Natural Language Engineering **9** (4): 421–424. © 2003 Cambridge University Press
DOI: 10.1017/S1351324903003280 Printed in the United Kingdom

Industry Watch

ROBERT DALE
Centre for Language Technology, Macquarie University, Sydney, Australia
e-mail: rdale@ics.mq.edu.au

A principal aim of this journal is to bridge the gap between traditional computational linguistics research and the implementation of practical applications with potential real-world use. This new column aims to help with that bridge-building, by providing a window on to what is happening with speech and language technologies in the industrial and commercial worlds, and a discussion forum for topics of interest in the technology transfer space.

The timelines for journal production mean that this cannot be an up-to-the-minute review of late-breaking news; for that kind of information, check out LT World at www.lt-world.org, or the HLT Central news pages at www.hltcentral.org. The intent of this column is a little different, turning the negative of traditional publishing delay into the positive of calmer retrospective. Mostly that means we'll review recent events; but to kick things off, we'll be *very* retrospective and consider where speech and language technologies have got to in the last 40 or 50 years. Forty to fifty years seems about right as an estimate of the age of the field: early research into machine translation was underway in both UK and US universities in the mid-1950s, and the first annual meeting of the Association for Computational Linguistics was in 1963.

The development of useful applications has been a motivation for research in natural language processing since the field was born. So how far we have we come in that time? An interesting intermediate stopping point is the rough halfway mark of 1985: that's the year that Ovum Ltd, a UK-based market research consultancy, published its first report into the commercial applications of natural language processing (Johnson 1985). Table 1, compiled from a number of tables in that report, shows that report's predictions for market size in a number of application areas in both the UK and US (the report did not explicitly consider other markets). The absolute values in this table are not interesting: predictions were considerably downgraded in Ovum's second report on the area in 1991, and market projections are always extremely variable, since they depend so much on varying assumptions and definitions. What's more interesting are the categories, their predicted *relative* growth, and how that compares with where we are now.

The categories have remained remarkably stable, although the fortunes of each are a little more varied. **Mainframe and micro interfaces** are what we know and love as natural language database interfaces (NLDBIs): in the Ovum report, NL interfaces on mainframes and PCs were treated separately, but this seems not so relevant now, so I've combined the two here. Taken together, this category was predicted

https://doi.org/10.1017/S1351324903003280 Published online by Cambridge University Press

Table 1. *Market projections from the 1985 Ovum Report*

Application Category	1985		1990		1995	
	UK	US	UK	US	UK	US
Talkwriters	$0m	$0m	$7m	$95m	$63m	$460m
Mainframe and micro interfaces	$0.4m	$12m	$17m	$179m	$57m	$434m
Text editing	$0m	$0m	$4m	$38m	$46m	$300m
Dialogue interfaces	$0m	$0m	$2m	$8m	$13m	$100m
Content scanning	$0m	$1m	$5m	$45m	$16m	$94m
Machine translation	$0.7m	$2m	$20m	$53m	$33m	$84m
Totals	$1.1m	$15m	$55m	$420m	$230m	$1500m

to be the second largest element of the overall NLP market; but NLDBIs are almost invisible these days. Interestingly, 1985 was also the year that Gary Hendrix brought out his Q&A database interface product, heralded soon after as the most widely-used natural language processing application. Calm retrospective suggests that the product was successful for its database and related back-end capabilities, rather than its natural language interface. The company Gary founded to sell the product was Symantec, now well known to most people for its anti-virus software and other utilities, with not an NLP application in sight.

But NLDBIs are the only category here that has died; evidence of the continued life of the other categories is all around us. **Talkwriters** are desktop applications like Dragon Naturally Speaking, IBM ViaVoice and Philips' FreeSpeech: the Ovum report predicted these to form the largest single category by 1995. Maybe they were indeed a winner, in that many offices have a copy of one or more on their shelves, but I've come across very few people who regularly use them. On the other hand, **dialog interfaces**—the report's term for the applications we now call spoken language dialog systems—is now, from a commercial perspective, the most clearly distinguished market sector; there's wide public awareness of the existence of the technology, and many people will by now have booked a taxi, checked a flight time, sold some shares or ordered a pizza by talking to a voice recognition system. Contrary to Ovum's expectations, this is also the area that appears to have demonstrated most market potential.

The other categories are still there, but you'd be hard-pushed to call any of them 'killer apps'. **Text editing** still exists as an application category, but a relatively dormant one, since the market space is all but monopolised by the grammar checker in Microsoft's Word. This is surely the natural language processing application with the largest installed base in the world, but people don't buy Word *because* of the grammar checker.

Content scanning, which in 1985 represented applications such as the Carnegie Group's processing of banking messages, would now include information extraction, text mining and related document processing technologies. Many crystal-ball gazers in the early 21st century would nominate this as the most likely future winner, although in 1985 it was near the bottom of Ovum's list.

https://doi.org/10.1017/S1351324903003280 Published online by Cambridge University Press

Machine translation has persisted as an application area, and one that has become much more publicly visible as a consequence of web search engine translation services. But it remains to be seen whether the inevitably limited quality of these services will help or hinder the development of that market.

Getting accurate assessments of current market size of any of these categories is difficult, but it would seem that, of all the applications Ovum surveyed, dialog interfaces came out on top, and the others are still waiting in the wings for their moment in the spotlight. So why didn't things quite work out as expected? The development of the graphical user interface in the 1980s may have played a role in putting paid to NLDBIs. The web, whose significance no one could have guessed in 1985, has also played a role in changing the landscape for language technologies: if you're looking for information, you can get a long way without relying on NLP just by using Google. And end-users don't seem to care enough about grammar checking to motivate much further development there. Most importantly, though, the other technologies currently have high enough error rates that their usefulness is limited: not surprisingly, people won't rely on a technology which they can't trust to get it right.

What does any of this tell us about the future? Well, mostly that crystal-ball gazing is hard, which we knew already; Ovum can't really be faulted, since they probably got their hot tips from people just like you, dear reader. A modern reappraisal of the technology space would be more likely to distinguish two broad categories of systems: spoken language systems and intelligent text processing capabilities. Of course, these broad domains have the convenient property of allowing us to hedge our bets on the substructure of the categories. So, enough of calm retrospective; here are two predictions aimed at generating some audience response.

- In the next five years, the market for dialog interfaces will rocket, with more phone calls to businesses being answered by voice recognition systems than by humans. But they'll almost all be pretty simple applications, and they'll probably use fairly simple linguistic interfaces, more along the lines of Rosenfeld's Universal Speech Interface (Rosenfeld *et al.* 2000). They won't be like the sophisticated, naturalistic dialog systems coming out of the research labs; recognition error rates will remain roughly where they are now, limiting the scope for complex applications to niche markets.
- XML, schmex-ML: Ken Church and his colleagues once pointed out (Church *et al.* 1994) that, much as our lives as language technologists would be easier if all we had to process was SGML, the world is using fax because it is more convenient for them. The moral of the story was that we ought to be developing techniques to do useful things with that kind of grungy data. Times have changed: now people are increasingly using PDFs rather than faxes, again because they are convenient. But times haven't changed all that much: many of those PDFs are scanned and OCRed, just like faxes, and that will remain the case for a few years to come. Text processing technologies won't gain a significant foothold until the ugliness of *real* real data is confronted head-on.

https://doi.org/10.1017/S1351324903003280 Published online by Cambridge University Press

Provoked? Do you have differing views on any of the above? Email me at rdale@acm.org, and I'll follow up in a future column.

References

Church, K., Gale, W., Helfman, J. and Lewis, D. D. (1994) Fax: An Alternative to SGML. In *Proceedings of Coling 1994*, Nantes, France.

Johnson, T. (1985) *Natural Language Computing: The Commercial Applications*, Ovum Ltd, London.

Rosenfeld, R., Zhu, X., Shriver, S., Toth, A., Lenzo, K., and Black, A. W. (2000) Towards a Universal Speech Interface. In *Proceedings of ICSLP 2000*, Beijing, China.

https://doi.org/10.1017/S1351324903003280 Published online by Cambridge University Press