

A CONVERSATIONAL EXPERT SYSTEM SUPPORTING BULLYING AND HARASSMENT POLICIES

Annabel Latham, Keeley Crockett, Zuhair Bandar

*Intelligent Systems Group, The Manchester Metropolitan University, Chester Street, Manchester M1 5GD, UK
a.latham@mmu.ac.uk, k.crockett@mmu.ac.uk, z.bandar@mmu.ac.uk*

Keywords: Expert system, conversational agent, knowledge engineering, knowledge tree, bullying, harassment.

Abstract: In the UK, several laws and regulations exist to protect employees from harassment. Organisations operating in the UK create comprehensive and carefully-worded bullying and harassment policies and procedures to cover the key aspects of each regulation. In large organisations, such policies often result in a high support cost, including specialist training for the management team and human resources (HR) advisors. This paper presents a novel conversational expert system which supports bullying and harassment policies in large organisations. Information about the bullying and harassment policies and their application within organisations was acquired using knowledge engineering techniques. A knowledge tree is used to represent the knowledge intuitively, and a dynamic graphical user interface (GUI) is proposed to enable the knowledge to be traversed graphically. Adam is a conversational agent allowing users to type in questions in natural language at any point and receive a simple and direct answer. An independent evaluation of the system has given promising results.

1 INTRODUCTION

Organisations in the UK have policies to enforce laws which protect employees from bullying and harassment. Designed to support several laws and regulations, such policies and procedures must be comprehensive and carefully-worded. These policies often result in a high support cost, requiring specialist training across the organisation. Victims of bullying and harassment may be reluctant to discuss such sensitive issues within the organisation, fearing lack of confidentiality, bias by support staff and repercussions. The cost to organisations of not dealing with bullying can be high – bullying at work is estimated to be costing the UK National Health Service more than £300m in sick days and recruitment costs (BBC.co.uk 2008). The rather inert nature and high support cost of bullying and harassment policies means they lend themselves to automation. An expert system would offer anonymous, 24-hour access to advice which was consistent, appropriate and valid (avoiding human bias), and would be generic, offering cross-sector application across the UK.

An ability to enter into a natural language dialogue with an expert system would be the

ultimate development in a user interface, giving users direct and non-linear access to expert knowledge. This is especially the case in heavily legislated areas where the language used may be difficult to understand. Human computer interface designers have used the metaphor of face-to-face conversation for some time, but have only recently attempted “*to design a computer that could hold up its end of the conversation*” (Cassell 2000). A conversational agent (CA) is a computer program which can interact using natural language (Cassell 2000). There are three main development approaches for CAs – natural language processing (NLP), pattern matching or artificial intelligence (AI). NLP studies the constructs and meaning of natural language, applying rules to process information contained within sentences (Khoury *et al* 2008). Pattern matching uses an algorithm to find the best match for key words and phrases within an utterance (Pudner *et al* 2007). The AI approach compares the *semantic similarity* of phrases to decide on the meaning of the input (O’Shea *et al* 2008; Li *et al* 2006), enabling a CA to cope with input which is not grammatically correct or complete. This paper presents a novel conversational expert system (CES) which models bullying and harassment policies and procedures in large

organisations. Knowledge engineering techniques are used to acquire information about the policies and how they are applied. A dynamic GUI is proposed for the CES, which changes according to the user's need, enabling the knowledge to be traversed graphically and which allows natural language questions to be asked at any point.

This paper is organised as follows: Section 2 outlines bullying and harassment, Section 3 describes CAs and their applications, Section 4 explains the methodology for constructing the CES, Section 5 summarises the evaluation of the CES and Sections 6 and 7 include the results, discussions and conclusions.

2 BULLYING & HARASSMENT

Bullying and harassment is widespread within the workplace (Burnes & Pope 2007, Lee 2000). In 2005, a national survey of the UK National Health Service (NHS) staff found that 15% of staff had experienced bullying, harassment or abuse from colleagues and 26% experienced these problems from patients (Healthcare Commission 2006). In addition to the health implications for victims, workplace bullying and harassment costs organisations lost productivity, sick leave, high staff turnover and the cost of replacing staff, as well as the possibility of legal actions by employees (Unite the Union 2007).

In the UK, no single piece of legislation covers bullying and harassment. Organisations must adhere to several different laws protecting employees from harassment due to diverse causes (ACAS 2007). Implementing this legislation has led organisations to develop comprehensive policies, often resulting in large and complex policy documents. These policies often require additional training and support for employees understanding and following the process of reporting bullying and harassment. The high support cost and fairly static nature of these policies suggests that automation would be advantageous. A CES would reduce costs and add benefits such as allowing anonymous, 24 hour access to information ensuring that the advice given is consistent, appropriate and valid.

3 CONVERSATIONAL AGENTS

Communicating with a computer using natural language has been a goal in artificial intelligence for

many decades, stimulated by the 'Turing Test' (Turing 1950). Conversational agents designed with the sole aim of holding a conversation are known as *chatbots* (Carpenter 2007). In the context of this paper, CAs are tools for addressing specific problems, such as ConvAgent's InfoChat, a goal-oriented CA (Michie & Sammut 2001).

CAs are ideally suited to simple question-answering systems (Sadek 1999, Convagent 2005, Owda *et al* 2007) as they are intuitive to use and allow users direct, non-linear access to information. A natural language questioning facility in an expert system would additionally widen access to the expertise. The role of a CA in the context of this paper is to accept natural language questions and produce an appropriate response. A pattern matching CA is suited to systems where conversations are restricted to specific knowledge areas, such as bullying and harassment policies. The pattern matching approach requires the development of conversation scripts, a similar idea to call centre scripts, which match key input words and phrases to suitable responses. CA scripts are often organised into contexts and may be linked in a tree or network structure, offering powerful pattern matching at various levels (e.g. a script level to respond to abusive language). The CA receives an input and searches scripts to find the best matching response. Different contexts and conversation histories are used to help find appropriate matches, for example, the meaning of a user utterance "Yes, please show me" can only be understood in relation to the current context and previous utterances of the conversation.

4 CONSTRUCTING THE CES

This section describes the methodology and key stages in constructing the CES. Figure 1 shows the overall structure of the CES. A *central controller* communicates with the knowledge base, dynamic GUI and CA to manage user interactions. The *knowledge base* contains expert knowledge and was constructed in two phases. First expert knowledge was gathered (knowledge engineering, described in Section 4.1) and then that knowledge was structured (explained in section 4.2). The CES user interface has two main components – the *dynamic GUI* displays a graphical representation of the knowledge which changes according to user choices (outlined in section 4.3), and the *CA* allows natural language questions to be asked (described in section 4.4).




Figure 1: CES system structure.

4.1 Knowledge Engineering

Knowledge engineering is the stage of expert system development concerned with the acquisition and elicitation of knowledge from experts, and was done in parallel with requirements gathering (Tuthill 1990). A domain review of current UK harassment legislation, policies and guidance was undertaken and a real-life large organisation's Bullying and Harassment policy was adopted. The documentation study was followed up by meetings with the organisation's management team to elicit further knowledge and clarification on points arising from the review. An analysis of knowledge acquired about the policy and its use by employees identified four key categories of user problem, forming the basis of the knowledge represented by the CES. Each category leads to a series of additional questions to further identify the problem. The four main questions to be addressed by the CES are:

- What counts as bullying and harassment?
- How do I report bullying or harassment?
- What can I do if I am not being taken seriously?
- What can I do if I am being victimised?

4.2 Structuring Knowledge

The knowledge base underpins a CES and therefore must accurately reflect the expert knowledge and its use in practice. A knowledge tree was the most suitable representation of the knowledge for this system as it follows a path of questions and feels natural to users. Knowledge trees are biological representations which consist of *nodes*, where values are assessed, *branches* to be followed according to node values, and *leaves* which terminate the branch and represent a conclusion. Information is gathered by the CES based on user choices, which determine the value of each node, and therefore the branch to follow and the next question to ask (the next node).

The knowledge tree was designed iteratively and modelled using Convagent's InfoRule TreeTool (Convagent 2005) which is an interactive system that graphically displays the tree. The TreeTool facility to simulate the running of the tree was used so that experts could verify its accuracy. Figure 2 shows a small section of the knowledge tree within the TreeTool. The entire tree has more than 1000 nodes over 11 levels, which illustrates the complexity of the problem. Within the Decision Tree folder, the main tree is called 'bullying', and the start node of the tree is called *att_query_type*. The start node has four branches, one for each of the main questions modelled. The *not_taken_seriously* branch leads to a leaf node, a CES response which advises the user to discuss their case with a specialist advisor. Also seen in Figure 2 is a subtree, *follow_up*, which is used to avoid repetition of nodes within the tree (similar to a macro).




Figure 2: CES knowledge tree.

4.3 Graphical User Interface

The CES users will range in experience from novice to expert so the user interface must enable two-way communication for all users (McGraw 1992); if users cannot use the CES as an effective consultant it will not be successful. Graphical interfaces are more intuitive to use for novice users. The CES is destined to be run on various web browsers over the Internet or within company intranets. The two main user interface elements will now be described.

4.3.1 Dynamic GUI

The GUI allows users to view a representation of the knowledge and select appropriate options. Depending on the user's position in the knowledge tree, the graphical representation changes to reflect the available options. This dynamic interaction was

intended to feel natural like a conversation rather than a linear procedure.




Figure 3: Dynamic user interaction.




Figure 4: CES response.

The screen was divided horizontally into three main areas (see Figure 3): the top area displays Adam’s instructions and advice, the centre shows a graphical representation of the knowledge tree and the bottom area allows natural language questions to be entered. Adam asks questions to gather information about the user’s query, and the user responds by clicking one of the highlighted answer boxes. The CES dynamically accesses the knowledge base and responds with a new question and set of answer boxes. Figure 3 shows an interaction where the user is exploring how to report bullying, having taken no action so far (previous answers are shown graphically using red triangles and lines). The user wants to take informal action, so clicks the ‘Informal’ answer box (Figure 4), and is given guidance by Adam, shown graphically using an ellipse. At any point the user can restart the session or go back and explore the knowledge further, using buttons at the foot of the screen. The dynamic GUI for the CES was designed with system reuse in mind as its display draws its information directly from the knowledge base. Amendments made to the knowledge tree are automatically shown

graphically on screen. This offers easy system maintenance and the option of applying new knowledge bases, e.g. finance policies.

4.3.2 Conversational Agent

The CA component of the CES integrates a natural language question-answering function. The ability to ask questions widens access to the system by allowing users to employ their own vocabulary. The CA also offers users direct access to knowledge without tracing through the CES questions.

Convagent’s InfoChat CA (Michie & Sammut 2001), a goal-oriented CA employing natural language pattern matching, was chosen. A list of frequently asked questions (FAQs) about bullying and harassment was compiled during the domain review. A number of scripts were developed which contained rules with patterns of keywords and phrases to match each FAQ to an appropriate response. A lower level of script was designed to respond to unknown or abusive language. An example rule from one of the InfoChat scripts is shown below. In the rule, *a* is the activation level used for conflict resolution (Sammut 2001); *p* is the pattern strength followed by the pattern and *r* is the response. Also seen in the example is the wildcard (*) and macros (<explain-0>) containing a number of standard patterns, each matched separately.

<Rule-01>

a:0.01

p:50 *<explain-0> *bullying*

p:50 *bullying* <explain-0>*

p:50 *<explain-0> *a bully*

p:50 *a bully* <explain-0>*

r: Bullying is persistent, threatening, abusive, malicious, intimidating or insulting behaviour, directed against an individual or series of individuals, or a group of people.




Figure 5: Asking a question.

At any stage within the CES, the user could type natural dialogue into an on-screen box to “chat” directly with Adam. Figure 5 shows an example interaction with Adam, continuing the previous interaction (Figures 3 and 4) by clarifying Adam’s response. The user types their question, ‘Which senior manager should I meet’, into the question box and presses ‘Ask Adam’. The CA has been scripted to understand this question, and Adam provides an appropriate answer at the top of the screen.

5 USABILITY EVALUATION

The knowledge tree was evaluated by the organisation’s management team, who simulated its execution using the TreeTool. The CES was evaluated by the HR policy experts within the organisation and an independent graphic designer specialising in websites. During the evaluation, a questionnaire was used to record feedback on system usability, appearance and functionality. Additionally a representative group of users were selected to evaluate the CES, all employees in different roles for different companies, aged between 25 and 45 and experienced in interacting with HR departments. The group included both men and women from administrative, academic and managerial backgrounds. A scenario of possible sexual harassment was developed along with a questionnaire to record feedback on the appearance, usability and functionality of the CES. The group were asked to read the scenario and then complete two tasks within the system – find out whether the behaviour described is sexual harassment and how to stop the behaviour without jeopardising your job. Users then completed the questionnaire anonymously and the results were analysed. Some of the users were unobtrusively observed during evaluation, giving qualitative feedback about facial expressions and speed of use. The results of the evaluation will now be discussed.

6 RESULTS AND DISCUSSION

Results of the independent evaluation show that the CES was generally well received, being understandable and intuitive to use. Table 1 shows the average results for each section of the questionnaire. The lowest score (of 5) was given to the typeface used within the GUI, although the average score was 7.8. The highest average score of

9.33 indicated that users had found the information they sought without difficulty (Table 1, Section 2, Question 6). One user commented that they “*did not need to ‘learn’ to use the advisor*”. The next highest average score of 9 was given to the clarity of instructions, ease of navigation and use of Restart and End buttons. 50% of users suggested differentiating Adam’s instructions from the advice given would improve the usability of the system, for example using a different coloured font to make it clear when users are expected to respond.

Table 1: User evaluation results.

Section 1 – Screen Design	Avg Score (out of 10)
1. Look and feel	8.5
a. Colour	8.17
b. Size	8.5
c. Typeface	7.83
2. Placing of buttons and boxes	8.5
3. Clarity of instructions	9
4. Central option buttons	
a. Size of buttons (width)	8.67
b. Positioning of buttons	8.67
c. Connecting lines (width)	7.83
d. Indication of current selection	8
Section 2 – Functionality	
1. Adam’s text for instructions	
a. Readable?	8.5
b. Noticeable?	8
2. Ease of navigation	9
3. Restart & End buttons	9
4. Clicking of option buttons	8.83
5. Typing of text and Ask Adam button	8.33
6. Did you find the information and advice you wanted?	9.33
7. How useful did you find the advice given to you?	7.83

Observation of the tests revealed two interesting issues. 33% of users attempted to click the response icon, and were expecting to be taken to an email or a form to arrange to see an advisor. Also, 33% of users pressed the browser’s Back button, which caused the tree buttons to behave unpredictably (causing an error in one case). Ease of use is critical to the success of a CES to support a Bullying and Harassment policy. Users felt that the system was intuitive and easy to use, giving high scores to ease of navigation and the ability to find information. The results have shown that the system allows users to access information about sensitive issues like bullying anonymously and at any time, aiding organisations in positively implementing harassment policies and improving workforce cohesion.

7 CONCLUSIONS AND FURTHER WORK

This paper has presented a novel conversational expert system which has been applied to a real world application to advise employees on bullying and harassment policies. The CES features a graphical representation of knowledge which may be explored dynamically by the user, and a facility to ask natural language questions and receive a response. During development, the CES was generalised, meaning it can adapt to a different, unknown knowledge base. As well as improving the ease of maintenance, similar applications can be developed more rapidly by reusing the GUI component and limiting new development to the knowledge base and conversational scripts. The system was evaluated by a representative group of users, who found the system intuitive to use and the information easy to access. Implementation of the system could benefit both employees and the organisation in providing anonymous access to information about the bullying and harassment policy at any time of day, leading to a positive implementation of the policy and a more cohesive group of employees.

In future a pilot evaluation study of the CES is required before conclusions may be drawn about the usability of the system compared to telephone advice or meeting advisors. It is then planned to develop new CES for different HR procedures. Future improvements include expanding the CA scripting to include further definitions and explanations, extending the system for multiple users, and investigating the use of different CAs which employ less labour-intensive techniques for developing dialogue scripts.

ACKNOWLEDGEMENTS

The authors thank Convagent Ltd for the use of their InfoChat and TreeTool software and Mr John Shiel of Krann Ltd for his evaluation of the GUI.

REFERENCES

ACAS 2007, *Advice leaflet - Bullying and harassment at work: guidance for employees*, ACAS, viewed 24 May 2007, <http://www.acas.org.uk/index.aspx?articleid=797>
BBC.co.uk 2008, 'NHS Bullying costing millions', BBC, viewed 10 March 2008, <http://search.bbc.co.uk/cgi->

[bin/search/results.pl?scope=all&tab=all&recipe=all&q=Bullying+NHS+2008&x=0&y=0](http://search.bbc.co.uk/cgi-bin/search/results.pl?scope=all&tab=all&recipe=all&q=Bullying+NHS+2008&x=0&y=0)
Burnes, B & Pope, R 2007, 'Negative behaviours in the workplace: A study of two Primary Care Trusts in the NHS', *International Journal of Public Sector Management*, vol. 20 no. 4, pp. 285-303.
Carpenter, R 2007, *jabberwacky – live chat bot*, viewed 20 June 2007, <http://www.jabberwacky.com/>
Cassell, J 2000, *Embodied Conversational Agents*, MIT Press, Cambridge MA.
Convagent Ltd 2005, *Convagent*, viewed: 23 May 2007, <http://www.convagent.com/demo.htm>
Healthcare Commission 2006, *National survey of NHS staff 2005 Summary of key findings*, Commission for Healthcare Audit and Inspection, London.
Khoury, R, Karray, F & Kamel, MS 2008, 'Keyword extraction rules based on a part-of-speech hierarchy', *International Journal of Advanced Media and Communication*, vol. 2 no. 2, pp. 138-153.
Lee, D 2000, 'An analysis of workplace bullying in the UK', *Personnel Review*, vol. 29 no. 5, pp. 593-612.
Li, Y, McLean, D, Bandar, Z, O'Shea, J & Crockett, K 2006, 'Sentence Similarity Based on Semantic Nets and Corpus Statistics', *IEEE Transactions on Knowledge and Data Engineering*, vol. 18, no. 8, pp. 1138-1150.
Michie, D & Sammut, C 2001, *Infochat Scriptor's Manual*, Convagent Ltd, Manchester, UK.
O'Shea, K, Bandar, Z & Crockett, K 2008, 'A Novel Approach for Constructing Conversational Agents using Sentence Similarity Measures', *World Congress on Engineering, International Conference on Data Mining and Knowledge Engineering*, pp. 321-326.
Owda, M, Bandar, Z & Crockett, K 2007, 'Conversation-Based Natural Language Interface to Relational Databases', *IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology*, pp. 363-367.
Pudner, K, Crockett, K & Bandar, Z 2007, 'An Intelligent Conversational Agent Approach to Extracting Queries from Natural Language', *World Congress on Engineering, International Conference of Data Mining and Knowledge Engineering*, pp 305-310.
Sadek, D 1999, 'Design Considerations on Dialogue Systems: from theory to technology – the case of Artemis', in *ESCA Workshop: Interactive Dialogue in Multi-modal Systems*, p. 173, Kloster, Irsee.
Sammut, C 2001, 'Managing Context in a Conversational Agent', *Electronic Transactions on Artificial Intelligence*, Vol. 3, No. 7, pp. 1-7.
Turing, A 1950, '[Computing machinery and intelligence](#)', *Mind*, vol. LIX, no. 236, p. 433.
Tuthill, GS 1990, *Knowledge Engineering: Concepts & Practices for Knowledge-Based Systems*, TAB Books Inc, Blue Ridge Summit, PA.
Unite the Union 2007, *Dignity at Work – tackling bullying in the workplace*, Unite the Union, viewed: 12 August 2008, <http://www.dignityatwork.org/>