

How are you? Introducing stress-based text tailoring

Simone Balloccu, Ehud Reiter

University Of Aberdeen
United Kingdom

simone.balloccu@abdn.ac.uk
e.reiter@abdn.ac.uk

Alexandra Johnstone, Claire Fyfe

Rowett Institute, University of Aberdeen
United Kingdom

alex.johnstone@abdn.ac.uk
c.fyfe@abdn.ac.uk

Abstract

Can stress affect not only your life but also how you read and interpret a text? Healthcare has shown evidence of such dynamics and in this short paper we discuss customising texts based on user stress level, as it could represent a critical factor when it comes to user engagement and behavioural change. We first show a real-world example in which user behaviour is influenced by stress, then, after discussing which tools can be employed to assess and measure it, we propose an initial method for tailoring the document by exploiting complexity reduction and affect enforcement. The result is a short and encouraging text which requires less commitment to be read and understood. We believe this work in progress can raise some interesting questions on a topic that is often overlooked in NLG.

1 Introduction

Healthcare can benefit greatly from NLG as communication plays a big role in therapy outcome (Stewart, 1995) and electronic documentation is useful (Ross et al., 2016) but heavy and time-consuming to be produced (Arndt et al., 2017). NLG can speed up the process and there is evidence of increasing adoption of data-to-text in healthcare (Pauws et al., 2019). Despite NLG being valid for producing standardised documentation, the situation is more complex when the targeted reader is a patient: people are all different and a single text is often not enough. Healthcare-NLG often addresses the problem by adopting user tailoring, which is the practice of generating a text customised to the user traits and preferences (Kukafka, 2005). In this work, we focus on user stress, a factor which is often not considered in NLG, as a feature for text-tailoring. We show that stress does matter when reading text and provide a real-world example (the dietary domain) along with some initial discussion and examples on how stress-based tailoring

could be carried out. Motivated by the PhilHumans project¹ and its goal of improving the interactions between personal health apps and users, we explore the benefits of personalising generated texts based on how much stressed the user is. Stress impacts people’s health, behaviour, receptivity to advice, and the ability to absorb complex information, so basing tailoring on this element could help health software support patients better. The structure of the paper is the following: in Section 2 we recap the state of user tailoring in NLG, emphasizing that no known system does properly address user stress; in Section 3-3.1 we provide a real-world example in the dietary domain, showing the results of statistical analysis that, in line with previous healthcare studies, hints at the correlation between stress and eating behaviour; Section 4.1 discusses possible methodologies to assess user stress, prioritising the non-intrusive ones; in Section 4.2 we propose two theoretical ways to tailor the text for a stressed user, along with an example; Section 5 closes the paper with our final considerations.

2 Related Work

Natural Language Generation has been used in the healthcare domain in a variety of ways, with different aims. Since the spread of e-health, various works have focused on automating the creation of documents. Some examples of these are clinical encounters (Finley et al., 2018), chief complaints (Lee, 2018) and handover reports (Schneider et al., 2013). There has also been some research on producing multiple texts from the same source, as in the BabyTalk project (Portet et al., 2008) which generated premature infants reports for doctors (Portet et al., 2009), nurses (Hunter et al., 2011) and parents (Mahamood and Reiter, 2011) in Neonatal Intensive Care Unit (NICU) environ-

¹PhilHumans Project page: <https://www.philhumans.eu/>

ment. Another scenario in which NLG contributed to healthcare is decision support (Binsted et al., 1995) (Hommes et al., 2019), where the text can help patients understand their clinical situation, and prepare them for shared decision making (Seminar, 2011) (Stiggelbout et al., 2012). The majority of these systems don't need any personal information about the reader because their final purpose is purely informative. It is, however, different when we move to behavioural change, as the text aims to trigger a lifestyle improvement in the reader. In this case we're dealing with persuasive communication, which is the act of communicating to make the reader perform certain actions or collaborate in various activities (Guerini et al., 2011). Persuasive communication systems can be distinguished in vertical and horizontal approaches (Maimone et al., 2018): vertical systems tends to be specialised, hence they are very effective but harder to move between different domains, while horizontal systems are often conceptual and focused on domain-independent knowledge and strategies. When persuasive communication has to be obtained through NLG, a popular approach is to profile the user to model the text accordingly. This practice is known as text tailoring (Kukafka, 2005) and has shown promising results (Kreuter and Wray, 2003). Moreover NLG tailoring is particularly helpful, since manual text customisation is often time-consuming (Pauws et al., 2019). NLG-driven tailoring was tested for lifestyle improvements and behavioural change in various domains. Vertical approaches includes solutions for smoking cessation (Reiter et al., 2003), driving conduct improvement (Braun et al., 2018), diet management (Anselma et al., 2018)(Anselma and Mazzei, 2018)(Anselma and Mazzei, 2017) and therapy recommendation (Skinner et al., 1994). Horizontal approaches deals with more general topics, like the key challenges and model design for behaviour change (Oinas-Kukkonen, 2013)(Kelders et al., 2016)(Oinas-Kukkonen and Harjumaa, 2009), argumentative persuasive communication (Zukerman et al., 2000) (Reed et al., 1996) or the inclusion of generic psychological traits and basic human values into the model (Ding and Pan, 2016). Finally, there have been some attempts to partially merge horizontal and vertical models (Maimone et al., 2018)(Dragoni et al., 2018)(Donadello et al., 2019). When the text must be built according to user traits, it is important to identify the variables which could

influence not only the behaviour but also user engagement. These include general psychological tailoring techniques (Hawkins et al., 2008); user numeracy, literacy (Williams and Reiter, 2008) and domain knowledge (McKeown et al., 1993)(Paris, 1988); affect (Mahamood and Reiter, 2011). User stress is a factor which, to our knowledge, is largely ignored in the existing literature. Some NLG works have considered stress, but not in a dynamic way. Some works in tactical NLG (text which tries to induce a certain emotional status) (Van Der Sluis and Mellish, 2008) considered stress, but only to assess the text effect. This makes sense since tactical NLG aims to induce an emotional status instead of detecting or managing it. BabyTalk-Family project (Mahamood and Reiter, 2011) tried to actively tailor the document by estimating stress value, but without detecting it from the user itself. Again, the choice is reasonable since having children in NICU can be a traumatic experience, therefore directly asking questions about stress could be intrusive and dangerous. Beside these boundary cases we argue that stress-based tailoring is of primary importance as stress has been shown to significantly impact individual reading skills (Rai et al., 2015)(Peng et al., 2018), potentially causing temporary difficulty in understanding the given text. Therefore stress data can be exploited to guide the tailoring process with precious hints on how to re-realise the text. This motivates researching how user stress assessment can be done in the least intrusive way.

3 A real world example: NeuroFAST

Before going into the details of how stress-based tailoring could be achieved, we want to introduce some preliminary evidence that stress does matter and can influence user behaviour. We exploit a dataset which is derived from the EU-funded project NeuroFAST², regarding the socio-psychological forces that could influence eating behaviour. Our dataset (Malone et al., 2015) contains a detailed food diary for 413 different workers, together with a "Daily Hassles" questionnaire in which the participants noted stressful events. Each entry is scored in a range between 0 and 4, with a higher score indicating a higher stress level. For example, we can find an individual that put "Caught in a traffic jam on the way to work" as a daily hassle. This resulted in a score of 2/4. We analysed the

²NeuroFAST project
<https://cordis.europa.eu/project/id/245009>

dataset to test the following hypothesis: does stress influence workers’ diet? This is not a novelty in healthcare, as different studies linked stress to eating disorders (Scott and Johnstone, 2012)(Torres and Nowson, 2007)(Puddephatt et al., 2020)(Riffer et al., 2019) but we still provide an additional, post hoc, analysis on said data to show a concrete example. Additionally, even if the original study was not meant to inspect such correlation, both data were readily available, making it possible to inspect it.

3.1 NeuroFAST stress analysis

To test our hypotheses we extracted the data, isolating the absolute calorie change for every couple of days: this was necessary since the dataset itself doesn’t provide any kind of dietary goal (such as the daily calorie target) and hence we could only calculate how each participant energy intake changed during the week. This gave use 1410 data points, each one represented as a pair in which the first element is the absolute energy shift (how much did the intake increase/decrease compared to the previous day?) and the second one is the daily stress score. Overall, the dataset contained only one week worth of data for each worker, so we couldn’t further inspect their eating behaviour. We also had to cut out some participants cause they didn’t fill in the “Daily Hassles” form. Now we proceed to show the results of our statistical analysis, to test whether stress does affect individual diet or not. It must be considered that a higher stress score doesn’t necessarily imply overeating: people react differently to stress (some will eat more, other will eat less) and that’s another reason why we calculated the absolute calorie change. For the first analysis, we looked for a correlation between the calorie shift and some stress trends which can occur in workers’ life. We identified four stress patterns by comparing the score for every given pair of days:

- **Stress variation:** the score changes (increases or decreases) while staying positive
- **Stress drop:** the score drops to zero from any given value
- **Stress stasis:** the score doesn’t change but stays positive
- **Stress absence:** the score stays equals to zero during the two days

Table 1: Stress trend analysis result (ONE-WAY ANOVA)

Trend	p-value
variation - absence	0.520
drop - absence	0.886
stasis - absence	0.0317
drop - variation	0.326
stasis - variation	0.722
stasis - drop	0.052

By defining these patterns we were able to check if a particular stress trend could be related to a calorie shift. We ran a one-way ANOVA which showed statistically significant difference in the distributions ($p \approx 0.008$; F-value = 3.9; Df = 3; $\eta^2 \approx 0.008$). An additional Tukey Honestly Significant Difference (HSD) test revealed that the only statistically significant difference was between the Stress Absence and Stress Stasis trends ($-tbfp \approx 0.031$). In other words, there was a significant difference in how people ate when they were consistently stressed during the weeks compared to when they weren’t stressed. The analysis results are summarized in Table 1.

For the second analysis we considered a simpler scenario: we aggregated every stress trend except for the absence, to inspect a link between being generally stressed in any possible way or calm and the calorie shift. In this regards we ran a Welch Two-Sample T-Test which gave us a significant result ($p \approx 0.046$). We want to remind that the analysis we ran is a post-hoc one: original experiment wasn’t mean to test our hypotheses. Hence our results can only be taken as an hint that stress could have influenced participants’ eating habits.

4 Stress-based tailoring

We now proceed to detail a hypothetical way to tailor a text-based on user stress. To do so we do stay in the dietary domain and consider producing a diet-coaching report by using NLG. To do so, we take into consideration our previously proposed architecture for dynamic tailoring in healthcare (Balloccu et al., 2020), which we summarise in Figure 1. The framework is designed to extract data regarding user diet, then produce a text which is tailored by taking into account both general tailoring techniques and user individual traits. What we do in here is to formulate a theoretical way to

enrich the tailoring logic with details about the user stress which will guide the size and content of the final text.

4.1 Assessing user’s stress

As a first step, we must determine whether the user is stressed and find a measure for this data: this is a major challenge, as it should be done in the most reliable and least intrusive possible way. Detecting stress is usually pursued through two different tools: body sensors and self-assessment tools (Plarre et al., 2011). Intrusion-wise, sensors are the least desirable approach: since our target is to deliver a report with a high frequency, it is unrealistic to expect the user to wear detection devices daily. Even if the high-frequency constraint is omitted, sensors tend to be reliable in controlled environments only (Healey and Picard, 2005)(Madan et al., 2011)(Plarre et al., 2011), as they tend to struggle to distinguish stressful events from normal activities which alter user’s psychophysical condition (like sports activities). On the other hand, self-assessment tools typically involve questionnaires like DASS-21 (Lovibond and Lovibond, 1995) or PANAS (Watson et al., 1988). Such tools can produce an estimation of users’ stress related to a relatively short period (a week for example). Moreover, it could be even possible to assess daily stress: this would make the system capable of tailoring the text “on the fly”, potentially fine-tuning the text details every single day. A possible way to address daily stress is, for example, the “Daily Hassles” form that was used in NeuroFAST project. However, measuring stress every day reduces self-assessment tools reliability, and reduces their validity to some specific scenarios (Adams et al., 2014) only. Finally, it must be considered that forcing the user to fill a form every single day could end up being a stressful operation as well. From a discussion with experts from NeuroFAST project, it turned out that the stress values were collected by using a mix of sensors and forms, an experience which the medical staff itself described as potentially stressful for the participants, especially when they were shift workers.

4.2 Enriching user tailoring with stress

Given what has been said until now, we proceed to formulate two key parameters on top of which the tailoring will take place:

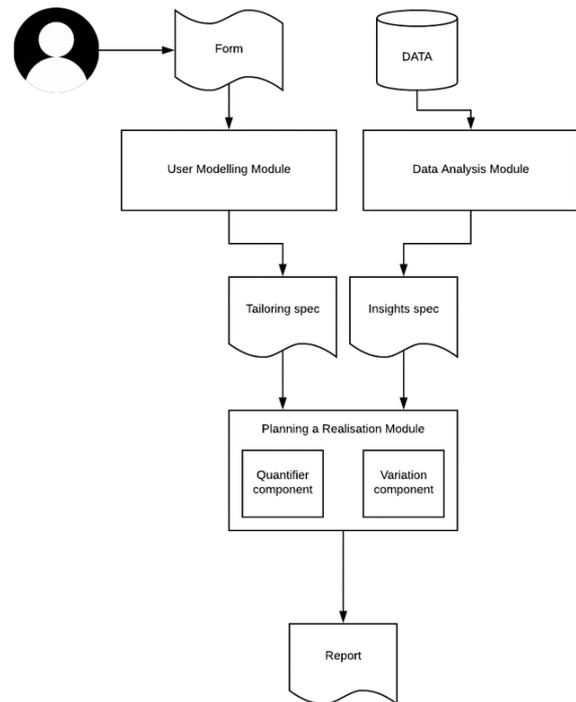


Figure 1: Original framework architecture

- Complexity:** this involves both the length and terminology of the given text. There have been proofs that stress does impact on working memory (Rai et al., 2015)(Peng et al., 2018). This means that being stressed potentially influences reading skills. Addressing this scenario can be done by delivering a shorter text and simplifying complex terminology, in line with previous work (McKeown et al., 1993)(Paris, 1988).
- Affect:** when dealing with stressful events, like work or personal problems, (Scott and Johnstone, 2012)(Torres and Nowson, 2007)(Puddephatt et al., 2020)(Riffer et al., 2019) users could feel demoralised after reading a report which states they didn’t meet their current goals. This applies well to the dietary domain but it might reasonably make sense in other areas. A possible way to address this problem is adopting an affective tone in the text, in line with previous studies (Mahamood and Reiter, 2011)(De Rosis and Grasso, 1999). Omitting certain information could be a valid option as well under some circumstances. For example (Van Deemter and Reiter, 2018), showed how “lying” could be useful in the NICU environment, since pre-

venting the child grandparents from knowing about negative development could have avoided potential health issues (i.e. heart attacks).

Regarding complexity, good tailoring should address both reducing text length and simplifying the content itself. This can be split into two complementary document planning (length) and micro-planning (terminology) steps. This is especially important given that, even without considering stress, patients typically struggle to understand complex medical language (Elhadad, 2006)(Zeng-Treitler et al., 2008). We leave terminology simplification as future work and focus, for now, on obtaining a shorter text. We do consider an output example from our system (Figure 2), which informs the user with some weekly insights about his/her diet. The text itself is lengthy, hence we try to shorten it to get an immediately understandable content. To do so we remove superfluous information which doesn't contribute to giving the user facts about their diet. A hypothetical result is showed in Figure 3. The text is now around a quarter of its original length but delivers the same information (i.e.: calories and nutrients). Potentially traumatic information was removed, like adverse effects, and some of the tailoring techniques got lost during the shortening phase: these include both psychological practices (Hawkins et al., 2008) and user choices like adverse effects. We argue that, by sacrificing a bit of customisation, we're gaining a lot of readability which is a critical parameter when the user is stressed. Moreover, the majority of user choices are still being kept. It is, however, fair to underline that the system should offer the user the chance to see the original uncut text, to access the omitted details. It is now necessary to re-introduce the affective bits inside the text as stated before: despite the new text being much faster to be read in terms of length, it sounds quite cold compared to the original one. Therefore we re-insert some motivational frames without being redundant and with the primary target of keeping the text short and simple. A mockup with an affect component can be seen in Figure 4: the improved strategy is pretty much the same we originally adopted (Balloccu et al., 2020) and focuses on lowering the weight of bad news and encouraging the user to pursue their goals (which were re-introduced into the text). Overall the text is still significantly shorter than the original one.

Hi Paul, this is your weekly diet report. You told us that you want to lose weight and gain self-confidence. So we wrote this to help you out.

You did a great job on calories; you had some problems with sugar and sodium.

Going a bit deeper, Monday was your best day! Your calories were about perfect then. Friday gave you some problems, as you ate about a third more than what you should.

Your sugar consumption is around twice more than what it should be, and much of it came from Coca Cola. We know it's hard to change what you eat, but sugar excess is bad for your health and can cause pale skin, anxiety and fatigue. Try having less sugary foods as it would turn in weight loss, less dental plaque and more energy.

Your sodium consumption is around half more than what it should be. A lot of it came from Pringles chips. Keep in mind that sodium excess can lead to nausea, headache and seizures. Less salty foods mean memory improvement, less bloating and lower blood pressure.

Paul, we came up with some suggestions based on your needs. Regarding sugar, you could replace sugary drinks with tea or coffee. Also, sodium will be lower if you avoid salty snacks and opt for fruits instead.

Figure 2: Starting text

Hi Paul, calories were good this week, especially Monday, but you had problems on Friday. Nutrient-wise, you had around twice more than your sugar target (cut a bit on Coca-Cola) and around a half more sodium than what you should get (try reducing chips).

Figure 3: Reduced example

Hi Paul, this week you did a great job on calories, especially Monday. You had some problem Friday: we're sure that next week will be better! Nutrient-wise, you had around twice more than your sugar target (cut a bit on Coca-Cola) and around a half more sodium than what you should get (try reducing chips). Changing what you eat is a long path, keep up will succeed in losing weight and gaining self-confidence!

Figure 4: Reduced example + affect

5 Conclusion and future developments

In this short paper, we highlighted the importance of addressing user stress when performing text tailoring. To the best of our knowledge there's not a single system which tries to use stress as a variable to tailor textual generation. We think that ignoring this factor prevents current systems from reaching a further step towards user engagement, especially in the domain of behavioural change, as stress does impact both reading capabilities and user motivation. We showed a real-world case in which stress does matter, the NeuroFAST project: our post hoc statistical analysis hints at the fact that the workers who took part into the experiment changed their diet significantly when stressed. This result is aligned with previous well-known evidence which state that diet and stress are linked. We then took our previously proposed framework, which dynamically tailors the user to communicate health data, and enriching its components to exploit stress value as an additional tailoring parameter. By doing so, we showed a few preliminary ideas on tailoring text based on the stress level. The new tailoring logic is mainly achieved by reducing the text size, hence drastically reducing the information delivery time, and enforcing its affective component to counterbalance the lack of motivation which stress can cause when following a diet.

5.1 Implementation

Given that this is a work-in-progress project and all of the given examples are mockup (no computational steps or formal definition of the problem were given, as this paper is purely conceptual), we intend to proceed to work on this thematic, including the implementation of a proper stress-based tailoring system to evaluate its effectiveness. Moreover, the proposed text lacked terminology simplification, a phase which is particularly important in healthcare, where the patient often ignores the meaning of medical terminology. Along with an actual system implementation, it is in our interest to design an evaluation framework, which could help us get a realistic idea of the tailoring strategy effects. This would also include considering and inspecting all of the ethical challenges that comes when working with stressed people.

5.2 Exploring different domains

We focused on the dietary domain, but we remind that behaviour change is a vast domain and we're

looking forward applying stress-based tailoring in different areas: at the moment we believe that sleeping apnea therapy (CPAP) could take great advantage of this technique as sleep deprivation is indeed a stressful event and tailored text showed to be a promising tool in increasing therapy adherence (Tatousek, 2016). Overall we do hope that this work raised some interesting research questions around a variable which is, as for now, not considered in NLG and user-tailoring.

5.3 Acknowledgments

This research was funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 812882.

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7-KBBE-2010-4 under grant agreement No: 266408. Collaborators from the University of Aberdeen, Rowett Institute gratefully acknowledge financial support from the Scottish Government as part of the Strategic Research Programme at the Rowett Institute.

References

- Phil Adams, Mashfiqui Rabbi, Tauhidur Rahman, Mark Matthews, Amy Volda, Geri Gay, Tanzeem Choudhury, and Stephen Volda. 2014. Towards personal stress informatics: Comparing minimally invasive techniques for measuring daily stress in the wild. In *Proceedings of the 8th International Conference on Pervasive Computing Technologies for Healthcare*, pages 72–79.
- Luca Anselma, Simone Donetti, Alessandro Mazzei, and Andrea Pirone. 2018. Checkyourmeal!: diet management with nlg. In *Proceedings of the Workshop on Intelligent Interactive Systems and Language Generation (2IS&NLG)*, pages 45–47.
- Luca Anselma and Alessandro Mazzei. 2017. An approach for explaining reasoning on the diet domain. In *1st Workshop on Natural Language for Artificial Intelligence co-located with 16th International Conference of the Italian Association for Artificial Intelligence (AI* IA 2017)*, volume 1983, pages 4–17. CEUR Workshop Proceedings (CEUR-WS.org).
- Luca Anselma and Alessandro Mazzei. 2018. Designing and testing the messages produced by a virtual dietitian. In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 244–253.

- Brian G Arndt, John W Beasley, Michelle D Watkinson, Jonathan L Temte, Wen-Jan Tuan, Christine A Sinsky, and Valerie J Gilchrist. 2017. Tethered to the ehr: primary care physician workload assessment using ehr event log data and time-motion observations. *The Annals of Family Medicine*, 15(5):419–426.
- Simone Balloccu, Steffen Pauws, and Ehud Reiter. 2020. A nlg framework for user tailoring and profiling in healthcare. In *SmartPhil@ IUI*, pages 13–32.
- Kim Binsted, Alison Cawsey, and Ray Jones. 1995. Generating personalised patient information using the medical record. In *Conference on Artificial Intelligence in Medicine in Europe*, pages 29–41. Springer.
- Daniel Braun, Ehud Reiter, and Advait Siddharthan. 2018. Saferdrive: An nlg-based behaviour change support system for drivers. *Natural Language Engineering*, 24(4):551–588.
- Fiorella De Rosis and Floriana Grasso. 1999. Affective natural language generation. In *International Workshop on Affective Interactions*, pages 204–218. Springer.
- Tao Ding and Shimei Pan. 2016. Personalized emphasis framing for persuasive message generation. *arXiv preprint arXiv:1607.08898*.
- Ivan Donadello, Mauro Dragoni, and Claudio Eccher. 2019. Persuasive explanation of reasoning inferences on dietary data.
- Mauro Dragoni, Tania Bailoni, Rosa Maimone, Michele Marchesoni, and Claudio Eccher. 2018. Horus. ai-a knowledge-based solution supporting health persuasive self-monitoring. In *International Semantic Web Conference (P&D/Industry/BlueSky)*.
- Noemie Elhadad. 2006. Comprehending technical texts: Predicting and defining unfamiliar terms. In *AMIA annual symposium proceedings*, volume 2006, page 239. American Medical Informatics Association.
- Gregory Finley, Erik Edwards, Amanda Robinson, Michael Brenndoerfer, Najmeh Sadoughi, James Fone, Nico Axtmann, Mark Miller, and David Suendermann-Oeft. 2018. An automated medical scribe for documenting clinical encounters. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 11–15.
- Marco Guerini, Oliviero Stock, Massimo Zancanaro, Daniel J O’Keefe, Irene Mazzotta, Fiorella de Rosis, Isabella Poggi, Meiyi Y Lim, and Ruth Aylett. 2011. Approaches to verbal persuasion in intelligent user interfaces. In *Emotion-Oriented Systems*, pages 559–584. Springer.
- Robert P Hawkins, Matthew Kreuter, Kenneth Resnicow, Martin Fishbein, and Arie Dijkstra. 2008. Understanding tailoring in communicating about health. *Health education research*, 23(3):454–466.
- Jennifer A Healey and Rosalind W Picard. 2005. Detecting stress during real-world driving tasks using physiological sensors. *IEEE Transactions on intelligent transportation systems*, 6(2):156–166.
- Saar Hommes, Chris van der Lee, Felix Clouth, Jeroen Vermunt, Xander Verbeek, and Emiel Kraemer. 2019. A personalized data-to-text support tool for cancer patients. In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 443–452.
- James Hunter, Yvonne Freer, Albert Gatt, Ehud Reiter, Somayajulu Sripada, Cindy Sykes, and Dave Westwater. 2011. Bt-nurse: computer generation of natural language shift summaries from complex heterogeneous medical data. *Journal of the American Medical Informatics Association*, 18(5):621–624.
- Saskia M Kelders, Harri Oinas-Kukkonen, Anssi Oörni, and Julia EWC van Gemert-Pijnen. 2016. Health behavior change support systems as a research discipline; a viewpoint. *International journal of medical informatics*, 96:3–10.
- Matthew W Kreuter and Ricardo J Wray. 2003. Tailored and targeted health communication: strategies for enhancing information relevance. *American journal of health behavior*, 27(1):S227–S232.
- Rita Kukafka. 2005. *Tailored Health Communication*, pages 22–33. Springer New York, New York, NY.
- Scott H Lee. 2018. Natural language generation for electronic health records. *NPJ digital medicine*, 1(1):1–7.
- Peter F Lovibond and Sydney H Lovibond. 1995. The structure of negative emotional states: Comparison of the depression anxiety stress scales (dass) with the beck depression and anxiety inventories. *Behaviour research and therapy*, 33(3):335–343.
- Anmol Madan, Manuel Cebrian, Sai Moturu, Katayoun Farrahi, et al. 2011. Sensing the “health state” of a community. *IEEE Pervasive Computing*, 11(4):36–45.
- Saad Mahamood and Ehud Reiter. 2011. Generating affective natural language for parents of neonatal infants. In *Proceedings of the 13th European Workshop on Natural Language Generation*, pages 12–21.
- Rosa Maimone, Marco Guerini, Mauro Dragoni, Tania Bailoni, and Claudio Eccher. 2018. Perkapp: A general purpose persuasion architecture for healthy lifestyles. *Journal of biomedical informatics*, 82:70–87.
- RM Malone, K Giles, NG Maloney, CL Fyfe, A Lorenzo-Arribas, DB O’Connor, and AM Johnstone. 2015. Effects of stress and mood on caffeine consumption in shift and non-shift workers. *Proceedings of the Nutrition Society*, 74(OCE1).

- Kathleen McKeown, Jacques Robin, and Michael Tanenblatt. 1993. Tailoring lexical choice to the user's vocabulary in multimedia explanation generation. In *31st Annual Meeting of the Association for Computational Linguistics*, pages 226–234.
- Harri Oinas-Kukkonen. 2013. A foundation for the study of behavior change support systems. *Personal and ubiquitous computing*, 17(6):1223–1235.
- Harri Oinas-Kukkonen and Marja Harjumaa. 2009. Persuasive systems design: Key issues, process model, and system features. *Communications of the Association for Information Systems*, 24(1):28.
- Cecile L Paris. 1988. Tailoring object descriptions to a user's level of expertise. *Computational Linguistics*, 14(3):64–78.
- Steffen Pauws, Albert Gatt, Emiel Krahmer, and Ehud Reiter. 2019. Making effective use of healthcare data using data-to-text technology. In *Data Science for Healthcare*, pages 119–145. Springer.
- Peng Peng, Marcia Barnes, CuiCui Wang, Wei Wang, Shan Li, H Lee Swanson, William Dardick, and Sha Tao. 2018. A meta-analysis on the relation between reading and working memory. *Psychological bulletin*, 144(1):48.
- Kurt Plarre, Andrew Raij, Syed Monowar Hossain, Amin Ahsan Ali, Motohiro Nakajima, Mustafa Al'Absi, Emre Ertin, Thomas Kamarck, Santosh Kumar, Marcia Scott, et al. 2011. Continuous inference of psychological stress from sensory measurements collected in the natural environment. In *Proceedings of the 10th ACM/IEEE international conference on information processing in sensor networks*, pages 97–108. IEEE.
- François Portet, Albert Gatt, Jim Hunter, Ehud Reiter, Somayajulu Sripada, and Feng Gao. 2008. Babytalk: A core architecture to summarise icu data as tailored text.
- François Portet, Ehud Reiter, Albert Gatt, Jim Hunter, Somayajulu Sripada, Yvonne Freer, and Cindy Sykes. 2009. Automatic generation of textual summaries from neonatal intensive care data. *Artificial Intelligence*, 173(7-8):789–816.
- Jo-Anne Puddephatt, Gregory S Keenan, Amy Fielden, Danielle L Reaves, Jason CG Halford, and Charlotte A Hardman. 2020. 'eating to survive': A qualitative analysis of factors influencing food choice and eating behaviour in a food-insecure population. *Appetite*, 147:104547.
- Manpreet K Rai, Lester C Loschky, and Richard Jackson Harris. 2015. The effects of stress on reading: A comparison of first-language versus intermediate second-language reading comprehension. *Journal of Educational Psychology*, 107(2):348.
- Chris Reed, Derek Long, and Maria Fox. 1996. An architecture for argumentative dialogue planning. In *International Conference on Formal and Applied Practical Reasoning*, pages 555–566. Springer.
- Ehud Reiter, Roma Robertson, and Liesl M Osman. 2003. Lessons from a failure: Generating tailored smoking cessation letters. *Artificial Intelligence*, 144(1-2):41–58.
- Friedrich Riffer, Manuel Sprung, Hannah Münch, Elmar Kaiser, Lore Streibl, Kathrin Heneis, and Alexandra Kautzky-Willer. 2019. Relationship between psychological stress and metabolism in morbidly obese individuals. *Wiener klinische Wochenschrift*, pages 1–11.
- Jamie Ross, Fiona Stevenson, Rosa Lau, and Elizabeth Murray. 2016. Factors that influence the implementation of e-health: a systematic review of systematic reviews (an update). *Implementation science*, 11(1):146.
- Anne Schneider, Alasdair Mort, Chris Mellish, Ehud Reiter, Phil Wilson, and Pierre-Luc Vaudry. 2013. Mime-nlg in pre-hospital care. In *Proceedings of the 14th European Workshop on Natural Language Generation*, pages 152–156.
- Clare Scott and Alexandra M Johnstone. 2012. Stress and eating behaviour: implications for obesity. *Obesity facts*, 5(2):277–287.
- Salzburg Global Seminar. 2011. Salzburg statement on shared decision making. *Bmj*, 342.
- Celette Sugg Skinner, Victor J Strecher, and Harm Hoppers. 1994. Physicians' recommendations for mammography: do tailored messages make a difference? *American Journal of Public Health*, 84(1):43–49.
- Moira A Stewart. 1995. Effective physician-patient communication and health outcomes: a review. *CMAJ: Canadian Medical Association Journal*, 152(9):1423.
- Anne M Stiggelbout, Trudy Van der Weijden, Maarten PT De Wit, Dominick Frosch, France Légaré, Victor M Montori, Lyndal Trevena, and Glenn Elwyn. 2012. Shared decision making: really putting patients at the centre of healthcare. *Bmj*, 344:e256.
- Lacroix J. Visser T. Den Teuling N. Tatousek, J. 2016. Promoting adherence to cpap with tailored education and feedback: A randomized controlled clinical trial. *Proc. Sleep 2015*.
- Susan J Torres and Caryl A Nowson. 2007. Relationship between stress, eating behavior, and obesity. *Nutrition*, 23(11-12):887–894.
- Kees Van Deemter and Ehud Reiter. 2018. Lying and computational linguistics. In *The Oxford Handbook of Lying*. Oxford University Press.

- Ielka Van Der Sluis and Chris Mellish. 2008. Towards affective natural language generation: Empirical investigations. In *Proceedings of the Symposium on Affective Language in Human and Machine, AISB*, pages 9–16.
- David Watson, Lee Anna Clark, and Auke Tellegen. 1988. Development and validation of brief measures of positive and negative affect: the panas scales. *Journal of personality and social psychology*, 54(6):1063.
- Sandra Williams and Ehud Reiter. 2008. Generating basic skills reports for low-skilled readers. *Natural Language Engineering*, 14(4):495–525.
- Qing Zeng-Treitler, Sergey Goryachev, Tony Tse, Alla Keselman, and Aziz Boxwala. 2008. Estimating consumer familiarity with health terminology: a context-based approach. *Journal of the American Medical Informatics Association*, 15(3):349–356.
- Ingrid Zukerman, Richard McConachy, and Sarah George. 2000. Using argumentation strategies in automated argument generation. In *INLG'2000 Proceedings of the First International Conference on Natural Language Generation*, pages 55–62.