

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/221483519>

Enhanced monitoring tools and online dialogue optimisation merged into a new Spoken Dialogue System design experience

Conference Paper · January 2010

Source: DBLP

CITATIONS

5

READS

64

3 authors, including:



Romain Laroche

Microsoft Maluuba

58 PUBLICATIONS 185 CITATIONS

SEE PROFILE



Ghislain Putois

Voxygen SAS

22 PUBLICATIONS 80 CITATIONS

SEE PROFILE

Enhanced Monitoring Tools and Online Dialogue Optimisation Merged into a New Spoken Dialogue System Design Experience

Romain Laroche¹, Philippe Bretier¹ and Ghislain Putois¹

¹Orange Labs, Issy-les-Moulineaux, France

romain.laroche@orange-ftgroup.com

Abstract

This paper shows how the convergence between design and monitoring tools, and the integration of a dedicated reinforcement learning can be complementary and offer a new design experience for Spoken Dialogue System (SDS) developers. Most industrial SDS developers use a graphical tool to implement the dialogue strategies. This article proposes first to integrate dialogue logs into this design tool, so that it constitutes a monitoring tool as well, by revealing call flows and their associated Key Performance Indicators (KPI). Second, the SDS developer is opened the possibility of designing several alternatives and of comparing visually his design choice performances. Third, a reinforcement learning algorithm is integrated to automatically optimise the SDS choices. The design/monitoring tool helps the SDS developers to understand and analyse the user behaviour, with the assistance of the learning algorithm. The SDS developers can then confront the different KPI and control the further SDS choices by removing or adding alternatives.

Index Terms : Dialogue Design, Online Learning, Spoken Dialogue Systems, Monitoring Tools

1. Introduction

Recent research in spoken dialogue systems (SDS) has called for a “synergistic convergence” between research and industry [1]. This call for convergence concerns architectures, abstractions and methods from both communities. Under this motivation, several research orientations have been proposed. This paper discusses three of them : dialogue design, dialogue management, and dialogue evaluation. Dialogue design and dialogue management reflect in this paper the respective paths that industry and research have followed for building their SDS. Dialogue evaluation is a concern for both communities, but remains hard to put into operational perspectives.

The spoken dialogue industry is structured around the architecture of the well known industrial standard VoiceXML¹. The underlying dialogue model of VoiceXML is a mapping of the simplistic turn-based linguistic model on the browser-server based Web architecture [2]. The browser controls the speech engines (recognition and text-to-speech) integrated into the voice platform according to the VoiceXML document served by an application server. A VoiceXML document contains a set of prompts to play and the list of the possible interactions the user is supposed to have at each point of the dialogue. The SDS developers², reusing Web standards and technologies (e.g. J2EE, JSP, XML...), are used to design directed dialogues modelled

by finite state automata. Such controlled and deterministic development process allows the spoken dialogue industry to reach a balance between usability and cost [3]. This paper argues that our tools are facilitators that improve the usability vs. cost trade-off thanks to an improved reliability of advanced technologies.

Spoken dialogue research has developed various models and abstractions for dialogue management : rational agency [4], Information State Update [5], functional models [6]. Only a very few number of these concepts have been transferred to industry. Since the late 90’s, the research has tackled the ambitious problem of automating the dialogue design [7], aiming at both reducing the development cost and optimising the dialogue efficiency and robustness. Recently, criticisms [8] have been formulated and novel approaches [9] have been proposed, both aiming at bridging the gap between research –focused on Markov-Decision-Process [10] based dialogue management– and industry –focused on dialogue design process, model, and tools. This paper contributes to extend this effort.

Regarding the dialogue evaluation topic, Paek [3] has pointed out that while research has exerted attention about “how best to evaluate a dialogue system?”, the industry has focused on “how best to design dialogue systems?”. This paper unifies those two approaches by merging system and design evaluation in a single graphical tool. We address all these convergence questions together as a way for research and industry to reach a technological breakthrough.

The second section is devoted to the presentation of the tools : the historical design tool, its adaptation to provide monitoring functionalities and the insertion of design alternatives. It is eventually concluded with an attempt to reassessing the dialogue evaluation. The third section describes the learning integration to the tool : the constraints we impose to the learning technique and the complementarity between the tools and the so-built learning capabilities. Finally, the last section concludes the paper.

2. The tools

Industry follows the VUI-completeness principle [1] : “*the behaviour of an application needs to be completely specified with respect to every possible situation that may arise during the interaction. No unpredictable user input should ever lead to unforeseeable behaviour*”. The SDS developers consider reliable the technologies, tools, and methodologies that help them to reach the VUI-completeness and to control it.

2.1. The Dialogue Design Tool

The graphical abstraction proposed by our dialogue design tool conforms to the general graph representation of finite state automata, with the difference that global and local variables en-

¹<http://www.w3c.org/TR/voicexml20/>

²In this paper, the term “SDS developers” denotes without any distinction VUI designers, application developers, and any industry engineers acting in SDS building.

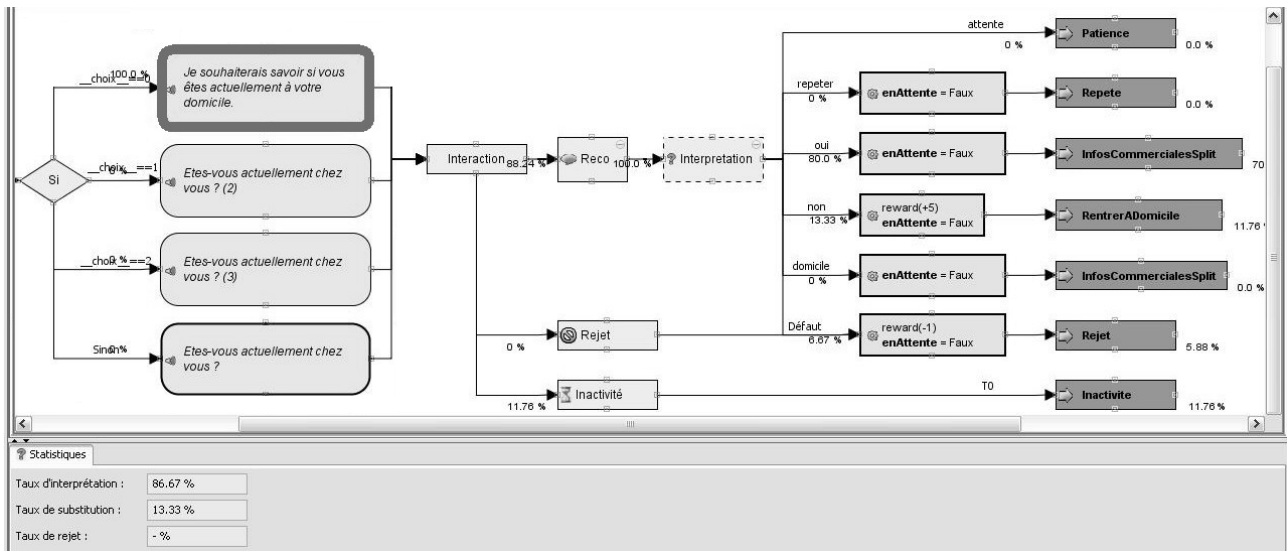


FIG. 1 – Some user experience feedbacks related to a selected prompt alternative.

able to factorise several dialogue states in a single node. Transitions relate to user inputs or to internal application events such as conditions based on internal information from the current dialogue state, from the back-end, or from the dialogue history. In that sense, dialogue design in the industry generally covers more than strict dialogue management, since its specification may indicate the type of spoken utterance expected from the user at each stage of the dialogue, up to the precise speech recognition model and parameter values to use, and the generation of the system utterance, from natural language generation to speech synthesis or audio recordings.

Our dialogue design tool offers to the SDS developers a graphical abstraction of the dialogue logic, sometimes also named the call flow. Thanks to a dynamic VoiceXML generation functionality, our dialogue design tool brings the SDS developers a complete guarantee that VUI-completeness at the design level automatically implies a similar completeness at the implementation level. During maintenance, If the SDS developers modify a specific part of the dialogue design, the tool guarantees that solely the corresponding code is impacted. This guarantee impacts positively VUI-completeness, reliability, and development cost.

2.2. Monitoring Functionalities inside the Design Tool

While researchers are focused on measuring the progress they incrementally reach, industry engineers have to deal with SDS tuning and upgrade. Their first dialogue evaluation criterion is task completion also called the automation rate because a SDS is deployed to automate specifically selected tasks. Most of the time, task completion is estimated thanks to Key Performance Indicators (KPI). KPI are difficult to exhaustively list and classify. Some are related to system measures, others are obtained thanks to dialogue annotations and the last ones are collected from users through questionnaires.

Some studies [11] investigated graphical monitoring tools to visualise a corpus of dialogue logs. Their tool aims at revealing how the system transits between its possible states. As a dialogue system is too complex to enumerate all its possible states, the dialogue logs are regarded as a set of variables that evolve during time and the tool proposes to make a projection

on a subset of these variables. This way, the generated graphs can either display the call flow, how the different steps are reached and where they lead, or display how different variables, as the number of errors evolve. This is mainly a tool for understanding how the users behave, because it does no connection with the way how the system was built. As consequence to this, it does not help to diagnose how to make it better. In other words, it does evaluate the system but does not meet one of our goal : the convergence between design and evaluation.

On the opposite, our graphical design tool provides an innovative functionality : local KPI projection on the original dialogue design thanks to an extensive logging. A large part of the KPI are automatically computed and displayed. As a consequence, it is possible to display percentage of which responses the system recognised, the users actually gave, and see how these numbers match the various KPI. It is one example among the numerous analysis views this graphical tool can provide.

2.3. Insertion of Alternatives

Figure 1 displays a monitoring view of an interaction implementation with alternatives (on the left). The recognition rate is the projected KPI on the graph at each branch. Other performance indicators are displayed at the bottom of the window : here, it is the actual rate of correct semantic decoding, the semantic substitution rate, and the semantic rejection rate. The selection of the highlighted box conditions the displayed logs.

Our design tool also provides a multivariate testing functionality. This method consists in testing multiple alternatives and selecting the best one on a fixed set of predetermined criteria. Regarding the VUI-completeness, presenting the complete automaton to the SDS developers is acceptable, as long as they can inspect and control every branch of the design. In general, they even come up with several competing designs or points of choice, which can only be properly validated in a statistical manner. The ability to compare all the dialogue design alternatives in the same test-field is a major factor to boost up SDS enhancement by drastically reducing the time needed. It increases the statistical relevance of the causal link between the tested alternatives and the differences in performance measures.

The KPI graphical projection into the dialogue design covers the dialogue alternatives : KPI computation just needs to be conditioned by the alternatives. Figure 1 illustrates the merge of several system prompt alternatives inside a single design : three speaking style versions of the same wording and one wording variant. The displayed KPI is conditioned by the selected alternative, here the first wording circled in bold grey. In some sense, embedding the user experience feedback inside the dialogue design forms a new material to touch and feel : the SDS developers can now sculpt a unique reactive material which contains the design and the KPI measures distribution. By looking at the influence of each alternative on the KPI when graphically selecting the alternatives, the SDS developers are given a reliable means to understand how to improve the system.

2.4. Reassessing Dialogue Evaluation

The traditional approaches to dialogue evaluation attempt to measure how best the SDS is adapted to the users. We remind that each interaction between the user and the SDS appears to be a unique performance. First, each new dialogue is co-built in a unique way according to both the very abilities of the user and the possibilities of the SDS. Second, the user adapts very quickly to new situations and accordingly changes her practices. The traditional approaches to dialogue evaluation are eventually based on the fragile reference frame of the user, not reliable enough for a scientific and an industrial approach.

This suggests for a shift in the reference frame used for dialogue evaluation : instead of trying to measure the adequacy between the SDS and the user in the user's reference frame, one can measure the adequacy between the user and the SDS in the design reference frame composed by the dialogue logic, the KPI and their expected values. Taking the design as the reference allows to reassess dialogue evaluation. The proposed basis for dialogue evaluation is reliable for the SDS developers because it is both stable and entirely under control. Deviations from the predicted situations are directly translated into anomalous values of measurable KPI that raise alerts. These automatically computable alerts warn the SDS developers about the presence of issues in their dialogue design.

3. Dialogue design learning

As presented in previous section, the alternative insertion is an enabler for the dialogue system analysis tools. It provides the SDS developers with a novel call flow visualisation experience. The further step to this approach is to automate at least a part of those analyses and improvements with learning capabilities.

3.1. Constraints

The objective is to automatically choose online the best alternative among those proposed in the design tool, and to report this choice to the SDS developers via the monitoring functionalities that are integrated to the design tool. This approach differs from the classical reinforcement learning methods used in the dialogue literature, which make their decisions at the dialogue turn level.

We use a technique from a previous work [12]. It does not need to declare the reachable states : they are automatically created when reached. This is also a parameter-free algorithm, which is very important when we consider that most dialogue application developers are not familiar with reinforcement learning theory. We keep the developer focussed on its main task. The two additional tasks required for the reinforcement learning

are to define the variable set on which the alternative choice should depend, and to implement a reward function based on the task completion, in order to get a fully automated optimisation with an online evaluation. The dialogue system automatic evaluation is a large problem that goes beyond the scope of this paper. However, sometimes, the dialogue application enables to have an explicit validation from the user. For instance, in an appointment scheduling application, the user is required to explicitly confirm the schedule he was proposed. This user performative act completes the task and provides a reliable automatic evaluation.

3.2. Learning and Monitoring Complementarity in the Design Optimisation

The learning algorithm and the SDS developers are two actors on the same object : the dialogue system. But, they work at a different time space. The learning algorithm updates its policy after each dialogue while the SDS developers monitor the system behaviour more occasionally. The same kind of opposition can be made on their action sets and their sight ranges. The learning algorithm can only change its policy among a limited amount of alternatives, while the SDS developers can make deeper changes, such as implementing a new dialogue branch, adding new alternatives, new alternative points, removing alternatives, etc... Their sight ranges vary a lot too. The learning algorithm is concentrated on the alternative sets and automatic evaluation and ignores the rest, while the SDS developers can apprehend the dialogue application as a whole, as a system or as a service. He can also have access to additional evaluations through annotations, or user subjective evaluations.

These functionality differences make their respective roles complementarity. The SDS developers have the responsibility for the whole application and the macro-strategic changes while the learning manager holds the real-time optimisation.

3.3. Control vs. Automation : the Trusting Threshold

As argued by Pieraccini and Huerta [1], finite state machine applied to dialogue management does not restrict the dialogue model to strictly directed dialogues. Finite state machines are easily extensible to powerful and flexible dialogue models. Our dialogue design tool offers various extensions : dialogue modules, hierarchical design, arbitrary function invocation at any point of the design, conditional statements to split the flow in different paths. All those extensions allow designing any topology of the finite state machine required to handle complex dialogue models like mixed-initiative interaction. Dialogue model is not the point where research and industry fail to converge.

The divergence point concerns the control aspect of VUI-completeness versus the automation of the dialogue design. As pointed out by recent works [8], MDP-based dialogue management aiming at automating the whole dialogue design is rejected by the SDS developers. Even more adaptive, it is seen as an uncontrollable black box sensitive to the tuning process. The SDS developers do not rely on systems that dynamically build their dialogue logic without a sufficient degree of monitoring and control.

Williams [9] has made a substantial effort to meet this industrial requirement. His system is a hybridisation of a conventional dialogue system following an industrial process, with a POMDP decision module, which is a MDP-based approach to dialogue management enhanced with dialogue state abstractions to model uncertainties. The responsibilities of each part of the system are shared as follows : the conventional system

elects several candidate dialogue moves and the POMDP decision module selects the most competitive one. This is a great step towards industry because the dialogue move chosen by the POMDP module has been first controlled by the conventional system design. Nevertheless, the so-built hybrid system is still not fully compliant with the industrial constraints for the following reasons.

First, contrarily to our approach, the SDS developer is called upon specific skills that cannot be demanded to a developer. As a consequence, the smoothness of the industrial process is deeply damaged.

Second, such a predictive module is not self-explanatory. Although the SDS developers have the control on the possible behaviour presented to the POMDP decision module, they are given no clue to understand how the choices are made. In fact, a learnt feature can never be exported to another context. At the opposite, our approach allows us to learn at the design level and consequently to report in the automaton the optimisation. The learning results are therefore understandable, analysable and generalisable.

4. Conclusion

4.1. A New Basis for Trusting Automatic Learning

This paper presents an original dialogue design tool that mixes dialogue design and dialogue evaluation in the same graphical interface. The design paradigm supported by the tool leads the SDS developers to predict value ranges of local KPI while designing the dialogue logic. It results a new evaluation paradigm using the system design as the reference and trying to measure deviations between the predicted and the measured values of the designed local KPI. The SDS developers rely on the tool to fulfil the VUI-completeness principle. Classically applied to dialogue design, the tool enables its application to the dialogue evaluation, leading to the comparison of dialogue design alternatives.

This places the SDS developers in a dialogue design improvement cycle close to the reinforcement learning decision process. Moreover, the inspectability offered by the user experience feedback functionality allows the SDS developers to understand, analyse and generalize all the decisions among the dialogue design alternatives. Combining the learning framework and the design tool guarantees the SDS developers keep control of the system. It preserves VUI-completeness and opens the way to a reliable learning based dialogue management.

4.2. Implementation

This approach to learning led us to deploy in October 2009 the first commercial spoken dialogue system with online learning. The system's task is to schedule an appointment between the customer and a technician. This service receives approximately 8k calls every month. At the time those lines are written, we are already in a virtuous circle of removing low-rated alternatives and replacing them with new ones, based on what the system learnt and what the designer understands from the data.

4.3. Future Work

On a social studies side, we are interested in collaborations to test advanced dialogue strategies and/or information presentation via generation. Indeed, we consider our system as a great tool for large scope experiments.

5. Acknowledgements

This research has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 216594 (CLASSIC project : www.classic-project.org).

6. References

- [1] R. Pieraccini and J. Huerta, "Where do we go from here ? research and commercial spoken dialog systems," in *Proceedings of the 6th SIGdial Workshop on Discourse and Dialogue*, L. Dybkjaer and W. Minker, Eds., 2005, pp. 1–10.
- [2] M. F. McTear, *Spoken Dialogue Technology : Toward the Conversational User Interface*. Springer, August 2004. [Online]. Available : <http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20&path=ASIN/1852336722>
- [3] T. Paek, "Toward evaluation that leads to best practices : Reconciling dialog evaluation in research and industry," in *Proceedings of the Workshop on Bridging the Gap : Academic and Industrial Research in Dialog Technologies*. Rochester, NY : Association for Computational Linguistics, April 2007, pp. 40–47. [Online]. Available : <http://www.aclweb.org/anthology/W/W07/W07-0306>
- [4] M. D. Sadek, P. Bretier, and F. Panaget, "Artimis : Natural dialogue meets rational agency," in *in Proceedings of IJCAI-97*. Morgan Kaufmann, 1997, pp. 1030–1035.
- [5] J. Bos, E. Klein, O. Lemon, and T. Oka, "Dipper : Description and formalisation of an information-state update dialogue system architecture," 2003. [Online]. Available : <http://citeseer.ist.psu.edu/bos03dipper.html>
- [6] R. Pieraccini, S. Caskey, K. Dayanidhi, B. Carpenter, and M. Phillips, "Etude, a recursive dialog manager with embedded user interface patterns," in *Automatic Speech Recognition and Understanding, 2001 IEEE Workshop on*, 2001, pp. 244–247.
- [7] O. Lemon and O. Pietquin, "Machine learning for spoken dialogue systems," in *Proceedings of the European Conference on Speech Communication and Technologies (Inter-speech'07)*, August 2007, pp. 2685–2688.
- [8] T. Paek and R. Pieraccini, "Automating spoken dialogue management design using machine learning : An industry perspective," *Speech Communication*, vol. 50, pp. 716–729, 2008.
- [9] J. D. Williams, "The best of both worlds : Unifying conventional dialog systems and POMDPs," in *International Conference on Speech and Language Processing*, 2008.
- [10] R. Bellman, "A markovian decision process," *Journal of Mathematics and Mechanics*, vol. 6, pp. 679–684, 1957.
- [11] A. Abella, J. Wright, and A. Gorin, "Dialog trajectory analysis," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, vol. 1, May 2004, pp. 441–444.
- [12] R. Laroche, G. Putois, P. Bretier, and B. Bouchon-Meunier, "Hybridisation of expertise and reinforcement learning in dialogue systems," in *Proceedings of Inter-speech. Special Session : Machine Learning for Adaptivity in Spoken Dialogue*, Brighton (United Kingdom), September 2009.