Institution: University of Aberdeen

Unit of Assessment: 11 (Computer Science and Informatics)

Title of case study: Promoting the Public Understanding of Artificial Intelligence

1. Summary of the impact

Computational Linguistics research at Aberdeen concentrates on Natural Language Generation (NLG), an area of Artificial Intelligence. NLG raises difficult issues about what makes a text effective. This case study explains how our Computing Science department has brought these issues to the general public, aiming to improve public understanding of, and enthusiasm for, Artificial Intelligence. This has been achieved through two projects, targeting different sections of the public. *The Joking Computer* project explains the mechanisms behind verbal *humour,* allowing non-scientists to engage with computer programs that generate puns. The online Joking Computer has received 200,000 hits, and allows users to contribute data to the research by rating its jokes. The second project targets a more sophisticated section of the public in a book entitled *Not Exactly: in Praise of Vagueness*. This book, which builds on years of research on Natural Language Generation, illuminates the role of *vagueness* in science and computing and has triggered worldwide debate on the role of precision in scientific, social, and political discourse, meeting with an enthusiastic reception. To date it has sold well over 6,000 copies.

2. Underpinning research

Natural Language Generation (NLG) is a well-represented research topic at Aberdeen, combining a mix of applied and theoretical work. This impact case study focuses on humour in generated text, and on the role of vagueness in communication.

a. **Computational Humour**. Research into humour is important both scientifically and from an engineering perspective. The cognitive processes involved in humour are not well understood, and computer modelling is a useful tool in exploring these. From a practical viewpoint, as computer programs become more 'intelligent' and human-like, it is desirable for software to detect human use of humour, and to create humour when appropriate within an interaction. This has led to a growing interest in computational models of humour [R3]. Much of this work focusses on verbal humour, in which humorous effects are created by algorithms which search through specially designed lexicons looking for words with particular characteristics, such as phonetic similarity or semantic relatedness. Ritchie (Senior Research Fellow, retired Sept. 2011) oversaw a sequence of pioneering projects, involving computer generation of punning riddles (such as *"What kind of tree is nauseated? A sick-amore").* The EPSRC-funded *STANDUP* project constructed the first large-scale, properly engineered, generally usable joke generator [R1]. STANDUP was originally awarded to Ritchie while at Edinburgh Informatics (GR/S15402/01); the project commenced on 01/10/2003. Ritchie moved to Aberdeen 10 months later (01/08/2004), and continued to lead the computational humour research in this project for the remaining 32 months of the project, following up with the EPSRC PPE project *The Joking Computer* (EP/G020280/1). Its interactive software helped children with communication disabilities (often a consequence of cerebral palsy) to explore language in an engaging way normally not available to them. The software was tested with users in a school for children with special needs, yielding positive results: children had no difficulty in using the software and gave it good ratings (as did teachers and parents), and teachers reported improvements in the children's demeanour and use of communication devices [R2].

b. **Vagueness**. Despite a venerable tradition of research by logicians, game theorists and linguists, the reasons why words often lack sharp definition (such words are called "vague") are still poorly understood, as are the logical and computational implications of vagueness. Research in Aberdeen over the last decade, by Van Deemter (Professor), Reiter (Professor) and Sripada (SL), has addressed these issues from a computational viewpoint. For example, we investigated how the preferred level of precision of an expression is determined by the data that give rise to the expression, by the context in which the expression is made, and by the audience for whom it is intended [R6]. Research in Aberdeen has increased our insight into the choices that are made by human speakers and/or preferred by hearers. This research has informed computer programs (e.g. VAGUE, described in [R5]; and SUMTIME-MOUSAM, in [R4]) that express quantitative information in a variety of ways, producing texts that resemble human-generated expressions. For example, instead of saying *"The expected temperature is between 12.6 and 13.1 degrees",* the

computer might generate *"It will be mild for the time of year"*. This work was brought to the wider public by the book *Not Exactly: in Praise of Vagueness.*

3. References to the research

a. Computational Humour. Published work underlying the STANDUP computer program:

[R1] R.Manurung, G. Ritchie, H. Pain, A. Waller, D. O'Mara, R. Black (2008). The construction of a pun generator for language skills development. *Applied Artificial Intelligence*, **22** (9): 841-869. http://dx.doi.org/10.1080/08839510802295962

** Describes the motivation, design and implementation of the STANDUP interactive joke generator. *Paper [R1] best indicates the quality of underpinning research.*

[R2] A.Waller, R.Black, D.A.O'Mara, H.Pain, G.Ritchie, R.Manurung (2009). Evaluating the STANDUP Pun Generating Software with Children with Cerebral Palsy. *ACM Transactions on Accessible Computing* (TACCESS) **1** (3), Art.16. http://dx.doi.org/10.1145/1497302.1497306

** On the evaluation of STANDUP in a special-needs school.

[R3] G.Ritchie (2009). Can computers create humor? *AI Magazine* **30** (3): 71-81. http://www.aaai.org/ojs/index.php/aimagazine/article/download/2251/2104

** In this paper the state of the art and the research issues within computational humour are discussed, including the STANDUP project.

b. Vagueness. The book is: Van Deemter (2010) *Not Exactly: In Praise of Vagueness*. Oxford University Press, Oxford, UK. *The book has a highly accessible web site, which contains reviews and background:* http://homepages.abdn.ac.uk/k.vdeemter/pages/NotExactly*.*

[R4] E.Reiter, S.Sripada, J.Hunter, J.Yu, and I.Davy (2005). Choosing Words in Computer-Generated Weather Forecasts. *Artificial Intelligence* **167**:137-169. http://dx.doi.org/10.1016/j.artint.2005.06.006

** This article discusses the borderlines of words like "evening", as used by human weather forecasters and NLG computer programs.

[R5] K.van Deemter (2006). Generating Referring Expressions that involve Gradable Properties. *Computational Linguistics* **32** (2): 195-222. http://dx.doi.org/10.1162/coli.2006.32.2.195

** This article details a computer program that uses gradable adjectives appropriately, and the research underlying this program. *Paper [R5] best indicates the quality of underpinning research.*

[R6] K.van Deemter (2009). Utility and Language Generation: the Case of Vagueness. *Journal of Philosophical Logic* **38** (6): 607-632. [vanDeemter3 in the REF2 for this unit.]

** This article asks why language generators should sometimes produce vague expressions, discussing previous answers to this question and proposing a new one. *Paper [R6] best indicates the quality of underpinning research.*

4. Details of the impact

a. Computational Humour. Non-academic impact in this area was achieved through an EPSRC-funded Partnerships for Public Engagement project, *The Joking Computer: An interactive language playground* [EP/G020280/1]. Led by Masthoff and Ritchie, this project developed a user-friendly, education-oriented implementation of computational humour algorithms (previously developed by Ritchie and colleagues), and closely related educational materials. The software allows users to compose their own jokes and to rate the jokes produced by the computer. It was installed, using interactive touch screen kiosks, in Glasgow Science Centre (Dec 2009 – August 2010), the Satrosphere Aberdeen (March 2010 – present), and Dundee Science Centre (May 2012 – present), where it took part in the 2012 headline exhibition "Robot: the Fantasy and the Reality" [S1]. The session logs kept by the software show that these exhibits have been used by over ten thousand members of the public. Workshops for schoolchildren were run at various festivals (Word, May 2010; Techfest, September 2010; Science in the City, October 2010), with more than 250 participants in total. Feedback was collected at the workshops, in the science centres, in

evaluation sessions in a primary school, and online. Teachers were very enthusiastic (e.g. "excellent teaching materials", "covered many literacy and ICT skills", "pupils thoroughly enjoyed it") [S2]. The data gathered during the public's interactions with the software will allow for more detailed analysis of the Joking Computer and the public understanding of humour. A public talk was also delivered to the general public (May 2010, 120 attendees) as part of the University of Aberdeen's Public Engagement programme. Materials created for these workshops have been made available to primary schools across Scotland. A child-friendly website has been created, with educational information and games, including a full-scale interactive version of the project's main software. This site has had many visitors; the online database shows that the interactive software has received nearly 200,000 hits. Users have been asked to supply feedback via the website. This feedback has indicated that they had gained insight in the mechanism behind jokes, and that their experiences with the Joking Computer had taught them what Artificial Intelligence can be used for. These activities have attracted much media attention, resulting in articles in the press, radio interviews (e.g., BBC Radio Scotland, BBC5 Radio, NorthSound2), and a TV appearance (STV), which have tended to applaud the beneficial impact on children with communication difficulties in particular. The media attention in this technology has reach across technical [S3], main stream [S4] and international [S5] media, and has sparked interest in using the technology commercially. Debenhams, for example, issued a press release noting their interest and engagement with the research team in using the software to generate puns for Christmas Crackers [S4].

Computational humour has proven to be a remarkably effective topic for explaining the challenges facing Artificial Intelligence, and NLG specifically. People with little background in computing often find it difficult to understand how a computer program can tackle new problems. The Joking Computer allows users to play with algorithms, to discover how a computer can itself invent a surprising new joke.

b**. Vagueness.** The book *Not Exactly: in Praise of Vagueness* presents the University's work on vagueness (e.g., [R5] & [R6]) to non-specialists. The book launch, organised by the KIM Center for Contemporary Arts (Riga) in collaboration with the University of Latvia's Centre for Cognitive Sciences and the Stockholm School of Economics, drew over 100 visitors. Other invited public presentations include the literary *Word* Festival (Aberdeen, May 2010), the *Science & Society* Forum at the European Molecular Biology Laboratory in Heidelberg (Germany, Dec 2010), and a public lecture at Sun Yat-Sen University, Guangzhou (China, Dec 2010).

In June 2013, a bit over a year after appearance of the paperback edition, the English–language version had sold 4,311 copies [S6]. Global interest in the text is substantial and growing: A translation into Complex Chinese, published by the China Times Publishing Company, Taiwan, has sold 1,950 copies (April 2013) and a licence for translation into Simplified Chinese has been signed in April 2013 with Beijing Time-Chinese Publishing House (covering 5,000 copies). An Arabic translation by the National Center for Translation, Cairo, is to appear in Winter 2013/14. Additionally, a pirate version of the English edition has appeared in China at Taobao (sales figures unknown).

Reviews in publications such as *Nature* [S7]; *Computational Linguistics*; *Minds and Machines*; and *Laval Theologique et Philosophique*, and in non-academic forums such as *The New Scientist*, *Wall Street Journal*, *The Sunday Times*, *Science News*, and *Guardian*, are in agreement that the book has contributed to civic debate on issues surrounding arbitrary borderlines and false precision, and to the public understanding of Artificial Intelligence. The *Times Higher Education* wrote:

> "the concept of vagueness is set to be of huge significance in the next few years (…). I suspect that the products and services that prove ultimately successful will be the ones that have fully accommodated vagueness as discussed here by van Deemter - thereby delivering a truly human-friendly interface between the virtual and the real world. The style throughout is amusing, persuasive, conversational and engaging, but this does not detract from the thoroughness with which van Deemter approaches his thesis." [S10]

There were prominent reviews in *Corriere della Sera* (one of Italy's main broadsheets), *Rigas Laiks* (Latvia's paramount culture magazine) and in Israel (*Globes* and *Calcalist*). *New Scientist* wrote:

> "the world turns out to be unexpectedly vague (see 'Kees van Deemter: The importance of being vague'). It's not just that things we think of as well-defined are actually a series of

> approximations, like the metre. More profoundly, vagueness is a key part of communication: unless we get to grips with it, robots will never 'talk' naturally to people (…). Just as well, then, that we have finally started to think precisely about vagueness." [S8]

Legal and political implications were highlighted repeatedly, for example in the *Wall Street Journal*:

> "The vague boundaries of political terms (…) are what allows for the differences in opinion that lend democracy its vibrancy. Vagueness, in Mr. van Deemter's view, is language's gift to civic culture." [S9]

The book is not without its own academic impact (with 40 citations under Google Scholar) and has achieved impact in pedagogic fields, being used in courses at the Ecole Normale Superieure in Paris (Philosophy, P. Egré); Northwestern University in the USA (Philosophy, P. Ludlow); Melbourne in Australia (Philosophy, G.Restall); and Leiden in The Netherlands (Archaeology, R. Corbey).

Jointly the two parts of this Case Study – The *Joking Computer* public engagement project and the popular science book *Not Exactly: in Praise of Vagueness* – are engaging significant sections of the public with Artificial Intelligence. In addition, they are giving a non-scientific audience an insight into the basics of Computing Science (e.g., What is an algorithm?) and are informing civic debate about the role of language in culture and society.

5. Sources to corroborate the impact

a. **Computational humour**

[S1] IT Manager, Dundee Science Centre – will corroborate the Joking Computer exhibit at the Science Centre and the numbers of visitors engaging with the exhibit.

[S2] A teacher, Aberdeen City Council – can corroborate the impact on school pupils' understanding of technology.

[S3] http://www.theengineer.co.uk/news/exhibit-features-joking-computer/1000320.article

[S4] The Debenhams press release is no longer available, but this is an example of an article reporting on it along with the core focus of the technology and impact: http://www.dailymail.co.uk/news/article-1337306/Heard-the-witty-It-creates-corny-Christmas-cracker-jokes.html

[S5] http://news.oneindia.in/2010/12/10/joking-computer-to-give-stand-up-comedians-run-for-their.html

b. **Vagueness**

[S6] Head of translations, reprints and permissions, Oxford University Press – will corroborate sales figures and translations for *Not Exactly*.

[S7] A review of *Not Exactly* in the journal *Nature*, focussing on arbitrary borderlines and false precision: http://homepages.abdn.ac.uk/k.vdeemter/pages/Nature-review-NE.pdf

[S8] Interview in *New Scientist*, 20-03-12, http://homepages.abdn.ac.uk/k.vdeemter/pages/NE-NewScientist.pdf Also the editorial: http://www.newscientist.com/article/mg20527512.700-we-need-to-think-precisely-about-vagueness.html

[S9] A review of Not Exactly in the *Wall Street Journal*, focussing on political implications: http://online.wsj.com/article/SB10001424052748703808904575025382998649088.html

[S10] Review of *Not Exactly* in the *Times Higher Education* (Book of the Week): http://www.timeshighereducation.co.uk/410239.article