Science and Technology in Machine Translation

Pius ten Hacken, Universität Basel

The major question to be addressed here is the conditions under which MT can be considered scientific. Before this question can be meaningfully discussed, a number of preliminary issues should be settled. They include the distinction between applied science and other kinds of science, the distinction between applied science and technology, and the epistemological and practical consequences of these distinctions (section 1). In section 2, the problem of MT is analysed in such a way that some general choices which have to be made stand out clearly. Each of these choices may influence the character of the approach to MT in terms of the distinction between (applied) science and technology. In section 3, one of the choices, determining the boundaries of the problem to be solved, is taken up for closer analysis. It has been chosen because its influence on the scientific nature of the approach, though pervasive, is not immediately obvious.

1. Applied Science

Applied science is found at the intersection of pure science and general problem solution. The former imposes epistemological conditions, the latter practical conditions. If a particular approach to MT is to count as applied science, it must simultaneously incorporate knowledge in such a way as to be scientific and constitute a working solution to a real-life problem.

The input and output of pure science can be represented as in Fig. 1.

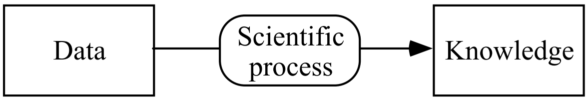

Fig. 1: Pure science.

The goal of pure science is to produce knowledge which is ultimately based on certain data. The interpretation of the elements in Fig. 1 is strikingly different for two types of science: empirical and formal sciences. In a formal science, the universe of the data is determined by a system of axioms controlled by the scientist and the knowledge consists of theorems for which the scientific process has been able to produce proofs. Examples of formal sciences are mathematics and logic. In empirical sciences, the system underlying the universe of data is not known in advance but given as a kind of black box. The knowledge consists of theories, i.e. hypotheses about this underlying system. The scientific process is governed by the empirical cycle in Fig. 2.

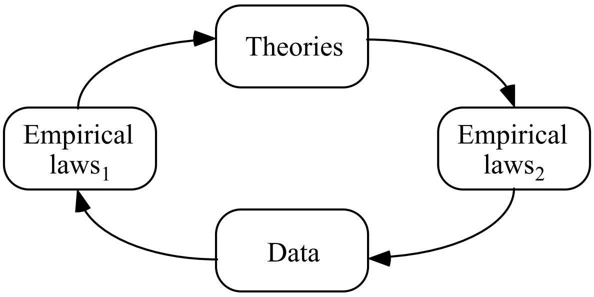

Fig. 2: The empirical cycle.

The data in Fig. 2 are results of individual experiments. Empirical laws are generalizations about relevant aspects of the data. They occur both as a result of the study of data (laws$_1$ in Fig. 2) and as predictions of the outcome of experiments (laws$_2$). A theory is a set of hypotheses providing an explanation of the laws$_1$ and, hence, of the data covered by these laws. It can be tested by confronting its predictions (laws$_2$) with the data. Prototypical examples of empirical sciences are astronomy and physics. A discussion of theoretical linguistics in terms of the empirical cycle can be found in ten Hacken (1997, 1998).

We started with the observation that applied science was at the intersection of pure science and problem-solving. Having said something on the former, let us now turn to the latter. Fig. 3 represents the input and output of problem-solving.

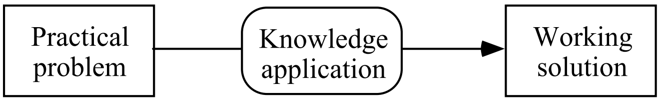

Fig. 3: Problem-solving.

The goal in problem-solving is to find a working solution to a real-life problem. Depending on the type of knowledge and the way it is applied, different types of problem-solving can be distinguished. A first distinction can be found in the history of agriculture. The problem to be solved here can be described as how to maximize the yield of a given plot of soil. Reijnders (1997:39ff.) describes the “first green revolution”, which took place in what is now Belgium and the Netherlands in the later Middle Ages. The major discovery underlying this revolution was that town refuse could be used as fertilizer, increasing the nutritious power of the soil. Though highly effective, this discovery had nothing to do with science. At the time no relevant scientific knowledge was available and the increased productivity was often referred to as a “miracle”. In the 19th century, the “second green revolution” introduced the use of artificial fertilizer. It could only be produced as a result of advances in biology and chemistry, such as Justus Liebig’s theories on the growth of plants (cf. Reijnders (1997:120f.)). It is clear that science has a role here, which can be characterized in terms of Fig. 3 as providing the knowledge for the practical solution.

There is a further distinction to be made here, however, between technology and applied science. Intuitively, this distinction is clear when we consider the history of chemistry. As described by Bensaude-Vincent & Stengers (1993), industrial production and scientific research are in close interaction in chemistry. Often production is far ahead of our understanding of why the product works. Science can ultimately come up with an explanation. On the other hand, science supplies ideas which may inspire the development of new products. Thus we can divide chemistry into three domains. First, there is chemistry as an empirical science, e.g. explaining what happens when a candle burns. It is this part which ends up being absorbed by physics with the advance of atomic theory. Second, there is chemistry as pure technology, e.g. developing a new dye stuff with a different colour. Third, there is chemistry as an applied science. It uses the explanatory nature of the corresponding empirical science to explain how, why, and to what extent problems can be solved.

The difference between technology and applied science can be expressed succinctly in terms of the elements of the empirical cycle. Whereas the knowledge applied in technology consists of empirical laws, the knowledge applied in applied science consists of theories of empirical science.

Perhaps the most prototypical example of an applied science is medicine. It can serve here to clarify the relations between different types of problem-solving approaches and empirical science. A first remark concerns the unit to be classified. As Bynum (1994) describes, the 19th century saw a ‘scientific revolution’ in medicine. Nevertheless, we still have applied science, technology, and commonsense problem-solving functioning next to each other in medicine. Some cures work well and are understood quite thoroughly, e.g. vaccination against smallpox. For other diseases, highly technological cures are used without properly understanding their working, e.g. current AIDS therapies. Finally, there are diseases on which science has little to say, e.g. allergies. Here the only effective treatment consists in common-sense precepts such as avoiding contact with certain substances. Only the first type of cases can be termed applied science. Therefore the unit of classification should be the approach to a particular problem.

A second remark concerns the use of such distinctions. Basically, the distinction between applied science and technology is of an epistemological nature. It determines our understanding, not necessarily the practical validity of a solution. Therefore it is correct to scrutinize a new cure by means of a large-scale practical test rather than by studying the underlying theory in detail. A chain of explanations may still increase our confidence, however.

A final remark concerns the distinction between empirical and applied science in practice. Much of medical research is actually highly specialized biology. It is empirical research aimed at producing the knowledge necessary to solve practical problems. Medicine as an applied science has its own body of knowledge. It consists of procedures to diagnose which disease a patient has, knowledge about which cure to apply etc. Therefore the interaction of empirical and applied science does not make it impossible to characterize the body of knowledge of an applied science.

2. Machine Translation

An MT system is a computer program which takes as input a text in one language and gives as output a corresponding text in another language. Here *text* may refer to written or spoken language. In Fig. 4 input and output of MT are represented.

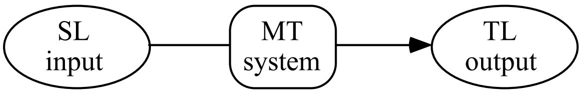

Fig. 4: Machine translation.

If we compare the representation of MT in Fig. 4 with that of problem-solving in Fig. 3, the SL text corresponds to the problem, the TL text to its solution, and the MT system to the process of knowledge application. In an alternative perspective, the MT system is the solution, achieved by the application of knowledge. In neither perspective is the problem stated in its most general form, i.e. the non-understanding of an SL-text. Certain properties of the solution are incorporated in the problem description. Although comparison of the performance of MT systems with human translation is a common way of evaluating MT systems (e.g. Van Slype (1982)), we would not accept human translation as a solution to the problem of MT. In an even broader perspective on the problem, ALPAC (1966:5) also consider second language learning as an alternative solution. This type of restriction is certainly not unique. A similar example is to be found in biological agriculture, which consciously excludes a number of solutions to the problem of increasing yield which are known to work.

The problem of MT is obviated rather than solved if there is no translation process involved whose basic organization depends on the use of computers. In the same way as biological agriculture, the success of MT can be measured in terms of its performance in comparison to alternatives consciously excluded.

In addressing the question of the conditions under which MT is applied science, two questions need to be answered first. Given that applied science distinguishes itself from technology by the availability of explanations, we have to determine what is explained in the case of MT and how. The second question is how the units to be classified as applied science can be found.

In distinguishing applied science and technology we found that the presence or absence of the explanatory component of the knowledge borrowed from empirical science plays a crucial role. In the case of MT the question arises of what the scope of explanation must be. In most approaches to linguistics as an empirical science the system of knowledge acquired by the native speaker explains the data. Data include (at least) grammaticality judgements.[1] If we want to take explanation in such a theory as a basis for an explanation in MT, it is the performance of the system which can be explained. Performance is measured in terms of evaluation criteria yielding precise observations which can be used as data. The explanation takes into account the original explanation of the theories borrowed from empirical science, the way these theories are employed in the MT system, and the relationship of the input of the MT system to the data explained by these theories.

In line with our observations on medicine in section 1, the unit of MT to be classified as applied science or technology is the approach to the problem of MT. An approach to MT differs from individual MT systems or projects by abstracting from a number of decisions. Typical examples are the choice of a programming language and of a computer platform. These choices influence performance in terms of the time required to produce an answer but do not affect the quality of the answer. The decisions listed in (1) – (4) characterize the approach.

(1) What exactly is the nature of the problem implicit in the SL text and what counts as a solution ?

(2) How can the problem be broken down into components ?

(3) How can it be determined whether or to what degree the problem has been solved adequately ?

(4) What is the scientific knowledge used in the MT system ?

Problem definition in (1) addresses the boundaries of the problem to be solved by the system and the relationship between SL and TL text. Its interaction with the applied science vs. technology question will be discussed in section 3 below.

Problem decomposition in (2) concerns such questions as the level of transfer. This question, with its traditional answers of zero (direct systems), infinite (interlingua systems), or anything in between (transfer systems), is often presented in introductory texts, e.g. Tucker (1987:22-25), Nirenburg (1989:5-7), as the major choice in the design of an MT system, or at least one of the most important criteria for the classification of MT systems, e.g. Slocum (1988:8), Hutchins & Somers (1992:4). In addition to the level of transfer, (2) includes such issues as multilinguality and reversibility. In general, (2) interacts with the applied science vs. technology question above all in the sense that the more general the solution aimed at, the more likely it is that the approach belongs to applied science. It is difficult to imagine what explanatory possibilities are associated with a strictly bilingual, non-reversible MT

---

[1] This description is vague enough to encompass Chomskyan linguistics, LFG, and HPSG. For a discussion of the similarities and differences between these approaches, cf. ten Hacken (1997, 1998). As outlined in ten Hacken (1998b), GPSG does not qualify as an empirical science because of its failure to explain. This is due to a conception of language as a set of sentences given in advance, as proposed by Quine (1972).

system, because explanatory theories in linguistics claim to be valid for any language and to cover knowledge of language independently of the direction it is used for (analysis or generation).

The evaluation criteria aimed at in (3) are of eminent importance to applied science because they provide the data to be explained. They are equivalent to the astronomer's telescope. The type of observations we get as data directly determines the type of explanatory theory required. Less direct, but not less important, is the effect on the expected success rate of explanation. In theoretical linguistics, the shift from the analysis of utterances by the Post-Bloomfieldians (e.g. Harris (1951) to grammaticality judgements in Chomskyan linguistics constituted a crucial component in the advance of explanatory success. In the case of MT, a pure cost-benefit analysis, such as discussed by Van Slype (1982b), yields data of a type comparable to Post-Bloomfieldian utterances in degree of explainability. Data resulting from a cost-benefit analysis do not provide a promising basis because they require an explanation formulated at least in part in economic terms, whereas the knowledge used in MT derives from linguistics rather than from economy.

The sources of knowledge in (4) directly influence the technological or scientific nature of an approach to MT in the sense that the use of non-explanatory knowledge can never lead to applied science. An example is the use of statistical techniques on an aligned corpus, as described by Brown et al. (1990). Although it could be claimed that the statistical formulae are scientific because they derive from mathematics, they do not belong to empirical science. Mathematics being a formal science, they are not meant to explain anything. An interesting parallel with medicine can be drawn here. As Bynum (1994) describes, at various points in the history of medicine researchers started collecting statistical data on diseases for which no other knowledge existed. In some cases interesting correlations were found, e.g. between the use of certain public water pumps in London and the spread of cholera (1994:76ff.). Whatever the merits of statistics, however, an explanatory account of cholera could only be found by a theory on the organism causing it and on the life cycle of this organism. This theory was a precondition for a cure which could be part of applied science. Statistics has thus at most heuristic value from the point of view of applied science. In an area where simple correlations such as the one mentioned for cholera are not to be expected, this heuristic value of statistics should not be overestimated.

For the two questions posed at the start of this section we found the following answers. MT as applied science is concerned with the explanation of performance of MT systems, as measured by evaluation criteria. An approach to MT is characterized by its answers to questions (1) – (4). For each question, certain answers may make an explanatory theory impossible to achieve.

3. Problem Definition

In the discussion of the questions characterizing an approach to MT, a detailed discussion of (1), the definition of the problem to be solved, was postponed. This question can be analysed into a number of more precise questions, listed in (5) – (7).

(5) What is the input the system is expected to accept ?

(6) What is the status of the output of the system ?

(7) What is the relationship between an element of the input domain and the output associated with this element by the system ?

Work in MT aims at a coverage of the full problem represented in Fig. 4, producing the best TL translation of any SL text. Everyone in the field agrees that this goal, if it can be reached in principle at all, is far beyond the scope of current realistic planning. Controversy exists, however, on what constitutes the most promising intermediate stage. Different opinions can be expressed as different answers to questions (5) – (7). In formulating these answers it will be convenient to refer to the diagram in Fig. 5.

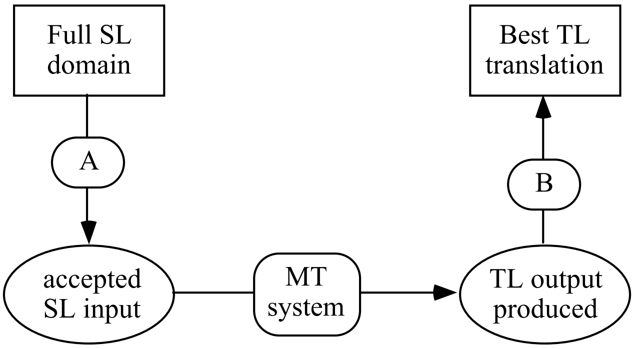

Fig. 5: MT and the real-life problem.

The two mappings labelled A and B in Fig. 5 represent the discrepancy between input and output as actually handled by the MT system and input and output as idealizations of the best possible solution. The input mapping A can be interpreted as either a selection from the full input domain or a problem reduction by pre-processing of the input. The output mapping B marks the difference between MT output and correct output for a given input. It should be noted therefore that Fig. 5 exhibits an asymmetry, in the sense that to the left of the MT system a set of possible inputs is represented, whereas to the right the output of a single element of this set is taken.

The discrepancy at the origin of mappings A and B can be related to the familiar distinction between competence and performance as used in linguistic theory. As defined by Chomsky (1965:4), competence is a system of knowledge and performance is the projection of this knowledge into a set of sentences, filtered through various cognitive and other constraints which may distort them. For a justification of the use of the terms competence and performance and their relevance in the context of computational linguistics, cf. ten Hacken (in press). The relevance of the opposition between them rests on the observation that whereas input text belongs to performance, explanatory linguistic theories tend to take competence as their object of study.

In order to overcome the discrepancy between general input and coverage of linguistic theories, three types of strategy exist. The first minimizes mapping A, thus determining the task of further development as progressively reducing mapping B. The second type of strategy starts from a well-defined but far-reaching initial restriction of the input, defining mapping A so as to minimize mapping B. Progress consists then in a gradual reduction of mapping A. Finally, one could devise a division of tasks such that mappings A and B are covered by different specialized modules, separate from the MT system proper. Each of these approaches will be discussed in turn.

3.1. Correction of the Output

In a project aimed at a general-purpose MT system as an immediate goal, the most common approach to question (5) is to identify the accepted input with the texts occurring in the SL, with as few restrictions as possible. A typical example is Systran. An answer to questions (6) and (7) is more problematic for such an approach. The output is the best approximation to the translation which could be achieved by the system. It can be used as an approximative translation or as input for post-editing. In both cases, unanalysed human intelligence is crucially involved.

There are two reasons why it is not possible to count approaches based on this answer to question (5) as applied science. The first concerns the use of linguistic knowledge. In theoretical linguistics, theories are accounts of linguistic competence. Their explanatory power, directed towards competence, cannot be used in a context where mapping A in Fig. 5 is eliminated. What we observe in practice is that the coverage of certain constructions by different theories is eclectically combined, with some constraints relaxed so as to avoid some of the problems arising from the difference between a direct projection of competence and actual performance. The coverage of individual constructions is exactly the linguistic embodiment of empirical laws as discussed in the context of the empirical cycle (Fig. 3). The use of empirical laws while ignoring their role in a theory explaining them is exactly our definition of technology as opposed to applied science.

The other reason concerns the status of the output. Reliance on human intelligence in post-editing implies that this status cannot be determined in relation to the desired output. On the other hand, the approximative relationship to the input, resulting from the goal of producing the best possible approximation to the translation, inhibits any characterization independent of (unanalysed) system performance.

Therefore any approach to MT taking the full domain of SL texts as input and aiming at the best possible approximation to their translation is purely technological.

3.2. Restriction of the Input

The second approach restricts the generality of the accepted SL input by defining mapping A in Fig. 5 so as to reduce mapping B. Mapping A can be interpreted as pre-processing or restriction of the domain. Under the name of pre-editing, the former is often presented as a natural counterpart to post-editing, e.g. Hutchins & Somers (1992:151). Its epistemological interpretation is radically different, however. In the way it is used in SUSY, for instance, as described by Maas (1978:53f.), pre-editing presupposes predictive knowledge of what constitutes a problem for the MT-system. Whereas post-editing relies on unanalysed human intelligence, successful pre-editing depends on knowledge of the performance of the MT-system. Of course, such an analysis only applies to restricted use of human pre-processing. If the pre-editor translates the input text, we gain nothing. Unfortunately, with such a restriction, pre-editing is not an alternative to post-editing, but only a supplement. In SUSY, it is actually optional and its scope is restricted to the assignment of individual words in the text to a category. In such cases, the same problems arise as for the approach dealt with in section 3.1.

Restriction of the domain of SL texts to a sublanguage is another way to reduce the problem of MT. It is proposed by Kittredge (1987) and by Lehrberger & Bourbeau (1988:51ff.) as the only viable strategy to achieve fully automatic MT, i.e. without post-editing or interactive human involvement. The idea behind sublanguage is that texts taken from a particular subject domain and text type only contain a fragment of language, which can be described with a more limited vocabulary and grammar than language in general. A partial solution of the general MT problem is then a full solution for the MT problem of a sublanguage.

The sublanguage approach certainly reduces the problem of lexical disambiguation. In view of the representation in Fig. 5 and the distinction between competence and performance, however, it does not seem to lead to an approach to MT which can count as applied science. The claim implicit in adducing this type of reduction of the input as a strategy towards applied science is that in a sublanguage we can still use general linguistic theories but the gap between actual text and competence is smaller. There does not seem to be any sign that this assumption is correct.

An example often quoted as a successful instance of a sublanguage approach is MÉTÉO (e.g. Kittredge (1987:60f.), Lehrberger & Bourbeau (1988:51)). Thouin (1982:40-42) and Isabelle (1987:261-265), however, describe the input of MÉTÉO as only that part of meteorological reports that conforms to a rigid style sheet. This excludes, for instance, the general reports on weather conditions. Therefore MÉTÉO is not a successful treatment of the sublanguage of weather reports, but an application of controlled language. As opposed to a sublanguage, a controlled language is an exhaustively defined subset, not an empirically motivated subset of a language. In a controlled language approach, translation is not an empirical but a formal question. It can be mathematically proven that every well-formed sentence of the controlled language is translated correctly. The empirical claim involved is that communication is possible given the controlled language. This is a very different problem to MT. I therefore do not take it to be an instantiation of MT as an applied science.

3.3. Explanatory Modules

The third type of approach takes the explanatory possibilities of linguistic theories as a starting point for the answer to questions (5) – (7). Such theories are concerned with explaining data in terms of competence and with explaining some chosen aspects of the competence in terms of a more general entity, common to all languages. In the context of MT, competence plays a role in SL and TL, but it does not cover the mapping from performance to competence and the translation. The former can be identified with mapping A in Fig. 5. The translation mapping is of course the central problem of MT.

A variant of this approach is the one advocated by Rosetta (1994). The translation theory adopted by Rosetta is based on linguistic meaning, defined as opposed to world knowledge meaning, stylistic meaning, etc., by its compositional basis. Input is restricted to individual sentences belonging to the language system, which answers (5). Output consists of the set of sentences with basically the same meaning in a different language. This answers (6). In view of possible ambiguity and vagueness in the SL sentence not mirrored in the TL, the TL set contains all sentences with at least one meaning in common with the SL sentence, the so-called *possible translations*. This answers (7). The Rosetta translation system leaves it to separate modules to cover mappings A (performance to competence) and B (choice of the best translation in context). Inherent in this approach is the assumption that the best translation is one of the possible translations. As shown in Stolze's (1994) overview, this assumption is contradicted in many theories of (human) translation. Mapping B can therefore be characterized more exactly as a ranking of possible translations.

At first sight, the modular approach adopted by Rosetta comes closer to applied science than to mere technology. Two properties of the Rosetta approach are problematic, however, if we intend to classify it as applied science.

First, Rosetta syntax is admittedly eclectic (1994:203), combining insights from various frameworks. As noted above, when empirical laws rather than the theories explaining them are used, this is a sign of technology rather than applied science. It should be pointed out, however, that borrowing of insights across frameworks is common in linguistics and need not by itself negatively affect explanatory power. The crucial issue is whether the choice of one mechanism or another is guided by a theory which aims at explanation or solely dependent on pragmatic criteria. Therefore this property, while for the moment making applied science unattainable, is not a necessary, inherent barrier to applied science for the approach as such.

The other property is more deeply entrenched in the core of the Rosetta approach. Throughout the presentation, Rosetta (1994) emphasizes its reliance on formal rather than empirical sciences. Concern with formal correctness is not restricted to the section on “Formal aspects” (Ch. 17-19), but is also pervasive in the presentation of the translation method (Ch. 2-9). Thus, as soon as the parser has been introduced, a measure condition for proving termination in finite time is discussed (1994:69). This way of representing the approach gives the impression that Rosetta aims more at proving formal correctness than explaining the degree to which it solves the MT problem.

This impression is reinforced by the extensive use of Montague grammar. Rosetta uses M-grammars described as “a computationally viable variant of *Montague grammar*” (1994:35). Montague assumed that a language is a set of sentences. As Thomason (1974:2) put it, “According to Montague the syntax, semantics, and pragmatics of natural languages are branches of mathematics”. Elaborating on this view, Quine (1972) draws the conclusion that, since there is an indefinite number of grammars generating a given set of sentences, the concept of a language as equivalent to a set of sentences makes it impossible to choose a single grammar. The absence of anything corresponding to competence to be described by the grammar excludes the possibility of explanation. Montague grammar is a formal rather than an empirical science. The absence of explanation makes it impossible to build an applied science on it.

The Rosetta approach does not qualify as applied science because it does not aim at explanation, as shown by its using non-explanatory knowledge and its concentrating rather on proving formal correctness wherever possible.

This conclusion does not affect the viability of the attempt to arrive at applied science in MT by a modular approach. In particular, the decision to have the MT system start from sentences accepted by grammatical competence rather than directly from text does not have a negative effect. In a modular approach to MT, it is perfectly conceivable to have a mixture of modules belonging to technology and to applied science. For the scientific modules it is essential that their input and output are characterized independently of the actual performance of the system and its modules. Furthermore, in order to belong to applied science, a module should be based on an explanatory theory, i.e. a coherent (non-eclectic) theory of empirical (not formal) science. For tasks for which no explanatory theory is available for use in a module, technological solutions may be adopted, making use of interaction with unanalysed human intelligence or statistical rules applied to a corpus. An interactive or statistical module for, for instance, lexical disambiguation does not affect the possibility of scientific modules preceding or following it, provided questions (5) – (7) are answered for each module independently of system performance. A similar solution may be considered for the performance-to-competence mapping.

4. Conclusion

The goal of treating MT as an applied science rather than mere technology imposes heavy constraints on the approach to the problem of MT as characterized in questions (1) – (4). The constraints following from the explanatory nature of applied science on the answers to (2) – (4) are more straightforward than those on the answer to (1). The only satisfactory approach to (1) turned out to be a modular analysis of the problem, such that at least the modules for which an explanatory theory can be used may belong to applied science.

An approach based on controlled language alters the original problem, solving (in the ideal case) a communication problem rather than MT. Although a formal proof of the translation of a controlled language is possible along the lines also followed in Rosetta, this does not make it applied science. An explanatory theory of communication is required to make the design of a controlled language into an applied science.

The fact that technological efficiency in MT is far ahead of explanation is not surprising. The history of chemistry and agriculture shows that in these cases too, technological solutions preceded the ability to explain them.

References

ALPAC [Automatic Language Processing Advisory Committee] (1966), *Language and Machines: Computers in Translation and Linguistics*, Washington (DC): National Academy of Sciences.

Bensaude-Vincent, Bernadette & Stengers, Isabelle (1993), *Histoire de la chimie*, Paris: La Découverte.

Brown, Peter; Cocke, John; Della Pietra, Stephen; Della Pietra, Vincent J.; Jelinek, Fredrick; Lafferty, John D.; Mercer, Robert L. & Roossin, Paul S. (1990), 'A Statistical Approach to Machine Translation', *Computational Linguistics* 16:79-85.

Bynum, W.F. (1994), *Science and the Practice of Medicine in the Nineteenth Century*, Cambridge: Cambridge University Press.

Chomsky, Noam (1965), *Aspects of the Theory of Syntax*, Cambridge (Mass.): MIT Press.

ten Hacken, Pius (1997), 'Progress and Incommensurability in Linguistics', *Beiträge zur Geschichte der Sprachwissenschaft* 7:287-310.

ten Hacken, Pius (1998), 'Chomskyan Linguistics and HPSG as Competing Research Programmes', to appear in Caron, Bernard (ed.), *Actes du 16^{è} Congrès International des Linguistes (Paris 20-25 juillet 1997)*, Oxford: Elsevier Sciences.

ten Hacken, Pius (1998b), 'Research Programmes in Linguistics', to appear in Jürg Strässler (ed.), *Tendenzen europäischer Linguistik, Akten des 31. Linguistischen Kolloquiums, Bern 1996*, Tübingen: Niemeyer, p. 67-70.

ten Hacken, Pius (in press), 'Performance vs. Competence in Computational Linguistics', to appear in Halter, Peter (ed.), *Performance*, Tübingen: Narr.

Harris, Zellig S. (1951), *Methods in Structural Linguistics*, Chicago: University of Chicago Press, (repr. as *Structural Linguistics*, 1960).

Hutchins, W. John & Somers, Harry L. (1992), *An Introduction to Machine Translation*, London: Academic Press.

Isabelle, Pierre (1987), ‘Machine Translation at the TAUM Group’, in King, Margaret (ed.), *Machine Translation Today: The State of the Art*, Edinburgh: Edinburgh University Press, p. 247-277.

Kittredge, Richard I. (1987), ‘The significance of sublanguage for automatic translation’, in Nirenburg, Sergei (ed.), *Machine Translation: Theoretical and Methodological Issues*, Cambridge: Cambridge University Press, p. 59-67.

Lehrberger, John & Bourbeau, Laurent (1988), *Machine Translation: Linguistic characteristics of MT systems and general methodology of evaluation*, Amsterdam: Benjamins.

Maas, Heinz-Dieter (1978), ‘Das Saarbrücker Übersetzungssystem SUSY’, *Sprache und Datenverarbeitung* 2:43-61.

Nirenburg, Sergei (1989), ‘Knowledge-Based Machine Translation’, *Machine Translation* 4:5-24.

Quine, W.V. (1972), ‘Methodological Reflections on Current Linguistic Theory’, in Davidson, Donald & Harman, Gilbert (eds.), *Semantics of Natural Language*, Dordrecht: Reidel, p. 442-454.

Reijnders, Lucas (1997), *Het boerenbedrijf in de Lage Landen: Geschiedenis en toekomst*, Amsterdam: Van Gennep.

Rosetta, M.T. (1994), *Compositional Translation*, Dordrecht: Kluwer Academic.

Slocum, Jonathan (1988), ‘A Survey of Machine Translation: its History, Current Status, and Future Prospects’, in Slocum, Jonathan (ed.), *Machine Translation Systems*, Cambridge: Cambridge University Press, p. 1-47.

Stolze, Radegundis (1994), *Übersetzungstheorien: Eine Einführung*, Tübingen: Narr.

Thomason, Richmond H. (ed.) (1974), *Formal Philosophy: Selected Papers of Richard Montague*, New Haven: Yale University Press.

Thouin, Benoît (1982), ‘The MÉTÉO System’, in Lawson, Veronica (ed.), *Practical Experience of Machine Translation*, Amsterdam: North-Holland, p. 39-44.

Tucker, Allen B. (1987), ‘Current Strategies in Machine Translation Research and Development’, in Nirenburg, Sergei (ed.), *Machine Translation: Theoretical and Methodological Issues*, Cambridge: Cambridge University Press, p. 22-41.

Van Slype, Georges (1982), ‘Conception d’une méthodique générale d’évaluation de la traduction automatique’, *Multilingua* 1:221-237.

Van Slype, Georges (1982b), ‘Economic Aspects of Machine Translation’, in Lawson, Veronica (ed.), *Practical Experience of Machine Translation*, Amsterdam: North-Holland, p. 79-93.