Towards A Separation of Pragmatic Knowledge and Contextual Information

Robert Porzel and **Hans-Peter Zorn** and **Berenike Loos**[1] and **Rainer Malaka**[2]

Abstract. In this paper we address the question of how *traditional* approaches to modeling world knowledge, i.e. to model shared conceptualizations of specific domains of interest via formal ontologies, can be enhanced by a pragmatic layer to solve the problem of explicating hitherto implicit information contained in the user's utterances and to further the assistance capabilities of dialog systems and how they can be connected to dedicated analyzers that observe topical contextual information. For this purpose, the notions of *context* and *pragmatics* are introduced as one of the central problems facing applications in artificial intelligence. We will argue that pragmatic inferences are impossible without contextual observations and introduce a model of context-adaptive processing using a combination of formal ontologies and analyzers for various types of context.

1 Introduction

In this paper two fundamental, but notoriously tricky, notions for mobile open-domain multimodal human-computer interface systems, such as SmartWeb [26], are discussed as one of the central problems facing both applications in artificial intelligence as well as in natural language processing. These, often conflated, notions are those of *context* and *pragmatics*. Indeed, in many ways both notions are inseparable from each other if one defines pragmatics to be about the encoding and decoding of meaning, which, as pointed out frequently [4, 28, 21], is always context-dependent. This, therefore, entails that pragmatic inferences (also called *pragmatic analyses* [4]) are impossible without recourse to contextual observations. In this paper, we will argue that the distinction between pragmatic knowledge - which is learned/acquired - and contextual information - which is observed/inferred - is of paramount importance in designing scalable context-adaptive systems, which seek to interact with human users and to collaborate intelligently with them. More specifically, we will focus on the use case of natural language understanding using ontology-based analyses of open-domain user utterances.

As the work presented here is part of a research undertaking that attempts to tie together semantic web technologies, natural language processing and assistance systems in an attempt to develop a mobile multimodal open-domain conversational question answering system , the central idea behind it is to employ ontological knowledge - if available - and revert to statistical processing in the absence thereof. In this paper we will focus on the ontology-based processing pipeline and examine how pragmatic knowledge and contextual infromation - needed to increase the conversational capabilities of dialogue systems - can be modeled and consequently employed. For this we give an overview of the state of the art in Section 2, followed by two motivating examples for distinguishing pragmatic knowledge from contextual information in Section 3. Thereafter, we will describe the ontological infrastructure as found in SmartWeb and our approach for modeling pragmatic knowledge as part of that infrastructure in Section 4. Finally, we will show how we *connected* this knowledge to contextual analyzers in Sections 5 and 6 followed by concluding remarks in Section 7.

2 State of the Art

In general, computational pragmatics can be defined as the attempt to enable artificial systems to encode meaning into a set of surface structures or to decode meaning from such forms In this given sense computational pragmatic resolution is equivalent to *decontextualization* in the sense of McCarthy [17]. While this work will, from now on, focus on the decoding processes it is theoretically quite possible to apply the same techniques to processes of encoding, but will not be the focus of this paper. As we will show herein, there are sound theoretical as well as practical reasons for modularizing and separating pragmatic knowledge, for which we propose an ontological model called PRONTO, from contextual information, which has to integrate numerous non-discrete, noisy and sub-symbolic sensor data in a robust fashion, for which dedicated analyzers and inference mechanisms for combining various observations can be employed.

In general terms, decoding meaning is *understanding*, however, no precise notions of where semantic processing ends and pragmatic processing begins exists, and might never be forthcoming. Various overviews describing the need for context-adaptiveness in natural language processing systems exist [4, 6, 21]. Given the goal of more intuitively usable and more conversational natural language interfaces that can someday be used in real world applications, the handling of pragmatic knowledge - needed for a felicitous decoding of the meaning encoded in user's utterances - is still one of the major challenges for understanding conversational utterances in dialogue systems, since a substantial part of that meaning is contained implicitly in the linguistic surface structures of the utterance, recourse to contextual information is needed for pragmatic analyses. The paramount importance of context for natural language understanding is frequently noted in the literature, albeit few dialogue systems take context explicitly into account and perform a corresponding context-dependent analysis of the given utterances at hand. We follow Porzel and Gurevych [21] and differentiate between four different types of contexts that contribute information relevant to natural language understanding, listed in Table 1. In dialogue systems these knowledge stores are commonly assigned to respective models: the situation model, dialogue model, user model and the domain model, e.g. represented in a formal ontology.

[1] European Media Laboratory, 69118 Heidelberg, Germany, email: firstname.lastname@eml-d.villa-bosch.de
[2] University of Bremen, 28359 Bremen, Germany, email: lastname@informatik.uni-bremen.de

Table 1. Context-types, content and their models

types of context	information observed	context model
situational context	time, place, etc	situation model
discourse context	what has been said	discourse model
interlocutionary context	user/system properties	user/system model
domain context	ontological knowledge	domain model

Recently developed multi-modal dialogue systems [27, 13, 23] equipped with the ability to understand and process natural language utterances from one ore more domains often employ ontologies as a formal, explicit specification of shared conceptualizations of their domains of interest [10]. At the same time the emerging Semantic Web [2] employs such formal conceptualizations to add semantic information to textual and other data available on the Internet. Efforts originating in various W3C and Semantic Web projects brought about several knowledge modeling standards: Resource Description Framework (RDF), DARPA Agent Mark-up Language (DAML), Ontology Interchange Language (OIL) and the Ontology Web Language standards (OWL (Lite, DL, Full)).[3]

Therefore, numerous mobile dialogue systems, such as MATCH, SmartKom or SmartWeb [27, 13, 26], employ ontologies to represent spatial and navigational knowledge; to support car, motorcycle and pedestrian navigation. Existing navigation ontologies [16, 12] describe route mereologies, which do not capture contextual dependencies. The same holds true for other domain ontologies used by the individual system(s), e.g. models of domains such as sports, entertainment and the like. Also, while ontologies commonly model a more or less static world, conceptual and common-sense knowledge [25, 11, 5] based on the standard combinations of frame- and description logics, contextual knowledge is induced in specific *instances* and highly dynamic states of affairs. In natural language processing many ambiguities arise, which can be resolved only by recourse to different contexts, e.g. discourse context has to be taken into account for reference resolution [9], domain context for hypothesis verification [22] or situational context for resolving pragmatic ambiguities [20].

Visible in all systems that are limited to an impoverished contextual analysis and precompilations, was their restrictedness in terms of their understanding capabilities, rendering them unscalable and in the case of more conversational input undeployable. This evidently shows up in the fragility of systems that fail when confronted with imperfect or unanticipated input, usually that also include perfectly unambiguous utterances that stray but a little from a scripted demo dialogue. Human conversations are between partners that share a rich background of pragmatic knowledge (involving topical observations of both more static & more dynamic contexts) without which natural language utterances become ambiguous, vague and incomplete. An interpreter with little contextual awareness and pragmatic reasoning will encounter problems and fail frequently; one which does not fail in unexpected or more complex situations is called *robust*. Several means have been used to increase robustness ranging from rules for grammatical relaxations, automatic acquisition of semantic grammars, automatic spelling correction to on-line lexical acquisition and out-of-vocabulary recognition. These so-called *low-level techniques* [4] have not solved the problem of enabling a system to react felicitously in dynamic contexts and for multiple domains. These techniques fail to assume a pragmatics-based approach where the fact that the user has an intention, communicated via a message, which has to be reconstructed by recourse to the current context, is explicitly taken into account. Therefore, today's systems using *pragmatics-free* ontologies face two options. One is to to restrict themselves to single applications with clearly defined application-specific contexts, e.g. offering single domain services - such as providing information about soccer scores - or guiding only pedestrians - always on foot and always on the shortest path. The other is to force the user to explicate each possible contextual parameter, which means reverting to controlled and restricted processing techniques.

However, if we wish to make use of (or combine) semantically described web services, which offer vast ensembles of tunable parameters, e.g. route, weather, and geo-services, or to employ semantic information extraction applications in a variety of domains, e.g. sports or news, we must provide the means to decode the appropriate meaning based on pragmatic knowledge and context-specific topical information. Moreover, we would like to do so in the least invasive way, i.e. minimizing the amount of information that needs to be obtained by asking the user in order to maximize dialogical efficiency and user satisfaction.In the following we motivate and describe how the ontologies used in the SmartWeb project were adapted to provide a principled approach for encoding pragmatic knowledge.

3 Contextual Information and Pragmatic Knowledge at Play

As mentioned above we apply our model of pragmatic knowledge and context-dependent processing to enhance the conversational understanding and ensuing assistance capabilities of dialog systems. While there exists quite a slippery slope where semantic processing ends and pragmatic assistance begins, we will try to motivate this distinction by means of two sample scenarios employed as running examples throughout this paper.

A question such as *How often was Brazil world champion?* poses a challenge to conversational open-domain dialog systems as the discourse domain of the utterance is not made explicit by the user. Since we regard the modeling of pragmatic knowledge as a major challenge for such systems and - in contrast to controlled systems - want the user to be able to make utterances in any domain of interest without placing the burden of explicating the exact context on him or her, we have to find a systematic and scalable way of modeling:

- that the pragmatic knowledge that a *correct* or felicitous answer to such a question (or many others for that matter) simply depends on what is talked about, and
- that any *intelligent* interlocutor has to know, keep track of or infer what is being talked about.

While these two statements may sound trivial, they are not. For one, the first statement expresses a fundamental bit of pragmatic knowledge that, to the best of our knowledge, has been proposed, implemented and evaluated in dialog systems only by Zorn et al. [15].[4]. This model explicetely and formally expresses such pragmatic knowledge, e.g. a bit that expresses that the *theme* of an utterance - what is new, unknown and asked about - depends on the given *rheme* - what is old, known and has been talked about. In Section 4 we show describe the corresponding ontological framework and in Section 5 how we integrate such knowledge with actual contextual observations, which as expressed in the second statement and can be regarded as an *observational* task assigned to the discourse model.

[3] See www.w3c.org/RDF, www.ontoknowledge.org/oil, www.daml.org, and www.w3.org/2004/OWL for the individual specifications.

[4] Of course, as shown in Section 2 most systems assume an implicitely given domain context or employ various shortcuts to deal with problems of underspecification.

That is to keep track and make inferences about what is being talked about or, in our terminology, to observe the given rheme at hand, which - as all contextual information - can change dynamically and even rapidly.

In a mobile dialog system contextual information is of high significance as a user expects the offer of topical services, while navigating through a dynamically changing environment (e.g. changing precipitation- and temperature levels and or traffic- and road conditions), which makes the adequate inclusion of extra-linguistic knowledge and context-sensitive processing inevitable for the task of felicitous navigational assistance. The necessity to couple extra-linguistic situative with pragmatic knowledge in the domain of spatial navigation has been demonstrated before [20, 14]. Some more obvious examples are given below:

- For instance, a pedestrian might prefer public transportation over walking when it is raining even for smaller distances.
- A motor bicyclist might prefer to use winding country roads over interstate highways when it is warm and sunny, but not, when road conditions are bad.
- A car driver might like to take a spatially longer route if shorter ones are blocked or perilous.

As mentioned above, existing navigation ontologies [16, 10] describe route mereologies, which do not capture contextual dependencies. Given a single application-specific context, e.g. guiding only pedestrians - always on foot and always on the shortest path, we can employ such a *context-free* ontology. However, if we wish to make use of the many tunable parameters offered by today's route planning and navigational systems one must provide the means to determine the right setting depending on the actual situation at hand in the least invasive way, i.e. minimizing the amount of parameters and settings obtained by bothering the user. In the following we motivate our ontological choices and describe the infrastructure employed in our approach to model the needed pragmatic knowledge for solving both sample use cases described above.

4 Pragmatic and other Ontologies in the SmartWeb Project

In order to allow systems such as the SmartWeb prototype [23] to employ a wide range of internal and external ontologies several ontological commitments and choices have to be made. The most relevant for our work are described below.

Foundational & Ground Knowledge: An important aspect in ontology engineering is the choice of a foundational layer, which is used to guarantee harmonious alignment of various independently crafted domain ontologies and their re-usability. The SmartWeb foundational ontology [5] is based on the highly axiomatized Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE) It features various extensions called *modules*, e.g. the Ontology of Plans and a module called *Descriptions & Situations* [8]. As the focus of our work lies on an application and elaboration of the latter module, it will be described more closely in the following section. Additional to the foundational ontology, a domain-independent layer is included which consists of a range of branches from the less axiomatic SUMO (Suggested Upper Merged Ontology ontology [18]), which is known for its intuitive and comprehensible structure. Currently, the SmartWeb Integrated Ontology (SWINTO) features, next to the foundation and domain-independent layers, several domain ontologies, i.e. a SportEvent-, a Navigation-, a WebCam-, a Media-, and a Discourse-Ontology.

Pragmatic Descriptions & Situations: The Descriptions & Situations framework is currently the sole ontological framework for representing a variety of reified contexts and states of affairs. In contrast to physical objects or events, the extensions of ontologies by non-physical objects pose a challenge to the ontology engineer. The reason for this lies in the fact that non-physical objects are taken to have meaning only in combination with some other *ground* entity. Accordingly, their logical representation is generally set at the level of theories or models and not at the level of concepts or relations. According to Gangemi and Mika [8] this is not generally true as recent work can address non-physical objects as first-order entities that can change, or that can be manipulated similarly to physical entities. So in many cases relations and axioms modeled and applied for physical entities are also valid for non-physical ones. Therefore, a modeling pattern was devised that connects:

- COURSES OF EVENTS sequenced by PERDURANTS, i.e. processes within the ground ontology, such as QUESTIONING,
- FUNCTIONAL ROLES played by ENDURANTS, i.e. objects within the ground ontology, such as a type of EVENT or BUILDING,
- PARAMETERS valued by REGIONS, i.e. scalar phenomena, such as TEMPERATURES or DOMAINS

For endowing the SmartWeb ontologies with a pragmatic layer, we, therefore, decided to employ the *Descriptions & Situations* (D&S) module and its modeling patterns. The central modeling choice that arises hereby concerns the question of how fine-grained such a description and relation hierarchy should be that links the corresponding courses, roles and parameters to elements of the ground ontology. Hereby the classic trade-off between modeling and axiomatization comes into play, i.e. if a corresponding axiomatization should bear the burden of associating the pragmatically grouped items of the ground (domain) ontologies, e.g. SOCCER DISCOURSE, WORLD CUP and QUESTIONING for describing the pragmatic context of a given question. In either case this elaboration of the *Descriptions & Situations* module extends the notion of deriving an instance (situation) from a description by modeling a more general pattern of pragmatic knowledge.

5 Connecting Pragmatic Knowledge with Contextual Observations

Our context model - used for observing contextual information - is implemented as a module, called Situation and Context Module (SitCoM) within `SmartWeb`'s dialog manager. It interacts with the dialog manager's iHUB middle-ware [24] . The internal communication format in SmartWeb is a RDFS adapted derivative of the EMMA w3c standard called SWEMMA. A SWEMMA document is a collection of instances, the actual interpretation is embedded within instances of a discourse and a special EMMA domain ontology. Within the dialog manager these EMMA documents are stored in an A-box. All dialog manager components access a common A-box per turn, the internal iHUB contains only pointers to the root instance of an interpretation within this A-box. Each dialog component then adds its own interpretation to the EMMA document.

SitCoM receives the semantic interpretation via the iHUB, which has been processed by the modality specific recognizers (e.g. for speech and gesture), parser and discourse model components before. The task for SitCoM is to change the semantic representation in such way that contextual information is semantically represented, as if the user would have done so explicitly. If no pragmatic descriptions are applicable the A-box is not modified and the message is sent back to the iHUB without any changes. For a pragmatic description to be applicable means that any of the ground entities contained in the SWEMMA document have been connected to COURSES OF EVENTS, FUNCTIONAL ROLES or PARAMETERS via the respecitve relations *sequenced by*, *played by* or *valued by*.

If SitCom can apply its pragmatic knowledge it will enhance the semantic representation of the user utterance. This is done either by specializing a concept or inserting missing instances into the interpretation. The necessary information stems from connections established to context providing services or sensors. Currently, we query web services for topical weather and road conditions, establish the user's current position via GPS build into the mobile device and communicate with other components of the system to obtain discourse and temporal information.

As stated above in a mobile dialogue system contextual information is of paramount importance as the user expects the offer of topical services. This alone makes the adequate inclusion of contextual factors intertwined with the corresponding pragmatic knowledge inevitable for the task of navigational assistance.

However, a closer examination shows that in a truly open domain system, such as SmartWeb, virtually every utterance becomes ambiguous in an open-domain context. Looking, again, at the question introduced above, i.e. *How often was Brazil world champion?*, we find that, without knowing the domain at hand, i.e. which type of sport - soccer, beachball or else - is talked about, it is not possible to answer these questions directly. Currently, this problem is handled by either restricting NLU systems to a pre-specified (hard-coded) domain or shifting the pragmatic disambiguation task back to the user, by asking him or her to specify the needed information, thereby producing less efficient and more cumbersome dialogues.

6 Adding Context to the System

Our context model - used for observing contextual information - is implemented as a module, called Situation and Context Module (SitCoM) within `SmartWeb`'s dialog manager. It interacts with the dialog manager's IHUB middle-ware [24] . The internal communication format in SmartWeb is a RDFS adapted derivative of the EMMA w3c standard called SWEMMA. A SWEMMA document is a collection of instances, the actual interpretation is embedded within instances of a discourse and a special EMMA domain ontology. Within the dialog manager these EMMA documents are stored in an A-box. All dialog manager components access a common A-box per turn, the internal IHUB contains only pointers to the root instance of an interpretation within this A-box. Each dialog component then adds its own interpretation to the EMMA document.

SitCoM receives the semantic interpretation via the IHUB, which has been processed by the modality specific recognizers (e.g. for speech and gesture), parser and discourse model components before. The task for SitCoM is to change the semantic representation in such way that contextual information is semantically represented, as if the user would have done so explicitly. If no pragmatic descriptions are applicable the A-box is not modified and the message is sent back to the IHUB without any changes. For a pragmatic description to be applicable means that any of the ground entities contained in the SWEMMA document has been connected to COURSES OF EVENTS, FUNCTIONAL ROLES or PARAMETERS via the respecitve relations *sequenced by*, *played by* or *valued by*. Additional inferencing mechanisms are needed for selecting appropriate descriptions, insertions of appropriate concepts and instances and combinations of observations, which have been proposed and are described in greater detail by Chang et al [3], Porzel et al. [20].

If SitCom can apply its pragmatic knowledge it will enhance the semantic representation of the user utterance. This is done either by specializing a concept or inserting missing instances into the interpretation. The necessary information stems from connections established to context providing services or sensors. Currently, we query web services for topical weather and road conditions, establish the user's current position via GPS build into the mobile device and communicate with other components of the system to obtain discourse and temporal information.

If `SitCoM` can apply its pragmatic knowledge it will enhance the semantic representation of the user utterance. This is done either by specializing a concept or inserting missing instances into the interpretation. The Situation and Context Module (SITCOM) is connected to other dialog processing modules, i.e. Speech Interpretation (SPIN), Fusion and Dialog Engine (FADE), Reaction and Presentation Manager (REAPR), the EMMA Unpacker/Packer that handles communication with the multimodal recognizer and the semantic mediator which manages access to the knowledge access services, within SmartWb's multimodal dialog processing architecture. In the following we will describe the processing steps undertaken by our module.

Collecting Pragmatic Descriptions: The `SitCoM` algorithm performs two passes over the instances contained in the SWEMMA documents found in the iHUB. These instances are part of the ground ontology and are bound via their respective properties to pragmatic description modelled in our pragmatic ontology (PrOnto). This way, the ground entities *evoke* certain description which describe contexts or situations in which the given concept may play a role. In the first pass, all these evoked descriptions are collected and put in an *active descriptions* pool.

Context Sources: The interface to the sensor data is encapsulated into so called context sources. These context sources are identified by a concept from the ground ontology and provide the context information as instance of this concept or a subclass of it. The context information can be a set of instances, in this case, the identifying concept is the anchor instance. Below, we describe a set of sources that are currently analyzed by our module.

- A GPS Receiver connected to the user device delivers current location data to the dialog manager which is passed as *external message* to `SitCoM` by the IHUB in small intervals. The GPS context source uses a web service to resolve the exact address using inverse geocoding. This information is cached and only updated if the location has changed significantly.
- The Weather Service context source polls a Web Service for current weather conditions depending on the current location.
- The Time context source encapsulates time information from the real time clock.
- This context source provides the current domain as recognized by a domain recognizer.

Context Insertion Step: These descriptions are matched against the context information and - if applicable - accordingly specialized. The parameter of the description is used to query the context source. If the resulting context information instance is some subclass of this parameter, the corresponding description-subclass is activated instead.

The last step is another iteration over all instances of the current interpretation. During this pass, all concepts are matched against the description within the active descriptions pool. If a description has been specialized in the previous pass, the ground entities corresponding to this more specific description are specialized as well.

For example: A `Tournament` instances evokes the "SportsTalk" description. This description is about talking about specific domains, e.g. sports. It consists of the functional Role `SportsRhema`, the parameter `SportsThema`. `SportsRhema` is connected to the `Tournament` ground entity and this way the description gets activated. `SportsThema` is linked to the `Domain` ground entity which is covered by the `Domain` context source. This context source returns an instance of `SoccerDomain` which is a subclass of `Domain`. This way a sub description "SoccerTalk", consisting of `SoccerRhema` and `SoccerThema` gets active. During the last step the Tournament instance is changed to a FIFAWorldCup instance to match the more specialized "SoccerTalk" description where the functional role is linked to.

7 Conclusion

In this paper we have argued that an inclusion of pragmatic knowledge is needed to scale context-adaptive systems and that this inclusion can be achieved by means of an ontological model based on an extension of the situations & descriptions framework. Additionally, we have pointed at the need to handle contextual information differently from pragmatic knowledge, as it is quite different in nature and requires other classification, inferencing and reasoning methods, for which ontologies are simply not suitable. As future work, a promising framework, called BayesOWL, originating in the work of Ding [7] constitutes a promising next step towards a better integration of symbolic and probabilistic reasoning. Additionally, the framework proposed by Porzel [19] can be employed to integrate the various contextual observations in probabilistic graphical models while keeping the conditional probability tables from exploding.

Acknowledgements

This work has been partially funded by the German Federal Ministry of Research and Technology (BMBF) as part of the SmartWeb project under Grant 01IMD01E and by the Klaus Tschira Foundation. We would like to thank the referees for their comments which helped improve this paper.

REFERENCES

[1] James Allen, Georga Ferguson, and Amanda Stent, 'An architecture for more realistic conversational system', in *Proceedings of Intelligent User Interfaces*, pp. 1–8, Santa Fe, NM, (2001).

[2] Tim Berners-Lee, James Hendler, and Ora Lassila, 'The semantic web', *Scientific American*, **May**, (2001).

[3] John Bryant, Nancy Chang, Robert Porzel, and Keith Sanders, 'Where is natural lnaguage understanding', in *Proceedings of AMLAP*, (2001).

[4] Harry Bunt, 'Dialogue pragmatics and context specification', in *Computational Pragmatics, Abduction, Belief and Context*, John Benjamins, (2000).

[5] Philipp Cimiano, Andreas Eberhart, Daniel Hitzler, Pascal Oberle, Steffen Staab, and Rudi Studer, 'The smartweb foundational ontology', *SmartWeb Project Report*, (2004).

[6] John H. Connolly, 'Context in the study of human languages and computer programming languages: A comparison', *Modeling and Using Context*, **Springer, LNCS**, 116–128, (2001).

[7] Zhongli Ding, *BayesOWL: A Probabilistic Framework for Semantic Web*, Ph.D. dissertation, University of Maryland, 2005.

[8] Aldo Gangemi and Peter Mika, 'Understanding the semantic web through descriptions and situations', in *Proceedings of the ODBASE Conference*. Springer, (2003).

[9] Peter C. Gordon and Randall Hendrick, 'The representation and processing of coreference in discourse', *Cognitive Science*, **22**(4), 389–424, (1998).

[10] Thomas Gruber, 'A translation approach to portable ontology specifications', *Knowledge Acquisition (5)*, (1993).

[11] Nicola Guarino, 'Some ontological principles for designing upper level lexical resources', in *Proceedings of the 1st International Conference on Language Resources and Evaluation,*, (1998).

[12] Iryna Gurevych, Robert Porzel, Elena Slinko, Norbert Pfleger, Jan Alexandersson, and Stefan Merten, 'Less is more: Using a single knowledge representation in dialogue systems', in *Proceedings of the 3rd IJCAI 2003 Workshop on Knowledge and Reasoning in Practical Dialogue Systems*, p. (in press), Acapulco, Mexico, (2003).

[13] Michael Johnston, Srinivas Bangalore, Gunaranjan Vasireddy, Amanda Stent, Patrick Ehlen, Marilyn Walker, Steve Whittaker, and Preetam Maloor, 'Match: An architecture for multimodal dialogue systems', in *Proceedings of ACL'02*, pp. 376–383, (2002).

[14] Berenike Loos and Robert Porzel, 'Towards ontology-based pragmatic analysis', in *In Processding of DIALOR'05*, Nancy, France, (2005).

[15] Hans-Peter Zorn Berenike Loos and Robert Porzel, 'Pragmatic knowledge and context-aware proccesing for dialogue systems', in *7th SIGdial Workshop on Discourse and Dialogue Sydney, July 15-16,*, (submitted).

[16] Raphael Malyankar, 'Creating a navigation ontology', in *Workshop on Ontology Management*, Menlo Park, CA, (1999). AAAI Press.

[17] John McCarthy, 'Notes on formalizing contexts', in *Proceedings of the Fifth National Conference on Artificial Intelligence*, eds., Tom Kehler and Stan Rosenschein, pp. 555–560, Los Altos, California, (1986). Morgan Kaufmann.

[18] Ian Niles and Adam Pease, 'Towards a standard upper ontology', in *Workshop on Ontology Management*, eds., Chris Welty and Barry Smith, Ogunquit, Maine, (2001). Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001).

[19] Robert Porzel, 'Baysian networks with expert nodes: An application in computational pragmatics', in *Master Thesis at the International University in Bruchsal*, (2006).

[20] Robert Porzel and Iryna Gurevych, 'Towards context-adaptive utterance interpretation', in *Proceedings of the 3rd SIGdial Workshop*, Philadelphia, USA, 2002, pp. 90–95, (2002).

[21] Robert Porzel and Iryna Gurevych, *Contextual Coherence in Natural Language Processing*, LNAI 2680, Sprigner, Berlin, 2003.

[22] Robert Porzel and Iryna Gurevych, 'Contextual coherence in natural language processing', in *Fourth International Conference on Modeling and Using Context*, eds., P. Blackburn, C. Ghidini, R. Turner, and F. Giunchiglia, Berlin, (2003). Springer (LNAI 2680).

[23] Norbert Reithinger, Simon Bergweiler, Ralf Engel, Gerd Herzog, Norbert Pfleger, Massimo Romanelli, and Daniel Sonntag, 'Look under the hood: Design and development of the first smartweb system demonstrator', in *Proceedings of IMCI*, Trento, Italy, (2005).

[24] Norbert Reithinger and Daniel Sonntag, 'An integration framework for a mobile multimodal dialogue system accessing the semantic web', in *Proceedings of Interspeech 2005*, Lissabon, Portugal, (2005).

[25] Stuart J. Russell and Peter Norvig, *Artificial Intelligence. A Modern Approach*, Prentice Hall, Englewood Cliffs, N.J., 1995.

[26] Wolfgang Wahlster, 'SmartWeb: Mobile applications of the semantic web', in *Proceedings of Informatik*, Ulm, Germany, (2004).

[27] Wolfgang Wahlster, Norbert Reithinger, and Anselm Blocher, 'Smartkom: Multimodal communication with a life-like character', in *Procedings of the 7th Eurospeech.*, pp. 1547–1550, (2001).

[28] Dominique Widdows, *A Mathematical Model of Context*, LNAI 2680, Sprigner, Berlin, 2003.