EPSRC First Grant Scheme Research Proposal
Case for Support

Project full title:

Coherent Dialogue Automatically Generated From Text

Project Acronym:

CODA

Contents

2 Description of the Research and its Context

2.1 Introduction

Dialogue as a means of information presentation goes back at least as far as the ancient Greeks – Plato conveyed his philosophy through fictitious conversations between Socrates and his contemporaries. In this kind of dialogue, henceforth *expository dialogue*, the main purpose is to inform the reader or audience; the information the characters convey to one another is subservient to this purpose. The popularity of expository dialogue in today's world is evidenced by the widespread use of dialogue in news bulletins (between presenters), commercials, educational entertainment, and games.

[You can take aspirin]$_1$ [if you have a headache.]$_2$ [Though aspirin does have side effects:]$_3$ [it can harm circulation.]$_4$

L: What if I have a headache?
E: You can take aspirin
L: But does it have side effects?
E: Yes, it can harm circulation.

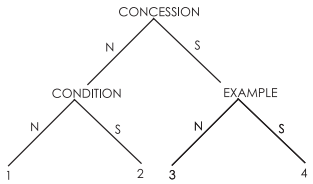

Figure 1: From left-to-right: monologue, corresponding dialogue, and the coherence relations underlying the monologue

Figure 1 shows how the information found in a Patient Information Leaflet, a monologue (on the left-hand side), can also be expressed as a dialogue between an expert (E), e.g., a doctor or pharmacist, and a layman (L) (middle column of the figure). According to a number of empirical studies, for educational and persuasive purposes, presentation in dialogue form is more effective than monologue; for example:

- Craig et al. [10] found that dialogue stimulates students to write more in a free recall test and ask twice as many deep-level reasoning questions in a subsequent tutor-guided task on a different topic,
- Lee et al. [19] report that there is more discussion between students and less banter,
- Cox et al. [9] established that student learning is at least as good as with monologue, and
- Suzuki and Yamada [36], comparing dialogue between two lifelike characters – a persuader and a persuadee agent in an online shopping scenario – with monologue, found that dialogue is more effective as a means for persuasion; in particular, dialogue increased the purchase likelihood when compared to monologue.

Despite the evidence that dialogue is more effective for education and persuasion than monologue, for the foreseeable future, *text* in monologue form is likely to remain the most common medium for the production of information content. It can be produced quickly and is used in leaflets, articles, books, web pages, etc. In order to nevertheless allow information to be presented as dialogue, **the aim of the current project** is *to develop the theory and technology for automatic transformation of text in monologue form to expository dialogue, specifically dialogues between a 'layman' (e.g., patient or student) and 'expert' (e.g., doctor or tutor) character.* The focus will be on:

- *Cohesion and coherence-preserving transformations* – a good text is not just a loose collection of sentences; rather, the sentences are connected: *cohesive* links [14] tie different parts, usually noun phrases, together through coreference relations, as in (*Harry . . . he . . . my friend . . .*), whereas *coherence* relations (e.g., [20]) connect larger units. For example, the tree in Figure 1 (right-hand side) shows the coherence relations (CONCESSION, CONDITION and EXAMPLE) underlying the text on left-hand side of Figure 1. Our text-to-dialogue transformations reorder and regroup information, but should not destroy the underlying cohesion and coherence relations.
- *Meaning-preserving transformations* – for many applications (healthcare, education, etc.) it is critical that transformations from one presentation format to another do not alter the content by adding, removing or changing information.
- *Robust, empirically motivated transformations* – on the input side, text, regardless of its topic and genre, should be transformable to dialogue, analogous to state-of-the-art speech synthesis systems, which can take almost any text and produce spoken language output. On the output side, the properties of the generated dialogue should be based on empirical studies of the properties of expository dialogue produced by expert human authors.

The proposed research will help realize the Government's Council for Science and Technology's (CST) aim of 'providing people with services and information when, where and how they need it' (quoted from the CST's 2007 report on 'Strategic decision making for technology policy'), specifically, 'how they need [the information]'. We have pointed out that there is empirical evidence that dialogue can be more effective than monologue for education and persuasion. Additionally, dialogue is extremely versatile: it can be presented as text – similar to frequently asked questions (FAQs) — but also, for

example, by a team of human-like computer-animated agents with automatic speech synthesis (see Figure 2). Information presented in this way can enhance the user's experience by making it more engaging and entertaining [11, 27]. Again, this is broadly in line with the CST's thinking: 'Interaction will be through next generation personal digital assistants, third and fourth generation mobile phones, set-top boxes, digital TV, clothing and doubtlessly a variety of other human-oriented methods as yet unforeseen.'

This Figure shows two animated characters, a client and a pharmacist, act out information contained in a patient information leaflet ('How should I take the tablet? You should first remove it from the foil and then swallow it with water'; see Piwek et al. [24]). The characters are from Dr. Helmut Prendinger's lab, who is proposed as visiting researcher for this project. Our collaboration with his team will ensure that the technologies resulting from the project will seamlessly interface with state-of-art computer-animated characters that can act out the expository dialogues.

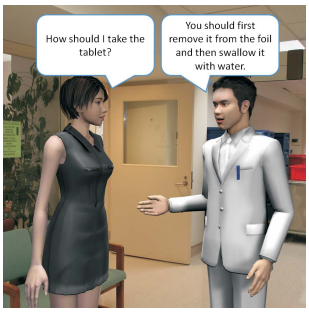

Figure 2: Presentation of expository dialogue with computer-animated agents

2.2 Prior work

In the field of Natural Language Generation (NLG), there has recently been a trend away from generation based on manually constructed inputs (usually in a knowledge representation language) to generation from widely available inputs. At least two complementary strands have emerged: data-to-text generation (D2T; e.g., [29]) and text-to-text generation (T2T/paraphrasing, e.g., [34, 8]). The current proposal is aimed at a specific kind of T2T: generation of a *text in dialogue form* from *a text in monologue form* (see Figure 1 for an example). Distinctive for the proposed work are the combination of: 1) generation from monologue, 2) meaning-preserving transformations to dialogue, 3) exploitation of cohesion and coherence relations in text that reach beyond sentence boundaries, and 4) robust and empirically motivated generation. All existing work on text-to-dialogue (T2D) differs in one or more respects from the current proposal.

Most research on text-to-text focuses on transforming a *single* sentence into another one (e.g., [34]). Within this type of research, from our dialogue generation perspective, work on mapping declarative sentences to questions is most relevant. In the US, Rus et al. [31] built a system that generates factual *who*, *what*, *where*, and *when*-questions from single sentences, whereas in the UK, Mitkov et al. [22] go one step further: they generate multiple choice items, i.e., a question regarding factual knowledge about a concept together with a correct answer and some 'distractor' answers. In Japan, several groups [23, 35] have developed multimedia presentation systems that generate entire dialogues from text. At the input side, these systems work, however, still with analyses of *individual* sentences (and, in particular, their information structure rather than underlying coherence relations, as we propose to do) that are mapped to dialogue fragments and then glued together. For example, Web2Talkshow [23] takes declarative sentences of the form *X of the Y did Z* and transforms them into dialogue fragments of the form *A: Who is X. B: I know. X is one of the Y. A: That's right! He did Z.* Moreover, Web2Talkshow aims to generate humorous dialogues by occassionally exaggerating the original content, whereas we aim at meaning-preserving transformations. Similar to the Web2Talkshow system, the systems developed in the Edinburgh-Stanford link CrAg project [15] generate entire dialogues as opposed to single questions or question-answer pairs. In CrAg, in contrast with our proposal, dialogue is generated from *existing dialogue* rather than monologue, with dialogue-to-dialogue transformations manipulating personality of dialogue participants.

The approach used in Intelligent Multimedia Presentation (IMMP) systems is complementary to the aforementioned T2T approaches. In IMMP systems, dialogue generation is automated by employing methods from artificial intelligence, knowledge representation, and planning (see André [1] for an overview). An IMMP system assumes a 'presentation goal' and uses planning methods to generate a sequence of presentation acts. The generation of a presentation is based on dedicated information sources that encode information about presentation content and objects [2]. The PI contributed to this line of research in the EU NECA project which integrated IMMP with techniques for full generation [38]. Such planning-based and knowledge intensive approaches have also been used to generate dialogue for *story telling* (e.g., [6, 21]), rather than information presentation. They are appropriate if the input consist of data or knowledge representations, and can result in dialogue presentations of high quality. A drawback is that they require domain-specific resources such as planning operators to produce the dialogue. This makes them labour intensive, since the resources typically have to be handcrafted by domain and knowledge representation experts. In contrast, the aim of the current proposal is develop a platform for domain-independent mappings from *text* to dialogue.

The current proposal derives from the PI's work [25] which describes how coherence relations that are implicit in text can be related to certain patterns in dialogue (e.g., the first question-answer pair in the dialogue of Fig. 1 expresses the CONDITION relation of the monologue). In Piwek et al. [25], the basic idea is introduced through a handcrafted example. We have only recently – in collaboration with Dr. Prendinger at the National Institute of Informatics (NII), Tokyo – carried out a pilot study into the feasibility of an *automated system*, with encouraging results – mappings involving the CONDITION relation achieved an accuracy $> 90\%$ [24, 28]. The proposed project will build on that preliminary work by extending and improving it along three dimensions:

1. *robustness/domain-independence*: The preliminary study focused on a single domain, i.e., Patient Information Leaflets, whereas the proposed research will aim to develop and evaluate transformations for different domains,
2. *empirical motivation*: like all work hitherto on monologue-to-dialogue generation, the preliminary study did not take into account systematically the techniques used by expert authors of expository dialogue to present information; in contrast, the proposed project will collect a *parallel* corpus of text–dialogue pairs, with dialogues authored by professional writers, and base the transformations on this corpus. We will investigate whether structured learning approaches for inducing and applying mappings in monologue-to-monologue generation (e.g., [8]) can be extended and applied to our monologue-to-dialogue task.
3. *scope and theoretical foundations of transformations*: The preliminary study only considered direct relations between adjacent spans of text, explored only a very limited subset of all possible text-to-dialogue transformations based on a small selection of coherence relations, and only evaluated the mappings for the CONDITION relation – in all these respects, the proposed project will go beyond the research reported in the preliminary study; it will deal with *nested relations*, the *full set of RST coherence relations*, and and *evaluation of all mappings*.

2.3 Objectives, Methodology and Research Questions

The overall aim of the proposed project is to *formulate a set of computable meaning-preserving transformations from text in monologue form to cohesive and coherent text in dialogue form*. This aim involves the following more specific objectives, each corresponding to a workpackage in the Diagrammatic Workplan: 1) constructing a parallel corpus of text-dialogue pairs, 2) formulating text-to-dialogue transformations based on the corpus, 3) implementing the transformations, and 4) evaluating the resulting system.

Objective 1 – Create a parallel text–dialogue corpus Work on dialogue generation from text in monologue form, including our own work, has so far relied on handcrafted transformation rules [31, 23, 35, 24], using the imagination and creativity of the system developers to come up with appropriate mappings. This means that much of the development of a text-to-dialogue system has hitherto been more of a craft than a science. As a step towards addressing this issue, the current project will develop the first parallel corpus consisting of monologue–dialogue pairs $\langle m, d \rangle$, with m and d conveying the same information. The problem will be to construct a corpus where the dialogues d are authored by professional writers and corpus construction relies on project staff, who are reasonably well-versed in writing monologue, but not experts at dialogue authoring. We solve this conundrum by starting with a set of professionally authored dialogues (about 80 pages of A4 of expository dialogues of varying lengths and on a variety of topics that are freely available in electronic form from online resources such as Project Gutenberg), and then leave it to the RA and PI to write one monologue m for each of the dialogues d that is equivalent in content. Moreover, the following annotations are added: 1) alignment of sentences in the monologue with spans of dialogue, for which we will use and extend alignment guidelines for parallel text, e.g., BAF.[1], 2) annotation of *Rhetorical Structure Theory* (RST) coherence relations [5] on the monologues and dialogues structure [4] on the dialogues, 3) annotation of cohesion relations on both monologues and dialogues based on the MUC-7 coreference task definition[2], and 4) annotation of syntactic structure of sentences in the monologue and dialogue, using an off-the-shelf parser [7].

We are aware that annotation schemes such as the aforementioned ones are still subject to discussion [37] – we will, however, employ schemes that are widely in use and have known scores for inter-annotator agreement. Additionally, will carry out our own validations by having a) annotators cross-check each other's work, and b) independent judges score informational equivalence for a sample of text-dialogue pairs.

Objective 2 – Formulate transformation rules The corpus, consisting of a collection of annotated $\langle m, d \rangle$ pairs, will allow us to formulate transformations from monologue to dialogue. We will collect instances of mappings from specific

[1] http://www.iro.umontreal.ca/ simardm/lrec98/

[2] See http://www-nlpir.nist.gov/related_projects/muc/proceedings/co_task.html.

rhetorical relations to dialogue acts/structures (such as Question-Answer pairs, but also, for instance, Acknowledge and Clarify moves [4]), and gather quantitative data using statistical methods and machine learning, specifically structured learning approaches such as the generic method for T2T generation/paraphrasing proposed by Cohn & Lapata [8]. The transformations will need to be sensitive to particular patterns of coherence relations in the input text, map these to dialogue and at the same ensure cohesion and syntactic well-formedness of the output dialogue. For this purpose, we will develop a Monologue-to-Dialogue Mark-up language (M2DML) which can capture regularities at these different levels of description. Our approach to cohesion preservation will be based on Siddharthan and Copestake [33] who propose a method for generating referring expressions in open domains. We will adapt it to dialogue, and aim to regenerate the referring expressions in the dialogue, since re-ordering as a result of the transformations can influence the set of properties that needs to be realized in referring expressions. For the theoretical foundations, we build on our own work on expressing meaning-preservation from monologue to dialogue with the help of the λ-calculus [24, 28].

Objective 3 – Implement the transformations Transformations will be stored declaratively in XML-compliant M2DML documents, with execution implemented in JAVA and *JavaDoc* documented . The other main task will be implementation of interfaces between M2DML and annotation schemes used in the system's input and output. Note that we will *not* develop technologies for analysing text or rendering dialogue text as spoken dialogue between computer-animated agents. On the input side this means relying on *off-the-shelf* modules for analysis and human-annotated gold standard inputs for evaluation. On the output side, we generate text, but suitably marked-up (e.g., using BML, the Behaviour Mark-up Language [18] to interface with state-of-art presentation systems. Thus, in Fig. 3, the primary path involves 9a and 10a and results in documents with a dialogue in text form; additionally, in collaboration with the visiting researcher, a version of the system will be delivered that produces dialogue for computer-animated characters (path 1 . . . 9b, 10b) via BML.

Objective 4 – Evaluate the implemented transformations Evaluation will address both content and organization of automatically generated dialogues: content evaluation will focus on whether the information from the input text is preserved in the output dialogue, whereas organization evaluation will focus on whether the organization of the dialogue acts in the output dialogues is coherent and whether cohesive referring expressions chains improve dialogue quality. We will follow the suggestions for evaluation of T2T in Rus et al. [32], which recommends avoiding automatic evaluation for open-ended tasks such as the current one, because it does not deal well with context-dependence and desirable variation. We will draw on a panel of 10 judges that will score both automatic and manually constructed text–dialogue pairs on informational equivalence, coherence and other measures of dialogue quality. We will investigate whether the cohesion processing step makes a noticeable difference through a defect experiment (*cf.* Callaway and Lester [3] for an application of this metholodogy in NLG). To investigate the robustness of the transformations, we will have the system generate the dialogues for evaluation from both monologues of the project's parallel corpus (ranging over a variety of topics) and other manually annotated text corpora such as the Wall Street Journal based RST Treebank [5] (extended with manual or automatic cohesion annotations [26]). Also, the use of human 'gold standard' annotated input is intended to *factor out of the evaluation* the reliability of the off-the-shelf analysis modules. Since we use human judges, we will seek approval from and comply with the the Open University's Human Participants and Materials Ethics Committee.

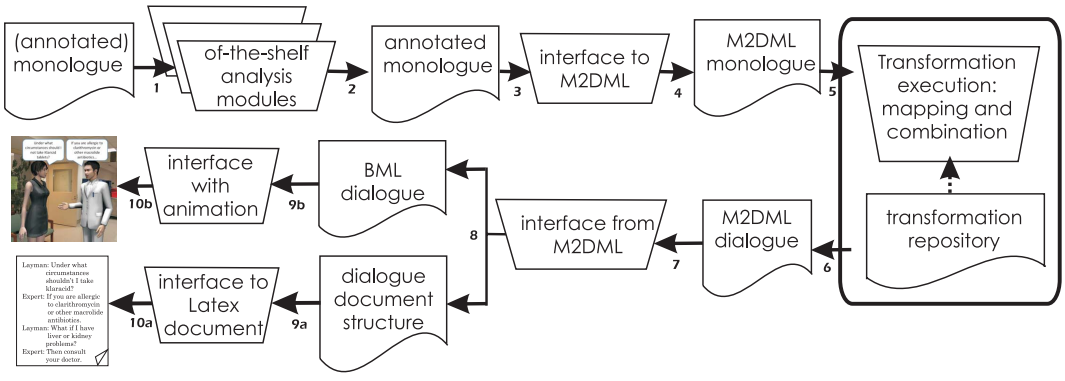

Figure 3: Architecture for transforming text to dialogue

Research questions Research questions that we hope to answer by addressing the aforementioned objectives include:

- Whether, and, if so, when and how are which coherence relations realized in professionally authored dialogue?
- Does the formal characterisation of meaning-preserving mappings [24, 28] apply to the mappings derived from the new parallel corpus, and if not, what characterises the new mappings? Are they author-specific?

- To what extent do monologue-to-dialogue transformations for one domain/text genre apply to another one?
- How important are cohesive, in addition to coherence, relations in mapping from text to dialogue?

2.4 Workplan

A diagrammatic workplan for this project is attached with this proposal. Each of the objectives is associated with a work package and milestones, with the latter indicating concrete measurable outcomes.

2.5 Project Management

The PI, responsible for project management, and RA will have weekly meetings to discuss progress and plan activities. Progress will also be reported at the weekly meetings of the OU's NLG group. The collaboration with Dr. Prendinger and his team at the National Institute of Informatics (Tokyo) will be maintained through mutual visits, email, skype conversations, a shared wiki on dialogue generation, and Access Grid meetings. The PI and RA's involvement in individual work packages is provided in attached Diagrammatic Workplan, with the PI in a supervisory role for all WPs.

2.6 Timeliness and Novelty

The field of Natural Language Generation is experiencing a shift from domain-specific approaches, that often face a knowledge bottleneck, to more generic approaches, such as text-to-text generation [32]. The current proposal extends this to the generation of expository dialogue. It will contribute a parallel monologue–dialogue corpus that is unique in its kind, and which will stimulate further empirical data-driven research. This approach also dovetails with the recent revival of work on automatic discourse parsing by various research groups – e.g., Webber at Edinburgh, Sporleder & Lascarides at Saarbr"ucken (Germany) and Edinburgh, Taboada at Simon Fraser University (Canada), Egg & Redeker at Groningen (Netherlands), and Georg & Jaulent (Paris). Finally, Knott [17] argues that applications of discourse parsers are the best way to gather solid evidence for and against specific approaches to discourse coherence. The proposed system promises to be a further test-bed for such comparisons.

2.7 Relevance to Beneficiaries

The proposed project will engage in basic scientific research. Its parallel corpus will provide a valuable resource for initiatives such as the one on Shared Task and Evaluation Challenges for Natural Language Generation (NLG) and T2T/paraphrasing [32]; the PI is co-organizer of a new STEC on Question Generation (which involves both the NLG and the Intelligent Tutoring Systems/AI in Education communities). Besides contributing to aforementioned communities through creation of resources and algorithms, the current project will benefit adjoining research areas such as the field of computer-animated Virtual Agents by providing automatically generated content for such agents – this will be fostered by collaboration with the visiting researcher, a leading expert in this area.

As mentioned in the introduction, the project is in line with the Government's Council for Science and Technology's aim of *providing people with information 'how they need it'*. We believe that if the research is successful it will have a strong impact on the R&D of methods for intelligent information presentation. We would expect the project to lead to potentially commercially exploitable middle-ware bridging content locked up in electronically stored text and multimedia presentation tools and devices, with applications in education (e.g., presenting textbook information in dialogue form), E-health (e.g., presenting patient information leaflets in dialogue form, see our preliminary study [24]) and gaming and serious/educational games (i.e., automatically generating dialogue content for non-player characters – with a tremendous potential impact on this industry witness the IGDA Game Writers Special Interest Group [13] who state that the amount of dialogue script for a character-driven computer game is usually many times that for the average film). We plan to follow up on exploitation through, for example, a Knowledge Transfer Partnership upon completion of the project.

2.8 Dissemination and Exploitation

The work will be disseminated through appropriate journal (e.g., CL, AIJ and NLE) and workshop/conference papers (see Justification of Resources for details). We will set up a demonstration website, the specification of M2DML will be made available, and the fully documented prototype will be released with a Creative Commons Attribution-Non-Commercial-Share Alike license. We also aim to organize a workshop on dialogue generation during the life of the project in collobaration with Dr. Prendinger, the visiting researcher, and researchers from elsewhere. In 2008, the PI was selected for participation by the National Endowment for Science, Technology and the Arts for the Crucible08 labs. These will provide him with training and opportunities for disseminating research to the general public.

References

[1] E. André. The generation of multimedia presentations. In R. Dale, H. Moisl, and H. Somers, editors, *Handbook of Natural Language Processing*, pages 305–327. Marcel Dekker, Inc, 2000.

[2] E. André, T. Rist, S. van Mulken, M. Klesen, and S. Baldes. The automated design of believable dialogue for animated presentation teams. In *Embodied Conversational Agents*, pages 220–255. The MIT Press, Cambridge, MA, 2000.

[3] Charles B. Callaway and James C. Lester. Narrative prose generation. *Artificial Intelligence*, 139(2):213–252, August 2002.

[4] J. Carletta, A. Isard, S. Isard, J. Kowtko, G. Doherty-Sneddon, and A. Anderson. The reliability of a dialogue structure coding scheme. *Computational Linguistics*, 23:13–31, 1997.

[5] L. Carlson and D. Marcu. Discourse tagging reference manual. Technical Report ISI-TR-545, ISI, September 2001.

[6] M. Cavazza and F. Charles. Dialogue Generation in Character-based Interactive Storytelling. In *Proceedings of the AAAI First Annual Artificial Intelligence and Interactive Digital Entertainment Conference*, Marina Del Rey, California, USA, 2005.

[7] S. Clark and J. Curran. Wide-Coverage Efficient Statistical Parsing with CCG and Log-Linear Models. *Computational Linguistics*, 33(4):493–552, 2007.

[8] T. Cohn and M. Lapata. Large margin synchronous generation and its application to sentence compression. In *Procs of EMNLP-CONLL*, pages 73–82, Prague, 2007.

[9] R. Cox, J. McKendree, R. Tobin, J. Lee, and T. Mayes. Vicarious learning from dialogue and discourse: A controlled comparison. *Instructional Science*, 27:431–458, 1999.

[10] S. Craig, B. Gholson, M. Ventura, A. Graesser, and the Tutoring Research Group. Overhearing dialogues and monologues in virtual tutoring sessions. *International Journal of Artificial Intelligence in Education*, 11:242–253, 2000.

[11] D. Dehn and S. van Mulken. The impact of animated interface agents: a review of empirical research. *Int. J. Human-Computer Studies*, 52:1–22, 2000.

[12] K. Echavarria, M. Genereux, D. Arnold, A. Day, and J. Glauert. Multilingual virtual city guides. In *Procs 15th International Conference on Computer Graphics and Applications*, Novosibirsk Akademgorodok, Russia, 2005.

[13] IGDA Game Writers Special Interest Group. International game developers association's (IGDA) guide to writing for games. IGDA White Paper, November 2003.

[14] M.A.K. Halliday and R. Hasan. *Cohesion in English*, volume 9 of *English Language Series*. Longman, 1976.

[15] A. Isard, C. Brockmann, and J. Oberlander. Re-creating Dialogues from a Corpus. In *Procs of the Corpus Linguistics 2005 ws on Using Corpora for Natural Language Generation*, Birmingham, UK, 2005.

[16] A. Isard, C. Brockmann, J. Oberlander, and M. White. The critical agent dialogue (CrAg) project. In *Procs of DiaBruck: 7th ws on the semantics and pragmatics of dialogue*, Saarbruecken, 2003.

[17] A. Knott. Book review of F. Wolf and E. Gibson's 'Coherence in Natural Language: Data, Structures and Applications'. *Computational Linguistics*, 33(4):591–595, December 2007.

[18] S. Kopp, B. Krenn, S. Marsella, A. Marshall, C. Pelachaud, H. Pirker, K. Thorisson, and H. Vilhjalmsson. Towards a Common Framework for Multimodal Generation in ECAs: The Behavior Markup Language. In *Intelligent Virtual Agents 2006*, LNAI 4133, pages 205–217, Berlin, 2006. Springer Verlag.

[19] J. Lee, F. Dinneen, and J. McKendree. Supporting student discussions: it isn't just talk. *Education and Information Technologies*, 3:217–229, 1998.

[20] William C. Mann and Sandra A. Thompson. Rhetorical structure theory: Toward a functional theory of text organization. *Text*, 8(3):243–281, 1988.

[21] M. Mateas and A. Stern. Structuring content in the faade interactive drama architecture. In *Proc. of Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, Marina del Rey, Los Angeles, June 2005.

[22] R. Mitkov, L. An Ha, and N. Karamanis. A computer-aided environment for generating multiple-choice test items. *Natural Language Engineering*, 12(2):177–194, 2006.

[23] A. Nadamoto and K. Tanaka. Complementing your TV-viewing by web content automatically-transformed into TV-program-type content. In *Proceedings 13th Annual ACM International Conference on Multimedia*, pages 41–50. ACM Press, 2005.

[24] P. Piwek, H. Hernault, H. Prendinger, and M. Ishizuka. T2D: Generating Dialogues between Virtual Agents Automatically from Text. In *Intelligent Virtual Agents: Proceedings of IVA07*, LNAI 4722, pages 161–174. Springer Verlag, 2007.

[25] P. Piwek, R. Power, D. Scott, and K. van Deemter. Generating Multimedia Presentations. In *Multimodal Intelligent Information Presentation*, volume 27 of *Text, Speech and Language Technology*, pages 203–225. Springer, Dordrecht, 2005.

[26] M. Poesio and M. Kabadjov. A General-Purpose, off-the-shelf Anaphora Resolution Module: Implementation and Preliminary Evaluation. In *Procs of LREC'04*, Lisbon, Portugal, 2004.

[27] H. Prendinger and M. Ishizuka, editors. *Life-Like Characters: Tools, Affective Functions, and Applications*. Cognitive Technologies Series. Springer, Berlin, 2004.

[28] H. Prendinger, P. Piwek, and M. Ishizuka. A novel method for automatically generating multi-modal dialogue from text. *International Journal of Semantic Computing*, 1(3):319–334, 2007.

[29] E. Reiter. An Architecture for Data to Text Systems. In *Procs of ENLG-2007*, pages 97–104, Schloss Dagstuhl, Germany, 2007.

[30] N. Reithinger, M. Loeckelt P. Gebhard, A. Ndiaye, N. Pfleger, and M. Klesen. VirtualHuman: dialogic and affective interaction with virtual characters. In *Procs of 8th intl conf on Multimodal interfaces*, pages 51–58, Banff, Alberta, Canada, 2006.

[31] V. Rus, Z. Cai, and A.C. Graesser. Experiments on Generating Questions About Facts. In *Procs. of CICLing 2007*, volume 4394 of *LNCS*, pages 444–455, Berlin, 2007. Springer Verlag.

[32] V. Rus, A.C. Graesser, A. Stent, M. Walker, and M. White. Text-to-text generation. In R. Dale and M. White, editors, *Shared Tasks and Comparative Evaluation in Natural Language Generation: Workshop Report*, 2007.

[33] A. Siddharthan and A. Copestake. Generating Referring Expersions in Open Domains. In *Procs. of 42nd ACL Conference*, Barcelona, Spain, 2004.

[34] R. Soricut. *Natural Language Generation using an Information-Slim Representation*. PhD thesis, University of Southern California, Department of Computer Science, 2006.

[35] K. Sumi and K. Tanaka. Transforming E-contents into a storybook world with animations and dialogues using semantic tags. In *Online Procs of WWW-05 Workshop on the Semantic Computing Initiative (SeC-05)*, 2005.

[36] S. V. Suzuki and S. Yamada. Persuasion through overheard communication by life-like agents. In *Procs of the 2004 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pages 225–231, Beijing, China, September 2004.

[37] K. van Deemter and R. Kibble. On coreferring: On coreference in MUC and related annotation schemes. *Computational Linguistics*, 26(4), 2000.

[38] K. van Deemter, B. Krenn, P. Piwek, M. Klesen, M. Schroeder, and S. Baumann. Fully Generated Scripted Dialogue for Embodied Agents. *Artificial Intelligence Journal*, to appear.