Industry watch
The commercial NLP landscape in 2017

R O B E R T   D A L E

Language Technology Group
e-mail: `rdale@language-technology.com`

(*Received 28 April 2017*)

Abstract

The commercialisation of natural language processing began over 35 years ago, but it's only in the last year or two that it's become substantially more visible, largely because of the intense popular interest in artificial intelligence. So what's the state of commercial NLP today? We survey the main industry categories of relevance, and offer comment on where the action is today.

---

1 The big picture

If you want to know what's happening in the commercial NLP space, and what the future of that space might be, a sensible place to start is by looking at the market research reports produced by technology analysts like Gartner, Forrester and IDC. These companies make a living out of talking to technology companies about their plans and prospects, and then writing up the results of that research in reports that typically sell for several thousand dollars a pop, identifying the major players and making predictions about future market size.

The first such report I'm aware of that addressed NLP technology was published in 1985 by Ovum.[1] Tim Johnson, the author of that report, identified the categories of NLP technology defined in Table 1.

At the time, the very first commercial applications in some of these categories were just appearing. In 1984, IBM had demonstrated a dictation system (then categorised by Johnson as a 'talkwriter') which could transcribe complete business letters, and that same year saw the introduction of a content-scanning application by Cognitive Systems that read money transfer telexes.[2] The early 1980s also saw the first commercially released NL interfaces to databases and the first commercial

---

[1] For a useful summary of the report, see Tim Johnson (1984), Natural Language Computing: The Commercial Applications, *The Knowledge Engineering Review*, Volume 1, Issue 3.

[2] You can read a conference paper about this early information extraction system at `http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.568.514&rep=rep1&type=pdf`.

Table 1. *NLP technology categories from Ovum's 1985 report*

Category	Definition
Database interfaces: mainframes and minis	Portable or integrated front-ends to DBMS and other applications software
Database interfaces: micros	PC or similar interfaces to application software or online services
Dialogue interfaces	Conversational interfaces to complete systems, such as expert systems or ICAl applications
Content scanning	Processing messages or other semi-formatted texts to decide action or routing
Text editing	Checking style and grammar of text and suggesting improvements
Machine translation	Computer translation of texts from source to target natural language
Talkwriters	Transcription of spoken input into computer text

MT systems. So it's not unreasonable to think of this period as the beginning of the commercialisation of NLP.

How have things progressed in the thirty-plus years since Johnson's analysis? A review of today's much larger universe of technology market research suggests some revisions to the set of categories, but they have remained remarkably robust:

- Johnson's first three categories, covering Database and Dialog Interfaces, have morphed into what today are called Conversational Systems, and have also had some of their functionality overtaken by the magic of web search—back in the 1980s, it was unthinkable that you would be able to type into a general search engine the query 'weather Sydney tomorrow' and get exactly what you were looking for displayed almost instantaneously.
- Context scanning has broadened out into the category we now know as text analytics.
- Text editing remains as a category, primarily thought of these days as grammar checking.
- Machine translation has remained as the most easily identifiable and persistent application category since the earliest days of NLP.
- Talkwriters, at least as defined by Johnson back in 1984, have become effectively a commodity, with speech recognition built into the leading operating systems, and widely available as a relatively inexpensive add-on with enhanced features.

So, with this slightly refreshed set of categories, what does today's commercial NLG landscape look like? This article provides a rundown, ordered in terms of where most activity appears to have been of late. We'll skip speech recognition, given its commodity status, the fact that it's perhaps more appropriately in the domain of publications other than this one, and because of space limitations.

2 Conversational systems

Gartner views conversational systems as one of the top ten strategic technology trends for 2017,[3] and predicts that, by 2020, the average person will have more conversations with bots than with their spouse.[4] This technological optimism has its detractors, with some commentators observing that the initial promise—that chatbots might replace apps and websites—doesn't quite seem to have been realised.[5] The massive numbers thrown around by the platform providers—such as Facebook's claim to be hosting 100,000 bots,[6] and Pandorabots' 'more than 285,000 chatbots created'[7]—hide the likelihood that many of those chatbots are trivially simple and probably have never been used by anyone other than their creators.

Most of the coverage of chatbots you'll find on the web focusses on what William Miesel calls 'specialized digital assistants',[8] these being very narrowly focussed interactive apps that let you achieve a very specific task, like buying flowers or tracking flights. The more interesting of these chatbots are conceptually similar to the much older category of telephony-based spoken language dialogue systems that let you book a taxi or order a pizza—incidently, an application category which itself has become much less significant given the capabilities on a typical smartphone. The fundamental differences are largely to do with modality, in that today's chatbots are predominantly deployed on text-based messaging platforms rather than by calling up a telephone number and communciating via an audio-only channel. That has two consequences. First, by removing or weakening the reliance on speech recognition, it's possible to avoid many of the issues around misrecognition and error handling that arise in voice-driven systems. Second, the fact that any phone on which you use a messaging app also provides supplementary modalities reduces the dependency on language. So, for example, when the app requires you to make a choice, you can select an item by touching an image on the screen, short-circuiting the need for you to actually describe what you want using language, and the consequent need for the machine to understand that description; and the app can use images to convey information far more efficiently than would ever be possible using language—try choosing a shirt via voice-only interaction.

So what's the current state of the chatbots world? It looks to me like the chatbot ecosystem is suffering the technical analogue of urban sprawl. There are now dozens of platforms for building chatbots for Kik, Twitter, Facebook Messenger, WeChat, WhatsApp, Slack, Skype and Line. As we've discussed in this column before, the Big Four—Apple, Amazon, Facebook and Google—provide SDKs that allow you to

---

[3] http://www.gartner.com/smarterwithgartner/gartners-top-10-technology-trends-2017/

[4] Gartner, *Top Strategic Predictions for 2017 and Beyond: Surviving the Storm Winds of Digital Disruption*, 14 October 2016.

[5] https://www.usatoday.com/story/tech/news/2017/04/18/facebook-messenger-takes-another-swipe-chat-bots/100596798/

[6] http://www.cnbc.com/2017/04/18/messenger-ceo-david-marcus-at-f8-more-bots-announced.html

[7] https://www.pandorabots.com

[8] http://www.tmaa.com/specializeddigitalassistantsandbots.html

extend the capabilities of their native interactive virtual assistants;[9] and in mid-2016, IBM entered the fray with its Conversation chatbot builder.[10] The barrier to entry is now very low, so that anyone can be a chatbot developer, and even the simplest interactive application appears to count as a chatbot.

Consequently, the chatbot landscape is like a Google Maps satellite view dense with thousands of indistinct red-roofed houses, many more garden sheds, and a few architecturally interesting buildings that are worth a visit. What marks out some of those interesting buildings is the enhancement of their utility via machine learning: while the simpler approaches to building chatbots rely on simple pattern matching against input utterances, the better frameworks provide learned classification of user utterances into a smaller set of intents that will be handled by the application. If there's a sensible way of organising the space so that chatbots with the same target functionality can be compared, then it's these smarter apps that are likely to win out. But in the interim, there's a significant risk of consumer disenchantment: it's hard to know what exists, and when you find something, there's a good chance it will be disappointing.

3 Machine translation

As noted above, machine translation has been a robust industry category for a long time. The major players in the space are in fact language services providers (LSPs) whose primary function is broader than just the application of machine translation: of note are Lionbridge Technologies (the largest publicly traded translation and localisation company in the US), Moravia IT, SDL, Systran International, Translations.com and Welocalize. While all of these companies make use of MT, really what they are about is managing projects that involve localisation and other translation services, which typically require higher quality than 'naked MT' will provide. So, an important focus for these businesses is their automation, support and post-editing tools.

For readers of this journal, the more visible MT players are of course Google and Microsoft, who have made non-curated automatic translation a commodity via their web-based translation services. But there's still a significant quality gap between fully automated MT and the human-assisted translations of the LSPs. It's of note that Microsoft and Google are also Lionbridge's two top clients, accounting for fifteen per cent and eleven per cent of total 2015 revenues, respectively.[11] As far as MT is concerned, eating your own dog food still remains something of an aspiration.

But perhaps that quality gap is due to shrink a bit. The big news this year is that the major players have moved towards adopting Neural MT as their base

---

[9] `https://www.cambridge.org/core/journals/natural-language-engineering/`
`article/the-return-of-the-chatbots/0ACB73CB66134BFCA8C1D55D20BE6392`
[10] `https://www.ibm.com/watson/developercloud/conversation.html`
[11] `https://slator.com/financial-results/what-we-learned-from-the-lionbridge-`
`10k-filing/`

technology. Google's NMT was announced in 2016[12] and made part of Google's standard API offering this year.[13] Not to be outdone, Microsoft Translator[14] and Systran[15] also announced the use of NMT in late 2016. This is definitely a space worth watching.

4 Text analytics

Text analytics, as we're using the term here, covers a wide range of technologies that aim to extract meaningful content from text, either in documents, emails or short-form communications such as tweets and SMS texts. The functionalities you'll typically find, although not all of these are provided by every vendor in the space, are named entity recognition, concept extraction, text classification, sentiment analysis, summarisation, and sometimes relation extraction and parsing.

Typical use cases are social media analysis, e-discovery and voice-of-the-customer analysis. In particular, a major focus for these technologies today is the determination of sentiment, and this has become such a focus for text analytics that some providers who used to see themselves as providing more generic functionality have zeroed in on this much narrower area. So, for example, two of the longest standing text analytics firms have reinvented themselves in this space: Attensity focussed in on social customer relationship management before being acquired by maker of contact-centre software InContact in 2016,[16] and Clarabridge has for some time branded itself as a provider of customer experience management software, rather than a pure text analytics company.

The rest of the today's text analytics offerings fall into two major clusters. First, every big IT company now claims a text analytics capability: IBM owns SPSS Text Analytics, SAS offers its Text Miner software, SAP has HANA Text Analytics, Oracle Data Miner incorporates text mining. The more interesting category, though, is the growing number of software-as-a-service text analytics APIs. In a previous Industry Watch column, we discussed five of these: AlchemyAPI (part of IBM's Watson Cloud offerings), Aylien, Lexalytics, Meaning Cloud and TextRazor.[17] AlchemyAPI has since been rebranded as IBM Natural Language Understanding.[18] Aylien seems to have gone relatively quiet: the most recent entry on their website's 'In the News' page is dated mid 2015, although it appears they obtained further funding in early

---

[12] https://research.google.com/pubs/pub45610.html
[13] https://cloud.google.com/translate/docs/premium
[14] https://blogs.msdn.microsoft.com/translation/2016/11/15/microsoft-translator-launching-neural-network-based-translations-for-all-its-speech-languages/
[15] http://www.systransoft.com/systran/translation-technology/pure-neural-machine-translation/
[16] http://odintext.com/blog/attensity-sells-ip-to-incontact-what-does-it-mean/
[17] https://www.cambridge.org/core/journals/natural-language-engineering/article/nlp-meets-the-cloud/EBD54F291A0BE29A60E1A15BD1DED736
[18] https://www.ibm.com/watson/developercloud/natural-language-understanding.html

2016, and around the same time they announced a News API.[19] Lexalytics, Meaning Cloud and TextRazor still appear to be quite active, with the last of these just announcing a Chinese language capability in March of this year.

Meanwhile, other text analytics players have come to the fore. Names that I've seen increasingly in recent months are Attivio, Cambridge Semantics, Expert System and OpenText; and old hands Basis Technology and Linguamatics seem to be alive and well. There's an increasing trend for the companies in this category to offer text analytics as one component in a wider package of technologies, which may suggest that consolidated and integrated solutions are where it's at today, and that it's a struggle to sell isolated text analytics capabilities. On reviewing the websites of many of these companies, one thing that strikes me is that we haven't seen much in the way of technical advances over the last few years. It will be interesting to see whether deep learning begins to make an impact here.

5 Text correction

This category primarily covers grammar checkers, but also includes spelling correction and style checking. Of all the areas of commercial NLP, this seems the most stagnant. Which is sort of weird, because you might have thought that the democratisation of publishing that social media has brought would have increased significantly the demand for checking the correctness of a text before hitting the 'Submit' button. Or if you're cynically of the view that millenials just don't care, surely at least there'd be increased demand from more professional writers as hard-pressed newspapers and publishers fire their subeditors and copy editors.

But if there is increased demand for this technology, it's not manifesting itself in the arrival of new entrants with new ideas. The grammar and style checking market seems somewhat dominated by two players, as it has been for a while. First, there's Microsoft, by virtue of the presence of the Word grammar checker being present on 1.2 billion desktops in 2016[20]—although I'm not suggesting that anything like this number pay attention to the red, green and blue underlining. Then, there's Grammarly (apparently with three million registered users as of August 2016[21]), which we've discussed at length in an earlier Industry Watch column.[22] There are a small number of other players here, but it's clearly not an area that's seeing much innovation. If you're in the mood for creating a disruptive start-up, maybe this is the one to go for.

---

[19] https://techcrunch.com/2016/04/07/aylien-launches-news-analysis-api-powered-by-its-deep-learning-tech/
[20] http://www.windowscentral.com/there-are-now-12-billion-office-users-60-million-office-365-commercial-customers
[21] https://www.grammarly.com/blog/3-million-registered-users-trust-grammarly-to-perfect/
[22] https://www.cambridge.org/core/journals/natural-language-engineering/article/checking-in-on-grammar-checking/EC45EFC337EA1EC1A4104D1E534B4D12

6 Summing up

So, that's where we're at in 2017. Chatbots are still the major focus of interest; machine translation appears to have received a jolt of energy via NMT; text analytics has settled down to become part of the furniture and the text correction market is looking like it needs a kick up the bum. Any takers?[23]

---

[23] This piece was written just before Grammarly's early May announcement of $110M funding: see `https://techcrunch.com/2017/05/08/grammarly-raises-110-million-for-a-better-spell-check/`.