially thousands of distinct forms. Word forms carry syntactic information in their suffixes and therefore word order is relatively free in Finnish sentences. Because Finnish is syntactically so different from most other Western languages, Finns face a higher language barrier than other Western Europeans do. Increasing foreign trade has forced major Finnish companies to systematically look for ways of making language translation more productive. Machine translation would of course seem to provide an ideal solution, but in practice both the state-of-the-art of MT research and the lack of computational models of Finnish have so far discouraged the companies in their attempts to apply MT software to alleviate the translation load.

SITRA Foundation in Finland is a public fund which allocates money for projects of notable national importance. In 1982 SITRA established the KIELIKONE project for the purpose of designing computational models of the Finnish language. The short term goals were to obtain concrete language technology products; the simultaneous long term goal was to build an infrastructure for MT research. During its period of activity so far the project has designed, implemented, and introduced to the market various software products for the Finnish language: a morphological analyzer (Jäppinen and Ylilammi 1986) and spelling checkers based on that model, a morphological synthesizer (Lassila 1988), a hyphenation algorithm, and dependency parsers (Nelimarkka et al. 1984; Jäppinen et al. 1986; Valkonen et al. 1987; Lassila 1989). Also, a synonym dictionary for Finnish has been produced both in a book (Jäppinen 1989) and electronic form.

As more direct steps toward MT, the project first developed an electronic bilingual Finnish-English dictionary. Later on, upon the request of a foreign customer, the project designed and implemented an MT workstation for a syntactically and semantically constrained sublanguage (Kulikov and Jäppinen 1989; Takala et al. 1991).

In 1986 it was decided that the project should concentrate on full-scale MT research in cooperation with two major Finnish companies. Telenokia OY exports telecommunication equipment, and all their products require extensive technical documentation. English is

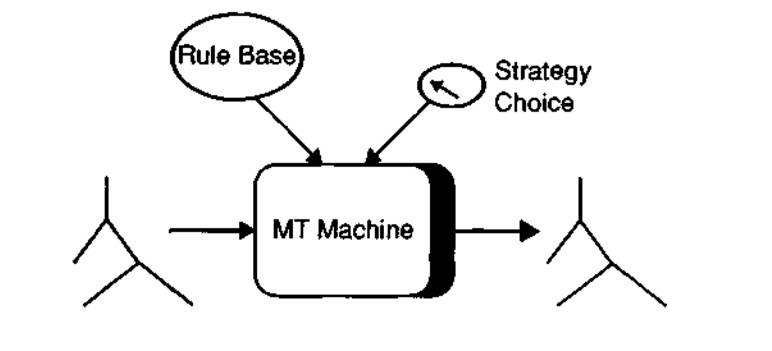

Figure 1: The MT Machine

their most important foreign language; this company is our pilot customer for the Finnish-English system. Finnair OY is the Finnish national air carrier company. Their problem is the translation of voluminous maintenance manuals from English into Finnish. This company is the pilot customer for our English-Finnish system.

The focus of our MT research has been the design of MT Workstations. By this term we mean personal computing systems which produce good quality raw translations and support post-editing with a user-friendly linguistic editor. To promote wide applicability the system architecture is designed to be maximally general (language independent), and the pans which hold language-dependent definitions are declarative. These principles have been realized in an MT Machine, which holds the algorithmic part of any given MT Workstation implementation. The MT Machine is totally language independent - it is not biased towards Finnish - and its execution is controlled by a declarative rule base.

At the moment of this writing, we have fully implemented and tested the MT Machine (in C under UNIX). Finnish-English Workstation has also been fully implemented and we are presently testing and tuning the system with real data.

1 The MT Machine

The MT Machine is a general tree-manipulation system with several built-in inference strategies. When a user applies the machine he/she writes a rule base to control the execution of the machine and chooses the appropriate inference strategy. The machine takes well-defined linguistic trees as input and produces as output trees which represent meaning-preserving transformations of the input trees (Fig. 1).

We will not discuss either the rule syntax or the inference strategies here. As for the linguistic trees, they are general feature trees (F-trees); the nodes of trees are represented by feature vectors.

Although the MT Machine is general, i.e. language independent, it does impose restrictions on what kinds of transformations are possible. The tree topology rules out, for instance, graph manipulation. The chosen rule syntax and the implemented inference strategies impose limitations of their own, but it is our belief that these restrictions are linguistically well-founded and do not constrain translations. The experience gathered with the Finnish-English Workstation system so far supports this conjecture.

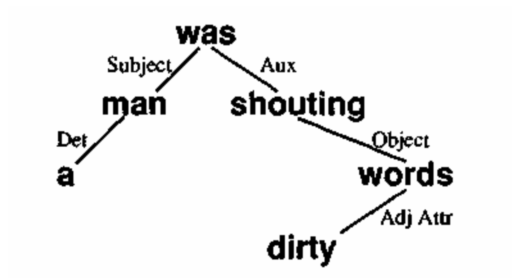

Figure 2: A dependency tree

It is important to notice in a positive sense how the MT Machine enables homogeneous processing. The data flow is in the form of F-trees throughout the process and descriptions of transformations are always rule bases (even lexicons are rule bases in our implementations). Processing corresponds to a monotonous application of F-tree transformations (Fig. 3). Homogeneity has many advantages; it means structural simplicity and thus advances clarity and maintainability.

2 Linguistic Commitments

MT Machine itself does not confine to the use of any specific linguistic theory. We have committed ourselves in our implementations to dependency theory as the model of sentence structure. We have studied dependency theory over the years and implemented parsers of Finnish based on that theory (Nelimarkka et al. 1984; Jäppinen et al. 1986; Valkonen et al. 1987; Lassila 1989). Dependency theory, we have argued, describes the sentence structure of so-called free-word-order languages better than constituent theories do.

Dependency trees do not explicitly show the constituent structure of a sentence. Instead, they exhibit the binary head-modifier relations between the words. The result of a parsing process is hence a tree whose nodes represent the words (more specifically, morpho-syntactic descriptions of the words) and whose branches represent binary dependency relations between the words of a sentence. The finite verb is the root of a full sentence. For example, the structure of "A man was shouting dirty words." is shown in Fig. 2.

It can be strongly argued that dependency theory is an advantageous representation model for MT. Dependency trees of sentences are close to their logical forms and hence closer to their meanings than the corresponding constituent trees. We do not delve into the matter here in more detail (see Schubert 1988 for a discussion on dependency theory and MT, and Mel'cuk 1988 and Starosta 1988 for general discussions on dependency theory). Dependency theory is applied in many other modern MT systems: DLT (Schubert 1988) and EUROTRA (Copeland et al. 1991) utilize it, and so do many Japanese MT systems.

Dependency structures have straightforward F-tree representations. If dependency relations are represented by their names in one feature in the dependent nodes, then an F- tree of a parsed sentence is a tree whose feature vector is a union of morphological, lexical, and relation features.

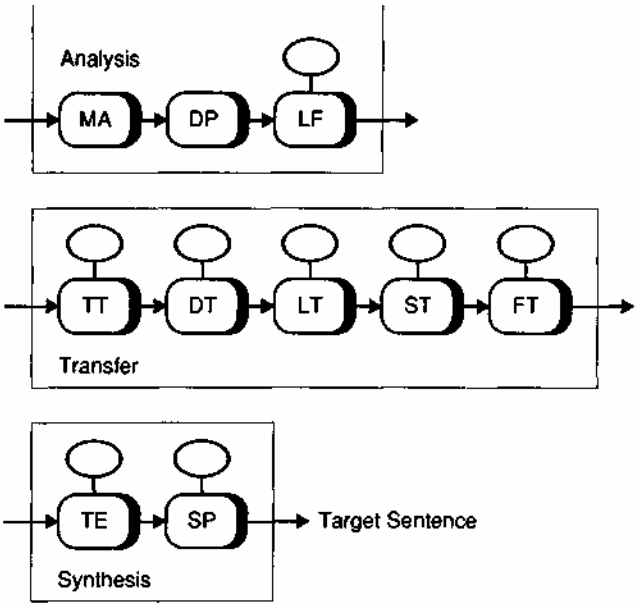

Figure 3: The translation process

3 Translation

The MT Machine and the dependency theory lend themselves naturally to a linear architecture of translation. When also each lexical transfer is described by a rule base, a possible system architecture has the simplicity of Fig. 3. That is in fact our implemented Finnish-English configuration. The MT Machine instances are marked with a special symbol.

The analysis phase includes morphological analysis (MA), dependency parsing (DP) and logical form reduction (LF). After DP and before LF data is converted into F-tree representation. Then the translation proceeds through several F-tree transformations: term and frozen phrase transfer (TT), domain- specific lexical transfer (DT), general lexical transfer (LT), structural transfer (ST), and feature transfer (FT). Then follows the synthesis phase which also utilizes the MT Machine: first the target tree expansion (TE) (inverse of logical form reduction) and then the target sentence production (SP). Each MT Machine application has its own rule base and each can choose its inference strategy independently from other phases. Notice how the sequence imposes hierarchy on the three lexical transfer phases.

The term "transfer" usually refers to projections between two languages that depend on both languages. Thus understood transfer is divided in our architecture in the subtasks shown in the figure. Transfer could of course be divided into subtasks in different ways. An administrative process, implemented on top of UNIX and not shown in the figure, controls the processes. It also includes tracing and debugging facilities.

4 MT Workstation

The translation architecture of the Finnish-English MT Workstations appears in Fig. 3. The workstation also has to provide an interface for the external world. Interaction with a user takes place through a graphic interface. The screen is divided into input and output windows which display source language and target language sentences, respectively.

The workstation concept takes post-editing seriously. One way of increasing translation quality in conjunction with positive user cooperation is to make editing and revising activities as convenient as possible. The user of the Workstation can edit the texts in the windows in different flexible ways. He/she can move text fragments around or delete or insert new words using similar services as offered by modern text editors. If necessary, he/she can also tag sentences for later scrutiny.

Another important editing function is lexical replacement. It is a well known fact that one of the greatest problems in MT is the correct lexical choice. The rules of the MT Machine permit quite elaborate contextual checks in the lexical transfer phase. However, some pragmatic factors outside the text affect translation, and these are not within the reach of any rule system. The Finnish-English system features a dictionary of translation equivalents: Finnish words with sets of possible translations (in some contexts). If the user is not satisfied with a given lexical choice in the target text, he/she can point at the word and a window with a list of alternative translations will appear on the screen. If an alternative is pointed, it will automatically replace the wrong word in the text - even in the correct form.

5 Knowledge Acquisition

The architecture of an MT system decomposes the translation task into subtasks. The architecture controls execution and imposes constraints on a system implementation. The architecture has the same function as the skeleton has for the human body: to create disposition for dexterity. Good architecture makes flexible and efficient systems possible, while bad architecture brings about rigidity and/or inefficiency. To move our bodies we need also muscles, and to translate the language into another we need lexicons and linguistic rules.

Descriptions of MT systems usually focus almost totally on architectural issues - and possibly on the syntax of linguistic and/or lexical rules - but they little pay attention to how linguistic knowledge has been acquired. Yet, languages are very complex and intricate communication systems, and to incorporate a sufficient amount of linguistic knowledge into a system so that it is capable of translating one language into another in a general fashion is a laborious task indeed, and well designed, systematic methods are in great demand.

In this short exposure we cannot discuss our knowledge acquisition method in detail but can give only a broad outline. In terms of translation theory, our transfer method is based on the hypothesis that translation is a decomposable task, that is, that a good-enough rough translation of a sentence results from the independent translations of its structural units. Having said this, the acquisition of linguistic knowledge centers around a document we call the Translation Map. The Translation Map features a contrastive analysis of the structural units of a given language pair from the viewpoint of the source language. More specifically, a Map is a depository of translation invariant structures (INTRA), showing what structures there are in the source language that can be translated into the target language in a general fashion, what the translations are, and how the translations progress through the various transfer phases.

INTRAs are extracted using an empirical method, which we visualize below by running through a simple example. For a given source language expression, which has been selected in a systematic manner, an accurate and closest possible target expression is defined. The dependency trees for both expressions are then written. Substitution tests are performed by replacing the lexical items with other items of the same type. Type similarity is a flexible notion, meaning that words belong to the same syntactic or semantic categories or subcategories. If the translation remains valid during the substitution tests, the typed pair is a. valid INTRA. If the translation is violated the “size” of the expression is decreased by either restricting the types of the lexical items or the topology of the trees. The procedure is then repeated, until a valid INTRA is found or no generalization holds. For example, this procedure zeroes in on the INTRA (1), which generalizes the translation of non-animate genitive attribute expressions from Finnish into English.

```
(1)  A:   (1, Reg=2, SRel=GenAttr, SCat=Noun, TSemCat#Animate)
     B:  (2, Reg=NIL, SRel=Head, SCat=Noun)
            >>INTRA: GenAttr10<<
     New:  (1, Reg=2, TRel=Det, TLex="the")
     B:  (2, Reg=0, TRel=Head)
     New:  (3, Reg=2, TRel=MaterAttr, TLex="of")
     A:       (4, Reg=3, TRel=PrepDep)
```

Trees in INTRAs are expressed by indented tuples, whose first three elements show the linear position of the node, its regent, and the relation of the node to its regent. The remaining elements describe the relevant features of the node; the input part shows the prerequisite attributes in the source tree, the output part shows the assigned attributes in the target tree. To the left of the tuples there is identification labels. (1) makes it known that the target equivalent unit has the premodifying genitive attribute relation transformed into a material attribute relation, which is a postmodifying OF- prepositional phrase.

An INTRA is hence a schematic mapping of a source language expression into the target language. At this point of the inquiry the units are independent of any particular translation method, save the decomposability assumption, and it is not yet clear whether the implied transformations can be carried out. The architecture, the rule syntax, and the inference strategies impose constraints on transformations. We say that an INTRA is well-defined if such rules can be written that carry out the transformation through the transfer phases, otherwise the INTRA is ill-defined. (1) is a well-defined one as (2) shows. Only well-defined INTRAs are documented in the Translation Map. Ill-defined INTRAs have to be modified by looking for other possible translations for the source expression until one with a well-defined INTRA is found. From now on, when we mention INTRAs we mean well-defined INTRAs recorded in a Map.

```
(2)  A:   (1, Reg=2, SRel=GenAttr, SCat=Noun, TSemCat\#Animate)
     B:  (2, Reg=NIL, SRel=Head, SCat=Noun)
            >>LF<<
            >>LT<<
            >>ST; GenAttr10<<
     A:   (1, Reg=2, TRel=MaterAttr, TPrep="of")
     B:  (2, Reg=NIL, TDef=Def)
            >>TE: PrepExp05<<
     New:   (1, Reg=2, TRel=Det, TLex="the")
     B:  (2, Reg=0, TRel=Head)
     New:   (3, Reg=2, TRel=MaterAttr, TLex="of")
     A:       (4, Reg=3, TRel=PrepDep)
```

INTRAs indicate the different translation phases as horizontal levels. If a phase participates in the translation its output tree is written below. Only the structural transfer phase (ST) participates in the translation of this particular INTRA, and the responsible rule is GenAttr10. (The lexical transfer phase, of course, produces the target words, but in this INTRA LT does not involve structural changes.) Term Expansion phase (TE) produces the target structure, where the preposition becomes the regent of the modifier noun and the head noun has a definite determiner as a dependent. Every INTRA shows also an example where the unit appears in the context of a whole sentence (3) and gives a few test sentences (4). The system automatically extracts the test sentences from the map and runs them through.

(3) Talon kattoa (maalataan)
>>DP<<
talo[GenAttr] katto[Head] (maalata)
>>LF<<
>>LT<<
house[GenAttr] roof[Head] (paint)
>>ST: GenAttr10<<
roof[Head, Def] house[MaterAttr, of] (paint)
>>FT<<
>>TE: PrepExp05<<
the [Det] roof [Head] of [MaterAttr] house [PrepDep] (paint)
>>SP<<
The roof of (the) house (is being painted)

(4) ##Genitive attribute, Non-animate noun
S: Talon katto pit{\"a}isi maalata.
T: The roof of the house should be painted.
S: Tunnetko maamme historiaa?
T: Do you know the history of our country?

Notice that INTRAs need not be translation invariant without any residue. Since in our architecture the lexical transfer phases have the full rule power available and these phases are invoked before the structural rules, INTRAs can be viewed as prototype translations which hold unless lexical rules override them. To continue our example, in Finnish noun phrases the genitive attribute has a broader usage than in English. Thus, the inanimate noun raha (money) can appear as a genitive attribute without any possessive association as in rahan vaihto (money exchange). This exception to the INTRA (2) is handled by the lexical rule (5).

(5) A: (1, Reg=2, SRel=GenAttr, SLex="raha")
B: (2, Reg=NIL, SRel=Head, SLex= "vaihto")
>>LT: raha<<
A: (1, Reg=2, TRel=NomAttr, TLex="money")
B: (2, Reg=NIL, TRel=Head, TLex="exchange")

The INTRAS in a Translation Map define the general transfer rules in a semi-formal manner. Since the INTRAs contain all relevant linguistic information on a high and proper level of abstraction, it would be advantageous to generate the rules automatically. So far we have written the rules by hand. A search for INTRAs should cover the source language as fully as possible in a systematic manner. The task is a very laborious one, unless there are documents available which classify syntactic structures of the source language one way or another. We have benefited from Hakulinen and Karlsson 1979, and Chesterman et al. 1979 in defining our Finnish-English Translation Map.

As a final point, we may reflect briefly upon translation quality. From the viewpoint of our theory, quality has three correlating variables. First, there is the decomposability hypothesis and its validity. Second, translation quality is adversely affected by the extent to which the implemented well-defined INTRAs deviate from high quality translations. Third, each INTRA which has no well-defined solution is a potential degrading factor. What is the overall quality of a translation is an empirical question, and an answer can be found only by testing real data.

6 Testing and Tuning

The INTRAs in the Translation Map indicate test sentences (4). The set of these sentences constitutes a simple test suite. The system extracts these sentences and runs them through. The sentences, if properly translated, give partial "proof" that the systems works as designed. The support is necessarily only a partial one because the INTRAs are general schemas and the written test sentences do not usually cover the domains of the schemas exhaustively.

A more important test of the system is its behavior when subjected to real data. If the purpose of a system is to translate texts of a given type, it is better to run the system on that text before it is installed in a production environment - regardless of other tests that have been performed until then. Testing and tuning has to us the following five objectives.

1. By running the test sentences written in the Map, we check and correct errors in the rules that realize the INTRAs.
2. By running Finnish sentences drawn from a real corpus, we test our hypothesis that translation of sentences can be decomposed into translation of the units.
3. By running tests on real data, we try to locate missing structures from our Translation Map and augment the Map accordingly with new INTRAs and the respective rules.
4. By running tests on real data, we increase the coverage of our lexicons.
5. By running tests on real data, we obtain statistical data about error frequencies in the translation rules and lexical entries.

We are currently testing and tuning the system using a general text type. As such we have chosen news items taken from the economic news section of a major newspaper in Finland. As the parser is functionally independent from the translation parts, we test them separately. The procedure for the testing of the translation part is as follows. A news item is randomly selected and the first ten sentences are typed in. The sentences are parsed and possible parsing errors are corrected by hand. The sentences are translated and errors or points of low quality are located in the translations. If an error appears in the existing rules of the current INTRAs or in the lexical entries, the errors are corrected immediately and statistical data about the corrections is recorded. If an error is due to insufficient coverage of the Translation Map, a new INTRA has to be generated, and the correction cycle is slower. The piece of text is translated again after the corrections, and if no errors remain, such that they can be treated by this method, the result is the raw translation of the text. Sometimes, however, translations have such "global" deficiencies that they must be left for a more careful analysis later on.

Figure 4: The sum of general rule errors in the test runs.

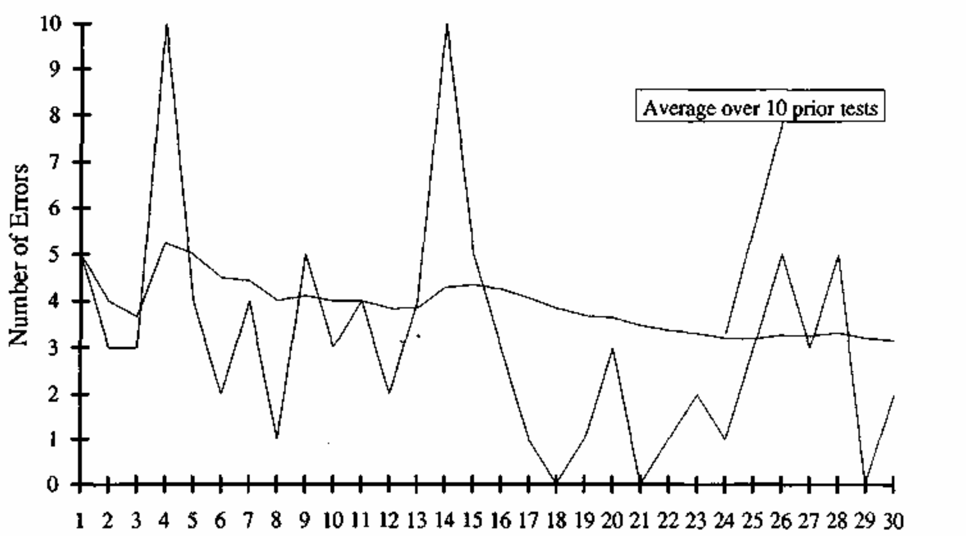

Examples of such deferred errors are the treatment of articles, insertion of commas, and the proper word order.

Appendix A shows the raw translation of one test after the correction cycle. Notice that, as just said, the articles, commas, and word order are not yet in their final shape. Fig. 4 shows the sums of the errors in the general rules (LF+ST+FT+TE+SP) found in the tests. They include both the number of errors in the existing rules and the number of new rules. The figure shows a clear downward tendency from about five errors per sentence on average at the beginning of the test to about three errors per sentence after 30 test runs.

We have also accumulated data about the respective lexical errors (LT) in the same tests. Again, the numbers include both errors in the existing entries and the new entries written. There has been about 0.8 modifications on lexical entries per sentence on average. As expected, the rate of lexical "errors" is not diminishing as clearly as the error rate of the general rules.

7 Evaluation

Figure 4 provides hard data about the behavior of the system in terms of error frequencies. Without going into precise probabilistic statements, the figure shows that the error rate of the general rules is approximately 0.3 errors per sentence on average after the 30 tests, and every ten test samples decreases the error rate by about 0.07 errors per sentence. That is, after about forty more tests of the same type the number of errors in the general rules should approach zero.

Yet, these figures say nothing about the quality of the raw translations produced by the system. It has been proposed that the quality of a translation should be judged on the basis of the intelligibility of the translations, on the one hand, and on their fidelity, on the other (Nagao 1989). Such measures rely on the inner workings of the human mind and are therefore subjective and difficult to measure precisely. A more objective measure would refer to the external behavior of the test subjects. Instead of asking a competent translator questions foreign to him/her, he/she is requested to do something natural. The most natural behavior vis-a-vis raw translations is correcting them. The quality measure of a raw translation should, then, correlate with the editing operations performed. The fingers do the talking, so to say. In addition to enhanced objectivity these measures have the additional advantage of being relative to the quality requirements set for the final texts.

We have not yet fully developed this idea. Some open questions still remain, such as should we observe “higher” level behavior by counting different high level editing operations, possibly weighing them, or should we rely on more elementary operations and weigh them equally? Presently we are experimenting with the most elementary operations - keyboard strokes and mouse clicks. A similar suggestion appears in Brown et al. 1990. We also measure the editing time. Appendix B shows an example in which the raw translation in Appendix A has been edited by one human editor. This example should not to be taken at its face value, as there was no training prior to the editing, and no quality requirements were set for the final translations. This tentative example shows that the editing intensity (the number of editing key strokes and mouse clicks in relation to the number of the characters in the text) is 18 %. Editing time rate (editing time in relation to the "standard" human translation time for Finnish-English = 1560 characters per hour) is 22 %. However, no final conclusions - one way or another - should be drawn from these preliminary figures.

8 System Status

The current Finnish-English system has about 35,000 different lexical entry words (but about twice as many entries, as some words - verbs in particular - have several entries) in the general bilingual lexicon. The coverage is quite good, as only about 0.8 additions to the lexicon are caused by an average sentence of a general news text type. The translation speed, when only the general lexical transfer is in use, is about 0.5 seconds/word in a VAXstation 3100 under ULTRIX (about 2 MIPS machine). That speed seems to be quite satisfactory because neither the MT Machine nor the rule bases have gone through optimizing efforts yet.

9 Conclusion

We have described a MT method which combats the complexity and openness of language translation problems by decomposing the task into well-defined subtasks and solving each using declarative, modular rules. The underlying computational abstraction is a general, language-independent MT Machine, designed specifically for the transformation of linguistic trees in MT. A full MT Workstation system is composed of sequential executions of that machine. We have discussed Finnish-English implementation. We have also explained the knowledge acquisition and testing methods employed.

Appendix A. The Raw Translation of One Test

- — ut3_32.src

Voin kulutus on lähtenyt tämän vuoden alkupuoliskolla kovaan nousuun.
Vuoden neljän ensimmäisen kuukauden aikana kulutus on kasvanut yli
kymmenen prosenttia. Voin suosion kasvu johtui vuodenvaihteessa

toteutetusta hinnanalennuksesta. Silloin voin kilohinta laski kuusi markkaa. Samaan aikaan margariinien kulutus on kääntynyt laskuun. Maaliskuussa margariineja kulutettiin viisi prosenttia edellisvuotista vähemmän ja huhtikuussa kulutuksen lasku oli jo yhdeksän prosenttia. Myös margariinien hinta aleni vuodenvaihteessa rasvaveron alentamisen vuoksi. Valion varatoimitusjohtaja ei halua ennustaa miten pysyvä ilmiö on. Hänen mukaansa varmempaa voidaan sanoa vasta puolen vuoden kuluttua. Voin vastainen terveyspropaganda tuntuu ihmisiltä unohtuneen.

- — ut3_32.raw

The consumption of butter has begun to rise during the first half of this year. During the first four months of the year , the consumption has increased by over ten per cent . The growth of the popularity of butter was caused by the price reduction carried out at the turn of the year . Then the price of butter per kilo fell by six marks . At the same time , the consumption of margarines has developed a downward trend . In March margarines were consumed five per cent less than the previous year and in April the decrease in the consumption was already nine per cent . Also the price of margarines fell at the turn of the year because of the reduction of the fat tax . The deputy managing director of VALIO does not want to predict how permanent the phenomenon is . According to him , something positive can be said only after half a year , . The people seem to have forgotten a health propaganda against butter .

Appendix B. The Edited Version of the Raw Translation

- — ut3_32.raw

The consumption of butter has begun to rise during the first half of this year. During the first four months of the year , the consumption has increased by over ten per cent . The growth of the popularity of butter was caused by the price reduction carried out at the turn of the year . Then the price of butter per kilo fell by six marks . At the same time , the consumption of margarines has developed a downward trend . In March margarines were consumed five per cent less than the previous year and in April the decrease in the consumption was already nine per cent . Also the price of margarines fell at the turn of the year because of the reduction of the fat tax . The deputy managing director of VALIO does not want to predict how permanent the phenomenon is . According to him , something positive can be said only after half a year , . The people seem to have forgotten a health propaganda against butter .

---- ut3_32.edt

The consumption of butter has begun to rise during the first half of the year . During the first four months of the year the consumption increased by over ten per cent . The growth of the popularity of butter was caused by the price reduction carried out at the turn of the year . At that time the price of butter per kilo fell by six marks . At the same time the consumption of margarine has developed a downward trend . In March margarines were consumed five per cent less than the previous year, and in April the decrease in consumption was already nine per cent . The price of margarines fell at the turn of the year , too, because of the reduction in the fat tax . The deputy managing director of VALIO does not want to predict how permanent the phenomenon is . According to him definite conclusions can be drawn only after six months . People seem to have forgotten the health propaganda against butter .

948 characters in the text.

Corrections: 123 keystrokes 48 mouse clicks together 171 (0.18 /character)

Time to edit: 472 s (7 min 52 s) (22% of the translation time)

References

Brown, P., J.Cocke, S.Della Pietra, V.Della Pietra, F. Jelinek, J.Lafferty, R.Mercer and P.Roossin. 1990. A statistical approach to machine translation. *Computational Linguistics*, **16**(2): 79-85

Chesterman, A., L.Korpimies, J.Lento, B.Sandlund, K.Varantola, P.Binham and A.Rantanen. 1979. *Finnish into English: An introduction to translation*. Gaudeamus.

Copeland, C., J.Durand, S.Krauwer and B. Maegaard. 1991. *The Eurotra Specification*. Studies in Machine Translation and Natural Language Processing, Volume 1. Office for Official Publications of the European Commission.

Hakulinen, A. and F.Karlsson. 1979. *Nykysuomen lauseoppia*. Suomalaisen kirjallisuuden seura.

Jäppinen, H. (ed.) 1989. *Synonyymisanakrja*. Werner Söderström Osakeyhtiö.

Jäppinen, H. and M.Ylilammi. 1986. Associative model of morphological analysis: an empirical inquiry. *Computational Linguistics*, **12**(4): 257-272.

Jäppinen, H., A.Lehtola and K.Valkonen. 1986. Functional structures for parsing dependency constraints. In *Proc. COLING86*, Bonn.

Kulikov, L.and H.Jäppinen. 1989. Automatic translation of a highly constrained language. In *Proc.SCAI'89*, Tampere.

Lassila, E. 1988. FORMA: program for Finnish word form synthesis. In *Proc. SteP-88*, Helsinki.

Lassila, E. 1989. Parsing Finnish sentences by performing functionally defined sequential subtasks. In *Proc.SCAI'89*, Tampere.

Mel'cuk, I. 1988. Dependency syntax. In: Mel'cuk, I. (ed.) *Studies in dependency syntax*. Kroma.

Nagao, M. 1989. *Machine translation: How far can it go?*, Oxford University press.

Nelimarkka, E., H. Jäppinen and A.Lehtola. 1984. Two-way finite automata and dependency theory: a parsing method for inflectional free-word-order languages. In *Proc. COLING84/ACL*, Stanford.

Schubert, K. 1988. The architecture of DLT – interlingua or double direct? In *Proc. New Directions in MT*, Budapest.

Starosta, S. 1988. *A case for lexicase*. Pinter Publishers.

Takala, J., J.Personen, L.Kulikov and H.Jäppinen. 1991. Machine translation for chemical safety information. *Journal of Hazardous Materials*, **27**: 213-230.

Valkonen, K., H.Jäppinen and A.Lehtola. 1987. Blackboard-based dependency parsing. In *Proc. 10th IJCAI*, Milan.