ADELE: Care and Companionship for Independent Aging

Brendan Spillane
ADAPT Centre
Trinity College Dublin
brendan.spillane@adaptcentre.ie

Emer Gilmartin
ADAPT Centre
Trinity College Dublin
gilmare@tcd.ie

Christian Saam
ADAPT Centre
Trinity College Dublin
saamc@cs.tcd.ie

Benjamin R. Cowan
University College Dublin
benjamin.cowan@ucd.ie

Vincent Wade
ADAPT Centre
Trinity College Dublin
vwade@adaptcentre.ie

Abstract

Dialogue system technology offers interesting prospects for services to seniors living independently. A spoken or text dialog system can support services such as monitoring medication adherence, fall prevention and reporting, exercise and wellbeing coaching, entertainment, companionship and the maintenance of social networks. Such systems would amalgamate several different types of conversation or speech-exchange systems, from well-defined tasks to engaging talk. Knowledge of how these types of talk function is essential to system design. The ADELE project aims to build an agent, ADELE, which can provide a range of services to seniors. Below, we outline plans for the system and describe challenges we anticipate in implementing the system.

1 Introduction

Spoken and text-based dialogue technology is the focus of increasing interest in the domains of geriatric care and coaching. These domains present interesting challenges, as they rely not only on the formulaic instrumental exchanges used in task-based systems, but also on an ability to perform human-like casual and social talk. Instrumental or task-based conversation is the medium for practical activities such as service encounters (shops, doctor's appointments), information transfer (lectures), or planning and execution of business (meetings). A large proportion of daily talk does not seem to contribute to a clear short-term task, but builds and maintains social bonds, and is described as 'interactional', social, or casual conversation. Early dialogue system researchers recognised the complexity of dealing with social talk (Allen et al., 2000), and initial prototypes concentrated on practical tasks such as travel bookings or logistics (Walker et al., 2001; Allen et al., 1995). Implementation of artificial task-based dialogues is facilitated by a number of factors. In these tasks, the lexical content of utterances drives successful completion of the task, conversation length is governed by task-completion, and participants are aware of the goals of the interaction. Such dialogues have been modelled as finite state and later slot-based systems, first using hand-written rules and later depending on data-driven stochastic methods to decide the next action. Task-based systems have proven invaluable in many practical domains. However, the creation of social companion application for healthcare or senior use entails the ability to wrap necessary tasks and recommendations in a matrix of social conversation. Although the aim of such agents is often described as conversational social companions, the ability for dialog systems to chat realistically lags their success in task based dialogue. Below, we describe a typical use case for our proposed system, ADELE. ADELE will be capable of monitoring medication, providing wellness advice and positive motivation, monitoring exercise and daily habits, storing and prompting general reminders, and engaging in companionable, social dialogue. We then outline relevant current approaches in dialogue systems and identify key challenges for companionable social talk, and highlight the challenges of supporting social talk interaction. We outline early progress on ADELE, and describe future work.

2 ADELE Use Case

The overall research goal of the ADELE project is to explore the use of personalisation in improving the efficacy of a digital companion that can communicate through informal, yet informed social dialogue, on a variety of topics of interest to a user over a prolonged time scale. A key point for the development of ADELE is that social spoken dialogue requires knowledge of the user to inform topic, style, and timing of conversation. This will be achieved by adapting the system persona and interaction based on the user's interests and profile to manage how and when the interactions occur, their content, and the means by which they are conveyed to the user to aid in comfortable, spoken delivery. The following is an example use case scenario between a future iteration of ADELE and Emma, a 77 year old woman.

ADELE: *"Emma, the next episode of that medical TV show you like should be on soon."*

Emma: *"Oh great, I'll check it out."*

ADELE: *"In the meantime, you have time to do your blood pressure daily check. I can then add the results to your file."*

Emma: *"Fine, let me just get the blood pressure monitor."*

ADELE: *"Your son Mark will also be here tomorrow morning at 10 to collect you for your doctors appointment."*

Emma: *"Oh great, I'd forgotten about that. It's his birthday next Thursday - can you remind me of that?"*

ADELE: *"Ok, would you like me to remind you to give him a call on the day?"*

Emma: *"Thanks ADELE, that would be great!"*

This interaction contains many elements which demonstrate what an everyday conversation entails, with topics flowing naturally. ADELE would use the time between the reminder of the TV show (a schedule reminder) and the actual start time to recommend to Emma to check her blood pressure, which must be done daily. This, like other user recommendations and their responses, could be marked as critical, resulting in regular reminders and/or notification of a third party. ADELE could also provide additional reminders such as a doctor's appointment. It could also reference events such as Mark's birthday, confirming with its calendar and asking if Emma wanted a reminder set. The use case scenario above raises a number of research challenges, which will be explored during the development of ADELE. There has been much progress in dialog technology and there is a body of work on the use of such systems in the elder care domain. Below, we briefly overview some existing systems.

3 Current Elder Care Systems and Applications

Academic work on companion applications and agents for the elderly is well established, and several commercial products have come to market. The work of Bickmore's Relational Agents Group, which includes several applications of hybrid social and task-based dialogue, is especially pertinent to the ADELE project. Their Senior Exercise Agent uses dialogue to encourage users to do more exercise, with moderate success with health literate adults (Bickmore et al., 2013). Their Virtual Nurse system takes patients through the transition from in-patient to discharge and aftercare through a dialogue. Users liked the system more than human doctors and nurses, citing the unrushed quality of interaction with the agent, and the possibility of re-checking every step without embarrassment. The group have also created agents which explain health documentation and provide counselling on a range of topics. Their early work on REA, a virtual estate agent which combined property viewing tasks with social talk, provided foundation research for hybrid task/social systems (Bickmore and Cassell, 2001). The Serroga system is an example of a non-speaking robot companion for domestic health care assistance which assists senior users in tasks from their day to day schedule and health care (Gross et al., 2015). The system emphasises social-emotional functions; in user trials participants accepted the non-speaking robot as a real social companion or relational agent. This was largely due to the successful establishment of co-presence. There is interesting recent work on multimodal systems providing basic care for the elderly and migrants (Wanner et al., 2016). Commercial examples of social care robots include ElliQ, Jibo, and GeriJoy[1]. The ADELE system will be based on dialog system and recommender system technology. Below, we briefly review dialog system design most relevant to our purposes and discuss challenges to the creation of casual talk.

4 Spoken Dialogue Systems

Dialogue systems have predominantly focussed on practical tasks. Classic dialogue systems are usu-

[1] www.elliq.com, www.jibo.com, and www.gerijoy.com

ally based on a division into several modules that handle the different problems of natural language dialogue (Jokinen and McTear, 2009). The Natural Language Understanding component converts input into an internal representation that can be reasoned on. The Dialogue Manager decides on a next action and supplies the Natural Language Generation with a specification of the next output which it converts to natural language. Dialogue management was initially based on handwritten rules, but this approach is severely limited in interactions other than simple question/answer sequences. More recently, research has concentrated on stochastic or machine learning methods, to better handle the uncertainty, noise, and variability inherent in spoken interaction. Stochastic dialogue systems have relied on the availability of large quantities of relevant data. While several dialogue corpora exist, these are generally collections of task-based dialogues, and not of casual talk (Serban et al., 2015a). Traditional approaches require this data to be labelled, adding significant cost to research as corpus annotation is a time and labour intensive undertaking. The ground truth provided by labelled data furnishes a good training signal to supervised learning for task-based interactions which tend to be short and relatively predictable in form and content. However, social talk can touch on virtually any topic and exhaustive annotation of corpora that capture realistic variation is prohibitively expensive. Therefore, approaches where labelling is not necessary are highly relevant. End-to-End systems can be trained directly from user input to system output without additional levels of training data annotation. This advantage comes at the price of requiring even greater amounts of training data. Sequence-to-Sequence (Sutskever et al., 2014) models are a popular Neural Network architecture which can achieve this goal. These models have proven effective in a variety of tasks. Early systems based on this approach essentially translated from a user turn to a system turn (Vinyals and Le, 2015). Some current versions use an additional level of hierarchy that allows the system to take into account longer histories by aggregating state information over previous turns (Serban et al., 2015b).

5 Challenges for Companionable Talk

Personalised systems are concerned with adapting and personalising services to individual users based on a range of models (Brusilovsky, 1998). Recent work on personalisation and dialogue agents includes a customer service agent (Verhagen et al., 2014), a social robot tutor (Gordona et al., 2015), and an eLearning agent (Peeters et al., 2016). Dialogue management in digital agents has long adopted personalisation strategies. Recent contributions include Ultes et al, who extended dialogue management to adapt to user satisfaction (Ultes et al., 2016), Litman and Pan, and San-Segundo et al. who developed methods to adapt the overall dialogue strategy based on the performance of the speech recognizer (Litman and Pan, 2002; San-Segundo et al., 2005), and Nothdurft et al. and Gnjatović and Rösner have developed approaches to adaptive dialogue (Nothdurft et al., 2012; Gnjatović and Rösner, 2008). An early approach to combine a personalised companion with adaptive dialogue was that of André and Rist who developed personalised information assistants for accessing information on the web (André and Rist, 2002). More recently, SARA was developed as a multifunctional conversational agent capable of personalised recommendations (Niculescu et al., 2014). To ensure a social and personalised interaction, the type and source of content that each user is recommended should be based on their history and content that they have explicitly expressed an interest in during speech interaction. This content may include news stories, conversational search terms or individuals from social feeds (Garcin et al., 2012). Each user also has preferred sources for such content. These may also include preferred social contacts such as close friends or family. A personalised intelligent companion should be capable of accessing and recommending content in a similar fashion. This is a complex process and includes topic extraction, weighting, mapping, longevity, context etc. The nuances of people's individual interests and preferences are difficult to model and require complex approaches to be accurately captured, if they are to allow personalised services to better meet the user's needs. Efforts have been made to include this ability in agents. Garcin et al. developed a personalised news recommendation system which demonstrated that collaborative filtering provided the best results for recommendations (Garcin et al., 2012). The content from these favoured sources could be used to select information to recommend when initiating or taking part in social talk. The time, location and context of delivery of interruptions are important in social agent interaction. A conversational social agent like ADELE needs to be able to instigate conversation which would involve interruption, but without excessively annoying or disturbing the user. It is also important to ensure that each interruption or conversation has a point and is delivered in a concise manner. Likewise, it is important that recommendations delivered within these interruptions are not overly repeated or become tiresome. The comparison by Bickmore et al. of strategies for interrupting the user to instigate exercise or take medication is highly relevant to this research (Bickmore et al., 2008). There are significant challenges around the time at which a social agent should interrupt a user's current task (to initiate conversation) as well as how to interleave with the user's utterances to provide an effective conversation. The urgency of the information or request to be delivered by the agent should also impact the agent's interruption strategy, especially in elderly care and social care contexts. Research is also needed on personal adaptation of this interruption pattern to suit individual tasks and user preferences. There are also significant challenges to be faced in the personalisation and recommendation of content based on granularity, word choice, comprehensibility etc. Much of the relevant work is in personalised eLearning where the delivery of content has received considerable attention (Daradoumis et al., 2013). It is also important to consider personality traits such as friendliness, chattiness and professionalism of ADELE, and the personality of the user. The importance of personality development for personal and virtual agents has been documented previously (Doce et al., 2010). Research has also highlighted the importance of matching a system's personality with that of its users (Lee and Nass, 2003). There are also several contributions on modelling personality traits, such as the framework of McQuiggan et al. for modelling empathy (McQuiggan et al., 2008). Further work is required to identify and understand how these traits may be perceived in social agent dialogue, the delivery of personalised recommendations, and the extent to which they need to be personalised to individuals. Personalised recommendation and dialogue. based on user personality has recently been investigated (Braunhofer et al., 2015; Vail and Boyer, 2014). ADELE's user model needs to account for physical and demographic attributes, usage preferences, context and temporal requirements, etc. It will also need to model additional attributes such as personality, conversational preferences, and critical care needs, such as medication requirements, and long term care needs such as memory monitoring. Many of these challenges and potential solutions, have been detailed by De Carolis et al. (Carolis et al., 2013). Current chat-oriented systems often exhibit a lack of variability in system output. One approach to counter this is the use of Reinforcement Learning to learn the production of utterances that are more beneficial to the long-term goal of the conversation (Li et al., 2016). A further step is to drive this learning by adversarial training which seeks to make system output indistinguishable from human-generated conversation (Li et al., 2017). Latent Variable models are another technique that can sample from a learnt stochastic variable to introduce randomness (Serban et al., 2017). An extension is to make this process controllable by making the distribution conditional on a variable (Sohn et al., 2015). This can facilitate the adjustment of a dimension that needs to be adapted such as friendliness. To work over longer contexts such as those found in casual interpersonal conversation which may lapse and restart over the course of hours or even days, memorization will need to be addressed. Research into Memory Networks has shown potential in the management of intra-session context (Bordes et al., 2016), and may prove applicable to the maintenance of longer contexts in ADELE. One of the goals of the ADELE system is to efficiently interleave different 'sub-dialogues' and sub-tasks - the system should be able to break off from story-telling or casual chat to perform a task such as medication checking and then return to the previous activity. Both the chat and task elements will vary between users depending on their care model and circumstances. A pertinent question is how to manage these different sub-dialogues to promptly intervene to perform subtasks? There has been some success with reinforcement learning in this context, which is being explored by the ADELE project (Yu et al., 2017).

6 Current Focus

Initially the ADELE project focused on investigating the greeting and leavetaking phases of an informal conversation. The project now aims to investigate how to generate the body of a friendly conversation (both interactive chat and longer chunks). To this end the project is focussing on topic shift and shading, the mechanisms which underpin the development of such conversations (Ries, 2001; Lambrecht, 1996). Consequently, it will be necessary for ADELE to be able to identify, strategise, render, and initiate topic shift and topic shading in a conversation. There are several reasons for this. Firstly, it will allow ADELE to change the course of a conversation. Secondly, it will enable ADELE to more easily and more naturally follow a conversation strategy, such as recommending a television show. Thirdly, it will enable ADELE to form more natural dialogue. The ADELE project is currently identifying and annotating examples of topic shift and topic shading in Switchboard (Godfrey et al., 1992), a corpus of 2,400 two-sided telephone conversations. A Wizard of Oz experiment based on a social care scenario is being conducted to generate additional social dialogues for training the Neural Network.

7 Conclusion

ADELE will be a virtual speech dialogue agent capable of informal, yet informed social dialogue. The importance of social dialogue to create social bonds cannot be underestimated to support independent living and companionship while providing a means for task reminders to promote beneficial self care. By treating the dialogue management as a personalisation task, the dialogue system can dynamically adapt to the users' preferences to enhance a more socially beneficial and comfortable conversational interaction. This will require an increased focus on personalisation strategies.

Acknowledgements

This research is supported by the Science Foundation Ireland (Grant 13/RC/2106) and the ADAPT Centre (www.adaptcentre.ie) at Trinity College, Dublin.

References

J. Allen, D. Byron, M. Dzikovska, G. Ferguson, L. Galescu, and A. Stent. 2000. An architecture for a generic dialogue shell. *Natural Language Engineering* 6, 3&4 (2000), 213—228.

James F. Allen, Lenhart K. Schubert, George Ferguson, Peter Heeman, Chung Hee Hwang, Tsuneaki Kato, Marc Light, Nathaniel Martin, Bradford Miller, Massimo Poesio, and David R. Traum. 1995. The TRAINS project: a case study in building a conversational planning agent. *Journal of Experimental & Theoretical Artificial Intelligence* 7, 1 (1995), 7–48.

Elisabeth André and Thomas Rist. 2002. From adaptive hypertext to personalized web companions. *Commun. ACM* 45, 5 (2002), 43—46.

Timothy Bickmore and Justine Cassell. 2001. Relational agents: a model and implementation of building user trust. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 396–403.

Timothy Bickmore, Daniel Mauer, Francisco Crespo, and Thomas Brown. 2008. *Negotiating Task Interruptions with Virtual Agents for Health Behavior Change*. International Foundation for Autonomous Agents and Multiagent Systems, 1241—1244. `http://dl.acm.org/citation.cfm?id=1402821.1402841`

Timothy W. Bickmore, Rebecca A. Silliman, Kerrie Nelson, Debbie M. Cheng, Michael Winter, Lori Henault, and Michael K. Paasche-Orlow. 2013. A Randomized Controlled Trial of an Automated Exercise Coach for Older Adults. *Journal of the American Geriatrics Society* 61, 10 (Oct 2013), 1676—1683. `DOI:http://dx.doi.org/10.1111/jgs.12449`

Antoine Bordes, Y.-Lan Boureau, and Jason Weston. 2016. Learning End-to-End Goal-Oriented Dialog. *arXiv:1605.07683 [cs]* (May 2016). `http://arxiv.org/abs/1605.07683` arXiv: 1605.07683.

Matthias Braunhofer, Mehdi Elahi, and Francesco Ricci. 2015. *User Personality and the New User Problem in a Context-Aware Point of Interest Recommender System*. Springer International Publishing, 537549. `DOI:http://dx.doi.org/10.1007/978-3-319-14343-9_39`

Peter Brusilovsky. 1998. Methods and techniques of adaptive hypermedia. In *Adaptive hypertext and hypermedia*. Springer, 1—43.

Berardina De Carolis, Irene Mazzotta, Nicole Novielli, and Sebastiano Pizzutilo. 2013. *User Modeling in Social Interaction with a Caring Agent*. Springer, London, 89116. `DOI:http://dx.doi.org/10.1007/978-1-4471-4778-7_4`

T. Daradoumis, R. Bassi, F. Xhafa, and S. Caballé. 2013. *A Review on Massive E-Learning (MOOC) Design, Delivery and Assessment*. 208—213. `DOI:http://dx.doi.org/10.1109/3PGCIC.2013.37`

Tiago Doce, Joäo Dias, Rui Prada, and Ana Paiva. 2010. *Creating Individual Agents through Personality Traits*. Springer Berlin Heidelberg, 257—264. `DOI:http://dx.doi.org/10.1007/978-3-642-15892-6_27` 00041.

F. Garcin, K. Zhou, B. Faltings, and V. Schickel. 2012. *Personalized News Recommendation Based on Collaborative Filtering*. Vol. 1. 437—441. DOI: http://dx.doi.org/10.1109/WI-IAT.2012.95

Milan Gnjatović and Dietmar Rösner. 2008. *Adaptive Dialogue Management in the NIMITEK Prototype System*. Springer, Berlin, Heidelberg, 14—25. DOI:http://dx.doi.org/10.1007/978-3-540-69369-7_3

John J Godfrey, Edward C Holliman, and Jane McDaniel. 1992. SWITCHBOARD: Telephone speech corpus for research and development. In *Acoustics, Speech, and Signal Processing, 1992. ICASSP-92., 1992 IEEE International Conference on*, Vol. 1. IEEE, 517—520.

Goren Gordona, Samuel Spauldinga, Jacqueline Kory Westlunda, Jin Joo Leea, Luke Plummera, Marayna Martineza, Madhurima Dasa, and Cynthia Breazeala. 2015. Affective Personalization of a Social Robot Tutor for Children's Second Language Skills. (2015).

H. M. Gross, S. Mueller, C. Schroeter, M. Volkhardt, A. Scheidig, K. Debes, K. Richter, and N. Doering. 2015. *Robot companion for domestic health assistance: Implementation, test and case study under everyday conditions in private apartments*. 5992—5999. DOI:http://dx.doi.org/10.1109/IROS.2015.7354230

Kristiina Jokinen and Michael McTear. 2009. Spoken Dialogue Systems. *Synthesis Lectures on Human Language Technologies* 2, 1 (2009). http://www.morganclaypool.com/doi/abs/10.2200/S00204ED1V01Y200910HLT005

Knud Lambrecht. 1996. *Information Structure and Sentence Form: Topic, Focus, and the Mental Representations of Discourse Referents*. Cambridge University Press. Google-Books-ID: bsXLCgAAQBAJ.

Kwan Min Lee and Clifford Nass. 2003. *Designing Social Presence of Social Actors in Human Computer Interaction*. ACM, 289—296. DOI:http://dx.doi.org/10.1145/642611.642662

Jiwei Li, Will Monroe, Alan Ritter, Dan Jurafsky, Michel Galley, and Jianfeng Gao. 2016. Deep Reinforcement Learning for Dialogue Generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, 1192–1202. https://aclweb.org/anthology/D16-1127

Jiwei Li, Will Monroe, Tianlin Shi, Sábastien Jean, Alan Ritter, and Dan Jurafsky. 2017. Adversarial Learning for Neural Dialogue Generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, 2157–2169. https://www.aclweb.org/anthology/D17-1230

Diane J. Litman and Shimei Pan. 2002. Designing and Evaluating an Adaptive Spoken Dialogue System. *User Modeling and User-Adapted Interaction* 12, 2—3 (Jun 2002), 111—137. DOI:http://dx.doi.org/10.1023/A:1015036910358

Scott W. McQuiggan, Jennifer L. Robison, Robert Phillips, and James C. Lester. 2008. *Modeling Parallel and Reactive Empathy in Virtual Agents: An Inductive Approach*. International Foundation for Autonomous Agents and Multiagent Systems, 167—174. http://dl.acm.org/citation.cfm?id=1402383.1402411

Andreea I. Niculescu, Ridong Jiang, Seokhwan Kim, Kheng Hui Yeo, Luis F. D'Haro, Arthur Niswar, and Rafael E. Banchs. 2014. SARA: Singapore's Automated Responsive Assistant, a multimodal dialogue system for touristic information. (2014).

Florian Nothdurft, Frank Honold, and Peter Kurzok. 2012. *Using Explanations for Runtime Dialogue Adaptation*. ACM, 63—64. DOI:http://dx.doi.org/10.1145/2388676.2388694

Marieke M.M. Peeters, Karel van den Bosch, John-Jules Ch. Meyer, and Mark A. Neerincx. 2016. *Agent-Based Personalisation and User Modeling for Personalised Educational Games*. ACM, 303—304. DOI:http://dx.doi.org/10.1145/2930238.2930273

Klaus Ries. 2001. *Segmenting Conversations by Topic, Initiative, and Style*. Springer, Berlin, Heidelberg, 51—66. DOI:http://dx.doi.org/10.1007/3-540-45637-6_5

Rubén San-Segundo, Juan M. Montero, Javier Macías-Guarasa, Javier Ferreiros, and José M. Pardo. 2005. Knowledge-combining methodology for dialogue design in spoken language systems. *International Journal of Speech Technology* 8, 1 (2005), 45—66.

Iulian Serban, Alessandro Sordoni, Ryan Lowe, Laurent Charlin, Joelle Pineau, Aaron Courville, and Yoshua Bengio. 2017. A Hierarchical Latent Variable Encoder-Decoder Model for Generating Dialogues. (2017). https://www.aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14567

Iulian Vlad Serban, Ryan Lowe, Peter Henderson, Laurent Charlin, and Joelle Pineau. 2015a. A Survey of Available Corpora for Building Data-Driven Dialogue Systems. *arXiv:1512.05742 [cs, stat]* (Dec. 2015). http://arxiv.org/abs/1512.05742 arXiv: 1512.05742.

Iulian Vlad Serban, Alessandro Sordoni, Yoshua Bengio, Aaron C. Courville, and Joelle Pineau. 2015b. Hierarchical Neural Network Generative Models for Movie Dialogues. *CoRR* abs/1507.04808 (2015). http://arxiv.org/abs/1507.04808

Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning Structured Output Representation using Deep Conditional Generative Models. In *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Eds.). Curran Associates, Inc., 3483–3491.

Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to Sequence Learning with Neural Networks. *CoRR* abs/1409.3215 (2014). `http://arxiv.org/abs/1409.3215`

Stefan Ultes, Hüseyin Dikme, and Wolfgang Minker. 2016. *Dialogue Management for User-Centered Adaptive Dialogue*. Springer, Cham, 51—61. `DOI:http://dx.doi.org/10.1007/978-3-319-21834-2_5`

Alexandria Vail and Kristy Boyer. 2014. *Adapting to Personality Over Time: Examining the Effectiveness of Dialogue Policy Progressions in Task-Oriented Interaction*. Association for Computational Linguistics, 4150. `http://www.aclweb.org/anthology/W14-4306`

Tibert Verhagen, Jaap van Nes, Frans Feldberg, and Willemijn van Dolen. 2014. Virtual Customer Service Agents: Using Social Presence and Personalization to Shape Online Service Encounters. *Journal of Computer-Mediated Communication* 19, 3 (Apr 2014), 529—545. `DOI:http://dx.doi.org/10.1111/jcc4.12066`

Oriol Vinyals and Quoc V. Le. 2015. A Neural Conversational Model. *CoRR* abs/1506.05869 (2015). `http://arxiv.org/abs/1506.05869`

Marilyn A. Walker, Rebecca Passonneau, and Julie E. Boland. 2001. Quantitative and Qualitative Evaluation of Darpa Communicator Spoken Dialogue Systems. In *Proceedings of the 39th Annual Meeting on Association for Computational Linguistics (ACL '01)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 515–522.

Leo Wanner, Josep Blat, Stamatia Dasiopoulou, Mónica Domínguez, Gerard Llorach, Simon Mille, Federico Sukno, Eleni Kamateri, Stefanos Vrochidis, Ioannis Kompatsiaris, and others. 2016. Towards a multimedia knowledge-based agent with social competence and human interaction capabilities. In *Proceedings of the 1st International Workshop on Multimedia Analysis and Retrieval for Multimodal Interaction*. ACM, 21–26.

Zhou Yu, Alexander Rudnicky, and Alan Black. 2017. Learning Conversational Systems that Interleave Task and Non-Task Content. (2017), 4214—4220. `https://www.ijcai.org/proceedings/2017/589`