*Industry watch**

R O B E R T   D A L E

Centre for Language Technology, Macquarie University, Sydney, Australia

A year ago in this column I wondered aloud whether 2007 was to be the year in which question-answering (QA) really took off in the commercial space. I was provoked to ask that question by the increasing number of Web-based QA systems that were portraying themselves as the Next Thing in search for the masses. There was, in particular, a lot of buzz around the $12.5 million funding deal announced by Powerset.[1] The San Francisco-based company had gained exclusive access to parsing technology from PARC,[2] but hadn't at that point displayed any of its wares to the general public. However, the company was inviting people to sign up for access to Powerset Labs,[3] where, we were told, we would get the opportunity to be the first to play with the technology and to provide feedback to make it better. Since then, there have been occasional screen images of the application seen in blog posts and other news items, and a small number of claimed sightings by bloggers who were granted privileged access. But Powerset Labs was finally launched at TechCrunch 40 in mid-September.[4] The company's Web site says that they've begun to let people have access to the technology, and that they'll be 'letting in the next wave of users as soon as possible'. My surf board is ready, but I'm not holding my breath. I signed up in June 2007, and haven't heard a thing since. Comments posted on the Powerset site suggest that there might be quite a few in the line before me.

Meanwhile, in a relatively brash New York kind of way, East Coast-based Hakia's Web-based QA system[5] is out on the street, up and running in beta form, so you can try it out. I asked Hakia: 'When will Powerset let people try its technology?' The results came back in standard search-engine layout as short snippets of text with key elements highlighted. The first couple of hits looked reasonably relevant, but, disappointingly, not much different from what Google returns in response to the same question. Hakia does offer some other interesting functionalities, though. They have a new ScoopBar tool, which, in response to a question, will highlight the appropriate content on a Web page, and let you 'Scoop 'n Save' [sic] a user-definable text window around the answer to a local file. This is quite neat; even without the

---

*Industry Watch is a semi-regular column that looks at commercial applications of natural language technology. The author can be contacted at `rdale@acm.org`.
[1] `www.powerset.com`
[2] `www.parc.xerox.com`
[3] `http://labs.powerset.com`
[4] `www.techcrunch20.com/2007/`
[5] `www.hakia.com`

QA capability, it's a useful productivity tool. But now I have so many toolbars and other stuff plugged into my browser that I can't see very much of the Web pages that I use them to find in the first place.

A couple of other QA systems that have achieved some visibility recently are Lexxe[6] and PowerAnswer.[7] I've been aware of Lexxe for a while, largely because it's based in my home town of Sydney; in recent months, it's received quite a few blog mentions. I asked Lexxe what will henceforth be known as 'the Powerset question': 'When will Powerset let people try its technology?' Lexxe was rather slow in providing a response, and the results seemed less relevant than Hakia's.

PowerAnswer comes from Lymba, a spin-off from Language Computer Corporation that's headed up by Dan Moldovan. The Powerset question took longer still here to process; I suspect speed of response will be a real showstopper for many of these applications. The results didn't look that promising either, until I got to the last of five hits: one extracted sentence in a box along with a URL: 'Powerset doesn't expect to launch a product until some time in 2008'. Question answered. Maybe they just got lucky, but I was impressed.

Still, Hakia seems to be the most inventive of the current crop when it comes to neat add-on functionalities. For some people of note, a query consisting of just that person's name returns results clustered under convenient headings complete with a table of contents. For example, a search for 'Tony Blair' returns links under the headings 'Official Website', 'Biography and Timeline', 'Awards and Accomplishments', and 'Speeches and Quotes'. A search for 'Britney Spears', on the other hand, provides clusters for 'Music Profile', 'Famous Songs', and 'Movies and Videos'. It's not clear exactly what kind of fame you need to have achieved to get this treatment: a search for 'John Howard', the Prime Minister of Australia,[8] returns a standard undifferentiated list of hits; but Stephen Harper, the Prime Minister of Canada, gets the full treatment. A search for 'Yorick Wilks' delivers clustered results too, but as we've noted in this column before, he is, after all, an advisor to the company.

More generally, people-search is a specific functionality that has seen quite a bit of activity in the last few months. I've reported previously in this column on ZoomInfo[9] and Spock[10]; in August, Spock released its first public beta.[11] The Magazine *Time* has an interesting overview article on people-search, under the headline 'Online Snooping Gets Creepy'[12]: not surprisingly, some people are concerned about the way that tools of this kind allow information gathering without the permission of the person concerned. And, if you haven't already, you have to check out Wonkosphere,[13] a kind of people-search-meets-sentiment-analysis site where you

---

[6] www.lexxe.com

[7] http://lymba.com/products/poweranswer.html

[8] Well, he was at the time of writing, in the run-up to an election. Perhaps not the Prime Minister of Australia when you read this. We live in hope.

[9] www.zoominfo.com

[10] www.spock.com

[11] http://blogs.business2.com/startups/2007/08/people-search-e.html

[12] www.time.com/time/business/article/0,8599,1649121,00.html

[13] http://wonkosphere.com

can keep track of the buzz around the US presidential candidates. Wonkosphere uses Crawdad Technologies' Listening Post technology,[14] which wraps up sentiment analysis in a neat commercial service.

All of this makes me think that there's considerable scope for a third party services market around the whole people-search area. To do this well, you need good algorithms for named entity recognition and cross-document coreference resolution; but it also helps if you have lots of data about people, and ways of integrating diverse data sources, both structured and unstructured. I foresee the Web 2.0 version of specialised *Who's Who*-like biographical resources in vertical markets, on-selling various quality grades of information wrapped up in RDF statements. Social networking sites, particularly of the professionally-oriented variety, are already in a good position to capitalise on this, of course, but I don't think it's too late to be thinking about compiling niche-market people data sets. Cost will be determined by quality: machine-constructed people profiles at a cent a time, maybe 25 cents for the human-curated versions, and so on. This is just the Internet-age equivalent of the market in address lists, of course. I can see a possible future where people package up their own profiles (verified correct, thus addressing a problem with many of the current services) with the last year's worth of clickstream data from their Web browsers; form corsortia of like-minded people; and auction off the resulting data sets to the highest-bidding marketing brokers on eBay.

───────────────

So much for text; what about the voice world? Most visibly, there's been a lot happening in the the voice-recognition-driven directory assistance area; see the *Wall Street Journal*'s comparison of a number of competing services.[15] Amonsgt the more notable events in the last few months was the announcement that Jingle,[16] who offer a free directory assistance service funded by advertising, have entered into a strategic parternship with Nuance.[17] Also worth a look in this area is vlingo,[18] a company started up by Mike Phillips, who was a cofounder of SpeechWorks. Their Web site explains the novel aspects of their approach to speech recognition, using what they call Hierarchical Language Models.

Another space that is starting to look rather interesting is the integration of speech and predictive text. Take a peek at TravellingWave's[19] VoicePredict. The idea here is that you use both the keypad and voice to provide mutual disambiguation and overcome imperfect speech recognition; for example you might input the word *technology* by typing *T* and saying 'technology'. But, as with so many good ideas, the big guys are likely to move in here and take over. In June, Nuance announced the

[14] www.crawdadtech.com
[15] http://online.wsj.com/article/SB118602200167285684.html?mod=googlenews_wsj
[16] www.free411.com
[17] www.nuance.com/news/pressreleases/2007/20070823_jingle.asp
[18] www.vlingo.com
[19] www.travellingwave.com

acquisition of Tegic (the people who developed the T9 predictive text technology) for \$265M.[20]

Speaker authentication and identification, as opposed to recognition, are also coming along. Voice Pay mobile[21] provides a service whereby you can buy things from shops, TV commercials or magazines by dialling a number, voice-authenticating yourself, and speaking the code that appears next to the product you want to buy. BeepCard[22] goes one step further: their ComTalk voice-command credit card contains a speaker, a microphone and a voice-recognition chip. To use the card, you speak into it; if your voice matches the recorded sample, the card sends 'high frequency, hard-to-fake beeps over the phone or a computer microphone, signaling the credit card company to authorize the transaction'.

All around the world, it appears, voice technologies are being used to prevent fraud and combat crime. In Brazil, an alleged leader of Colombia's biggest drug cartel, who had had his features radically altered by plastic surgery, was apparently identified by anti-drug agents using voice recognition technology.[23] Lambeth Council in the United Kingdom intends to use Voice Risk Analysis technology to catch welfare-benefit cheaters[24]: those who contact the call centre will be asked a number of questions, with the technology scanning what they say to detect 'micro-tremors' that indicate when stress is generated by an attempt to deceive. Perhaps with this in mind, a proposal to use speech recognition in the Scottish Parliament led to a suggestion that a lie detector be integrated into the software.[25]

Of course, if you happen to be familiar with the kind of Scottish accent my sister has, you might think the Scottish Parliament would also benefit from speech-to-speech translation. But worry not, it looks like Ectaco may have it covered. Their iTRAVL NTL-9C Talking Two-Way Multilingual Language Communicator and Electronic Dictionary,[26] which sells for £347.95, attempts to replace the foreign language phrase book:

By simply speaking into the iTRAVL, you can have your phrases instantly translated and spoken aloud using a sophisticated combination of speech recognition modules, real human voice narration and speech synthesis.

The device is available for English to Russian, Spanish, Polish, German, Portuguese, French, Italian and Chinese. But at the top of the list of language pairs mentioned on the company's Web site is English to English. Most likely a mistake, but maybe not. Perhaps Alex Salmond (First Minister for Scotland; doesn't get the clustering treatment from Hakia) might be interested in some dialect translation for parliamentary debates.

---

[20] www.nuance.com/news/pressreleases/2007/20070621_tegic.asp
[21] www.voice-pay.com
[22] http://beepcard.com
[23] www.cnn.com/2007/WORLD/americas/08/10/brazil.drugs.ap/
[24] www.theregister.com/2007/09/11/lambeth_council_voice_stress_test/
[25] http://scotlandonsunday.scotsman.com/scotland.cfm?id=1395012007
[26] www.ectaco.co.uk/ECTACO-NTL-9C/