

A Personalized Data-to-Text Support Tool for Cancer Patients

Saar Hommes

TiCC, Tilburg University
s.hommes@uvt.nl

Chris van der Lee

TiCC, Tilburg University
c.vdrlee@uvt.nl

Felix Clouth

Tilburg University
f.j.clouth@uvt.nl

Jeroen Vermunt

Tilburg University
j.k.vermunt@uvt.nl

Xander Verbeek

Netherlands Comprehensive
Cancer Organization (IKNL)
x.verbeek@iknl.nl

Emiel Krahmer

TiCC, Tilburg University
e.j.krahmer@uvt.nl

Abstract

In this paper, we present a novel data-to-text system for cancer patients, providing information on quality of life implications after treatment, which can be embedded in the context of shared decision making. Currently, information on quality of life implications is often not discussed, partly because (until recently) data has been lacking. In our work, we rely on a newly developed prediction model, which assigns patients to scenarios. Furthermore, we use data-to-text techniques to explain these scenario-based predictions in personalized and understandable language. We highlight the possibilities of NLG for personalization, discuss ethical implications and also present the outcomes of a first evaluation with clinicians.

1 Introduction

Data-to-text generation systems are increasingly used in the health domain (Pauws et al., 2019). They can, for example, be used for automation of health reports, clinical decision support, encourage behavioural change, ensure patient engagement or assist patients with making health decisions (Pauws et al., 2019). The tool we present here focuses on the latter two in the context of shared decision making (SDM) (Elwyn et al., 2017) for colorectal cancer patients. Since patients are increasingly encouraged to have an active role in treatment decision making (Pieterse et al., 2008), patients need to be accurately informed about their treatment options. Next to information on incidence and survival, patients also want to consider how a treatment is going to affect their quality of life (QoL) (Zafar et al., 2009). Since survival rates for colorectal cancer patients are increasing (Mols et al., 2013), the relevance of QoL becomes more prominent and patients are more likely to consider, for example, how treatments will impact their social life, ability to go

to work, or emotional well-being. Importantly, in recent years, a dedicated effort has started to collect data on such QoL dimensions. However, this information is often not communicated to the patient, or is generic and difficult to understand (Brundage et al., 2005). In this paper, we describe the design and implementation of a new patient support tool that is able to communicate QoL information to individual patients in an understandable and personalized way. Additionally, ethical issues are considered and results of an initial evaluation with clinicians are discussed.

2 Background

To explain health data in a clear way and assist patients with making treatment decisions, so called ‘decision aids’ are developed. These are tools that explain health information and lay out the benefits and risks of treatments to patients. A recent systematic review concluded that “[...] people exposed to decision aids feel more knowledgeable, better informed, and clearer about their values, and they probably have a more active role in decision making and more accurate risk perceptions” (Stacey et al., 2017). Vromans et al. (2019a,b) looked at the content and communication styles of such decision aids in more detail. The authors found that information communicated in such tools is often generic (based on the general population instead of the individual patient), which seems to undermine the potential effects of decision aids since understanding personalized risks (rather than generic risks) is easier and more relevant for patients (Thorne et al., 2005).

Many patients are unable to benefit from this generic health information, because these documents fail to communicate crucial information that influences the patient’s understanding of these materials (e.g. the patient’s individual risks, con-

cerns and values) (Acharya et al., 2019). Tailoring health communication to individual patients seems a fruitful solution. And although such personalized health information can be effective (Kreuter and Wray, 2003), it is not without challenges. Personalizing health information manually is time consuming, costly and the outputs are often inconsistent (Pauws et al., 2019). Natural language generation (NLG) techniques can tackle these problems, and are therefore increasingly used in the health domain (Di Eugenio and Green, 2010; Pauws et al., 2019). A leading example of tailored health information using NLG is the BabyTalk-system developed by Gatt et al. (2009). This system generates personalized hospital stay summaries for parents of babies in a neonatal intensive care unit. Not all tailored NLG information is successful. Reiter et al. (2003) created ‘STOP’, a system aimed at generating tailored smoking cessation letters. However, the non-tailored letters were just as effective as the tailored letters. For our system, we do not need to change health behavior, but rather inform patients in the best way in order to assist them with decision making.

Some NLG health applications have already been developed to facilitate shared decision making to some degree. PIGLIT (Binsted et al., 1995; Cawsey et al., 2000), for example, was developed to generate explanations of patient records to help patients make sense of their prognosis. Additionally, Gkatzia et al. (2014) use NLG techniques to generate textual summaries of medical sensory data and personalize the presentation format of these summaries.

To date however, no tool has been developed to communicate personalized Quality of Life outcomes, which is the aim of the current tool. By deploying data-to-text techniques, the tool can communicate QoL data in a personal, relevant, understandable and consistent way. What follows is a description of how the tool is able to do so.

3 The system

3.1 Clinical setting

For the current tool, we focus on colorectal cancer as the health subject. Colorectal cancer is the third most common cancer in the world (American Cancer Society, 2019; World Cancer Research Fund, 2019). It is expected that in 2020, 17.000 people will suffer from colorectal cancer in the Nether-




Figure 1: The support tool within the three-talk shared decision model (Elwyn et al., 2017).

lands alone (KWF Kankerbestrijding, 2019). At the same time, survival rates are improving (Mols et al., 2013). This means that increasingly, people with colorectal cancer will have to face the long-term effects of the treatments they underwent. Partly because of this, patients want to be involved in treatment decision making (Shay and Elston Lafata, 2015). In order to take part in these treatment decisions, however, patients need to have access to the relevant quality of life data. Providing patients this access is the aim of the system we developed.

3.2 Shared decision making

Figure 1 shows the system’s ecosystem within the context of shared decision making. Elwyn et al. (2017) identified three different kinds of talks in the SDM process, which can be supported by a decision support tool. Shared decision making begins with the ‘Team Talk’. Here, patient and clinician work together as a team and explore the choices and goals of different treatment options.

Secondly, the clinician and patient discuss the different treatment options in the ‘Option Talk’. Within this talk, data become important since patient and clinician need to discuss risks and benefits of the options. Therefore, they can decide to use decision aid support (or directly go to the decision talk without interference of a decision support tool).

When a decision support tool is used, the SDM-support stage is entered (see Figure 1). Within this stage, it is essential that treatments are explained and in order to do so, data are required. For our support tool, there are two different kinds of input data. The first is a registry data set called “PROFILES” (Patient Reported Outcomes Following Initial treatment and Long term Evaluation of Survivorship (van de Poll-Franse et al., 2011)). This is a dataset that consists of over 21.000 Dutch cancer patients who have reported on health-related quality of life measures within the Netherlands Cancer Registry (NCR). Of those patients, 1.631 were colorectal cancer patients. The questionnaire-based data are acquired via the European Organisation for Research and Treatment of Cancer (EORTC) Quality of Life Questionnaire (QLQ) C30 (version 3.0) which assesses health-related quality of life (HRQOL) (Mols et al., 2013). The QoL measures include: physical functioning, role functioning, emotional functioning, cognitive functioning, social functioning, global health status, fatigue, nausea/vomiting, pain, dyspnea, insomnia, appetite loss, constipation, diarrhea, and financial impact.

The second form of input data are based on the individual patient prognosis data (gender, age, tumor stage, comorbidities, examined lymph nodes, differentiation grade, topography and histology). This information is put into the system by the patient’s clinician.

Both kinds of data feed into the prediction model that assigns patients to different outcome scenarios. The model has been developed by statisticians (Clouth et al., 2019), and uses latent class analysis (Vermunt and Magidson, 2002). The model clusters patients, from the PROFILES registry, into five latent classes (or scenarios as we call them here). That is, patients in the same outcome class have a comparable combination of answers on EORTC items. Based on the clinical information from the NCR, the model predicts class membership. This way, we can predict QoL outcome scenarios for new patients with unknown




Figure 2: Module selection of the system.

scores on the EORTC questionnaire. More precisely, we estimate the probability of an individual patient as belonging to one of the five latent classes assuming treatment is known. See 4.1 for more.

In turn, the different data inputs and the prediction model output feed the support tool. See 3.3 for a more detailed explanation of the current support tool. The decision support tool can assist patients and clinicians within the ‘Decision Talk’, where patient and clinician discuss the preferences for treatment.

3.3 Module selection

Figure 2 is an overview of the system’s module selection. The system is rule- and template-based. Contrary to neural approaches, rule- and template-based systems ensure precise and consistent communication without room for new or incorrect interpretations of the data. Since the current system communicates sensitive health data, this approach is essential. The module selection framework is based on van der Lee et al. (2017) who developed a system called ‘PASS’. PASS automatically generates soccer reports (from soccer data) and tailors these texts based on the receiver’s club preference. Similarly, the current support system translates medical data into texts and personalizes that information based on relevance.

As van der Lee et al. (2017) argue, a major advantage of such a modular approach is the flexibility of such a system. One can easily add modules. For the current system, this means that we plan on adding other cancer types and treatments, but also expand the personalization techniques used.

The modules used for the system are derived from PASS, although some modules were modified or omitted. For instance, modules that present references and information in a varied way were suitable for PASS, since the system’s goal was to generate enjoyable and varied output, but these goals are not appropriate for the current system for which effective communication of sensitive health data is most salient.

After data is obtained using the prediction

model and the two input data types (PROFILES and individual patient prognosis data), the data-to-text system is initiated. The system starts with loading the patient’s prognosis data and activating the *lookup module*. This module opens the template database and retrieves all the different templates that can be used to describe the medical information for the patient.

After collecting these templates within the *template selection module*, the *governing module* will walk through all topics that need to be discussed one by one, and send all the templates pertaining to the topic to the template selection module. Since patients have the option to choose if they want to receive (more) information on a QoL measure (see 4.1), this governing module will continuously update the topics based on the patient’s information preference.

In the template selection module, a template is selected for a topic based on the *ruleset module*, which checks which template is most appropriate based on the input data as well as personalization condition (see 4.2). After choosing a template, the empty slots in the template are filled out with the corresponding data using the *template filler module*. Finally, the filled out templates are ordered within the *text collection module*, and a personalized text is generated. Next, we will discuss which personalization techniques are used specifically.

4 Personalization techniques

The most salient personalization technique used is based on communicating the relevant health information to the individual patient by generating personalized predictions. Although it may seem evident that patients need information based on their own personal information, most patient education materials or decision aids are not tailored towards specific patient outcomes (Vromans et al., 2019a,b). By using the output of the prediction model as a starting point, the current system is able to do so since it can identify different quality of life outcome scenarios. What follows is a brief description of these scenarios and how textual output is tailored towards them.

4.1 Scenarios

Based on latent class analysis, patients are predicted to belong to one of five different scenarios. Broadly speaking, patients can (1) have good outcomes on all quality of life outcomes, (2) have

good outcomes for physical dimensions, but relatively poor outcomes for mental dimensions, (3) perform average across all dimensions, (4) perform well on mental dimension, but poorly on physical outcomes and (5) perform poorly across all dimensions.

When patients have a better indication of how treatment will impact their specific lives, they can make better treatment decisions. For patients belonging to scenario 5 for example (who perform badly across all dimensions), it seems essential to know that treatment is going to affect your quality of life greatly. For patients in scenario 1 however, knowing that a particular treatment is not going to have a huge impact on your quality of life might be a relief. All in all, knowing which scenario is most applicable can help the patient to plan better for their life with and after treatment.

In total, patients can view 15 QoL outcomes (see the QoL outcomes mentioned in 3.2). [Text 1](#) and [Text 2](#) give an overview of how patients within the different scenarios are informed about their general personal QoL outcomes.

[Fixed text]:

This tool calculated how chemotherapy will probably affect your quality of life. It did this by comparing how your personal medical information relates to the medical information of other colorectal cancer patients who also underwent chemotherapy.

[Scenario 1]:

[s1] Overall, patients like you have good to very good outcomes on quality of life. [s2] This means that chemotherapy will probably not have a major effect on your quality of life afterwards. [s3] Do keep in mind that these outcomes are probabilities, so real life can be different.

[Scenario 2]:

[s4] Overall, patients like you experience few to very few physical problems. [s5] Overall, mental problems are quite common for patients like you. [s6] You might benefit from talking to a therapist or going to meetings with other colorectal cancer patients. [s7] Also, keep in mind that these results are probabilities, so your real life can be different from these outcomes.

[Scenario 3]:

[s8] Overall, patients like you could experience several problems relating to quality of life. [s9] There are a lot of things that could help you with possible problems you might experience. [s7].

Text 1: Patients in different scenarios (1, 2 and 3) receive different explanations of QoL outcomes. All patients receive the fixed texts. Some sentences overlap between scenarios (sentence s7).

[Scenario 4]:

[s10] Overall, patients like you experience few to very few mental problems. [s11] Overall, physical problems are quite common for patients like you. [s12] You might benefit from going to a physical therapist or taking medication. [s7]

[Scenario 5]:

[s13] Overall, patients like you do experience problems relating to quality of life. [s14] This means that chemotherapy could have a big impact on your quality of life afterwards. [s9] [s7]

[Fixed text]:

Click on the different outcomes to see a more detailed description of your risk, examples of how it can affect your life and the things you can do to help you manage.

Text 2: Patients in different scenarios (4 and 5) receive different explanations of QoL outcomes. All patients receive the fixed texts. Some sentences overlap between scenarios (sentence s7 and s9).

After viewing the general explanation of the personal QoL outcome scenario (Text 1 and Text 2), patients can click on specific QoL outcomes they would like to view. For example, if a patient clicks on physical functioning (one of the QoL measures), Text 3 shows an example of how information presentation is done for a patient within scenario 2 (good on physical dimensions), and Text 4 shows this for a patient within scenario 5 (poor outcomes overall).

Most patients like you have **excellent** physical functioning.

This means that **most patients** like you have **no difficulty** with dressing themselves, eating or washing up. Also, patients like you can **easily** lift a heavy suitcase.

If you **do** have physical problems, a physical therapist could help you.

Text 3: Information on physical functioning for patients in scenario 2. Bold text is personalized based on scenario outcomes.

A vast majority of patients like you have **difficulty with** physical functioning.

This means that **most patients** like you have **difficulty** with dressing themselves, eating or washing up. Also, patients like you **might have difficulty with** lifting a heavy suitcase.

If you have physical problems, a physical therapist could help you.

Text 4: Information on physical functioning for patients in scenario 5. Bold text is personalized based on scenario outcomes.

Note that individual patients only get information based on their own scenario. This way, patients in scenario 5 are not confronted with the better outcomes in other scenarios (which is the case for health materials not tailored towards the individual patient). Also, patients can opt for the QoL outcomes they want to view, so they do not have to view all the information. Finally, patients can choose whether they want the personal information based on scenarios, or general information based on the whole colorectal cancer population (e.g. “In general, about half of the colorectal cancer patients have difficulty with physical functioning. This means that half of the patients have difficulty with dressing themselves, eating or washing up. Also, half of the patients have difficulty with lifting a heavy suitcase. If you have physical problems, a physical therapist could help you.”).

4.2 Experimentation with personalization

Understanding statistical information about the risks and benefits of treatment, or QoL outcomes, is difficult to interpret for many patients (Gigerenzer et al., 2007). Gigerenzer et al. (2007) even state that there is a “collective statistical illiteracy” (p.53). At the same time, when communicating health information, statistical information is inevitable. That is why within the current tool a ‘statistical module’ is included. Within this module, textual output can vary based on the statistic that is used, for example:

“Most patients like you/ about 80% of patients like you/ 8 out of 10 patients like you experience X”

One of the strengths of NLG is that the textual output is consistent. Since research has shown that patients have a strong preference for the presentation of statistical information (e.g. Brundage et al. (2005)), but studies also show that patients vary in their preferences (e.g. Hagerty et al. (2004)), the current system could help investigate which patterns underlie these statistical preferences.

Another module that is built in in the current system is the ‘framing-module’. Some research indicates that “the framing effect” occurs when delivering messages (Akl et al., 2011). That is, that the intended message depends on how it is formulated. For example, “there is a 10% chance of dying” statistically means the same as “there

is a 90% chance of surviving”. Arguably though, the information and acceptance of that information can be interpreted differently by the patient. Since a neutral way of presenting health related quality of life information is essential in this case, the tool can assist with determining the most appropriate way of phrasing the information. This may depend on personal preferences, but might also have to do with the scenarios patients are in.

4.3 Affective NLG

Although clinicians sometimes fear that patients cannot handle poor outcomes (de Haes and Koe-doot, 2003), other research indicates that patients prefer to know all the information, even when this information is upsetting (Kehl et al., 2015). In any case, when dealing with delivering this sensitive information, it is crucial to think about the context in which such information is delivered. Earlier evaluation studies of NLG-systems, such as Mahamood and Reiter (2011), showed that all patients – regardless of their stress level – prefer affective texts over neutral ones. That is why affective language is used in all scenarios. For example, possible solutions such as (“if you experience physical problems, a physical therapist might help you”) are included in all scenarios. Also, fixed texts convey messages such as “getting a cancer diagnosis may be overwhelming”. Patient evaluation will reveal how such texts are received within this context.

5 Ethical considerations

Since sensitive data are being used, some ethical considerations need to be discussed. More so than systems such as PASS, NLG-systems within the health domain need to be accurate since generation mistakes can have severe implications for patients. That is why, while designing the support tool, the ethical checklist of Smiley et al. (2017) is kept in mind. Smiley et al. (2017) give a 12-item checklist relating to human consequences, data issues, generation issues and provenance. Next, these items are discussed in light of the current support tool.

5.1 Human consequences

The questions relating to human consequences are (1) “are there any ethical objections to building the application”, (2) “how could a user be disadvantaged by the system” and (3) “does the system use any personally identifiable information”. With re-

spect to the first two questions, it should be noted that patients should always have the choice not to receive the information, if they prefer not to be informed. For the current system, we think it is best if the clinician and patient jointly decide whether they want to use the support system at all. Also, patients can choose the option to only receive general information (not based on their own medical situation). Additionally, personal identifiers such as names, residence or specific birth dates are not asked so data is anonymized (question 3).

5.2 Data issues

For data issues, Smiley et al. (2017) identified four questions: (1) “how accurate is the underlying data”, (2) “are there any misleading rankings given”, (3) “are there (automatic) checks for missing data” and (4) “does the data contain any outliers”. The quality of the data used for the support tool is excellent. The information on QoL outcomes of patients is based on a representative sample of colorectal cancer survivors in the Netherlands (van de Poll-Franse et al., 2011) and all clinical information is registry based data from the NCR. With regard to the second question, patients need to be informed about the probabilities of belonging to a certain outcome scenario. The system provides patients with a personalized text that explains what the scenarios are, how the tool is able to get these scenarios and that the information communicated are probabilities so that real life outcomes can still differ. For patients that have a relatively negative outcome (scenario 5), the system generates a more detailed explanation of how probabilities can differ from reality (see Text 2). Patient evaluation will demonstrate if people are satisfied with these explanations.

Questions 3 and 4 refer to the quality of the data. Both the PROFILES and NCR data are rigorously and manually checked for missing of implausible values to ensure high quality of the data. Further, latent class analysis is particularly suited for dealing with missing data and robust methods are used to ensure accurate predictions. Mistakes could be made with the clinical information the clinician puts in. That is why the system prints the patient data so both patient and clinician can check this during the Decision Talk (Figure 1). When a mistake was made, the clinician can adjust the data accordingly. The patient does not have this access to ensure that he or she only receives accurate in-

formation and cannot inadvertently introduce mistakes.

5.3 Generation issues

For generation issues, three questions are formulated: (1) “can you defend how the story is written”, (2) “does the style of the automated report match your style” and (3) “who is watching the machines”. Questions 1 and 2 will be kept in mind while doing patient evaluation. Because the underlying data is NCR data, governance of the data is ensured.

5.4 Provenance

Finally, (1) “will you disclose your methods” and (2) “will you disclose the underlying data sources” concludes the provenance section. Both of these questions are answered with a ‘yes’ for the current system. Patients have the option to click on “more information” sections, which will explain how the data and probabilities are acquired (in fixed text). Patient evaluation will reveal how this will be communicated specifically to ensure that patients understand these explanations.

6 Evaluation with clinicians

As [Mahamood and Reiter \(2012\)](#) point out, involving clinicians in an early stage of the development of NLG systems can “significantly enhance the quality of many NLG systems” (p.100). Therefore, as an initial iteration, the tool was evaluated with two clinicians (both colorectal cancer surgeons). Two individual semi-structured interviews were conducted and both conversations lasted about 30 minutes. The goals of the interviews were to assess (1) what clinicians thought of the clinical information that was used, (2) how the tool should be implemented, (3) whether they thought the outcomes of the tool are relevant, (4) what they thought of the interpretation of the outcomes and (5) what the general aim of the tool is. Outcomes of the evaluation are discussed next.

6.1 Clinical information

Initially, names of the patients were also inserted into the system. Both clinicians agreed however, that working with anonymized data would be best to ensure safety. This way, clinicians could also access the tool without a password, as they both agreed remembering login information is not ideal. As for all the other clinical information

(gender, age, tumour stage, comorbidities, examined lymph nodes, differentiation grade, topography and histology), the clinicians agreed that this information is readily available for them and does not take up too much time to enter.

6.2 Implementation of the tool

When asked about where in the consultation the tool could be implemented best, both clinicians agreed that they were able to reserve limited time for this during consultation. The patient can then access his or her information at home and results can be discussed during the next consultation with an oncologist (cf. the Option Talk and the Decision Talk, see [Figure 1](#)).

6.3 Relevance of outcomes

Both clinicians agreed that knowing how treatment is going to impact the quality of life for specific patients is valuable information and that this information is currently lacking in clinical practice. They agreed that when communicating QoL outcomes, patients may know better what to expect and clinicians can provide better care. However, one clinician noted that not all quality of life measures should be communicated in the same way. For example, the effect of chemotherapy on financial issues was communicated as in [Text 5](#).

About half of the patients like you (5 out of 10) experience financial issues.

Financial issues mean for example that you cannot pay the bills on time.

When you experience financial issues, you can contact a financial expert.

Text 5: Financial issues communicated before evaluation with clinicians.

One clinician noted however that predicting how treatment is going to affect your financial status is dependent on many external factors (such as current financial status, whether or not patients are (self-)employed, what sort of job they have, et cetera). In order to make more valid predictions, patients would have to answer a lot of questions beforehand and even then, predictions could easily be wrong. That is why we choose to communicate financial issues in a more general way ([Text 6](#)).

Depending on your own financial situation, there are some patients who experience financial issues.

Financial issues mean for example that you cannot pay the bills on time.

When you experience financial issues, you can contact a financial expert.

Text 6: Financial issues communicated after evaluation with clinicians.

Furthermore, both clinicians agreed that including examples of how quality of life is going to affect the lives of patients is very useful (e.g. “you will probably have difficulty lifting a heavy suitcase”, “you might feel lonely at times”, “most patients like you do not have difficulty falling asleep”).

6.4 Interpretation issues

Before evaluation, the different scenarios were called “profiles”. Both clinicians agreed however, that interpretation of these profiles is hard. That is why, after evaluation, profiles are not mentioned and patients only receive information on scenarios with the phrase “patients like you” instead of “patients belonging to your profile”.

6.5 Aim of the tool

Both clinicians noted that it would be nice to also include information on survival outcomes and how this is affected by chemotherapy. In a next phase, we plan on incorporating such data if our data-set permits this. Also, data on QoL outcomes for patients who haven’t had chemotherapy will be included, as both clinicians noted that this is important information for the decision making process. Incorporating both survival outcomes and QoL outcomes for non-chemotherapy patients, the tool would be able to support patients and clinicians with shared decision making even more.

7 Conclusios & future work

In the current paper, a personalized data-to-text support system for colorectal cancer patients was described. The tool is aimed at communicating quality of life outcomes in an understandable, personalized, relevant and consistent way. As a starting point, the system was evaluated with two clinicians and yielded positive results. Both clinicians were convinced of the relevance of communicating quality of life information in a more person-

alized manner. As the support tool will be implemented in a clinical setting, several ethical issues should be kept in mind when developing the system further. The next step is therefore to evaluate the tool with patients.

For the development of this tool, we build on the PASS data-to-text system (van der Lee et al., 2017), which was originally developed for the tailored generation of soccer reports. Building on an analysis of existing decision aids (Vromans et al., 2019a,b), templates were defined with different framing policies relating to different outcome scenarios. Dedicated modules were developed for, among other things, generating descriptions of personalized probabilities. In this way, it becomes possible to generate descriptions for personalized, individual tailored predictions in many different ways, something which would never be possible for a traditional decision aid (which typically is generic).

The main aim of this prototype was to come up with medically relevant information for patients in a personalized manner. We are aware that this undermines the linguistic variation of the textual output. However, as we plan on adding more modules (for example include relevant patient testimonials), linguistic variability will increase. Furthermore, decision support tools in current practice are very static because they almost never take into account personalized information from the patient (Vromans et al., 2019a,b). Our support tool is therefore an important contribution to ensure personalized medicine. We hope to have shown that personalized treatment decision aids are both a new and interesting practical application of NLG-techniques, as well as an exciting testbed for the development of new NLG techniques.

Acknowledgements

We received support from The Netherlands Organization for Scientific Research (NWO Data2Person 628.011.030), for which we are grateful. We would also like to thank the clinicians for their expert opinions during the evaluation of the data-to-text support tool, and the anonymous reviewers for their helpful comments.

References

Sabita Acharya, Andrew Boyd, Richard Cameron, Karen Lopez, Pamela Martyn-Nemeth, Carolyn Dickens, Amer Ardati, and Barbara Di Eugenio.

2019. [Incorporating personalization features in a hospital-stay summary generation system](#). In *Proceedings of the 52nd Hawaii International Conference on System Sciences*, pages 4175–4184.
- Elie Akl, Andrew Oxman, Jeph Herrin, Gunn Vist, Irene Terrenato, Francesca Sperati, Cecilia Costiniuk, Diana Blank, and Holger Schünemann. 2011. [Framing of health information messages](#). *Cochrane Database of Systematic Reviews*, (12).
- American Cancer Society. 2019. [Key statistics for colorectal cancer](#).
- 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.
- Michael Brundage, Deb Feldman-Stewart, Ann Leis, Andrea Bezjak, Lesley Degner, Karima Velji, Lisa Zetes-Zanatta, Dongsheng Tu, Paul Ritvo, and Joseph Pater. 2005. Communicating quality of life information to cancer patients: A study of six presentation formats. *Journal of Clinical Oncology*, 23(28):6949–6956.
- Alison Cawsey, Ray Jones, and Janne Pearson. 2000. The evaluation of a personalised health information system for patients with cancer. *User Modeling and User-Adapted Interaction*, 10(1):47–72.
- Felix Clouth, Gijs Geleijnse, Lonneke van de Poll-Franse, Steffen Pauws, and Jeroen Vermunt. 2019. Quality of life profiles of colon cancer survivors: A three-step latent class analysis. In *Abstract from IFCS, Thessaloniki, Greece*.
- Barbara Di Eugenio and Nancy Green. 2010. Emerging applications of natural language generation in information visualization, education and health-care. In Nitin Indurkha and Fred J. Damerau, editors, *Handbook of Natural Language Processing*, volume 2, chapter 23, pages 557–575. CRC Press, Taylor and Francis Group.
- Glyn Elwyn, Nan Cochran, and Michael Pignone. 2017. Shared decision making—the importance of diagnosing preferences. *JAMA internal medicine*, 177(9):1239–1240.
- Albert Gatt, Francois Portet, Ehud Reiter, Jim Hunter, Saad Mahamood, Wendy Moncur, and Somayajulu Sripada. 2009. From data to text in the neonatal intensive care unit: Using NLG technology for decision support and information management. *AI Communications*, 22(3):153–186.
- Gerd Gigerenzer, Wolfgang Gaissmaier, Elke Kurz-Milcke, Lisa Schwartz, and Steven Woloshin. 2007. Helping doctors and patients make sense of health statistics. *Psychological science in the public interest*, 8(2):53–96.
- Dimitra Gkatzia, Verena Rieser, Alexander McSporran, Alistair McGowan, Alasdair Mort, and Michaela Dewar. 2014. Generating verbal descriptions from medical sensor data: A corpus study on user preferences. *BCS Health Informatics Scotland. Glasgow, UK*.
- Hanneke de Haes and Nelleke Koedoot. 2003. Patient centered decision making in palliative cancer treatment: a world of paradoxes. *Patient education and counseling*, 50(1):43–49.
- Rebecca Hagerty, Phyllis Butow, Peter Ellis, Elizabeth Lobb, Susan Pendlebury, Natasha Leighl, David Goldstein, Sing Kai Lo, and Martin Tattersall. 2004. Cancer patient preferences for communication of prognosis in the metastatic setting. *Journal of clinical oncology*, 22(9):1721–1730.
- Kenneth Kehl, Mary Landrum, Neeraj Arora, Patricia Ganz, Michelle van Ryn, Jennifer Mack, and Nancy Keating. 2015. Association of actual and preferred decision roles with patient-reported quality of care: shared decision making in cancer care. *JAMA oncology*, 1(1):50–58.
- Matthew Kreuter and Ricardo Wray. 2003. Tailored and targeted health communication: Strategies for enhancing information relevance. *American journal of health behavior*, 27(1):227–232.
- KWF Kankerbestrijding. 2019. [Kanker in nederland tot 2020. Trends en prognoses](#).
- Chris van der Lee, Emiel Krahmer, and Sander Wubben. 2017. Pass: A dutch data-to-text system for soccer, targeted towards specific audiences. In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 95–104.
- 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.
- Saad Mahamood and Ehud Reiter. 2012. Working with clinicians to improve a patient-information nlg system. In *Proceedings of the Seventh International Natural Language Generation Conference*, pages 100–104. Association for Computational Linguistics.
- Floortje Mols, Tonneke Beijers, Valery Lemmens, Corina van den Hurk, Gerard Vreugdenhil, and Lonneke van de Poll-Franse. 2013. Chemotherapy-induced neuropathy and its association with quality of life among 2-to 11-year colorectal cancer survivors: Results from the population-based PROFILES registry. *Journal of Clinical Oncology*, 31(21):2699–2707.
- 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.

- Arwen Pieterse, Monique Baas-Thijssen, Corrie Mar-
ijnen, and Anne Stiggelbout. 2008. Clinician and
cancer patient views on patient participation in treat-
ment decision-making: A quantitative and qual-
itative exploration. *British Journal of Cancer*,
99(6):875–882.
- Lonneke van de Poll-Franse, Nicole Horevoorts, Mies
van Eenbergen, Johan Denollet, Jan Anne Roukema,
Neil Aaronson, Ad Vingerhoets, Jan Willem Co-
ebergh, Jolanda de Vries, and Marie-Louise Essink-
Bot. 2011. The patient reported outcomes follow-
ing initial treatment and long term evaluation of sur-
vivorship registry: scope, rationale and design of
an infrastructure for the study of physical and psy-
chosocial outcomes in cancer survivorship cohorts.
European Journal of Cancer, 47(14):2188–2194.
- Ehud Reiter, Roma Robertson, and Liesl Osman. 2003.
Lessons from a failure: Generating tailored smok-
ing cessation letters. *Artificial Intelligence*, 144(1-
2):41–58.
- Aubree Shay and Jennifer Elston Lafata. 2015. Where
is the evidence? A systematic review of shared de-
cision making and patient outcomes. *Medical Deci-
sion Making*, 35(1):114–131.
- Charese Smiley, Frank Schilder, Vassilis Plachouras,
and Jochen Leidner. 2017. Say the right thing right:
Ethics issues in natural language generation sys-
tems. In *Proceedings of the First ACL Workshop on
Ethics in Natural Language Processing*, pages 103–
108.
- Dawn Stacey, France Légaré, Krystina Lewis, Michael
Barry, Carol Bennett, Karen Eden, Margaret
Holmes-Rovner, Hilary Llewellyn-Thomas, Anne
Lyddiatt, Richard Thomson, and Lyndal Trevena.
2017. [Decision aids for people facing health treat-
ment or screening decisions](#). *Cochrane Database of
Systematic Reviews*, (4).
- Sally Thorne, Barry Bultz, and Walter Baile. 2005. Is
there a cost to poor communication in cancer care?
a critical review of the literature. *Psycho-oncology*,
(14):875–884.
- Jeroen Vermunt and Jay Magidson. 2002. Latent
class cluster analysis. *Applied latent class analysis*,
11:89–106.
- Ruben Vromans, Mies van Eenbergen, Steffen Pauws,
Gijs Geleijnse, Henk van der Poel, Lonneke van de
Poll-Franse, and Emiel Kraemer. 2019a. Com-
municative aspects of decision aids for localized
prostate cancer treatment—A systematic review. In
*Urologic Oncology: Seminars and Original Investi-
gations*, 37, pages 409–429.
- Ruben Vromans, Kim Tenfelde, Steffen Pauws, Mies
van Eenbergen, Ingeborg Mares-Engelberts, Galina
Velikova, Lonneke van de Poll-Franse, and Emiel
Kraemer. 2019b. Assessing the quality and com-
municative aspects of patient decision aids for early-
stage breast cancer treatment: A systematic review.
In *Breast Cancer Research and Treatment*.
- World Cancer Research Fund. 2019. [Colorectal cancer
statistics](#).
- Yousuf Zafar, Stewart Alexander, Kevin Weinfurt,
Kevin Schulman, and Amy Abernethy. 2009. De-
cision making and quality of life in the treatment
of cancer: a review. *Supportive Care in Cancer*,
17(2):117–127.