Revolution in Computational Linguistics

Towards a Genuinely Applied Science

Pius ten Hacken

Abteilung für Geisteswissenschaftliche Informatik, Universität Basel

Abstract

Among people working in Computational Linguistics (CL) around 1990 and still active in the field now, there is a widely shared feeling that they have witnessed a revolution. This paper shows which developments are responsible for this perception and which elements are central in the actual revolution. In order to avoid terminological confusion, the concept of revolution as it is used here is clarified first. Then the development in the subfield of Machine Translation is studied in some detail. It is argued that the actual revolution consists in a shift of attention from the application of theoretical knowledge to the solution of practical problems. To the extent that this shift is representative of more general developments in the field, the conclusions can be generalized to CL as a whole.

Among people working in Computational Linguistics (CL) around 1990 and still active in the field now, there is a widely shared feeling that they have witnessed a revolution. The purpose of this paper is to show which developments are responsible for this perceived revolution and how the consequences of these developments for the field can be evaluated. Section 1 creates a basis by establishing which sense of revolution is meant in the title and how such a phenomenon can be recognized in general. The position of Machine Translation (MT) in CL and the general orientation of CL before the revolution are sketched in section 2. Some aspects of four MT systems representative of the development in the 1980s and 1990s are outlined in section 3 and analysed in section 4. From this analysis, a more general characterization of the revolution in CL is derived in section 5.

1 Revolutions

The concept of *revolution* is one loaded with a large connotative value, which tends to hamper its successful application in a specific context such as the question of whether and how exactly CL has undergone a revolution in the past decade.

A revolution can be described generally as a radical change in the orientation of a system. In a sufficiently broad interpretation of *system*, this definition even covers the sense of 'rotating movement', as in a motor. More to the point are the cases in which 'system' refers to a sociologically organized complex. In this context an interesting parallel can be observed between political revolutions and scientific revolutions. The interest resides in the possibility of identifying some key concepts of revolutions which can be used successfully in both contexts.

While political revolutions are always radical changes in the orientation of the political system of a country, not every radical change amounts to a revolution. When Mitterand won the French presidential elections in 1981 and his socialist party won a landslide victory in the subsequent parliamentary elections, this certainly caused a radical reorientation of the French political landscape. Nevertheless, it was not a revolution. By contrast, when in 1830 King Charles X was overthrown and replaced by King Louis Philippe, the change in political orientation was probably much less radical, but it was certainly a revolution (cf. Louessard, 1990). The main difference between these events is that the system in place in 1981 catered for a changeover of power of the type which took place, whereas in 1830 the changeover was not foreseen by the rules of the system.

In the context of science, the term was originally used in a definite sense, '*the* scientific revolution', referring to the emergence of science in the modern sense after the Middle Ages (cf. Rossi 1997). This is certainly not the sense in which CL was revolutionized in the 1990s. More relevant is the concept of revolution as discussed by Kuhn (1970). This sense of (scientific) revolution has more in common with the notion of (political) revolution sketched above. Kuhn showed that a science at a given point in its existence has a specific system of rules, which regulates what counts as normal, accepted behaviour, more or less likeits political counterpart. This system determines how science is done, what counts as progress in the development of science, and what is good science. A scientific revolution is the change of this reference system for a science in a sense not foreseen by the rules of the system itself.

In politics and in science, the underlying system is not a purpose in itself, but rather a background for doing something else. It requires strong motivation to spend a lot of time and energy reworking this background rather than making progress along better known lines. Therefore a revolution is always caused by a crisis. It is not necessary that this crisis is felt by everyone involved in the system, nor can a definite threshold be set when a revolution will start (much less when it will be successful), but revolutions do not occur at random.

Only when a group of scientists perceive their way of science as not promising sufficient progress any more will they start investing time in a revision of the underlying system of assumptions. As a consequence, a revolution is often accompanied by a general spirit of optimism among this group, because they believe that their radical measures will bring a faster road to progress. By itself, observing such a spirit need not be a sign of a revolution, however. Here a parallel with politics can be drawn. A similar type of optimism could be observed in France after the 1981 elections (at least in one part of society), for which we established previously that it did not constitute a revolution.

As a scientific revolution involves a change in reference system which determines how scientific practice should be evaluated, it is not possible to make a fair comparison of theories developed before and after the revolution by applying the standards of one reference system to the theories formulated on the basis of the other. Kuhn (1970) calls this the incommensurability of theories inherent in the comparison of different paradigms or disciplinary matrices. It should be emphasized that this means not only that the 'old' theory is judged unfairly with the new evaluation criteria, but also that the 'new' theory will probably be bad in terms of the evaluation criteria of the old reference system.

Mutual misunderstanding is often a sign of incommensurability. Thus, the following statements from the discussion between Householder (1965) and Chomsky & Halle (1965) provide a strong argument for the analysis of the emergence of Chomskyan linguistics as a scientific revolution. After a review of different aspects of Chomsky & Halle's phonological theory, Householder's conclusion is (1). In a detailed reaction to Householder's argument we find such remarks as (2).

> (1) "on matters of phonology, their claims and assertions, if all wholly true, would tend to make all phonological work impossible *on any known lines*." [Householder (1965:13), my emphasis].

> (2) "We have no idea what this comment means, and therefore make no attempt to discuss it." [Chomsky & Halle (1965:118)]

A consequence of incommensurability is that in a scientific revolution the newly emerging framework cannot be expected to encompass everything the old framework was able to deal with. A revolution always involves both progress and losses (cf. Chen 1997). If the revolution is successful the scientific community apparently considers the progress promised by the new framework greater than the losses it induces. Such a judgement cannot be fully rational, however, because it involves the choice between two standards of evaluation. There does not exist a common set of criteria to choose between two scientific paradigms.

Returning to CL, in order to show that a revolution took place, we should look for a radical change in the system of assumptions underlying scientific work. Properties of such a change include the incommensurability of results of work before and after the revolution.

2 Computational Linguistics before the Revolution

CL is special as a science because it requires more than just a theory. Whatever is produced in CL should work. In this sense, the opposition between CL and theoretical linguistics can be compared to the one between electrical engineering and theoretical physics. Despite the opposition, linguistics is usually understood as including CL, and physics as including electrical engineering. In addition, fields such as CL and electrical engineering share some important goals and properties with science at large. Science is about the extension of knowledge. In the list of priorities for this extension, there is a tension between two types of argument. Scientists are primarily looking for interesting knowledge, e.g. knowledge which extends the scope or precision of a theory. Their financial supporters are first of all attracted by the possibility of useful knowledge, e.g. knowledge which is beneficial for society or commercially attractive. Therefore, scientists usually try to convince potential financial supporters of the usefulness of what they find interesting.

This fundamental dilemma governs the relationship between empirical science, which is concerned with explaining data, and applied science, which is concerned with the solution of practical problems. The track record of past successes in producing useful knowledge determines to a large extent how directly the immediate applicability of the envisaged results influences the research agenda of a particular science. Physics has in the course of the centuries established a solid reputation that its results are useful in the long run, even if no immediate application is in sight. Computational Linguistics (CL) is in a more difficult position, not only because of its shorter history, but also because it has repeatedly disappointed the expectations of its financial supporters.

The early history of Machine Translation (MT) offers a number of clear examples. The upsurge of MT in the USA in the 1950s enabled research teams to get a rather large amount of money, but when the ALPAC (1966) report concluded that the promised results were not forthcoming, financial resources were withdrawn. In the 1970s and 1980s, the availability of linguistic theories of a new kind constituted a new promise. However, the TAUM group in Montral only managed to secure funding for its celebrated Mto project by making a compromise in which much of the linguistic knowledge was sacrificed to a useful application. When the subsequent Aviation project, in which more of the linguistic knowledge could be applied, did not produce results fast enough, funding stopped (cf. Isabelle 1987).

For the 1970s and 1980s it is not too much to claim that the overall goal of CL was perceived by people working in the field as the modelling of human language on a computer. This is the picture arising from conferences such as Coling (1984) and textbooks such as Grishman (1986) and Pereira & Shieber (1987). Thus, at Coling, Shieber (1984:362) remarked that ”The goal of natural-language processing research can be stated quite simply: to endow computers with human language capacity.”, and, in the introduction to his widely used textbook, Grishman (1986:4) states that ”Computational linguistics is the study of computer systems for understanding and generating natural language.” Characteristic of the general approach of this period is Kay's (1973) Mind system. This system consists of a number of modules for different aspects of language, which can be connected in different ways so as to perform different tasks. Pereira & Shieber (1987) take as their point of departure the task of ”natural language analysis”, which can indeed be considered the central task of CL in this approach.

Scientific work can be analysed in general as solving problems. The scope of the work is then determined by the type of problem chosen, the type of knowledge used for the solution, and the type of solution aimed at. In CL as it functioned in the 1970s and 1980s, these parameters were set as in (3):

(3)	a.	Problem:	Understanding human language processing
	b.	Knowledge:	Contemporary linguistic theories
	c.	Solution:	A running program on a computer

MT had a relatively important position in CL in this period. There are at least two reasons for this. First, whereas it is difficult to check how much a computer has understood after processing a sentence, it is much easier to check whether a particular task can be carried out. Translation is a task which uses a large part of the input and results in a representation which is open to easy human inspection.

Second, translation is a task which can easily be explained to potential financial supporters and which has economic value.

This position of MT within CL explains why translation theory as it had developed in the 1960s and 1970s was largely ignored in MT. The main focus of interest was not translation as such. Rather translation was used as an excuse for syntactic and semantic research. Thus, Lehrberger and Bourbeau (1988:1) claim that ”The obstacles to translating by means of the computer are mainly linguistic.” Translation was modelled (in most cases implicitly) as a fairly simple meaning equivalence.

By the end of the 1980s, this type of MT research had entered a crisis. A sign of this crisis is the fact that financial support for a number of projects ran out. Around 1990 all three of the major MT projects represented in the Netherlands - Eurotra, Rosetta, and DLT - were phased out.

3 Four MT Systems

In order to illustrate how the general trend in MT is reflected in individual systems, let us consider some relevant aspects of four systems in more detail. Eurotra, Rosetta, the statistical system developed at IBM, and Verbmobil represent different attitudes to the crisis in MT as it arose at the end of the 1980s.

Eurotra had its roots in the late 1970s. By the time the signs of a crisis were apparent, the project had entered its final phase. Major design decisions had been taken and also much of the linguistic theory had been established and ’frozen’ so as not to hamper development by undermining this basis. At the time of its inception, Eurotra had been innovative in being a large-scale, truly multilingual transfer system. The scale of the organization contributed to a certain inertia in view of the emerging crisis. In fact, the organization was designed to optimize the coordination of results of interesting linguistic research rather than foundational issues of translation. Even so, the project leadership felt obliged to mention the problem of the choice of a translation theory in an overview of the project, as in (4):

> (4) ”Any machine translation system, be it a research prototype or an operational system, builds upon some notion of theory of translation, implicitly or explicitly. [...] Unfortunately, we have found no theory of human translation which could be taken over and formalized.” [Maegaard & Perschke (1991:75)]

While at the start of the project the question of what constitutes a translation had been taken for granted, (4) shows that towards the end of the project the lack of an answer had come to be felt as a drawback.

As Kuhn (1970) states, within established circles in a science there is a large resistance to fundamental changes. As long as the need for changes is not strongly imposed by a crisis, scientists will go on with business as usual.

Rosetta existed at around the same time as Eurotra, but it was much smaller. In addition, the director of the project, Jan Landsbergen, put a particular emphasis on formal correctness from the beginning. Rosetta (1994) describes and motivates decisions in the design and theoretical development of the system systematically and in detail. As discussed more extensively in ten Hacken (to appear), the major design decisions such as the choice of an interlingua model and the tuning of grammars of different languages to each other were the subject of controversial discussion in the field of MT, but the explicit model of the translation problem was not. Rosetta assumes sentences as the input to the translation process, sets of possible translations as output, and compositional translation as the relation between an element of the input and an element of the output. In Rosetta a formal model of natural languages and their semantic correspondence is constructed. For the evaluation of such a model, internal and external criteria can be used. Internal evaluation criteria justified by the approach in Rosetta concern the formal consistency of the models of the individual languages and the formal correctness of the mapping between them. External evaluation criteria concentrate on the correspondence between the model and the actual language or the actual input to the system.

The translation theory of Rosetta is characteristic of the approach adopted more or less generally in the 1970s and 1980s. The most remarkable feature was its explicitness. Although other MT systems did not necessarily accept every detail of it, it was not attacked in its outline and no alternative statement was proposed. Proposals for comparative system evaluation such as Melby (1988) are based on the same presuppositions about the correct (external) evaluation criteria.

In terms of Kuhn (1970) the lack of discussion is less remarkable than the fact that such an explicit, detailed statement of the underlying assumptions of an MT system was made at all. It can well be interpreted as a sign of a crisis. As Kuhn describes it, a crisis brings about an atmosphere where we find on the one hand a group of people exploring new avenues, on the other hand a group which defends the traditional approach, for example by explicitly stating its foundational principles. Rosetta (1994) definitely belongs to the latter group.

An example of the former reaction to the crisis is found in the statistical MT system developed at IBM. Deliberately provocative in their presentation, Brown et al. (1988) at first outlined the possibility of doing MT without linguistic analysis. In later presentations, they suggest a type of collaboration with more traditional approaches (cf. below). Despite its resolutely innovative choice of a source of knowledge, the IBM system is rather traditional in other respects. In particular, translation is still considered as a matter of producing a target language sentence corresponding to the source language sentence in the input. This correspondence is still understood in terms of equivalence of meaning. The central evaluation criterion proposed is the percentage of correctly translated sentences, i.e. sentences which have the desired meaning. Only the calculation of the correspondence is meaning-independent, in accordance with Shannon's (1948) information theory. The external evaluation criteria for Rosetta can be used successfully for the IBM system as well, resulting in a measure of correctness of the output which can be compared by the same standards.

Verbmobil offers us an example of yet another type of reaction to the crisis.

In their study, which sketches the goals of the prototype to be pursued, Kay et al. (1994) devote a rather long section (one eighth of the entire text) to the argument that the fact that language is situated makes translation difficult. To professional translators this point is obvious. The reason why it has to be argued in so much detail is that it contradicts one of the central assumptions of what MT is as commonly understood among people working in this field at the time. In their overview of MT systems, the main criticism of Rosetta focuses on their use of possible translations, stating that ”this whole line of argument carries to an extreme that fallacy of thinking of translation as a function from a source to a target text.” (Kay et al. 1994:85). By contrast, they propose that the intention of a message rather than the meaning of a sentence should be rendered in translation.

Instead of the traditional transfer, interlingua, or even direct architectures for MT, which all assume a unidirectional mapping from source to target language, Kay et al. (1994:93) propose a negotiation model. Such a model, should offer the possibility to compare different analyses and different target language renderings before deciding on the best one in the given context. Moreover, they propose to exploit the situational context of a dialogue system. The fact that language is situated, so that the intention underlying a message cannot be fully retrieved without considering the situational context, complicates the analysis from a traditional point of view, but it can also be used as a strategy to support the solution of a real-life problem. Human communicative behaviour in a regular dialogue among people is marked by questions for clarification and confirmations of conclusions, communicative mechanisms which can also be used by an MT system to solve ambiguities and confirm analyses.

In his introduction to the volume presenting the results of the project, Wahlster (2000) represents the multi-engine architecture as the principal innovation reflecting this proposed new approach to translation. The negotiation model is reflected in packed representations of ambiguities at all stages, which can be unpacked when the negotiation process requires it. Perhaps the change in understanding of what constitutes successful translation is even more fundamental as an innovation. Once accepted, however, it is hardly spectacular enough to present it as such.

4 The Nature of the Revolution in Computational Linguistics

Let us now return to the question of what the revolution in CL in the 1990s amounts to. As a starting point, let us consider how the four MT systems whose general approach is described in section 3 relate to the description of pre-revolutionary CL in section 2, and in particular to the parameters in (3), repeated here for convenience.

(3)	a.	Problem:	Understanding human language processing
	b.	Knowledge:	Contemporary linguistic theories
	c.	Solution:	A running program on a computer

The problem of MT in the sense in which it is taken by Eurotra and Rosetta can be seen as an instantiation of the problem of understanding human language processing. These systems select the knowledge to be used from among linguistic theories, supplemented by contrastive studies for the translation part. Insofar as these theories say anything about processing, they are concerned with human processing. The goal of producing a running program on a computer has to be stated explicitly in order to distinguish research in Eurotra and Rosetta from purely theoretical linguistic research.

The emergence of probabilistic approaches in CL is often felt as a revolution. In MT this development is instantiated by the IBM system of Brown et al. (1988, 1990, 1993). In terms of the parameters in (3) the central difference from Eurotra and Rosetta is the type of knowledge used. Although the statistical modelling of natural language can be seen as a linguistic theory in the broadest sense, it does not reflect the theories which are the subject of debate in linguistics. In particular, it is not intended as a valid model for human language processing. Therefore, in at least one sense, the selection of knowledge in probabilistic approaches implies a change in the problem of CL as well.

In a narrower sense, however, the problem remains the same: the same external evaluation criteria can be applied to Eurotra, Rosetta, and the IBM system. In fact, the percentage of correctly translated sentences is used as a basis for the claims about the performance of statistical MT by Brown et al. (1990:83f.) in the same way as it had long been done for conventional systems based on linguistic theories. The acceptance of these figures by both sides indicates that there is a level of problem statement on which they agree. The main difference is that in the probabilistic approach the problem of matching corresponding sentences from two languages is taken in a literal sense, rather than as an instantiation of the problem of understanding and modelling human language processing.

Another indication showing whether a revolution separates two approaches is the way they behave towards each other. Clearly, ideologically oriented antagonism is not sufficient as a proof of a revolution. We also find it in French politics around 1981, although we established in section 1 that Mitterand's election as president was not a revolution. The same goes for the general spirit of optimism. A further sign might be taken from the type of collaboration between people working in the two frameworks. Here some care is needed to interpret the observations correctly.

Even in generally accepted cases of scientific revolutions the new approach will try to take over results from the old one. The large extent to which this took place in the Copernican Revolution is described in detail by Kuhn (1957). Thus Copernicus continued to use epicycles developed by Ptolemy, and Kepler based his hypothesis of elliptical orbits around the sun on Tycho Brahe's observations within a geocentric framework. In a linguistic context, Chomsky took over a significant part of the analyses and rule types used in the Post-Bloomfieldian framework, despite the separation by the Chomskyan revolution illustrated by (1) and (2). A more extensive justification of the analysis of the emergence of Chomskyan linguistics as a revolution is given in ten Hacken (1997, 2000). Laudan (1977) explains in more detail why new frameworks after a revolution can and even must take over certain insights from the preceding ones.

Against this background, the fact that Brown et al. (1993:294f.) explain the need to include morphological knowledge in a probabilistic system and Brown et al. (1990:84) even suggest including syntactic trees does not provide decisive evidence against the analysis of the emergence of probabilistic approaches as a revolution. More pertinent indications are the analysis of this emergence in Gazdar (1996), who observes a ”merger” of the two frameworks, and the work exemplified by Klavans & Resnik (1996), which illustrates this merger. The central difference can be described as follows. When a new framework after a revolution takes over knowledge from the old one, the new framework is dominant and what is taken over is a mixture of data, formal mechanisms, and low-level analyses, which are generally reinterpreted quite drastically. In the case of linguistic and probabilistic approaches exemplified by Eurotra, Rosetta, and the IBM system, there is enough agreement on the basic question to be answered that a competition of approaches can take place in the components of a common framework. The typical kind of collaboration is that for components such as morphological analysis, syntactic analysis, semantic analysis, and translational disambiguation probabilistic and linguistic approaches compete in order to produce an overall system which is better than one which is entirely probabilistic or entirely linguistic in its choice of knowledge.

In conclusion, the emergence of probabilistic approaches in CL, exemplified by the IBM system for statistical MT, does not constitute a revolution, but only a radical change in the choice of knowledge.

Finally let us turn to Verbmobil. As noted in section 3 above, a strong attack at the foundations of traditional MT is at the root of this project. As opposed to the IBM project, the focus of the attack is not primarily the type of knowledge used but the very problem as it is modelled. Thus, Verbmobil’s alternative to (3b) is not dramatically different, but (3a) is. Instead of language processing in isolation, it is a particular problem with its practical context which is the starting point. As a result, (3c) is also changed quite radically. It is no longer sufficient to have a running program as a proof that the proposed model of language processing is possible. Instead, only a fully usable, practical solution is an acceptable outcome.

The contrast in the position of the practical application stands out most clearly in comparison with the treatment by Rosetta (1994). Rosetta suggests that given the core solution of the translation problem in the project, consisting of a set of possible translations for each input sentence, the context-dependent choice of the best translation might be supported by an interactive interface or approximated by whatever technique turns out best. For Verbmobil there cannot be a separation between core solution and practical application, because the practical application motivates not only the enterprise as such but also the structure of the solution.

As a result of the radical break with the traditional choices in (3a) and (3c), it is no longer possible to use the same evaluation criteria for a fair comparison of Verbmobil and any of the other systems discussed. An evaluation of Verbmobil on the basis of the percentage of correctly translated sentences would be entirely beside the point. For the evaluation of systems such as Eurotra and Rosetta in a practical context, Krauwer (1993) uses the metaphor of a gearbox. It is not possible to evaluate gearboxes without building cars around them. In this metaphor, Verbmobil is intended as a car.

We can therefore conclude that the revolution in CL is instantiated in MT by the emergence of systems such as Verbmobil. What remains to be explained is the fact that, in the perception of many people working in MT, the revolution is rather associated with the emergence of probabilistic methods. The reason for this is that the emergence of probabilistic methods changed the tasks of these people and the tools for performing them quite radically. However, the term *revolution* is used here for an epistemological change, not a practical change. While we can say informally that the invention of the telescope revolutionized astronomy, its role in the Copernican revolution as an epistemological process is minor. It only served to confirm certain predictions, when the actual change of mind had already taken place (cf. Kuhn 1957). Similarly, while the emergence of probabilistic techniques revolutionized work in MT only in an informal sense, the revolution as an epistemological process is associated with the reorientation from a principled solution to the abstract problem of translation to the practical solution of a practical problem involving translation.

5 Computational Linguistics after the Revolution

After this analysis of the repercussions of the revolution in CL in the subfield of MT, we are now in a position to broaden our scope and consider which general change in CL the revolution in MT reflects. On the basis of our observations in the field of MT, we can hypothesize the following set of assumptions for the new orientation of CL corresponding to (3).

(5)	a.	Problem:	A practical problem occurring in real life
	b.	Knowledge:	Whatever turns out to be helpful in a solution
	c.	Solution:	A system or program in practical use

The most striking difference between (3) and (5) is the relative importance of the choice of knowledge. In the older type of CL, based on (3), ideological oppositions were most often associated with the preference for one theory or another. In the new type of CL, based on (5), there is hardly any room for ideological oppositions on this basis, because what counts is the practical use of the solution.

At the same time, in connection with this reduced importance of the choice of knowledge, we find a change in the status of practical applications. In the older type of CL, a suggestion of how the program might be used in practice was considered enough in most cases. If the question was addressed in more depth, this was typically because financial support required it. In the new type of CL, the practical use of the product is part of the conception of the problem from the start.

As noted in section 1, a scientific revolution is always associated with progress in the field, but certain possibilities are also lost compared to the older framework. If progress is measured in terms of the amount of money invested and the number of people working in the field, the 1990s certainly featured substantial progress in CL. While ten Hacken (to appear) observes a relative decrease in the number of papers on MT at Coling conferences from 1988 to 1998, this is largely (and perhaps entirely) due to the diversification of the field, in which a greater variety of problems are taken on than ever before. The 1990s also saw the emergence of specialized conferences on a scale never seen before. Finally, the practical use of solutions produced by CL has increased. This is not surprising if the formulation of problems is guided not by the theoretical knowledge to be applied but directly by the observation of the problem in real life. What is lost in this development is the close connection between computational and theoretical approaches to language. The determinism which depends on the validity of formal correctness as a goal is lost.

In conclusion, the revolution in CL in the 1990s consisted in the reorientation of the field from knowledge to problems. Before the revolution, CL was performing theoretical linguistics by slightly different means. If practical applications played a role at all, as in Météo, it was a compromise imposed by financial constraints. Since the revolution, researchers in CL have discovered how interesting practical problems are. Taking into account what works in practice needs no longer to be imposed after the fact, but guides the research from the start. Many of the tasks concerning details have remained the same, but the overall design of CL projects has changed.

It is interesting to compare this development to what happened in electrical engineering. The Leyden jar, invented by Pieter van Musschenbroek in 1746, was a device which can be said to work in the sense that it is not merely theoretical. Its only use at the time, however, was in classroom demonstrations showing that there is such a thing as electricity. By contrast, when Thomas Edison invented the carbon-filament lamp in 1879, it was the conclusion of a search for a solution to the problem of producing electric light. This problem had occupied a host of researchers for fifty years or so. CL before the revolution can be compared to electrical engineering in the 18th century. It produced working systems which like the Leyden jar did not have an immediate practical function but could be used to demonstrate certain theoretical notions. In MT, the emphasis was on linguistic theory rather than on the practical problem of translation. Since the revolution, CL has become more like electrical engineering in the 19th century. It is no longer sufficient that a system works, it must also be usable in practice. Therefore, the starting point of CL research has shifted from applying linguistic theory to solving practical problems. Arguably, only since the revolution has CL really become an applied science.

References

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

Brown, Peter; Cocke, John; Della Pietra, Stephen; Della Pietra, Vincent J.; Jelinek, Fredrick; Mercer, Robert L. & Roossin, Paul S. (1988), 'A Statistical Approach to Language Translation', in Vargha (ed.), p. 71-76.

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.

Brown, Peter F.; Della Pietra, Stephen A.; Della Pietra, Vincent J. & Mercer, Robert L. (1993), 'The Mathematics of Statistical Machine Translation: Parameter Estimation', *Computational Linguistics* 19:263-311.

Chen, Xiang (1997), 'Thomas Kuhn's Latest Notion of Incommensurability', *Journal for General Philosophy of Science* 28:257-273.

Chomsky, Noam & Halle, Morris (1965), 'Some controversial questions in phonological theory', *Journal of Linguistics* 1:97-138.

Coling (1984), *Proceedings of Coling84, 10th International Conference on Computational Linguistics, 22nd Annual Meeting of the Association of Computational Linguistics*, Stanford CA: Stanford University.

Gazdar, Gerald (1996), 'Paradigm Merger in Natural Language Processing', in Wand, Ian & Milner, Robin (eds.), *Computing Tomorrow: Future research directions in computer science*, Cambridge: Cambridge University Press, p. 88-109.

Grishman, Ralph (1986), *Computational Linguistics: An Introduction*, Cambridge: Cambridge University Press.

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

Hacken, Pius ten (2000), *Reusability and Research Programmes in Computational Linguistics*, Habilitationsschrift, Universitt Basel (Switzerland).

Hacken, Pius ten (to appear), 'Has There Been a Revolution in Machine Translation ?', to appear in *Machine Translation*.

Householder, Fred W. (1965), 'On some recent claims in phonological theory', *Journal of Linguistics* 1:13-34.

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.

Kay, Martin (1973), 'The Mind System', in Rustin, Randall (ed.), *Natural Language Processing: Courant Computer Science Symposium 8, December 20-21, 1971*, New York: Algorithmics Press, p. 155-188.

Kay, Martin; Gawron, Jean Mark & Norvig, Peter (1994), *Verbmobil: A Translation System for Face-to-Face Dialog*, Stanford (Calif.): CSLI.

Klavans, Judith L. & Resnik, Philip (eds.) (1996), *The Balancing Act: Combining Symbolic and Statistical Approaches to Language*, Cambridge (Mass.): MIT Press.

Krauwer, Steven (1993), 'Evaluation of MT Systems: A Programmatic View', *Machine Translation* 8:59-66.

Kuhn, Thomas S. (1957), *The Copernican Revolution: Planetary Astronomy in the Development of Western Thought*, Harvard University Press, Cambridge (Mass.).

Kuhn, Thomas S. (1970), *The Structure of Scientific Revolutions*, Second Edition, Enlarged, Chicago: University of Chicago Press.

Laudan, Larry (1977), *Progress and Its Problems: Towards a Theory of Scientific Growth*, Berkeley: University of California Press.

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

Louessard, Laurent (1990), *La rvolution de juillet 1830*, Spartacus.

Maegaard, Bente & Perschke, Sergei (1991), 'Eurotra: General System Design', *Machine Translation* 6:73-82.

Melby, Alan K. (1988), 'Lexical Transfer: Between a Source Rock and a Hard Target', in Vargha, (ed.), p. 411-413.

Pereira, Fernando C.N. & Shieber, Stuart M. (1987), *Prolog and Natural-Language Analysis*, Stanford: CSLI.

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

Rossi, Paolo (1997), *La nascita della scienza moderna in Europa*, Roma / Bari: Laterza.

Shannon, Claude S. (1948) 'The mathematical theory of communication', *Bell Systems Technical Journal* 27:379-423 and 27:623-656.

Shieber, Stuart M. (1984) 'The Design of a Computer Language for Linguistic Information', in Coling (1984), p. 362-366.

Vargha, Dnes (ed.) (1988) *Coling Budapest: Proceedings of the 12th International Conference on Computational Linguistics*, (2 vol.), Budapest: John von Neumann Society for Computing Sciences.

Wahlster, Wolfgang (2000). Mobile Speech-to-Speech Translation of Spontaneous Dialogs: An Overview of the Final Verbmobil System, *in* Wahlster, Wolfgang (ed.), *Verbmobil: Foundations of Speech-to-Speech Translation*, Berlin: Springer.