Natural Language Engineering **14** (2): 283–288. © 2008 Cambridge University Press
doi:10.1017/S1351324908004683 Printed in the United Kingdom

Industry Watch: Language technology, meet social networking

ROBERT DALE
Centre for Language Technology, Macquarie University, Sydney, Australia
e-mail: rdale@acm.org

(*Received 9 January 2008*)

Here's a round-up of notable events in the commercial language technology space in the last quarter of 2007, organized by broad application category. A common thread that pops up throughout many of these is the integration of language technology into social networking applications and other related Web 2.0 themes. I'd put my money on this being a hot direction in 2008.

1 Speech

The range of things you can do with speech recognition via your mobile phone just keeps growing.

In October, **SpinVox** announced 'Social Networks through SpinVox',[1] adding to its range of voice-to-text options a new functionality: now you can post on your FaceBook, Jaiku or Twitter accounts by making a phone call, rather than having to mess with one of those pesky keyboard things. Is this a good thing? Those web genres that are traditionally populated with poorly written text will now also have to suffer speech recognition errors. On the other hand, it will probably have the effect of improving the spelling overall.

Also in October, **Microsoft** joined the voice search space that Google and V-Enable, amongst others, have been developing, with Live Search for Windows Mobiles now accepting voice input.[2] The speech recognition is server-based, using the Tellme hosting infrastructure Microsoft acquired back in March.[3]

In December, **ReQall**[4] announced that its voice-to-text reminder system was being made available in the United Kingdom and Canada. With a functionality similar to Jott, you call up the system on your phone, verbalise what you don't want to forget and a transcribed reminder is made available via your reQall and iGoogle pages, by RSS, through a standard iCal feed, a text message alert, or in a daily email agenda.

[1] www.spinvox.com/social-networks-facebook-twitter-and-jaiku-all-gain-voice.html
[2] www.livesearchmobile.com
[3] See the 'Industry Watch' column in Volume 13, Issue 2 of this journal.
[4] www.reqall.com

https://doi.org/10.1017/S1351324908004683 Published online by Cambridge University Press

At least for the moment, trial subscriptions are free. Their website has a pointer to a page of YouTube demo videos.[5]

Before long, virtually every voice communication will be subjected to speech recognition. **Pudding Media**[6] is set to monetize this ability. When you call any telephone number in the United States or Canada for free via their web app, they use speech recognition to analyse what you are saying, and then place helpful relevant links—and of course advertisements—on the page during your call; see the screenshot on their website.[7]

Doing something useful with this on a mobile phone screen may be more challenging, but if these guys get together with SpinVox, I bet there's a fortune to be made off the back of the inane activity-tracking comments that populate FaceBook and Twitter; imagine a Starbucks' ad on your phone every time you tweet 'Just getting a cup of coffee'.

2 Named entity recognition

To continue on this social networking theme for a bit: Social networking is all about communities, and let's not forget that communities were made up of physically co-located individuals long before they were part of the virtual landscape. So you might suspect geospatial named entity recognition could have a role to play here, and it does. **YourStreet**[8] collects news stories and places them on its map-based interface, down to the nearest street corner. The site collects 30,000 to 40,000 articles a day from more than 10,000 RSS feeds, mostly from community newspapers and blogs. Given a location (as far as I can tell, limited to the United States, like most of these things), the site shows a map of that area, stuck with pins that represent the locations of news stories, user-generated content called conversations and people who have added themselves to the map: inevitably, the site includes social-networking features, such as the ability to log in, meet neighbours, start conversations and leave comments as annotations on stories. YourStreet says it uses the **MetaCarta** service to 'assist in location identification', but it's not clear whether they have added anything to MetaCarta's geospatial NER.[9]

More text analytics vendors are making it easier to see how well their technology actually works. **Lexalytics**[10] have made a free version of their text analytics software available for download on their website. The trial installation is limited to 50 documents at a time, but handles various file types, including HTML, Word, PowerPoint, PDF and text. It can also take input from a list of full-text RSS feeds. The application identifies people, companies, places and other entities; it

[5] Some of which are not about ReQall at all, but have been mislabelled with a misspelling of *recall*, which probably wouldn't have happened if the providers had uploaded and annotated their videos using Spinvox instead of relying on their own spelling ability.
[6] `http://puddingmedia.com`
[7] `www.thepudding.com/media/app/thepudding_screenshot.html`
[8] `www.yourstreet.com`
[9] See `www.metacarta.com` for a two-minute demo video of what MetaCarta does.
[10] `www.lexalytics.com`

https://doi.org/10.1017/S1351324908004683 Published online by Cambridge University Press

also attempts to assess the sentiment expressed towards each entity and provides summaries of the documents. None of this is rocket science by today's standards, but the application is a very straightforward install, and gives an idea of what commercial capabilities in this space are like. The company's website also provides a number of white papers that give a little bit of background on their technology.

Even better value for money is **IBM**'s Omnifind Personal Email Search, a free fully-functioning download.[11] The application integrates with Lotus Notes and Microsoft Outlook (2003 or 2007 only); it works by extracting and organizing concepts and relationships from your email, and then allowing you to carry out 'semantic searches' by posing queries using simple keywords: so, for example, you might ask for 'PPT from Rob' or 'Bill phone', and what you're looking for will be found even if expressed in other ways. Apparently you can extend the range of concepts the application identifies using a tagger utility, and the taggers you produce can be shared with others.

3 Sentiment analysis

Sentiment analysis now seems to be a 'must-have' for any generic text analytics vendor, but there are also a fair number of companies focussing just on this area. To get an idea of what they sell, check out Cymfony (`www.cymfony.com`; said to be the market leader in this category), Biz360 (`www.biz360.com`), Brandimensions (`www.brandimensions.com`), MotiveQuest (`www.motivequest.com`), Nielsen BuzzMetrics (`www.nielsenbuzzmetrics.com`) and Umbria (`www.umbrialistens.com`), all of which have been around for a while. Some newer players are also worth a look:

- **SentimentMetrics** (`www.sentimentmetrics.com`) mines huge numbers of weblogs for information about celebrities, trademarks and so on. It uses a web service from Scientio to determine the sentiment of each weblog about the trademark or celebrity and aggregates these as a daily updated popularity measure. See the website for examples of the reports they produce and a free 14-day trial if you have a product or service whose buzz you'd like to track.
- **Scout Labs** (`www.scoutlabs.com`) 'helps companies make sense of the mass of positive and negative things that are said about their brands in blogs, user-generated videos, and images'; to use the application, you create 'scouts' for each brand you want to track, and the application then searches through archived material from the beginning of 2007. See their demo page for a video that shows the various elements of their dashboard.

Browsing just a few of these sites demonstrates the importance of good visualizations. If you don't have a marketing budget to expend on these kinds of services, check out **QuarkRank** (`www.quarkrank.com`), an engine which crawls the web for opinions on various products and services and automatically summarizes them feature by feature,

[11] From `www.alphaworks.ibm.com/tech/emailsearch`; you have to register for an IBM ID if you don't already have one.

https://doi.org/10.1017/S1351324908004683 Published online by Cambridge University Press

and provides an API for retrieving this information; use of the engine is free for commercial as well as non-commercial purposes.

4 Search and question-answering

Another Web 2.0 theme that surfaces increasingly in language technology applications on the web is user-contribution. This is not surprising, really: technologies like natural language processing which find it so hard to be 100 percent correct are obvious candidates for occasional helpful nudges and clicks as the wise crowd of users moves past. **Digger** (`www.textdigger.com`) provides a relatively simple example of this.[12] This application aims to improve search results through sense disambiguation; amongst other things, the application shows the results of its disambiguation explicitly, so you can modify its decisions if it got things wrong, with the feedback no doubt being squirrelled away for future system tuning.

But the most impressive thing I've seen recently is the question-answering application being developed by TrueKnowledge.[13] Natural language questions are answered by reasoning over a knowledge base made up of known facts. I know that sounds very 1970s, but this really looks like it has a chance of working. What might make it feasible is that the application displays the steps in its reasoning, so you can tell it where it has gone wrong. You can also explicitly add new knowledge to the fact base. There's a very interesting demo on the website. This is definitely one to watch.

But Peter Norvig, Director of Research at Google, remains unconvinced that question-answering is the way to go for natural language processing in search.[14] If not QA on its own, then maybe in conjunction with ... social networking? Never one to sit still for long, **Hakia**[15] continues to explore the addition of interesting features around their QA system: now social networking is part of the mix. Their new 'Meet Others' feature lets you connect with other people who are searching for the same things you are. What's interesting here is Hakia's use of their underlying NLP capability to match your query with other similar queries; for example, if your question was "What drug treats a headache?", Hakia might match this with someone whose query was "'What medicine relieves migraines?". When you do a search on Hakia, a 'Meet Others' icon appears above the search result; clicking on that takes you to a chat room where you can converse with those doing similar searches.

5 Grammar checking

I really thought this was a moribund field in terms of new commercial developments,[16] but that appears not to be the case. In November, **WhiteSmoke**,[17] mentioned

[12] See the video at `http://www.demo.com/demonstrators/demo2007/91364.php`.
[13] `www.trueknowledge.com`
[14] See the interview at `www.technologyreview.com/Infotech/19868/`.
[15] `www.hakia.com`
[16] See the 'Industry Watch' column in Volume 10 Issue 1 of this journal.
[17] `www.whitesmoke.com`

https://doi.org/10.1017/S1351324908004683 Published online by Cambridge University Press

in an earlier column, released a new version of their writing assistant, WhiteSmoke 2008. According to the blog entry where I saw this mentioned:[18]

> Focused research and development efforts resulted in significant upgrades improving this revolutionary online text enhancement tool, which now takes the user's experience to a whole new dimension. The optimized features for better usability now dynamically pop-up when needed in a cleaner state-of-the-art interface.

But it's hard to find any details at all on WhiteSmoke's updated website: everything you click seems to lead to a purchasing page. Be wary of entering your email address anywhere here in your search for information; I made that mistake a few months back, and have since received a constant stream of sales emails trying to entice me with special offers on the WhiteSmoke software.

Another recent grammar checking application is CorrectEnglish[19] from Vantage. This application claims to use artificial intelligence and natural language processing to evaluate document organization, focus, development and style as you type. The application suggests improvements in various areas, including vocabulary, grammar, and document formatting. The company's website includes a comparison showing how, in a no-doubt carefully constructed document, MS Word's grammar checker finds only one error but CorrectEnglish catches a whopping 22.

You will have noticed that there's no social networking to be seen in this application area. Which just means there's a niche just waiting to be filled. 'Meet people who spell as badly as you do!' 'Share your grammar errors with others!'

By the time you read this it will be rather late for Christmas, but the last quarter of the year did throw up a number of notable cute speech-based applications that (in some cases at least) could serve as stocking fillers, so you might want to save the valuable (and surely timeless) information below for next year.

First, for US$29.95, you can buy the Intela Voice Dimmer.[20] Plug this into a power outlet, then plug in your Christmas tree lights and you're all set to control your greenhouse gas production by voice. You can train the device to respond to any phrases you choose; for example, you might create strobe effects by repeatedly reciting 'Save the planet! Kill the planet!' at a fast pace.

Or how about a voice-recognition-enabled safe? This item allows you to set your own password which, when spoken, opens the lock. Originally listed on the UK Science Museum site at £19.99 (US$34.67), Argos will sell you one at half price.[21] Comes complete with sound effects when the door is opened, a money bank slot and built-in LCD clock; requires 3 AA batteries. And in case you were worried about the risk of your money being locked away forever by virtue of speech recognition errors, you can also open it by typing in a number on the device's keypad.

[18] http://write-better.blogspot.com/2007/11/whitesmoke-releases-its-new-2008.html
[19] www.correctenglish.com
[20] http://trends.move.com/voice-recognition-gadgets-make-life-easier
[21] www.househacker.com/permanent/Voice-Recognition-Safe

https://doi.org/10.1017/S1351324908004683 Published online by Cambridge University Press

My favourite, though, for only £17.95, is the Nobby Voice Recognition Clock. Nobby is a blue, long-necked and single-eyed alien-looking object.[22] He responds to voice commands like 'How are you feeling?', 'Sing me a song', 'Make a decision', 'How do I look?', 'Are you hungry?' and 'Go to sleep'. But most usefully, before you go to bed, you tell Nobby what time you want to get up, and in the morning he will start beeping and chatting at you. You can either say 'Alarm off', or grab him by the neck and shake him until he shuts up.

By this point you'll be thinking of all kinds of devices you could embed speech in. To help you, Sensory has announced an update to its FluentChip firmware that includes what it calls TimeSet technology. This allows any RSC-4x[23] enabled device with a clock to have its time set using natural spoken phrases. Blinking clock radios, video recorders and microwaves are set to be a thing of the past; you just *say* the time to any device that has one of these chips in it, and you're done. A video demonstration can be seen on YouTube.[24]

All the above are relatively cheap items, as befits stocking fillers. If you'd rather go for something a bit more upmarket, you might put yourself on the waiting list for the Lady Bird.[25] This is a 1-meter tall, 1.35-meter-long prototype robot whose principal purpose in life is to clean public toilets at highway rest areas. It comes complete with water tanks, brushes and all the other tools needed for heavy-duty scrubbing. It has obstacle detection sensors so that it can safely perform its duties without running into people (or trying to clean them, presumably). But most importantly, in addition to cleaning, Lady Bird can engage in simple conversations with restroom users, thanks to its on-board speech recognition and synthesis capabilities. The robot has access to the latest information about traffic conditions on nearby roads, which it will relay to anyone comfortable enough to ask. When ready in March 2009, Lady Bird robots are expected to sell for about 3.5 million yen ($30,000) each. I expect the Lady Bird will soon have a FaceBook page and many friends.

[22] See `www.gizoo.co.uk/Products/HomeGarden/Home/NobbyClock.htm` for a photograph.
[23] Google is your friend.
[24] `http://www.youtube.com/watch?v=GtEBO34iQfY`
[25] `http://www.nextnature.net/?p=1738`

https://doi.org/10.1017/S1351324908004683 Published online by Cambridge University Press