Conversational Agents in a Virtual World

Peter Adolphs[1], Anton Benz[2], Nuria Bertomeu Castello[2], Xiwen Cheng[1], Tina Klüwer[1], Manfred Krifka[2], Alexandra Strekalova[2] and Hans Uszkoreit[1], and Feiyu Xu[1]

[1] Deutsches Forschungszentrum für Künstliche Intelligenz (DFKI) GmbH, Alt-Moabit 91c, 10559 Berlin, Germany
[2] Zentrum für Allgemeine Sprachwissenschaft (ZAS), Schützenstrasse 18, 10117 Berlin, Germany

Abstract. This paper presents a system that builds on theoretical and experimental insights from linguistic pragmatics, uses novel techniques from computational linguistics and combines them with robust baseline technologies to provide intelligent Non Player Characters (NPCs), which naturally act and talk in a virtual world. Current NPCs still lack the necessary linguistic knowledge and methods to apply them to the numerous conversational application areas in virtual worlds. The system presented in this paper manages two NPCs, a barkeeper and a furniture sales agent, which highly depend on conversational abilities.

1 Introduction

In the latter decade, Massive Multiplayer Online Role Playing Games and virtual reality environments, such as *Warcraft* and *Second Life*, are proliferating. In such games, players cohabit a 3D-environment with other players. To populate such 3D-worlds in the early stages of the game's existence and to fill roles which humans do not want to play, such as waiter, shop-assistant, etc. non-player-characters (NPC) are created. NPCs are virtual characters, who can serve different purposes in the game, ranging from providing information to helping the user to carry out some task. However, the application areas of NPCs are currently very limited, mainly due to their marginal linguistic capabilities. There are very few commercial games which handle linguistic input[3].

Recently, there has been some research on providing virtual characters in 3D-worlds with conversational capabilities. The NICE fairy-tale game [9] and the Mission Rehearsal Exercise System [10] are some of the most sophisticated resulting prototypes. Moreover, several research projects develop stand-alone embodied conversational agents (ECA) which are not part of a virtual game, for example the Companion project [4], Justina [12], Max [14] or relational agents projects [3]. Nevertheless, a lot of research still needs to be done. Especially the processing of unrestricted natural language input and the integration of

[3] *Lifeline*, released by SCEI and Konami, is an example of a game which allows for spoken commands

human-like dialogue capabilities enhance the naturalness of the NPCs. These capabilities may include e.g. the understanding of implicatures, the ability of producing pragmatically adequate responses as well as the interpretation of the user's utterances and the spatial situation. Additionally, NPC systems always need to preserve the robustness of real-world applications.

The system described in this paper builds on theoretical and experimental insights from linguistic pragmatics, uses novel techniques from computational linguistics and combines them with robust baseline technologies to provide NPCs which possess the necessary intelligence to act and talk in a virtual world. Our NPCs offer various services pertinent to virtual worlds through conversation with game users. Furthermore, the utilisation of detailed knowledge representations enables semantic search and inference.

The paper is organised as follows: The next section describes the virtual world we use as a basis for our system as well as the application scenarios. Section 3 details the human-system interaction corpus which was used for an initial analysis of the domain and the training of our methods. Section 4 presents the underlying knowledge of the NPCs and how it was acquired. Important aspects of dialogue management within a virtual world are explained in section 5: the interpretation of the user's input, the selection of an appropriate reaction, as well as the ability of handling spatially situated dialogue. Finally, section 6 summarises the paper and suggests future work.

2 Application Scenarios

We use the virtual world *Twinity*[4] as a testbed for our system. *Twinity* provides a digital mirror of an urban part of the real world. At the time of this writing,

[4] `http://www.twinity.com/`, accessed 1 May 2011

Fig. 1. Our two NPCs in interaction with human customers

the simulated section of reality already contains 3D models of the cities of Berlin, Singapore, London, Miami and New York and it keeps growing. Users can log into the virtual world, where they can meet other users and communicate with them using the integrated text chat function. They can style their virtual appearance, rent or buy their own flats and decorate them as to their preferences and tastes.

We model two specific characters with different facets, a furniture sales agent and a barkeeper (cf. Fig. 1). The task of the furniture seller is to help users with the interior design of their virtual apartments. Users can buy pieces of furniture and room decoration from the NPC by describing their demands and wishes. The barkeeper, on the other hand, sells cocktails to the users, but moreover, he can entertain his guests by providing trivia-type information about pop stars, movie actors and other celebrities.

We chose these two characters not only because of their value for the *Twinity* application but also for our research goals. Furniture sales conversations are governed by a complex task model involving rich knowledge models about objects to be sold and pragmatic strategies of goal negotiation that guide the conversation. The barkeeper agent on the other hand has to be able to handle less constrained situations. The main role of the barkeeper is to be a conversation companion, which allows us to study and model small talk strategies.

3 Interaction Data

Although there exist some corpora of interactions between humans, mostly children, and virtual agents, e.g. [15, 9], data about human-NPC interactions in 3D-environments are still rather scarce. In order to better understand the nature of this type of interaction, we collected two corpora of human-NPC interactions by means of Wizard-of-Oz (WoZ) experiments.

The interactions of our WoZ-experiments took place in the virtual Twinity environment through a chat-interface by means of written typed language. A furniture-sales scenario was simulated, in which the NPC played the role of an interior designer/furniture saleswoman, whose task was to help the subjects furnishing a living-room. She had to determine the subject's preferences for the different types of furniture, show objects according to these preferences and place them in the room following the subject's instructions. The experimental setting is thoroughly described in [2]. An example dialogue is shown in (1).

The second experiment differed from the first one only in that subjects were made aware of the possibility to do small talk with the NPC and the Wizard was told to initiate small talk when possible.

(1)
USR.1: And do we have a little side table for the TV?
NPC.1: I could offer you another small table or a sideboard.
USR.2: Then I'll take a sideboard thats similar to my shelf.
NPC.2: Let me check if we have something like that.

In a third experiment, the NPC played the role of a barkeeper, whose task was to serve a drink and carry out a small talk conversation with the subject. The subject was told to order a drink and have a conversation with the barkeeper. The latter experiments were carried out in German language.

As a result of the first two experiments, we obtained two mixed-initiative human-NPC corpora consisting of eighteen interactions with one hour of duration each. The first corpus with 23.015 alpha-numeric strings, 4.313 utterances and 3.171 turns was fully annotated. The third experiment resulted in a corpus of 12 dialogues with 1477 utterances. We developed an annotation scheme and a dialogue state representation format with *minimal joint projects* [5] as annotation unit. A minimal joint project has a purpose and is usually realised by an adjacency pair. An adjacency pair consists of two ordered actions carried out by different agents. The first action initiates the joint project, by raising an issue, and the second action completes it. For each joint project we annotated its function, goal, existence of embedded projects, information state at the time of initiating the project, and the initiating and completing actions. The annotation scheme and representation format are presented in detail in [2].

4 Knowledge Representation and Acquisition

4.1 Furniture Ontology

The interior designer has to be able to give guidance and advice. Hence, we need structured knowledge about both decoration, e.g. styles, colours, materials etc. and the catalogue of furniture items available to the user. For this purpose we constructed a detailed ontology specified in the OWL 1.1 language. Currently the ontology consists of 975 classes, 54 properties, 327 instances and 1712 facts (property and class assertions). The ontology contains general concepts (e.g. Room, Object), information about furniture objects (e.g. Sofa, Table) and particular catalogue items from the virtual environment (e.g. Sofa_Isadora). Each item is characterised by several properties, such as *hasStyle*, *madeOf*.

Among concepts modelled in the ontology colour and style are particularly interesting. All instances of the general class *Colour* are defined through properties *hue*, *saturation* and *value* (according to the HSV colour model) and *relative luminance*. Luminance values allow us to absolutely order colours according to their paleness. This enables us to answer comparative queries, e.g. a sofa in a paler or darker colour than the one currently shown. The HSV values allow to answer queries about bright (100% saturation) or dull (unsaturated) colours. Abstract classes, e.g. *Pale Colours*, *Dark Colours*are defined through numeric intervals for the HSV or luminance values. Reasoners can then identify all instances that belong to an abstract class and classify them accordingly.

Some further information about colours is specified, for example, feelings or abstract elements that a particular colour is associated with, and the temperature of a colour (warm vs. cold). All this information allows us to find adequate objects for unspecific wishes like the following: *I want a calm atmosphere*In such cases, reasoning can be applied to find the colour that can achieve the desired effect and use it in the query to retrieve appropriate objects, without asking the user to directly specify a colour preference.

Our style modelling accounts for the fact that users might have different knowledge degrees about architectural styles. A user with little knowledge about styles may ask for something modern or antique, whereas an expert can ask for a Modern Classic or Art Deco furniture piece. For this purpose styles like *ArtDeco* or *ArtNouveau* can be classified into more coarse-grained categories, such as *Traditional* or *Modern* These more general classes were defined in our ontology based on numeric (year) intervals for the value of the *stylePeriodBeginning* and *stylePeriodEnd* properties, e.g. *Retro = 1950-1980.*

4.2 Information Extraction and Data Merging

Knowledge bases that are large enough to satisfy user's information requests are laborious to create and often incomplete. We use Semantic Web technology for building the knowledge base for our agents. This knowledge base is in particular important for the barkeeper scenario, where the NPC has to be able to interpret user statements and queries about the world. We created a biographical ontology, the "gossip ontology", defining biographical and career-specific concepts for people. This ontology is accompanied by a huge data resource of celebrities by combining i) existing Semantic Web resources, ii) Semantic Web resources which have been created from semi-structured textual data in the Web and iii) with relation information extracted from free natural language texts [19], [18]. This resource covers nearly 600,000 persons and relations between them like family relationships, marriages and professional relations.

When combining different knowledge sources with overlapping domains, it is crucial to solve the object identity resolution task [7]. This involves finding records in both data sources that refer to the same entities in the world. The resources we use have overlapping concepts and instances for people, groups and locations. Instances were merged by applying a set of heuristics based on common sense as well as on cultural-specific knowledge. For example, we assumed that two instances of type *Person* belong to the same entity when they share at least some features such as their name descriptions, birthdays or birthplaces. However, it occurs in many cases that some important features are unavailable, or that their property values do not match exactly each other. In such cases, we decided to soften the matching constraints. For instance, we have to utilise a normalisation rewriting grammar for names in order to cope with spelling variants or we have to make use of the existence of common relatives.

5 Dialogue Management

The core of the dialogue system is the dialogue manager whose main tasks are the interpretation of the input from the user as well as the selection of an appropriate reaction. Interpretation focuses on the recognition of the dialogue acts which the user performs by his utterances. Due to the mixed-initiative approach used in the agent, the dialogue system is also responsible for the control of the system initiative on the basis of dialogue context.

5.1 Dialogue Act Recognition using Syntactic and Semantic Relations

Dialogue Acts (DAs) represent the functional level of an utterance, such as a greeting, a request or a statement. Dialogue acts are verbal or nonverbal actions that incorporate participant's intentions.

The dialogue act recognition in the described system is a hybrid component, which uses patterns, statistical methods and rules to detect the appropriate intention belonging to the input. Patterns contain assignments of regular expressions matching the input to a special dialogue act. This is useful for highly predictable, idiosyncratic or one word input such as “hi” or “you are welcome”. The rules and the statistical model use a cue-based method for dialogue act classification with various features from multi-level knowledge sources. Features include information extracted from the incoming utterance as well as information about the previous dialogue. In contrast to existing systems using bag-of words representations of the utterances [6], lexical features, i.e. words, single markers such as punctuation [17] or combinations of various features [16], the results of our interpretation component are based on syntactic relations and a minimal dialogue context [13]. Relations are extracted from the utterance through a predicate argument analysis. Context features are: the last preceding dialogue act, equality between the last preceding topic and the actual topic, and sentence mood. Syntactic relation features are: syntactic predicate class, arguments, and negation.

The rules are generalised assignments of feature combinations to dialogue acts. The rules are learned from the same feature lists which serve as input for the statistical model. Learning is supervised since we decided to name some feature values explicitly (e.g. personal pronouns) whereas others are treated anonymously. Rules contain conditions regarding the existence or non-existence of elements as well as the equality of element's values and are organised according to their specificity. The following example shows the rule covering questions such as “Can we add a table?” in XML.

(2) <rule pred=”add” subj=”we” dobj=”+”>
<da>r_ar</da>
</meaning>
</rule>

The statistical model is generated by a Bayesian classifier on the basis of the corpus annotated with dialogue acts and relational information. The model is integrated into the interpretation component and deals with input which does not match a pattern or a rule.

The dialogue act recognition was evaluated using 10-folded cross validation with the annotated corpus data described in section 3.

(3)

Method	Accuracy
AODEsr	68,7%
Rules	49,5%

The table presents the results of the dialogue act classification methods. The Bayesian Classifier is the most reliable method with 68,7% accuracy. The best result we could achieve with the rules is 49,5% accuracy without the information about the preceding dialogue act. If the preceding dialogue act is included in the classification process the accuracy of the rules drops to 34,4%.

5.2 Question Answering

Interaction data show that users ask NPCs about all sorts of things, as for instance their personal information and their doing, the surrounding, available choices in selling situations, or the state of affairs in the world. Question-answering (QA) technology forms a central pillar for handling many of these information-seeking requests. Following a general design paradigm for building our system, we use detailed knowledge representations for solving this task (cf. section 3). On top of this knowledge base, we built a question answering module, which allows users to access information in a smooth natural language dialog.

The QA module is embedded into our input processing pipeline: each user input is first linguistically analysed and interpreted with respect to the current dialog context. The result of the input interpretation is a dialogue act and a more fine-grained semantic representation of the users input in case of information-seeking requests. This question semantics is turned into a query to the knowledge base. The query results are then turned into an abstract representation of the answer and spelt out by a natural-language generator [1]. Our module is able to provide the QA functionality in a smooth and connected dialog. A dialog memory and the dialog state allow the module to resolve pronouns from previous entity mentions as well as to pose and react to clarification questions in case of a possible misunderstanding.

In order to achieve robustness and accuracy for question processing, we take a hybrid approach by combining two strategies. One is a fuzzy pattern matching algorithm which utilises regular lexico-syntactic patterns based on surface strings and recognised named entities, while another makes use of dependency tree structures as patterns. The lexico-syntactic patterns are very robust and not dependent on the performance of a parser. However, the enhancement of the pattern set with the dependency trees allowed us to reduce the 1067 lexico-syntactic patterns to 212 dependency tree patterns, with almost the same linguistic coverage. Thus the utilisation of syntactic parsing results eases maintenance of the pattern set in a significant way.

While tree-based rules help to abstract from surface variations and thus to reduce the number of rules that have to be coded, relevant expressions for the domain still have to be found and formalised for a specific domain. In order to discover all the different means to express a certain fact or question, we drastically extended our initial pattern base by using crowd-sourcing methods in a further effort. We built a platform for acquiring paraphrases of seed questions for biographical facts from multiple human annotators. The seed questions are associated with their semantic arguments and functions. As before, the users input is mapped against this set of paraphrases using fuzzy pattern matching.

The resulting resource can be both used for deriving further modular pieces of expression for a compositional input analysis or it can serve as a gold resource for pattern acquisition and system evaluation.

5.3 Generation of Optimal Answers

A problem often faced by recommender systems is that there is no product in the catalogue exactly matching the user's requirements, that is, the user's query is over-constrained. In such situations recommender systems typically inform the user about the failure to retrieve a valid match and provide some kind of cooperative response. In our two corpora for the sales scenario, we find several types of cooperative responses to retrieval failures; see e.g. (4).

(4) USR: *Let me see a modern one ... If it's possible a yellow one, please.*
SYS: *If you would like something modern it is going to be in white or grey. But I can offer you a yellow vintage sideboard.*

A pre-requisite for the generation of cooperative answers to retrieval failures is a judgement about how adequate the available alternatives are with respect to the user's preferences. Preference statements are qualitative in nature and in many cases they do not reveal the relative strength of customer's preferences over the different attributes and values. For example, the preference statement '*I want a purple leather sofa*' translates into the attribute–value matrix [TYPE = sofa, COLOUR = purple, MATERIAL = leather].

In decision theory, preferences are represented by utility functions which map the possible outcomes of decisions, in our case the objects of the catalogue, to real values. As is often done in applied decision theory [11], we make the simplifying assumption that the customer's preferences can be represented by an *additive multi–attribute utility function.* This means that each database object a can be identified with a sequence of attribute–values $\langle a_1, \ldots, a_n \rangle$ such that the customer's utility function F can be decomposed into the sum of his preferences over the different attributes which in turn can be represented by a non–negative real valued function F_i for the i'th attribute:

$$F(a) = F_1(a_1) + F_2(a_2) + \ldots + F_n(a_n). \quad (1)$$

The utility function F can be further constrained by dividing the attributes into hard and soft attributes. For example, when the customer states that he wants a *purple leather sofa*, we can assume that [TYPE = sofa] is a hard constraint, and that COLOUR and MATERIAL are soft attributes. Hence, searching for optimal alternatives is equivalent to a constraint optimisation problem for which a database object a has to be found which satisfies all hard constraints and optimises $F(a)$, where F is a sum of the utilities $F_i(a_i)$ for soft attributes i. The main problem is to optimise $F(a)$ without actually knowing the objective function F.

This can be solved by assuming that the domains of values of the different attributes have the geometrical structure of *conceptual spaces* as argued for by

Gärdenfors [8]. A conceptual space consists of geometrical representations corresponding to the different quality dimensions of the concept. The customer's preference statements define a *target* t in a conceptual space. d_i measures the distance between two values for attribute i. The goal is to minimise the distance to the target value, so that $d_i(t, a_i) < d_i(t, a'_i) \Rightarrow F_i(a_i) > F_i(a'_i)$. We can then order all values of the i'th dimension according to increasing distance. This allows us to retrieve Pareto-optimal objects, that is, objects that, when compared to any other object, are better than this in at least one dimension. Whatever the strength of the different preferences are, there is at least one object in the retrieval set that optimally fulfils them.

Let's look at an example. Consider the target [TYPE = sofa, COLOUR = purple, MATERIAL = leather] and the following catalogue of items:

(5) (a) [Sofa_Alatea: COLOUR = red, MATERIAL = fabric],
(b) [Sofa_Nadia: COLOUR = black, MATERIAL = leather],
(c) [Sofa_Isadora: COLOUR = amethyst, MATERIAL = fabric]

Figure 2 shows the catalogue items with respect to the target in the conceptual space. The points (3,0) and (1,1) correspond to Sofa_Nadia(b) and

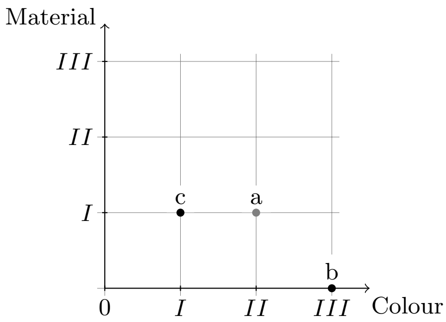

Fig. 2. Geometric representation of the search space for optimal candidates.

Sofa_Isadora(c), respectively. They dominate the other four sofas and are therefore best candidates. If we consider the properties of these objects, it is apparent that Sofa_Nadia is the only sofa made of leather, the requested material, and Sofa_Isadora, though made of fabric, is of colour amethyst, a shade of purple and the closest to the requested colour.

5.4 Spatial Model and Spatial Expressions

Human perception and cognition is spatially grounded. Users move in space and refer to space, also when being in a virtual reality. In order to give our NPCs an understanding of the surrounding space, we use a three-layered spatial representation introduced for natural-language dialogue components in mobile robots [20]. Following this approach, spatial information is represented in the following layers: (1) *Metrical representation:* Rooms and physical objects are represented by their coordinates. (2) *Topological representation:* Spatial objects (rooms and items) are topologically ordered. The topological model represents which objects contain which other objects, which objects are next to each other, etc. (3) *Conceptual representation:* Objects from the ontological representation are connected to object representations in the knowledge base, which provides class information, specific properties and information about similar concepts.

The metrical representation is fed to our system by the game engine. For each spatial object (room or item) in the game, we create an RDF representation in our dynamic knowledge base, which represents its spatial position and dimensions as well as basic object identifiers. We mimic human perception in that we limit the perceived objects to a certain radius and to the view angle. The resulting metrical representation is constantly topologically ordered, and the topological relations are as well added to our dymamic knowledge base. Finally, the dynamic item representations are linked with the ontological models by mapping the item identifiers provided by the game to the classes modelled in our rich furniture ontology. This link is established in form of a type assertion about the dynamic item representation. The employed reasoner will then ensure that all supertypes and properties associated with the class will be available for the item.

Spatially referring expressions are then generated and resolved with respect to the topological and conceptual spatial representations. Generation and resolution are realised similarly; for the sake of brevity, we only explain the generation part. In order to generate a spatially referring expression, we need an anchor, from which the spatial expression ought to be understood, and a reference, i.e. the object for which the expression should be generated. In order to establish the referring expression, we search for a shortest path between anchor and reference within the topological model. The natural-language reference is then generated by traversing the path from anchor to reference. Each node and edge on the path is realised verbally by looking up the names of nodes from the corresponding classes in the conceptual model and by mapping topological relations to language templates (usually prepositional constructions) for edges. In case of possible ambiguities arising from a node fragment, additional distinguishing properties such as the colour may be used. Finally, the node and edge fragments are concatenated. The generated spatial expression neither contains too less nor too much information and shall allow the addressee to dereference the expression unambiguously.

6 Conclusion & Future Work

This paper describes a system for conversational agents which are used for non-player characters in a virtual world. The system makes use of sophisticated, novel methods from computational linguistics on the one hand as well as robust state-of-the-art approaches on the other hand. Novel technologies include pragmatically-motivated modules for dialogue management and input interpretation as well as the processing of situated dialogue and knowledge acquisition and organisation. Due to this combination we can provide NPC agents for the conversational application areas of sales assistance and information provision. The paper describes the main technologies used in the dialogue system as well as the used corpus and knowledge sources.

As the system is still work in progress, evaluation results can only be given for some units of the system. However, we plan to perform a user evaluation of the complete system with regard to the two scenarios in the near future.

Other future work includes the further personalisation of the agents, mainly through the handling of personal opinions and opinion mining. In a first step, the agent will be able to browse a huge database of user's reviews, extract the predominant opinion and provide a summarisation to the user. The next step will include on-line extraction of the user's opinion from the input.

Acknowledgement The project KomParse is funded by the ProFIT programme of the Federal State of Berlin, cofunded by the EFRE programme of the European Union. The research presented here is additionally supported through a grant to the project TAKE, funded by the German Ministry for Education and Research (BMBF, FKZ: 01IW08003) and the German DFG Cluster of Excellence on "Multimodal Computing and Interaction". Many thanks go to the supporting company Metaversum.

References

1. Adolphs, P., Cheng, X., Klwer, T., Uszkoreit, H., Xu, F.: Question answering biographic information and social network powered by the semantic web. In: Calzolari, N., Choukri, K., Maegaard, B., Mariani, J., Odjik, J., Piperidis, S., Rosner, M., Tapias, D. (eds.) Proceedings of the 7th International Conference on Language Resources and Evaluation. European Language Resources Association (ELRA) (5 2010)
2. Bertomeu, N., Benz, A.: Annotation of joint projects and information states in human-NPC dialogues. In: Proceedings of the First International Conference on Corpus Linguistics CILC-09. pp. 723–740 (2009)
3. Bickmore, T., Cassell, J.: Social Dialogue with Embodied Conversational Agents. In: van Kuppevelt, J., Dybkjaer, L., Bernsen, N.O. (eds.) Advances in Natural Multimodal Dialogue Systems, Text, Speech and Language Technology, vol. 30, chap. 2, pp. 23–54. Springer, Berlin/Heidelberg (2005)
4. Cavazza, M., de la Cámara, R.S., Turunen, M., Gil, J.R.n., Hakulinen, J., Crook, N., Field, D.: 'How was your day?': an affective companion ECA prototype. In: Proceedings of the 11th Annual Meeting of the Special Interest Group on Discourse and Dialogue. pp. 277–280. SIGDIAL '10 (2010)
5. Clark, H.H.: Using Language. Cambridge University Press, Cambridge (1996)
6. Crook, N., Granell, R., Pulman, S.G.: Unsupervised classification of dialogue acts using a dirichlet process mixture model. In: Annual Meeting of the Special Interest Group on Discourse and Dialogue. pp. 341–348 (2009)

7. Elmagarmid, A., Panagiotis, I., Verykios, V.: Duplicate record detection: A survey. IEEE Transactions on Knowledge and Data Engineering (TKDE) 19(1) (2007)
8. Gärdenfors, P.: Conceptual Spaces: The Geometry of Thought. The MIT Press, Cambridge, MA (2000)
9. Gustafson, J., Boye, J., Fredriksson, M., Johannesson, L., Knigsmann, J.: Providing computer game characters with conversational abilities. In: Proceedings of Intelligent Virtual Agent (IVA05). Kos, Greece (2005)
10. Hill, A.W., Gratch, J., Marsella, S., Rickel, J., Swartout, W., Traum, D.: Virtual humans in the mission rehearsal exercise system. In KI Embodied Conversational Agents 17, 32–38 (2003)
11. Keeney, R., Raiffa, H.: Decisions with Multiple Objectives - Preferences and Value Tradeoffs. Cambridge University Press, Cambridge (1993)
12. Kenny, P.G., Parsons, T.D., Gratch, J., Rizzo, A.A.: Evaluation of Justina: A Virtual Patient with PTSD. In: Prendinger, H., Lester, J.C., Ishizuka, M. (eds.) IVA. Lecture Notes in Computer Science, vol. 5208, pp. 394–408. Springer (2008)
13. Klüwer, T., Uszkoreit, H., Xu, F.: Using syntactic and semantic based relations for dialogue act recognition. In: Proceedings of the 23rd International Conference on Computational Linguistics (COLING 2010). Association for Computational Linguistics, o.A. (2010)
14. Kopp, S., Gesellensetter, L., Krämer, N., Wachsmuth, I.: A conversational agent as museum guide – design and evaluation of a real-world application. In: Proc. of Intelligent Virtual Agents (IVA 2005). vol. 3661, pp. 329–343. Springer (2005)
15. Narayanan, S., Potamianos, A.: Creating conversational interfaces for children. IEEE Transactions on Speech and Audio Processing 10, 65–78 (2002)
16. Stolcke, A., Ries, K., Coccaro, N., Shriberg, E., Bates, R., Jurafsky, D., Taylor, P., Martin, R., Van, C., dykema Marie Meteer, E.: Dialogue act modeling for automatic tagging and recognition of conversational speech. Computational Linguistics 26, 339–373 (2000)
17. Verbree, A., Rienks, R., Heylen, D.: Dialogue-act tagging using smart feature selection: results on multiple corpora. In: Raorke, B. (ed.) First International IEEE Workshop on Spoken Language Technology SLT 2006. IEEE Computer Society, Palm Beach (2006)
18. Xu, F., Uszkoreit, H., Krause, S., Li, H.: Boosting relation extraction with limited closed-world knowledge. In: Proceedings of the 23rd International Conference on Computational Linguistics. Association for Computational Linguistics, o.A. (2010)
19. Xu, F., Uszkoreit, H., Li, H.: A seed-driven bottom-up machine learning framework for extracting relations of various complexity. In: Proceedings of ACL 2007, 45th Annual Meeting of the Association for Computational Linguistics. Prague, Czech Republic (6 2007)
20. Zender, H., Kruijff, G.J.: Multi-Layered conceptual spatial mapping for autonomous mobile robots. In: Schultheis, H., Barkowsky, T., Kuipers, B., Hommel, B. (eds.) Papers from the AAAI Spring Symposium on Control Mechanisms for Spatial Knowledge Processing in Cognitive / Intelligent Systems. pp. 62–66. Papers from the AAAI Spring Symposium, AAAI Press, Menlo Park, CA, USA (2007)