

Pilot Data Collection Survey and Analytical Techniques for Persuasion Engineering Systems

Annye Braca^{*}, Brendan Spillane[†], Vincent Wade[†] and Pierpaolo Dondio^{*}
Technological University Dublin^{*}, Trinity College Dublin[†],
Ireland

Abstract

Persuasive communication is learned from childhood and can further be refined through study and experience. The ability to persuade through effective argumentation is regarded as a sign of intelligence, enticing others to form social bonds. Persuasive dialogue can be used to encourage interlocutors to take up self-care and beneficial habits or help them stop harmful routines or addictions. This paper details a pilot survey, and data analysis techniques, designed to inform the development of a larger follow on experiment. The aim is to predict the most effective form of argumentation based on user characteristics.

1 Introduction

Industry and academia have increasingly been studying how to influence attitudes and behaviours (Orji and Mandryk, 2014; Grimes et al., 2010; van Leer and Connor, 2012; Mutsuddi and Connelly, 2012). Persuasive experiences are rendered to users and consumers every day in the form of advertising, public relations and media strategies, and more recently, through the use of interactive technologies such as chat bots.

Effective persuasion is accomplished by understanding the bidirectional relationship between human psyche and persuasion (Hunter, 2018). Persuasion influences people’s attitudes and behaviours - therefore, human behaviour can be triggered and predicted (Fogg, 2009). Argumentation theory studies how conclusions can be reached through logical reasoning. It centres on relational support structures and evidence to make claims (Baroni et al., 2009). Persuasive argumentation concentrates on the type of language targeted at re-framing peoples’ opinions and beliefs (Besnard

and Hunter, 2008). It does this by focusing on: “*language cues aimed at shaping, reinforcing and changing people’s opinions and beliefs.*” (Yang et al., 2019).

Dialogue systems, whether speech or text based, are becoming increasingly ubiquitous in homes (Porcheron et al., 2017), health care (Laranjo et al., 2018; Richards and Caldwell, 2018; Spillane et al., 2017, 2018) and banking (Trivedi, 2019). Many websites now provide dialogue agents to help users find and purchase products and services (Majumder et al., 2018; Chattaraman et al., 2012). Users increasingly interact with their phones, search agents (Schalkwyk et al., 2010), entertainment systems (Bernhaupt et al., 2012), personal computers, smart home devices (Hoy, 2018) and other technologies through text based dialogue and voice. Persuasive dialogue systems have the potential to engage users in discussions on topics to help them explore different options and formulate their own thoughts and decide on a position or a course of action. To do this, they often gather data about a user.

This paper focuses on collecting and preparing a dataset of *Input Features* for a learning algorithm, containing features such as personality traits and decision making styles. This will enable us to predict which Principles of Persuasion, via user provided *Influence Scores*, are most effective for the given *Input Features*. The motivation behind this work, was to conduct a pilot test of this data collection process and analysis, and to demonstrate an initial version of the model which will be used to identify the most effective type of argument for users. The lessons learned from this pilot study will be incorporated into a larger follow on study.

The remainder of this paper is structured as follows. Section 2 highlights relevant work. Section 3 details the research aim and hypothesis. Section 4 provides the methodologies employed to design

the survey, collect and analyse the data, and build the models. Section 5 is the evaluation. Section 6 provides the discussion. Section 7 discusses future work and Section 8 concludes the paper.

2 Related Work

Recent years have seen several novel persuasive dialogue agents (Corrége et al., 2017; Schulman and Bickmore, 2009; Wells and Reed, 2008; Delecroix et al., 2013). The early work of Andrews focused on the impact of system personality on persuasion in human computer dialogue (Andrews, 2012). His system introduced a personality utterance generator designed to control the systems extraversion personality trait to increase the persuasiveness of the dialogue. He found that users' own level of extraversion influences the perception of the dialogue, particularly its trustworthiness and persuasiveness.

Rosenfeld and Kraus's contribution adopted a more technical approach (Rosenfeld and Kraus, 2016). They believed that many of the existing efforts in the domain: "*use unnatural assumptions regarding persuasive interaction, which creates doubt regarding their applicability for real world deployment with people.*" Their work focused on combining theoretical argumentation modelling (Baroni et al., 2011), machine learning (ML) and Markovian optimization techniques. Two studies demonstrated that it significantly improved on baseline systems and that it was no worse than humans at persuading other humans.

De Carolis and Mazzotta developed PORTIA, a computational model for persuasion (De Carolis and Mazzotta, 2017). It combined emotional and rational modes and was based on rational theories of persuasion (Miceli et al., 2006). It adapted to characteristics of the user to choose between rational strategies and emotional ones and has been tested in well-being and health related domains. It employed a User Model to implicitly learn characteristics of a user, including their personality traits and living habits. It also tried to learn users' goals to ascertain if there is a link between certain personality traits and user goals.

This pilot study and the subsequent larger follow on study is the next stage of evolution. Unlike previous work, it aims to identify the most persuasive forms of arguments based on users' characteristics such as their personality, demographics, and their belief systems and values. The eventual

aim is to mine online knowledge bases to automatically generate the most persuasive argument for different users.

3 Research Question

The overall aim of this pilot study and the eventual follow on study is to predict the most effective Principle of Persuasion argument via the combination of the users' *Input Features* and their corresponding *Influence Scores*. This will allow us to construct the syntactic structure of a persuasive argument for each user. The most effective syntax is decided by which Principles of Persuasion *Influence Scores* are the most effective for a user. They include the use of communication channels such as visual and kinaesthetic and psychological approaches such as authority and social proof. The aim is to discover if there is a link between user *Input Features* and these Principles of Persuasion *Influence Scores* with a view to predicting the most effective argument for persuading a given user. Long term, The aim is to automatically build such arguments from a knowledge base and other available resources.

3.1 Hypothesis

H_A Features based on Personality traits and Principles of Persuasion can be used to identify the most persuasive argument for a user

4 Methodology

To gather the *Input Features* required for our analysis, a survey was designed that captures: demographics, personality traits, decision making styles, and beliefs and values. To measure the perceived persuasiveness of the strategies present in the arguments, a star ratings scale methodology was adopted. This will suggest a guideline (Sparling and Sen, 2011) to better understand individual preferences and predict which argument is most effective for a given user i.e. the *influence scores*.

4.1 Survey: Input Features

4.1.1 Demographics

The first section of the survey had four questions which identify the users age, gender, education and religion.

4.1.2 Personality Traits

The interaction between personality traits and the effectiveness of personalized persuasive strategies

has been confirmed in marketing and advertising research (Gerber et al., 2013; Ho et al., 2008; Halko and Kientz, 2010; Hirsh et al., 2012; Chen and Lee, 2008). Subjects were asked to complete the Ten Item Personality Measure (TIPI) to measure their Big Five personality traits (Gosling et al., 2003). The TIPI asks respondents to report the extent to which: “I see myself as:” with a series of 10 personality trait pairs (e.g. Extroverted, Anxious, Sympathetic). This was measured using a seven-point Likert scale ranging from Disagree Strongly to Agree Strongly.

4.1.3 Beliefs and Values

It is appropriate to remark that personality traits do not present a comprehensive view of individual personalities (Gerber et al., 2013). People also vary in their attitudes, values, identities and other attributes that are often considered to be components of personality. Thus, elemental personality traits and beliefs are examined to understand how they shape thoughts and behaviour (McAdams, 1995). Individuals have different attitudes, values, identities and other attributes that are often considered to be elements of character. This section of the survey used the Life Value Inventory and had 44 questions (Crace and Brown, 1996).

4.1.4 Decision Making Styles

The Decision Scale Style (DSS) questionnaire was used to measure the decision-making style of participants (Hamilton et al., 2016). This aimed to determine patterns of habitual response in decision-making situations such as: rational or intuitive and central route or peripheral route thinking.

4.2 Survey: Principles of Persuasion

The second stage of the survey asked the user how influenced they were by 39 separate arguments. In his seminal work, Cialdini defined: *Reciprocity, Commitment and Consistency, Social Proof, Liking, Authority and Scarcity* as the six Principles of Persuasion. His work has proved influential in domains such as: marketing and advertising (Griskevicius et al., 2009; Sonnemans and Schilperoord, 2014), business (Hoy and Smith, 2007), negotiation (Guthrie, 2001; Bülow-Møller, 2005) and social psychology (Guadagno et al., 2013). More recently, they have been adopted in persuasive technologies (Kaptein et al., 2015; Oyibo et al., 2017; Josekutty Thomas et al., 2017). Argument variations are: ‘As a socially responsible mem-

ber of society, you should get vaccinated against disease as this benefits the community.’(Follow the rules - personality), ‘You should get vaccinated against disease as this benefits the community and the majority of people do so.’ (Follow the crowd - personality), ‘You wish to protect your own health? You wish to protect your family’s health? Then you should get vaccinated against disease.’(Closed questions eliciting Yes).

4.3 Argumentation schemes and Cialdini’s six Principles of Persuasion

Recently, efforts have been made to map Cialdini’s six Principles of Persuasion to argumentation schemes (Thomas et al., 2018b,a, 2019). However, this is still early work and has yet to be sufficiently tested or adopted. Consequently, it was decided to use Walton’s argumentation scheme (Walton, 2007). According to Walton, an audience can be: “*persuaded by means of arguments that are (at least typically) not syllogisms, but kinds of arguments they are familiar with in everyday thinking and discourse.*” (Walton, 2007, p. 23). He further states: “*An audience will be more effectively persuaded by arguments that they think are reasonable*”. De Carolis and Mazzotta, citing O’Keefe, differentiate between argumentation and persuasion. They maintain that: “*argumentation means to induce a belief, persuasion means to induce an intention to do something*” (De Carolis and Mazzotta, 2017; O’Keefe, 2015). The data resulting from the 39 arguments in the survey on the Principles of Persuasion will be mapped into Walton’s argumentation schema (Walton, 2007). This will then be combined with the users *Input Features* in a ML model to determine the most effective type of argument to persuade each user.

4.4 Data: Collection

The survey was deployed using Qualtrics¹ in early July 2019. Participants were recruited via the crowd-sourcing marketplace Prolific Academic², and paid £6.00 for their time. The survey took approximately fifteen minutes to complete. Three attention questions were included in the survey to ensure that participants were engaged. Participants were provided with instructions requesting them to answer all questions as honestly as possible. They were also given the instruction: “*Please*

¹www.qualtrics.com

²www.prolific.ac

answer each of the questions as you feel right now”.

In total, 30 participants completed the experiment. 14 submissions were not included in the analysis either because they failed one of the three attention question or because they were missing significant data. This could be because the participant did not provide answers or because there was a problem with the survey tool. Thus, the final survey dataset consisted of N=16 participant submissions. The data was largely homogeneous with no extreme outliers. A visual inspection revealed no abnormal or vexatious submissions. The survey collected the user’s profile, presented them with argument variants (differing persuasion structure) and recorded their inputted *Influence Score*. The design matrix consisted of two main parts: the independent *Input Features* (user profile data) and the dependent variable (argument *Influence Score*). The topics were: arguments about visiting the dentist, avoiding smoking, changing toothbrushes, vaccination and going to bed early.

4.5 Data: Preprocessing

Before fitting the ML model, two important processes were performed. The first was to reverse score calculations on two of the survey elements - *Personality Traits* and *Beliefs and Values*. Certain questions are natural pairs such as a pair of questions relating to extroversion and introversion. Given the opposite nature of the questions, reverse scoring on one of the questions allows differentiation of the two attributes. As all other data in the survey was in the form of categorical variables and some ML algorithms cannot handle text data, one hot-encoding (Pedregosa et al., 2011) was used to convert the data to a numerical form. This process is a data transformation that yields a new design matrix with new columns that correspond to one of the unique values of each original feature. The final design matrix had 112 input variables.

4.6 The Persuasion Model

The purpose of the model is to predict the user *Influence Scores* for each argument variant. This would allow the selection of the most effective arguments to persuade a given user. To predict the *Influence Scores*, the model was built at the argument level. The design matrix consisted of a range of x variables i.e. the *Input Features* and a single y variable(*Influence Scores*). The user rated the

39 arguments on a scale of 1-10: 1-3 (low), 3-6 (medium), 7-10 (high confidence) as shown in Fig 1. A multinomial logistic regression ML model (Pedregosa et al., 2011) was used in the classification of statement *Influence Scores*. The three class logistic regression model was fit to 16 observations for each of the 39 arguments i.e. 39 models. A training and testing split of 75%:25% (12:4) observations was implemented. The classifier (logistic regression) was applied on each of the 39 user/argument datasets. A later model with more respondents would attempt to classify *Influence Score* over 10 classes (1-10) instead of 3. In place of a multinomial logistic regression, other algorithms such as deep learning neural networks and random forests would be investigated.




Figure 1: Class distribution

5 Evaluation

The dataset collected from this study does not have enough samples to identify relationships between *Input Features* and responsiveness to persuasive arguments (*Influence Scores*). However, this was expected. The purpose of this study was to pilot test data collection and analytical techniques for a larger subsequent study. The elements tested included: the amount, type and format of questions, the survey instrument, crowd-sourcing and data cleaning and processing strategies. The analytical techniques pilot tested include: Multinomial Logistic Regression and Principle Components Analysis. To test these techniques on the data collected, the analysis was performed on the limited dataset without expectation of robust and definitive findings. However, the knowledge gained from running them is invaluable to a planned later experiment. Thus, the following results are presented as pilot test results. When built on the complete

set of 102 input features the model yielded an average accuracy of 43.91% over the 39 argument models. A baseline model which guesses the user responses would expect to return an accuracy of 33%, given that the problem is a trinary classification task. Principal Components Analysis (PCA) was investigated as a feature selection approach in the hope of improving the discriminative power of the model. After transforming the original data into the space of the principal components and fitting the model with this new dataset, it was found that 90% of the variance was explained by the first 8 principal components. Fig 2. shows a plot of the 1st and 2nd principal components (37.5% and 15.6% of variance explained respectively). There is a suggestion of possible user clusters by statement score but more data would be necessary to investigate this robustly. The average accuracy did not improve at 43.24% using PCA. A robust model would require a considerably larger dataset which is planned in the follow on experiment.




Figure 2: Cluster at the argument level

6 Discussion

Research into persuasive dialog systems is important as it can be used to increase the acceptance of important information such as health advice to improve people's lives. The persuasion model at its core is a statistical pattern recognition system. Over the course of the research, user attributes were discovered that are useful in predicting attitudes and behaviour, in addition to designing argument variations that should map to these attributes. The size of the initial dataset is insufficient for an effective pattern classification system. Predictive models are only as good as the data (volume and quality) that they are trained on. Clearly, a larger

dataset is needed to reasonably capture the relationships that may exist between *Input Features* and *Influence Scores*. The design and deployment of this larger experiment is ongoing. It is uncertain how large a dataset will be required for training to achieve a high level of learning accuracy. We will address this problem through empirical investigation. A potential weak spot is domain expertise on persuasion. A domain expert for designing the syntax of the arguments under the different persuasive techniques would be of great assistance.

7 Future Work

Future work includes the procurement of a larger dataset which would greatly support robust validation. The aim is to further investigate the formation of pattern classes, this would involve deep investigation of the principal components of the dataset. It is our intention to design and deploy an experiment aimed at examining the effects of the predictive power of the persuasion engineering model. To do so, a web based persuasive agent will be used to obtain A/B testing statistics. Two versions of the system will be tested, one that uses the persuasion engineering model and a control system that does not. Participants will be recruited to interact with the pilot system and measure the persuasiveness of the agent. In both systems, the user will be prompted with questions about their own opinions regarding the topics in order to collect more arguments. Attitudes towards the topics will be measured pre-experiment and post-experiment: A Likert format instrument application form can serve as the attitudinal measure of persuasiveness. This will help to determine significant effects between the treatment and control groups.

8 Conclusion

The purpose of this paper was to pilot test each stage of a larger, more detailed experiment. As such it proved invaluable in terms of the lessons learned. This includes how to formulate a survey, including methods of data collection. It also included how to crowd source participants. Once data was gathered, several important lessons were learned on data analysis techniques. We outline the theoretical and methodological approach to the design of a persuasion based argumentation system that could be used to potentially deliver enhanced communication of information to users.

An example of such a system might be a health bot that helps users to understand the importance of certain beneficial practices e.g. regular dental checks. By speaking the user's language, such a system could help to improve their quality of life. Potentially, a persuasion system can learn how to be more persuasive by selecting optimal arguments given the users profile. To this end, it is important that data is of sufficient volume and quality to allow the discovery of deep relationships between user attributes and behaviour. The study of the cuing effect in argument syntax needs special attention and the model would benefit from consultation with persuasion psychology experts. On these grounds, it is hoped that future research will be better placed to explore the relationships among user characteristics, argument quality and attitude change and formation.

9 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 (www.tcd.ie) and Technological University Dublin (www.tudublin.ie).

References

- Pierre Y. Andrews. 2012. [System personality and persuasion in human-computer dialogue](#). *ACM Transactions on Interactive Intelligent Systems*, 2(2):1–27.
- Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. 2011. An introduction to argumentation semantics. *The knowledge engineering review*, 26(4):365–410.
- Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Giovanni Guida. 2009. Encompassing attacks to attacks in abstract argumentation frameworks. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 83–94. Springer.
- Regina Bernhaupt, Mael Boutonnet, Bastien Gattelier, Yannik Gimenez, Christian Pouchepanadin, and Latifa Souiba. 2012. A set of recommendations for the control of iptv-systems via smart phones based on the understanding of users practices and needs. In *Proceedings of the 10th European conference on Interactive tv and video*, pages 143–152. ACM.
- Philippe Besnard and Anthony Hunter. 2008. *Elements of argumentation*, volume 47. MIT press Cambridge.
- Anne Marie Bülow-Møller. 2005. *Persuasion in Business Negotiations*. John Benjamins Publishing. Google-Books-ID: 6Zs6AAAAQBAJ.
- Veena Chattaraman, Wi-Suk Kwon, and Juan E Gilbert. 2012. Virtual agents in retail web sites: Benefits of simulated social interaction for older users. *Computers in Human Behavior*, 28(6):2055–2066.
- Shu-Hui Chen and Kuan-Ping Lee. 2008. The role of personality traits and perceived values in persuasion: An elaboration likelihood model perspective on online shopping. *Social Behavior and Personality: an international journal*, 36(10):1379–1399.
- Jean-Baptiste Corrége, Céline Clavel, Nicolas Sabouret, Emmanuel Hadoux, Anthony Hunter, and Mehdi Ammi. 2017. Persuasive dialogue system for energy conservation. In *Adjunct Proceedings of the 12th International Conference Persuasive Technology: Development and Implementation of Personalized Technologies to Change Attitudes and Behaviours (PERSUASIVE 2017)*, Amsterdam, The Netherlands, pages 4–6.
- R Kelly Crace and Duane Brown. 1996. Life values inventory. *Minneapolis, MN: National Computer Systems*.
- Berardina De Carolis and Irene Mazzotta. 2017. A user-adaptive persuasive system based on 'arational' theory. *International Journal of Human-Computer Studies*, 108:70–88.
- Fabien Delecroix, Maxime Morge, and Jean-Christophe Routier. 2013. A virtual selling agent which is persuasive and adaptive. In *Agreement Technologies*, pages 625–645. Springer.
- Bj Fogg. 2009. [A behavior model for persuasive design](#). In *Proceedings of the 4th International Conference on Persuasive Technology - Persuasive '09*, page 1. ACM Press.
- Alan S Gerber, Gregory A Huber, David Doherty, Conor M Dowling, and Costas Panagopoulos. 2013. Big five personality traits and responses to persuasive appeals: Results from voter turnout experiments. *Political Behavior*, 35(4):687–728.
- Samuel D Gosling, Peter J Rentfrow, and William B Swann Jr. 2003. A very brief measure of the big-five personality domains. *Journal of Research in personality*, 37(6):504–528.
- Andrea Grimes, Vasudhara Kantroo, and Rebecca E Grinter. 2010. Let's play!: mobile health games for adults. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, pages 241–250. ACM.
- Vladas Griskevicius, Noah J Goldstein, Chad R Mortensen, Jill M Sundie, Robert B Cialdini, and Douglas T Kenrick. 2009. [Fear and loving in las vegas: Evolution, emotion, and persuasion](#). *Journal of Marketing Research (JMR)*, 46(3):384–395.

- Rosanna E. Guadagno, Nicole L. Muscanell, Lindsay M. Rice, and Nicole Roberts. 2013. [Social influence online: The impact of social validation and likability on compliance](#). *Psychology of Popular Media Culture*, 2(1):51–60.
- Chris Guthrie. 2001. Principles of influence in negotiation. *Marquette Law Review*, page 20.
- Sajaneer Halko and Julie A Kientz. 2010. Personality and persuasive technology: an exploratory study on health-promoting mobile applications. In *International conference on persuasive technology*, pages 150–161. Springer.
- Katherine Hamilton, Shin-I Shih, and Susan Mohammed. 2016. The development and validation of the rational and intuitive decision styles scale. *Journal of personality assessment*, 98(5):523–535.
- Jacob B Hirsh, Sonia K Kang, and Galen V Bodenhausen. 2012. Personalized persuasion: Tailoring persuasive appeals to recipients’ personality traits. *Psychological science*, 23(6):578–581.
- Shuk Ying Ho, Michael J Davern, and Kar Yan Tam. 2008. Personalization and choice behavior: the role of personality traits. *ACM SIGMIS Database: the DATABASE for Advances in Information Systems*, 39(4):31–47.
- Matthew B Hoy. 2018. Alexa, siri, cortana, and more: an introduction to voice assistants. *Medical reference services quarterly*, 37(1):81–88.
- Wayne K. Hoy and Page A. Smith. 2007. [Influence: a key to successful leadership](#). *International Journal of Educational Management*, 21(2):158–167.
- Anthony Hunter. 2018. [Towards a framework for computational persuasion with applications in behaviour change](#). *Argument Computation*, 9(1):15–40.
- Rosemary Josekutty Thomas, Judith Masthoff, and Nir Oren. 2017. [Personalising healthy eating messages to age, gender and personality: Using cialdini’s principles and framing](#). In *Proceedings of the 22nd International Conference on Intelligent User Interfaces Companion*, IUI ’17 Companion, pages 81–84. ACM. Event-place: Limassol, Cyprus.
- Maurits Kaptein, Panos Markopoulos, Boris de Ruyter, and Emile Aarts. 2015. [Personalizing persuasive technologies: Explicit and implicit personalization using persuasion profiles](#). *International Journal of Human-Computer Studies*, 77:38–51.
- Liliana Laranjo, Adam G. Dunn, Huong Ly Tong, Ahmet Baki Kocaballi, Jessica Chen, Rabia Bashir, Didi Surian, Blanca Gallego, Farah Magrabi, Annie Y. S. Lau, and et al. 2018. [Conversational agents in healthcare: a systematic review](#). *Journal of the American Medical Informatics Association*, 25(9):1248–1258.
- Eva van Leer and Nadine P Connor. 2012. Use of portable digital media players increases patient motivation and practice in voice therapy. *Journal of Voice*, 26(4):447–453.
- Anirban Majumder, Abhay Pande, Kondalarao Vonturu, Abhishek Gangwar, Subhadeep Maji, Pankaj Bhatia, and Pawan Goyal. 2018. Automated assistance in e-commerce: An approach based on category-sensitive retrieval. In *European Conference on Information Retrieval*, pages 604–610. Springer.
- Dan P McAdams. 1995. What do we know when we know a person? *Journal of personality*, 63(3):365–396.
- Maria Miceli, Fiorella de Rosis, and Isabella Poggi. 2006. [Emotional and non-emotional persuasion](#). *Applied Artificial Intelligence*, 20(10):849–879.
- Aditya U Mutsuddi and Kay Connelly. 2012. Text messages for encouraging physical activity are they effective after the novelty effect wears off? In *2012 6th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth) and Workshops*, pages 33–40. IEEE.
- Daniel J. O’Keefe. 2015. *Persuasion: Theory and Research*. SAGE Publications.
- Rita Orji and Regan L Mandryk. 2014. Developing culturally relevant design guidelines for encouraging healthy eating behavior. *International Journal of Human-Computer Studies*, 72(2):207–223.
- Kiemute Oyibo, Rita Orji, and Julita Vassileva. 2017. Investigation of the influence of personality traits on cialdini’s persuasive strategies. page 13.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Martin Porcheron, Joel E. Fischer, Stuart Reeves, and Sarah Sharples. 2017. [Voice interfaces in everyday life](#). In *Proceedings of the 2018 ACM Conference on Human Factors in Computing Systems (CHI ’18)*. ACM.
- D. Richards and P. Caldwell. 2018. [Improving health outcomes sooner rather than later via an interactive website and virtual specialist](#). *IEEE Journal of Biomedical and Health Informatics*, 22(5):1699–1706.
- Ariel Rosenfeld and Sarit Kraus. 2016. [Strategical argumentative agent for human persuasion](#). In *Proceedings of the Twenty-second European Conference on Artificial Intelligence, ECAI’16*, pages 320–328. IOS Press. Event-place: The Hague, The Netherlands.

- Johan Schalkwyk, Doug Beeferman, Françoise Beaufays, Bill Byrne, Ciprian Chelba, Mike Cohen, Maryam Kamvar, and Brian Strope. 2010. "your word is my command": Google search by voice: A case study. In *Advances in speech recognition*, pages 61–90. Springer.
- Daniel Schulman and Timothy Bickmore. 2009. Persuading users through counseling dialogue with a conversational agent. In *Proceedings of the 4th international conference on persuasive technology*, page 25. ACM.
- Merel Sonnemans and J Schilperoord. 2014. The influence of previous behavior and personality traits on the effect of personalized marketing communication for non-profit purposes.
- E Isaac Sparling and Shilad Sen. 2011. Rating: how difficult is it? In *Proceedings of the fifth ACM conference on Recommender systems*, pages 149–156. ACM.
- Brendan Spillane, Emer Gilmartin, Christian Saam, Benjamin R Cowan, and Vincent Wade. 2018. Adele: Care and companionship for independent aging. In *ICAHGCA@ AAMAS*, pages 18–24.
- Brendan Spillane, Emer Gilmartin, Christian Saam, Ketong Su, Benjamin R Cowan, Séamus Lawless, and Vincent Wade. 2017. Introducing adele: a personalized intelligent companion. In *Proceedings of the 1st ACM SIGCHI International Workshop on Investigating Social Interactions with Artificial Agents*, pages 43–44. ACM.
- Rosemary J Thomas, Matthew Collinson, and Judith Masthoff. 2018a. Caught by phishing emails? how can argumentation schemes be used to protect users? page 4.
- Rosemary J. Thomas, Judith Masthoff, and Nir Oren. 2019. Is argumessage effective? a critical evaluation of the persuasive message generation system. In *Persuasive Technology: Development of Persuasive and Behavior Change Support Systems*, Lecture Notes in Computer Science, pages 87–99. Springer International Publishing.
- Rosemary J Thomas, Nir Oren, and Judith Masthoff. 2018b. Argumessage: A system for automation of message generation using argumentation schemes. page 5.
- Jay Trivedi. 2019. Examining the customer experience of using banking chatbots and its impact on brand love: The moderating role of perceived risk. *Journal of Internet Commerce*, 18(1):91–111.
- Douglas N. Walton. 2007. *Media argumentation: dialectic, persuasion, and rhetoric*. Cambridge University Press.
- Simon Wells and Chris Reed. 2008. Magtalo: an agent-based system for persuasive online interaction. In *AISB 2008 Convention Communication, Interaction and Social Intelligence*, volume 1, page 29. Citeseer.
- Diyi Yang, Jiaao Chen, Zichao Yang, Dan Jurafsky, and Eduard Hovy. 2019. Let's make your request more persuasive: Modeling persuasive strategies via semi-supervised neural nets on crowdfunding platforms. page 11.