Helping Cancer Patients to Choose the Best Treatment: Towards Automated Data-Driven and Personalized Information Presentation of Cancer Treatment Options

Emiel Krahmer[1] ✉ 🆔
Department of Communication and Cognition, Tilburg University, The Netherlands

Felix Clouth 🆔
Department of Methodology and Statistics, Tilburg University, The Netherlands
Netherlands Comprehensive Cancer Organisation (IKNL), Utrecht, The Netherlands

Saar Hommes 🆔
Department of Communication and Cognition, Tilburg University, The Netherlands
Netherlands Comprehensive Cancer Organisation (IKNL), Utrecht, The Netherlands

Ruben Vromans 🆔
Department of Communication and Cognition, Tilburg University, The Netherlands
Netherlands Comprehensive Cancer Organisation (IKNL), Utrecht, The Netherlands

Steffen Pauws 🆔
Department of Communication and Cognition, Tilburg University, The Netherlands
Innovation Excellence, Innovation & Strategy, Philips, Eindhoven, The Netherlands

Jeroen Vermunt 🆔
Department of Methodology and Statistics, Tilburg University, The Netherlands

Lonneke van de Poll-Franse 🆔
Netherlands Comprehensive Cancer Organisation (IKNL), Utrecht, The Netherlands
Department of Medical and Clinical Psychology, Tilburg University, The Netherlands
Division of Psychosocial Research and Epidemiology, The Netherlands Cancer Institute, Amsterdam, The Netherlands

Xander Verbeek
Netherlands Comprehensive Cancer Organisation (IKNL), Utrecht, The Netherlands

───── **Abstract** ─────

When a person is diagnosed with cancer, difficult decisions about treatments need to be made. In this chapter, we describe an interdisciplinary research project which aims to automatically generate personalized descriptions of treatment options for patients. We relied on two large databases provided by the Netherlands Comprehensive Cancer Organisation (IKNL): The Netherlands Cancer Registry and the PROFILES dataset. Combining these datasets allowed us to extract personalized information about treatment options for different types of cancer. In a next step we provided personalized context to these numbers, both in verbal statements and in narratives, with the aim to facilitate shared decision making about treatments. We discuss strengths and limitations of our approach, illustrate how it generalizes to other health domains, and reflect on the overall research project.

---

[1] Corresponding author

1 Introduction

In his book *Mortality*, the British author Christopher Hitchens memorably describes how following his cancer diagnosis he found himself transported "from the country of the well across the stark frontier that marks off the land of malady" [34]. Every year many other people experience the same transition, and have to adjust to the medical, emotional, and cognitive challenges of their new habitat. One of the first of these challenges involves making informed decisions about possible treatments. Depending on the form and stage of the cancer, there may be a range of treatment options, each resulting in different chances for survival, risks of side effects and general impact on quality of life during and after treatment.

For example, in the case of prostate cancer, treatment options include (but are not limited to) radiotherapy, prostatectomy (removal of part or all of the prostate gland) and active surveillance (monitoring the cancer closely, hoping to avoid invasive treatment) [28]. Interestingly, these treatments have highly comparable survival rates. Hamdy et al. [28] report about 1 death per 1000 person-years for each of these treatments, and also find no significant differences between the treatment groups in the number of deaths from any cause. However, the treatments differ markedly in terms of potential impact on quality of life and side effects, which may range from incontinence or erectile dysfunction, to fear or anxiety [17].

In general, and in the ideal case, doctor and patient (often in concordance with loved ones) jointly decide upon the treatment plan.[2] However, such "shared decision-making" [19] is not always easy, and crucially relies on patients being adequately informed about favorable outcomes (such as long-term survival) and risks of adverse effects (such as side effecs) of potential treatments.

In short, patients can only make an informed choice about treatment and thereby participate in shared decision-making, when they are properly informed. This, however, is not easy. Oncologists generally face difficulties in how to inform their patients, because decisions typically hinge on weighing chances, which are not always easy to access, understand, and explain to patients [26]. Often the benefits and risks of different options for an individual patient are not precisely known in advance, and the statistics are derived from general prediction models, which are based on existing data of earlier cases (e.g., [20]). When doctors want to communicate this kind of information to patients, they can make use of information leaflets or websites, and sometimes also of decision aids. The latter are tools that can help patients to play a role in decision making, together with their doctor, by making the decision points explicit, and linking them to options and possible outcomes [7].

---

[2] A note about terminology: we use the term 'cancer survivor' for anyone who has had a cancer diagnosis and is still alive, while we reserve the term "patient" for anyone who is currently preparing for or receiving treatment.

However, there are three major challenges with this general scenario:

1. There is a focus on survival, with less attention to other factors. Obviously, long term (preferably recurrence-free) survival is a crucial factor for cancer treatment decision making. But, arguably, it is not the only one that should be taken into account. Side-effects of treatment, quality of life after treatment, and general well-being are example factors that patients increasingly consider to be very important as well. These factors will only become more relevant when cancer is increasingly likely to be viewed as a chronic disease [55].

2. Information is often collected and presented on a population level, rather than personalized (e.g., [38, 60, 67]). Still, it is intuitively clear that personalizing such information increases the relevance. Consider the case of prostate cancer again: it makes a big difference whether one receives a first diagnosis at 50 or at 80 years old, both in terms of the progression of the disease and of the potential impact of possible treatment side effects. In this situation, personalized information is clearly much more relevant.

3. Good and clear communication is a *conditio sine qua non* for information presentation about treatment options. However, it is difficult to get this right, especially when information is personalized, which creates two new challenges: (a) It is not obvious how personalized risks and opportunities should be explained to patients, nor to what extent they will be able to understand them. (b) Moreover, when information is personalized, this also means that communication should be different for each individual person, and this is clearly not something that can be obtained with a static patient information leaflet. Hence, this calls for a new approach.

In our research project, we tackled these three related challenges in an interdisciplinary approach, in which statisticians and epidemiologists collaborated with AI-researchers, computational linguists and communication scientists. Additionally, we intensively collaborated with the Netherlands Comprehensive Cancer Organisation (IKNL). IKNL has been actively involved from the very initial stages of the project, helping write the proposal, co-defining the research questions and generally helping to shape and conduct the research we undertook. Crucially, IKNL also provided two data-sets, which we relied on as the foundation of the project: The Netherlands Cancer Registry [44] and the PROFILES dataset [56]. Since 1989, the NCR tracks all new cancer diagnoses in The Netherlands, with data about diagnosis and treatment collected directly from hospitals. It currently contains data of over 2,500,000 persons. PROFILES is a registry for the study of physical and psychosocial impact of cancer and its treatment from a population-based cohort of cancer survivors. Since 2004, over 25,000 people have participated, reporting the consequences of cancer and its treatment on perceived quality of life, symptoms, and societal participation. Combining these datasets allowed us to extract personalized information about treatment options for different types of cancer (as we describe in Section 2 below). This enables us to determine, say, that a patient with cancer-type W and demographic factors X has a probability Y of experiencing treatment side-effect Z. But how do we communicate this to the patient? Do we use words or numbers? Do patients actually want and understand this kind of personalized information? This is addressed in Section 3. Besides describing statistics, it is important to also give personalized context to these numbers, since we can anticipate that merely communicating numbers will not be perceived as helpful and persuasive by everyone. Hence we also explore, in Section 4, whether providing personalized information in the form of tailored narratives is feasible and appreciated. We end this paper, in Section 5, with a general discussion of lessons learned and outstanding issues.

2   Computing personalized information

Computing personalized predictions on quality of life (QoL) for cancer survivors poses a multitude of interesting challenges. Many prediction models have been developed to support decision-making in oncology [30, 57] and there is a large amount of resources and methodological guidance to support the development of new prediction models. However, computing personalized predictions on QoL poses a rather unique problem compared to existing prediction models. Most prediction models have a clearly defined, directly observable outcome, e.g., additional time of survival [68], recurrence-free survival [8], or recurrence of the tumor [9]. In contrast, QoL is neither clearly defined nor directly observable [40]. If we were to ask ten people what their QoL is we would get ten different answers. Not necessarily because their QoL is very different but because they were thinking about different aspects of QoL when providing their answers. While the first person might have been thinking about the physical struggles to carry home a bag of groceries, the second person was happy to feel less pain than last month, and the third person was thankful for still being able to play with his or her grandchildren.

Ideally, we want to account for all aspects of QoL when constructing a prediction model. For this to be possible, we need high-quality data that contains information about all of the aspects of QoL. This requires tedious data collection using questionnaires, often at a very difficult time in the survivors' lives. Thankfully, the PROFILES study [56] offers data of such patient-reported outcome measures (PROMs) for cancer survivors with a wide range of diagnoses and treatments. Furthermore, the data from the PROFILES study can be directly linked to the Netherlands Cancer Registry (NCR [44]). The rich information from this register allows access to all relevant medical information including treatments of the survivors for whom we have PROMs data. The PROFILES study uses validated questionnaires such as the EORTC QLQ-C30, a comprehensive tool for assessing QoL that accounts for its multidimensionality, and its data has been widely used for research on cancer survivorship (for an overview, see [49]).

Subsequently, the model chosen to predict QoL should reflect the multidimensionality of this construct [11, 40]. Rather than developing a prediction model for each dimension separately, our goal was to predict one single outcome that entails all the information of the multiple dimensions of QoL. Keeping in mind the aim of using the predictions in decision aids, instead of providing explanations for predicted scores on each (for some patients potentially uninteresting) dimension of QoL, we aimed for a single outcome that could more easily be communicated to patients (described in Section 3), or even be translated into narratives (Section 4 further elaborates on this use).

To achieve this, we chose to use latent class analysis (LCA), a model-based clustering algorithm. Based on a set of indicators, in this case the dimension scores of QoL, LCA identifies distinct latent classes of cancer survivors with similar patterns of QoL. These QoL classes can then be interpreted based on the (average) values on the QoL items for the survivors in the respective classes. For example, in Clouth et al. [11], one of the five identified classes was characterized by very good scores on all QoL dimensions (class 1; Figure 1), while another class was characterized by very good scores on the social, cognitive, and emotional functioning scales but showed limitations on the physical and role functioning, fatigue, and pain dimensions (class 4), and a third class had moderate to good scores on the (short-term) side effect symptoms but low scores on the more long-term functioning dimensions (class 5).[3]

---

[3] Note that the identification of the best model in terms of number of classes is difficult. Several fit

LCA then estimates the probabilities of membership in these classes for each individual. The more comprehensive interpretations of the classes are more graspable and (potentially) more interesting for patients. Furthermore, these interpretations can directly be used as input for narratives in decision aids.
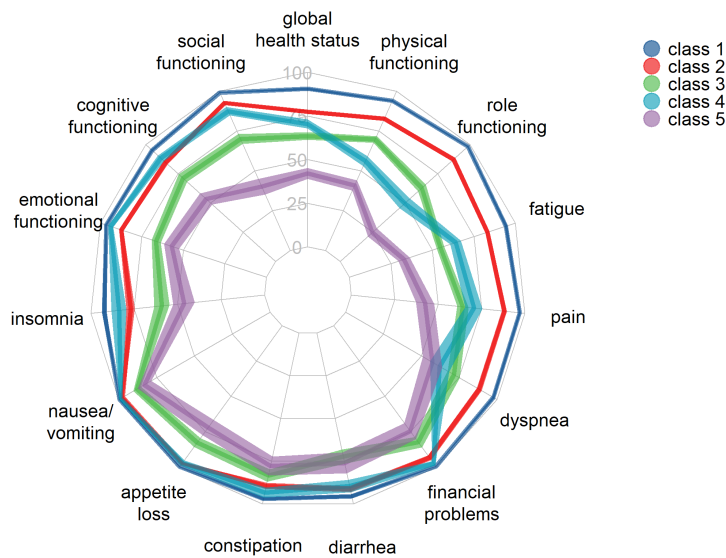
Figure 1 Means and confidence intervals of all five health-related quality of life classes for the 15 European Organisation for Research and Treatment of Cancer Quality of Life Questionnaire C30 dimensions. On the upper half are the functioning scales and on the bottom half are the symptom scales of (short-term) side effects related to treatment. The symptom and single-item scales were reversed.

After identifying the outcome of interest, decision-making requires relevant and person-alized information [19]. When choosing a specific treatment option, organizing follow-up care, and thinking about lifestyle arrangements, predictions based on average QoL can give some indication. However, the real interest lies in personalized predictions, that is, what is the expected QoL for "someone like me".[4] Statistically speaking, "someone like me" refers to individuals of similar age, gender, socioeconomic status, general (physical) health, and similar diagnosis in terms of location, size, and aggressiveness of the tumor. These predictors were included in a regression model to estimate the expected probabilities for membership in the QoL classes conditional on the values of the predictors [11]. With the estimated model

---

indices are available to support decision making but interpretability of the obtained solution is often regarded as equally important.

[4] Additional to personalized predictions, benchmarks might also be important to patients as they often want to know how they are doing compared to others. However, such benchmarks are most often well known and readily available to the treating doctors. We therefore did not focus on this aspect in our project.

parameters, we can then in turn calculate the expected score of class membership probability for every new patient for whom we have the relevant information on the predictors. For instance, imagine a new patient X, who is a 47 years old male with stage 2 colon cancer, no comorbidities, and who received chemotherapy additional to surgery. For this patient, we would then compute a large probability (68%) of being part of class 1 (excellent QoL), and a smaller, but still substantial probability (25%) to be in class 2 (good QoL with prevalence of insomnia). Noticeably, this patient would have a very small probability (2%) of being in class 5 (poor QoL with severe limitations). We deliberately chose for LCA in combination with regression techniques for estimating the personalized predictions to account for the uncertainty around our predictions. While there are many algorithms for clustering and prediction available, parametric approaches such as LCA and regression have the advantage that information about uncertainty, in particular the standard errors of the estimates, are obtained implicitly.

Predictions like the ones described above can be greatly beneficial for patients to support decision-making after diagnosis. However, one aspect of this decision-making, the decision about treatment, requires additional considerations. The question of which treatment option is preferable is a classical "what if" question, that is, a question concerning the causal effect of the treatment [32]. While prediction is aimed at forecasting observations under known conditions, causal inference is aimed at making predictions under unknown conditions. That is, the prediction of the effect of a treatment option A for a specific patient is only valid if that specific patient would actually receive treatment A. The problem is that if there is a good reason for this patient to receive treatment B, only the prediction of the effect of treatment option B (but not of treatment option A) is valid. Since we will only ever observe one of these scenarios, we cannot predict how much better treatment option A is for this specific patient. To arrive at the answer for this, we need to ensure that the patients under treatment A are equal (or similar enough) to the patients under treatment B (exchangeability assumption [32]). While this is usually achieved using randomized controlled trials, it is a major challenge when using observational data (such as registry data). Several causal inference techniques have been developed to overcome these challenges [32]. Unfortunately, these techniques are not directly applicable to be used in combination with clustering algorithms such as LCA. In Clouth et al. [12] and Clouth et al. [13], we developed new methods that combine these established causal inference techniques with LCA. In particular, we show how inverse propensity weighting (IPW) can be included in LCA to estimate the causal effect of surgery compared to an active surveillance strategy for low-stage prostate cancer patients. After estimating a measurement model for the latent QoL classes, the IPW weights can be combined with weights for the measurement errors from this first model. These new weights can then be included in a regression model to estimate the causal effect. In Clouth et al. [12], we show that there is no causal effect of surgery compared to an active surveillance strategy on the probability of class memberships in either good or bad QoL classes.

The impact of the methodological work in this project is beyond the domain of oncology and survivorship research. We validated the methodology of using a combination of LCA and regression techniques as well as causal inference techniques on several other data sets, both cross-sectional and longitudinal. The examples range from Australian data on breast cancer survivors [15], to data on trauma patients admitted to the ICU [29], cohort data from the US on drug abuse [13], Dutch panel data on mental health [10], and data on self-awareness and autonomy [42].

3   Explaining personalized information

A critical question that needs to be addressed is whether patients really want to receive personalized statistics about treatment outcomes obtained via prediction models. Should we disclose such statistics, and if so, how and to whom? By conducting both qualitative [37, 36, 66] and quantitative [61] studies, we found that the vast majority of patients and survivors, especially those with higher subjective numeracy skills and an information-seeking coping style, have a desire to receive personalized statistics about risks of treatment side effects, survival, and quality of life. Patients' considerations for wanting personalized cancer statistics related to feelings of being in control or making better-informed decisions, while considerations for not wanting statistics were about the unpredictability of future events for individual patients or negative experiences with statistics in the past. Personalized statistics were also considered more useful and relevant compared to generic, non-personalized statistics. Importantly, and in line with related research (e.g., [23, 58, 70]), whether personalized risks should be communicated to patients strongly depends on individual differences. For instance, those who are not actively seeking for detailed information and who are subjectively less numerate may have less desire for received personalized statistics information about treatment outcomes [48].

Once we established that patients have a need for receiving personalized statistical information, we wanted to explore what this would add to clinical practice. In two qualitative studies, we interviewed patients (N=35) and clinicians (N=6) about the role of personalized statistics in clinical practice. What is missing now and how could clinical practice benefit from personalized statistical information? We found that patients have a general "gist" knowledge about personalized predictions, they understand for example that their outcomes would improve over time and that some outcomes might take a while before they improve [6, 37, 36]. It seemed more difficult however, to understand the prediction model, the input data, and the statistical procedures that were responsible for the personalized predictions. Figure 2 shows an example of how the (input to the) prediction model was visualized. Even without precise knowledge on the working of the prediction model, receiving predictions reassured patients and most liked knowing what life would look like after their disease – especially if their prospects were not too positive. Some patients were not that worried about life after their diagnosis so they experienced a lesser need for detailed information. The most important implication for clinical practice would be that patients indicated they would initiate conversations with their doctor more about how they were doing and they would think about what to say during a consultation beforehand.

However, merely providing patients with personalized data is not enough. These numbers need to be effectively and adequately communicated to help patients make sense of their personalized health information. Therefore, we conducted several pre-registered experiments among a sample of cancer patients and survivors (N=141) and a Dutch representative sample of healthy participants (N=1807) with the aim to examine how varying message formats and contextual strategies for communicating personalized risks can influence people's cognitive, emotional, and behavioral responses. When presenting personalized risks to patients, they can best be communicated numerically using natural frequencies (e.g., "40 out of 100 men like you will experience this side-effect") combined with verbal descriptors (e.g., "this risk is common") and/or visual displays such as icon arrays [62, 66]. As illustrated in Figure 3, verbal-only formats (without any numerical data) should be avoided, as they may lead to variable and inaccurate risk perceptions (for similar reasoning, see [5]).

Figure 2 Example of the way the input data for the prediction model could be visualized. This example was used for trauma patients (people who had been in an accident) and was based on the prediction model developed in [29].

Moreover, even single risk statistics that are personalized and presented in isolation may be difficult to evaluate by patients and are sometimes even ignored by them [21, 72]. Results from both needs assessments and an experimental study show that people generally want comparative data about the average person's risk or survival rate when being provided with their personalized risk. By doing so, patients are better able to interpret even unfamiliar personalized risks because both types of risks serve as a reference for each other, which helps patients determine the "goodness" and the "badness" of their personalized health data [71]. Our studies show that providing such comparative data (Figure 4) increases information evaluability and can help patients derive meaning from their personalized risk data, without negatively influencing their risk perceptions, emotional responses or behavioral treatment intentions [63]. In addition, providing simple explanations about how personalized risks are determined (e.g., by discussing risk factors) may help patients better understand and recall those risks, and further suggest that they can play an important role in shared decision-making about treatment [64].

4 Narrating personalized information

One of the challenges in communicating statistics to patients is that people might not always correctly interpret numbers [26]. This lack of numeracy skills might in turn influence their perceived usefulness or negatively impact health-related decisions [31]. One way of
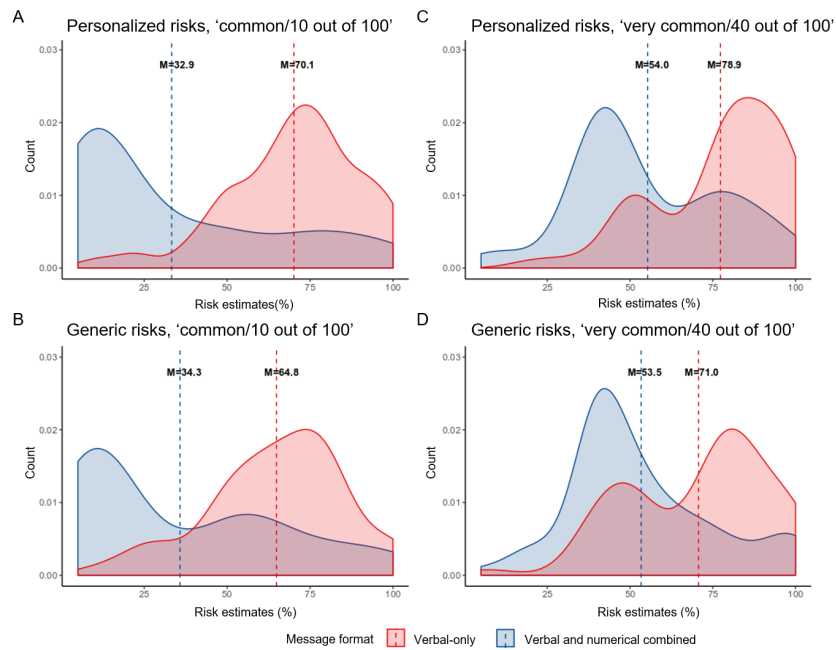
Figure 3 Variation in risk estimates given by patients when being presented with personalized and generic risks presented via verbal descriptors (e.g., "this side effect is common") with or without numbers ("occurs in 10 out of 100 patients like you").

overcoming this challenge might be by contextualizing numbers in stories, often called narratives. Narratives can be defined as "a representation of connected events and characters that has an identifiable structure, is bounded in space and time, and contains implicit or explicit messages about the topic being addressed" ([41], p.222). In health communication, narratives have been extensively studied in persuasive contexts where readers need to be persuaded to make healthier choices. They have proven effective tools in trying to persuade people to eat healthier, attend screening or quit smoking for example [54]. Narratives might also lead to better recall, fewer counterarguments and experiencing more positive emotions [52].

In informative settings, the setting of our project, the role of narratives is less straightforward. Whether or not narratives are helpful in informing readers might be measured by a range of variables including for example better recall, knowledge increase or risk estimation. In these informative settings, where narratives need not persuade readers, there is worry for the narrative bias effect [14, 39, 53]. Narrative bias entails that a narrative is so powerful that it outweighs statistical or factual information, even when the narrative counters the statistical information. It seems that this is especially the case when narratives are about adverse effects of health decisions, such as a narrative portraying severe side effects after a vaccine [27]. To combat narrative bias and to still make use of its positive effects, it is arguably more fruitful to develop narratives that are in line with the statistical information rather than narratives that counter the facts. That is why in our final studies, we used balanced narratives that are in line with the statistical information.

One way of making narratives more appealing and relevant for readers is by personalizing them. Personalizing (or tailoring) health information has shown promising results. Personalization means that the narrative or the character in the narrative is matched to the reader of

PERSONALIZED RISK

GENERIC RISK

Risk for men like you

For men like you (with the same age and the same type of colon cancer) 38 out of 100 experience neuropathy.

Risk for all people

For all people (of all ages with all types of colon cancer) 18 out of 100 experience neuropathy.

Side effect   No side effect

Side effect   No side effect

■ **Figure 4** Providing comparative data about the average person's risk may help patients better evaluate their personalized risk data.

that narrative. To give an example, Baezconde-Garbanati et al. [2] showed that a narrative personalized to the cultural background of readers was more effective in persuading readers to undergo cervical cancer screening than a narrative that was not personalized to the reader. In addition, a systematic review by Noar et al. [45] found that in general, personalized health information leads to better health outcomes than generic information. Additionally, studies found that the first-person perspective leads to better effects than other perspectives [51].

A downside to personalized health information however may be that it is a time-consuming and labor-intensive process if all information should be personalized based on individual characteristics of the reader. That is why we call for automatic personalization rather than manually changing health information [4, 47]. By applying basic data-to-text techniques, we are now able to construct a narrative that is personalized based on individual patient characteristics. Figure 5 shows examples of narratives we used in our study and how they were personalized depending on the reader.

Figure 5 The first paragraphs of two personalized narratives from fictional patients 'David' and 'Hanneke'. Narratives are personalized based on the reader's characteristics. These included: gender (male/female), age (old, middle, young), diagnosis (skin, bladder, breast cancer), physical functioning score (better or worse than critical value) and role functioning (better or worse than critical value). Additionally, patients could view their personal scores in a table at the end of the narrative. The photos were automatically generated from a website (`https://generated.photos/`). Narratives were constructed based on previously conducted interviews about patient experiences with cancer and were developed together with an expert team including clinicians, communication and data experts.

5 General Discussion and Conclusion

In the foregoing, we have described our novel interdisciplinary framework with the aim of better informing cancer patients about treatment options. By combining different data sources, such as the Netherlands Cancer Registry and the PROFILES dataset, we were able to compute personalized profiles on treatment side-effects and general quality of life after treatment, for patients newly diagnosed with cancer, using latent class analysis. On the basis of these profiles, we could generate personalized information about a range of associated outcomes of different treatment options to patients. These were communicated via various formats, ranging from words and numbers to personalized narratives, in which case both the information and its presentation can be tailored towards the reader. Evaluation studies revealed that which formats are preferred and work best for information presentation is dependent on personal and contextual features. Our approach towards personalized information on treatment options has been evaluated in the oncology domain, but can be generally applied to other healthcare contexts (as we have shown for trauma care). In this final section, we discuss a number of challenges and outstanding issues.

5.1 Challenges when using registry data

Personalized predictions are based on data and the quality of these predictions is inherently linked to the quality of the data. The NCR collects data about diagnosis and treatment directly from hospitals for all new cancer diagnoses in The Netherlands since 1989. Furthermore, the PROFILES study complements the information in the NCR with PROMs data on quality of life for a subset of the patients and survivors in the registry. The direct linkage between these two data sources allows for extracting all the information from the NCR when developing prediction models for QoL. On the other hand, using retrospective registry data is also restrictive.

In particular, we encountered three challenges.

1. Registry data is observational, that is, this data does not contain experimental interventions of any kind. Information is observed for treatments that are administered based on the survivors' characteristics (rather than randomly) and outcomes are observed as a consequence of all of these influencing factors rather than as a consequence of treatment alone. This indication bias, otherwise known as confounding by indication or self-selection into treatment bias, prohibits direct conclusions about the effectiveness of certain treatment options. As discussed in Section 2, several causal inference methods have been developed to estimate causal effects from observational data. Using these methods allows us to give answers to the "what if" questions about which treatment option is preferable.

2. Information from registry data might be outdated to address decisions today. Registry data is retrospective; information on patients is collected over many years and patients might have been treated years ago with methods that might be considered outdated by today's standards. For example, in many cases, lymph node dissections are no longer performed or chemotherapy is replaced by immune therapy. Consequentially, outcome information of recent treatments is often limited in registry data. Evidence-based predictions about such recent treatments are then not possible using registry data.

3. Registry data contains survivorship bias (or immortal time bias). Patients for whom a treatment is registered survived at least until they received such treatment. The information about treatments and their effectiveness is thus based on the subset of survivors among all patients. This problem intensifies when considering outcomes such as QoL as these outcomes are of interest and registered some time (often years) after treatment. The information about QoL is thus based on an even smaller subset of patients that survived treatment and some time after. Put differently, we will never observe QoL for patients who did not survive. However, as this set of patients is still present at the time of decision-making about the treatment, predictions based on the subset of survivors will be biased (survivorship bias). A solution for this problem might be the use of baseline assessments of QoL. Baseline QoL has been shown to often be predictive of QoL after treatment and such information is increasingly often being collected. As QoL is always conditional on survival, this information is still very valuable, for example, to describe the current state of cancer survivorship in the Netherlands. However, as input for decision-making about the treatment, the potential of this data is limited.

As a consequence, we draw two conclusions: (1) Predictions for QoL based on data from long-term survivors are often inadequate for decision-making about treatment. And (2): we need to base predictions for long-term QoL on the observed treatment and these predictions are always conditional on survival. Facing these limitations, we need to stress that these predictions (predictions conditional on survival and based on the observed treatment) can still be tremendously useful for cancer survivors. A cancer diagnosis changes a person's life

fundamentally and even after important decisions about the treatment have been made, there will still be a large amount of uncertainty about the future. Tools like the ones described above can play an important role in understanding what a survivor's life will look like.

5.2 The role of text generation for personalized treatment information

Obviously, when we move from generic to personalized information about treatment options, this implies that static information presentation (on a leaflet, say) is no longer feasible. Fortunately, tools for automatic information presentation based on data have been under development for many years (e.g., [50]). Currently, there are two broad classes of systems [24]: rule-based and statistical ones, both existing in many variants. Of the latter category, generative large language models (or LLMs) are the most prominent manifestation, including, for example, OpenAI's ChatGPT system which took the world's fancy late 2022. While LLMs are capable of producing text with impressive depth and fluency, they suffer from a number of problems. There are general concerns about, for example, the inherent biases and stereotyping of LLMs, the ecological footprint of training them and the poor quality for low-resource languages, see e.g., [3] for further discussion.

Two specific problems for current purposes are the following: (1) The texts generated by LLMs need not be factually true and correct, for example, saying that a treatment has a relatively low risk of a particular side-effect while in truth the risk is rather high or unknown; (2) They may hallucinate information, that is: generate output that does not correspond to anything in the input, for example discussing a non-existing or dispreferred, outdated treatment.

For these reasons, it is currently assumed that LLMs are not practically useful in the kind of health contexts we consider in this project, and could even be harmful [1]. Some LLM model-developers therefore explicitly rule out such applications. For example, the HuggingFace website for BLOOM (an open-access LLM with 176B parameters; [69]) explicitly lists the biomedical domain as well as "critical automated decisions" as out-of-scope uses of their model. For all these reasons, in this project we relied on more traditional rule-based text generation models, which are capable of generating a wide-variety of texts, following a number of fixed patterns [35].

As a result, it can be guaranteed that generated output is always factually correct and hallucination-free. The price to be paid for this is that the generated output is not as fluently flexible as LLM-generated text. In future research, more work needs to be done on harnessing the strengths of both methods, paving the way for fluid and flexible generated output that can be controlled for factualness and lack of hallucinations.

5.3 Investigating other modalities during personalized risk communication

A critical note on our approach has been its predominant focus on developing communication methods for representing personalized statistical information in the written, non-interactive communication domain (e.g., presenting statistics and probabilities in web-based patient decision aids). Relatively underexplored has been the effect of non-verbal communication of personalized information on patients' responses during spoken risk communication [25]. For instance, when doctors are using clinical prediction models during consultations with their patients, they can non-verbally signal uncertainty around individualized risk estimates by using vocal cues such as a rising, question-like intonation contour or preceded by filled pauses (e.g., "uhm...") [16].

So far, these non-verbal or prosodic cues have largely been ignored by risk communication researchers, while other studies show that they may have a substantially larger impact on people's risk and uncertainty perceptions than the risk estimates themselves [65]. As a next step for future research, we should discover which non-verbal communication strategies of disclosing risk information during patient-doctor interactions are most effective while also mitigating potential negative psychological responses such as fear or losing trust. Understanding these effects will help doctors to communicate individualized outcome information obtained via prediction models to their patients in a transparent and understandable way.

5.4 Communicating to patients with lower health literacy

When dealing with health information provision to patients, it is crucial that this information is understandable for all. However, patients might have lower health literacy skills, come from diverse socio-economic backgrounds and have different education levels. Especially when designing eHealth solutions, it is therefore crucial to take into account the (digital) health literacy of patients. A definition of (digital) health literacy is "the degree to which individuals have the capacity to obtain, process, and understand basic [digital] health information and services needed to make appropriate health decisions" [46, 18]. Patients with lower health literacy skills are less likely to participate in research than patients with higher health literacy skills, better socio-economic background and higher education levels [33, 22]. We recommend researchers to work closely together with organizations that have ample experience with communicating to this specific group and advise to measure health literacy skills so that samples can be described and recruitment strategies may be adjusted when few patients with low health literacy have been recruited. Additionally, we also conducted some studies with a representative sample of non-patients as some effects (e.g., effects on knowledge, utility) can be tested without burdening vulnerable patient populations.

5.5 Lessons learned about the collaborations

Besides involving patients in our project, we collaborated with many other stakeholders. The starting point of our project was the collaboration between Tilburg University and the Netherlands Comprehensive Cancer Organization (IKNL). This collaboration ensured accessibility to data and a mix of researchers from different disciplines (e.g., statistics, communication, data science, epidemiology, linguistics, (clinical) psychology). Although each discipline often had their own terminology, we found that bringing together ideas from different fields strengthened the feasibility of explaining personalized predictions in practice.

The close collaboration with IKNL was unusual, and a key-factor to the success of the project. The collaboration started with the preparation of the project proposal, and continued throughout. We could not have carried out this project without access to and support for using the NCR and the PROFILES registry. The PhD researchers working on this project also spend part of their time at IKNL, facilitating this process. Additionally, monthly consortium meeting helped in making sure that research ideas were grounded in practice, and the feasibility and applicability of the approach were a constant point of attention in these meetings. Finally, IKNL provided a bridge to actual hospital practice, which was instrumental in the evaluation studies we conducted.

During the project we expanded the scope to different health domains in oncology and physical trauma by working closely together with the Antoni van Leeuwenhoek Hospital (skin, bladder, breast cancer; data set: quality of life data obtained through the EORTC-QLQ-C30), the Catharina Hospital (colorectal cancer; data set: PROFILES, NCR) and the Elisabeth

TweeSteden Hospital (physical trauma; data set Brabant Injury Outcome Surveillance). This gave us the rare opportunity to work on a solution that would generalize to different disease domains, and provided insights into different patient communication traditions across health areas, which are not usually documented. Finally, we also collaborated with different information providers (Kanker.nl, ZorgkeuzeLab, Interactive Studios) that are experienced with communicating complex health information in an understandable way to patients.

5.6 Implications for (clinical) practice

The idea that patients should receive personalized information might seem simple, the clinical reality is unruly. Estimates show that it takes around 17 years for research evidence to be implemented in clinical practice [43] and although eHealth solutions might seem promising, uptake can be low [59]. Nevertheless, the findings of our project also have immediate implications for real-world clinical practice, including patient decision aid developers, health-care professionals, and general websites about cancer such as online cancer communities. For instance, as a result of the findings presented in this paper, a real-life web-based tool "Specifieke Cijfers" (Specific Numbers) will be launched in 2023 on the Dutch online cancer community website Kanker.nl (`https://kanker.nl`), which will be able to communicate personalized cancer statistics such as (conditional) survival to patients and their relatives. In this tool, visitors will have the opportunity to enter a number of personal and disease related characteristics, for which in return they will see the number of similar patients that are still alive 5 years after diagnosis. Similar to the statistical prediction models outlined above, these personalized cancer statistics are derived from a real-life connection with NCR. Moreover, in line with our project's findings, the most recent version of this tool discloses personalized statistical information while keeping the information and data entry characteristics short and concise, and also provides contextual information by explaining the gist (or bottom-line) meaning of conditional survival outcomes. Finally, insights from our project have already been acknowledged, valued, and adopted outside the oncology domain, such as the trauma setting where patients need to be informed about their recovery after an injury. For this, our data-driven approach for determining personalized narratives will soon be developed and implemented in the Patient Journey App, a patient-centered application for patients used by the Elisabeth Tweesteden Hospital in the Netherlands.

── **References** ──

1 A. Arora and A. Arora. The promise of large language models in health care. *The Lancet*, 401:641, 2023.

2 L. A. Baezconde-Garbanati, J. S. Chatterjee, L. B. Frank, S. T. Murphy, M. B. Moran, L. N. Zhao N. Werth, P. Amezola de Herrera, D. Mayer, J. Kagan, and D. O'Brien. Tamale lesson: A case study of a narrative health communication intervention. *Journal of Communication in Healthcare*, 7(2):82–92, 2014.

3 E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623, 2021.

4 N. Bol, E. Smit, and M. Lustria. Tailored health communication: Opportunities and challenges in the digital era. *Digital Health*, 6:1–3, 2020.

5 C. Bonner, L. J. Trevena, W. Gaissmaier, P. K. J. Han, Y. Okan, E. Ozanne, E. Peters, D. Timmermans, and B. J. Zikmund-Fisher. Current best practice for presenting probabilities in patient decision aids: fundamental principles. *Medical Decision Making*, 41(7):821–833, 2021.

6   E. Boomstra, S. Hommes, R. D. Vromans, A. M. van den Burg, A. M. Schrijver, M. W. J. M. Wouters, I. M. C. van der Ploeg, M. W. van de Kamp, E. J. Krahmer, L. V. van de Poll-Franse, and K. de Ligt. Numbers call for action, personalized narratives provide support and recognition": A qualitative assessment of cancer patients' perspectives on patient reported outcome measures (proms) feedback with narratives. *Journal of Cancer Survivorship*, 2024.

7   C. Charles, A. Gafni, T. Whelan, and M. A. O'Brien. Treatment decision aids: conceptual issues and future directions. *Health Expectations*, 8(2):114–125, 2005.

8   S. S. Chauhan, J. Kaur, M. Kumar, A. Matta, G. Srivastava, A. Alyass, J. Assi, I. Leong, C. MacMillan, I. Witterick, T. J. Colgan, N. K. Shukla, A. Thakar, M. C. Sharma, K. W. M. Siu, P. G. Walfish, and R. Ralhan. Prediction of recurrence-free survival using a protein expression-based risk classifier for head and neck cancer. *Oncogenesis*, 4(4):e147, 2015.

9   P. C. Chen, Y. M. Yeh, B. W. Lin, R. H. Chan, P. F. Su, Y. C. Liu, C. T. Lee, S. H. Chen, and P. C. Lin. A prediction model for tumor recurrence in stage ii-iii colorectal cancer patients: From a machine learning model to genomic profiling. *Biomedicines*, 10(2), 2022.

10   F. J. Clouth, M. J. Bijlsma, S. Pauws, and J. K. Vermunt. The parametric g-formula for latent Markov models, 2024. PsyArXiv preprint.

11   F. J. Clouth, A. Moncada-Torres, G. Geleijnse, F. Mols, F. N. van Erning, I. H. J. T. de Hingh, S. C. Pauws, L. V. van de Poll-Franse, and J. K. Vermunt. Heterogeneity in quality of life of long-term colon cancer survivors: A latent class analysis of the population-based profiles registry. *The Oncologist*, 26(3):e492–e499, 2021.

12   F. J. Clouth, S. Pauws, F. Mols, and J. K. Vermunt. A new three-step method for using inverse propensity weighting with latent class analysis. *Advances in Data Analysis and Classification*, 16(2):351–371, 2022. `doi:10.1007/S11634-021-00456-5`.

13   F. J. Clouth, S. Pauws, and J. K. Vermunt. Three-step latent class analysis with inverse propensity weighting in the presence of differential item functioning. *Structural Equation Modeling: A Multidisciplinary Journal*, 30(5):737–748, 2023. `doi:10.1080/10705511.2022.2161384`.

14   A. De Graaf, J. Sanders, and H. Hoeken. Characteristics of narrative interventions and health effects: A review of the content, form, and context of narratives in health-related narrative persuasion research. *Review of Communication Research*, 4:88–131, 2016.

15   B. H. de Rooij, I. Ramsey, F. J. Clouth, N. Corsini, J. S. Heyworth, B. M. Lynch, J. K. Vallance, and T. Boyle. The association of circadian parameters and the clustering of fatigue, depression, and sleep problems in breast cancer survivors: a latent class analysis. *Journal of Cancer Survivorship*, 17(5):1405–1415, 2022.

16   C. Dijkstra, E. Krahmer, and M. Swerts. Manipulating uncertainty: The contribution of different audiovisual prosodic cues to the perception of confidence. In *Proceedings of the Speech Prosody Conference*, pages 1–4, 2006.

17   J. L. Donovan, F. C. Hamdy, J. A. Lane, M. Mason, C. Metcalfe, E. Walsh, ..., and ProtecT Study Group. Patient-reported outcomes after monitoring, surgery, or radiotherapy for prostate cancer. *New England Journal of Medicine*, 375(15):1425–1437, 2016.

18   P. Dunn and E. Hazard. Technology approaches to digital health literacy. *International Journal of Cardiology*, 293:294–296, 2019.

19   G. Elwyn, D. Frosch, R. Thomson, N. Joseph-Williams, A. Lloyd, P. Kinnersley, Cording E, Tomson D, Dodd C, Rollnick S, Edwards A, and M. Barry. Shared decision making: a model for clinical practice. *Journal of General Internal Medicine*, 27:1361–1367, 2012.

20   E. G. Engelhardt, A. H. Pieterse, N. van Duijn-Bakker, J. R. Kroep, H. C. de Haes, E. M. Smets, and A. M. Stiggelbout. Breast cancer specialists' views on and use of risk prediction models in clinical practice: a mixed methods approach. *Acta Oncologica*, 54(3):361–367, 2015.

21   A. Fagerlin, B. J. Zikmund-Fisher, and P. A. Ubel. If i'm better than average, then i'm ok? comparative information influences beliefs about risk and benefits. *Patient Education and Counseling*, 69(1-3):140–144, 2007.

22 J. G. Ford, M. W. Howerton, G. Y. Lai, T. L. Gary, S. Bolen, M. C. Gibbons, J. Tilburt, C. Baffi, T. P. Tanpitukpongse, R. F. Wilson, N. R. Powe, and E. B. Bass. Barriers to recruiting underrepresented populations to cancer clinical trials: A systematic review. *Cancer*, 112:228–242, 2008.

23 S. J. Franssen, S. M. Lagarde, J. R. van Werven, E. M. Smets, K. T. Tran, J. T. M. Plukker, J. J. van Lanschot, and H. C. de Haes. Psychological factors and preferences for communicating prognosis in esophageal cancer patients. *Psycho-Oncology: Journal of the Psychological, Social and Behavioral Dimensions of Cancer*, 18(11):1199–1207, 2009.

24 A. Gatt and E. Krahmer. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *Journal of Artificial Intelligence Research*, 61:65–170, 2018. `doi:10.1613/JAIR.5477`.

25 J. Geipel, C. Hadjichristidis, L. Savadori, and B. Keysar. Language modality influences risk perception: Innovations read well but sound even better. *Risk Analysis*, 43:558–570, 2022.

26 G. Gigerenzer, W. Gaissmaier, E. Kurz-Milcke, L. M. Schwartz, and S. Woloshin. Helping doctors and patients make sense of health statistics. *Psychological Science in the Public Interest*, 8(2):53–96, 2008.

27 N. Haase, C. Betsch, and F. Renkewitz. Source credibility and the biasing effect of narrative information on the perception of vaccination risks. *Journal of Health Communication*, 20(8):920–929, 2015.

28 F. C. Hamdy, J. L. Donovan, J. A. Lane, M. Mason, C. Metcalfe, P. Holding, ..., and ProtecT Study Group. 10-year outcomes after monitoring, surgery, or radiotherapy for localized prostate cancer. *New England Journal of Medicine*, 375(15):1415–1424, 2016.

29 R. J. M. Havermans, F. J. Clouth, K. W. W. Lansink, J. K. Vermunt, M. A. C. de Jongh, and L. de Munter. Prediction of recovery in trauma patients using latent Markov models. *European Journal of Trauma and Emergency Surgery*, 48(3):2059–2080, 2022.

30 B. He, W. Chen, L. Liu, Z. Hou, H. Zhu, H. Cheng, Y. Zhang, S. Zhan, and S. Wang. Prediction models for prognosis of cervical cancer: Systematic review and critical appraisal. *Frontiers in Public Health*, 9:654454, 2021.

31 L. Heilmann. Health and numeracy: the role of numeracy skills in health satisfaction and health-related behaviour. *ZDM Mathematics Education*, 52:407–418, 2020.

32 M. A. Hernan and J. M. Robins. Estimating causal effects from epidemiological data. *Journal of Epidemiology and Community Health*, 60:578–586, 2006.

33 R. Hernandez-Ramos, A. Aguilera, F. Garcia, J. Miramontes-Gomez, L. E. Pathak, C. A. Figueroa, and C. R. Lyles. Conducting internet-based visits for onboarding populations with limited digital literacy to an mhealth intervention: Development of a patient-centered approach. *JMIR Formative Research*, 5(4):25299, 2021.

34 C. Hitchens. *Mortality*. Hachette, UK, 2012.

35 S. Hommes, C. van der Lee, F. J. Clouth, X. Verbeek, and E. J. Krahmer. A personalized data-to-text support tool for cancer patients. In *Proceedings of the 12th International Conference on Natural Language Generation, Tokio, Japan*, pages 443–452, 2019.

36 S. Hommes, R. Vromans, M. de Jongh, T. Houwen, N. Bol, E. de Groot, M. Antheunis, and E. Krahmer. Knowing what life has in store for me is better than being in the dark: Trauma patients' perspectives on receiving personalized predictions for life after injury, 2023.

37 S. Hommes, R. D. Vromans, J. Bloemen, S. van Gent, G. Schipper, X. Verbeek, and E. J. Krahmer. Personalized predictions about quality of life after chemotherapy: Development of a data-driven support tool and initial evaluation of colorectal cancer patients and clinicians, 2023.

38 S. Hommes, R. D. Vromans, F. J. Clouth, X. Verbeek, I. de Hingh, and E. J. Krahmer. Communication in decision aids for stage i-iii colorectal cancer patients: A systematic review. *BMJ Open*, 11:e044472, 2021.

39   A. Kalch, C. Küchler, A. Albani, H. Bilandzic, S. Fischer, and I. Kirchberger. On the need for narratives in patient information: Differentiating types and functions of narratives from pulmonary embolism patients' point of view. *Frontiers in Communication*, 7:990819, 2022.

40   K. Kenzik, M. Y. Martin, M. N. Fouad, and M. Pisu. Health-related quality of life in lung cancer survivors: Latent class and latent transition analysis. *Cancer*, 121(9):1520–1528, 2015.

41   M. W. Kreuter, M. Green, J. Cappella, M. Slater, M. Wise, D. Storey, E. M. Clark, D. J. O'Keefe, D. O. Erwin, K. Holmes, L. J. Hinyard, T. Houston, and S. Woolley. Narrative communication in cancer prevention and control: A framework to guide research and application. *Annals of Behavioral Medicine*, 33(3):221–235, 2007.

42   L. E. Kunst, M. A. L. M. van Assen, F. J. Clouth, C. Hunt, M. Abbott, J. Maas, and M. H. J. Bekker. Who cares what other people think? a longitudinal investigation of long-term self-esteem stability and its associations with autonomy-connectedness, 2024. Manuscript submitted for publication.

43   Z. S. Morris, S. Wooding, and J. Grant. The answer is 17 years, what is the question: understanding time lags in translational research. *Journal of the Royal Society of Medicine*, 104(12):510–520, 2011.

44   Netherlands Cancer Registry (ncr) and Netherlands Comprehensive Cancer Organisation. (iknl), derived via www.iknl.nl/en/ncr/ncr-data-figures [03/31/2023].

45   S. M. Noar, C. N. Benac, and M. S. Harris. Does tailoring matter? meta-analytic review of tailored print health behavior change interventions. *Psychological Bulletin*, 133(4):673–693, 2007.

46   R. Parker and S. C. Ratzan. Health literacy: A second decade of distinction for americans. *Journal of Health Communication*, 15(2):20–33, 2010.

47   S. Pauws, A. Gatt, E. Krahmer, and E. Reiter. Making effective use of healthcare data using data-to-text technology. In *Data Science for Healthcare*, pages 119–145. Springer, 2019. `doi:10.1007/978-3-030-05249-2_4`.

48   E. Peters. *Innumeracy in the wild: Misunderstanding and misusing numbers*. Oxford University Press, 2020.

49   PROFILES.   derived   via.   URL: `https://www.dataarchive.profilesregistry.nl/publications`.

50   E. Reiter and R. Dale. *Building Natural Language Generation Systems*. Cambridge University Press, 2000.

51   S. Salem, T. Weskott, and A. Holler. Does narrative perspective influence readers' perspective-taking? an empirical study on free indirect discourse, psycho-narration and first-person narration. *Glossa: A Journal of General Linguistics*, 2(1):61, 2017.

52   V. A. Shaffer, E. S. Focella, A. Hathaway, L. D. Scherer, and B. J. Zikmund-Fisher. On the usefulness of narratives: An interdisciplinary review and theoretical model. *Annals of Behavioral Medicine*, 52(5):429–442, 2018.

53   V. A. Shaffer and B. Zikmund-Fisher. All stories are not alike: A purpose-, content- and valance-based taxonomy of patient narratives in decision aids. *Medical Decision Making*, 33(1):4–13, 2013.

54   F. Shen, V. C. Sheer, and R. Li. Impact of narratives on persuasion in health communication: A meta-analysis. *Journal of Advertising*, 44(2):105–113, 2015.

55   R. L. Siegel, K. D. Miller, and A. Jemal. Cancer statistics. *2016. CA: A Cancer Journal for Clinicians*, 66(1):7–30, 2016.

56   L. V. van de Poll-Franse, N. Horevoorts, M. van Eenbergen, J. Denollet, J. A. Roukema, N. K. Aaronson, A. Vingerhoets, J. W. Coebergh, J. de Vries, M. L. Essink-Bot, and F. Mols. The patient reported outcomes following initial treatment and long term evaluation of survivorship registry: scope, rationale and design of an infrastructure for the study of physical and psychosocial outcomes in cancer survivorship cohorts. *European Journal of Cancer*, 47(14):2188–2194, 2011.

57 H. G. van den Boorn, E. G. Engelhardt, J. van Kleef, M. A. G. Sprangers, M. G. H. van Oijen, A. Abu-Hanna, A. H. Zwinderman, V. M. H. Coupe, and H. W. M. van Laarhoven. Prediction models for patients with esophageal or gastric cancer: A systematic review and meta-analysis. *PLoS One*, 13(2):e0192310, 2018.

58 N. C. van der Velden, H. W. van Laarhoven, S. A. Burgers, L. E. Hendriks, F. Y. de Vos, A. M. C. Dingemans, J. Jansen, J. W. van Haarst, J. Dits, E. M. Smets, and I. Henselmans. Characteristics of patients with advanced cancer preferring not to know prognosis: A multicenter survey study. *BMC Cancer*, 22(1):1–12, 2022.

59 J. E. van Gemert-Pijnen, N. Nijland, M. van Limburg, H. C. Ossebaard, S. M. Kelders, G. Eysenbach, and E. R. Seydel. A holistic framework to improve the uptake and impact of ehealth technologies. *Journal of Medical Internet Research*, 13(4):e111, 2011.

60 R. Vromans, K. Tenfelde, S. Pauws, M. Van Eenbergen, I. Mares-Engelberts, G. Velikova, L. van de Poll-Franse, and E. Krahmer. Assessing the quality and communicative aspects of patient decision aids for early-stage breast cancer treatment: a systematic review. *Breast Cancer Research and Treatment*, 178:1–15, 2019.

61 R. D. Vromans, S. Hommes, F. J. Clouth, D. N. N. Lo-Fo-Wong, X. A. A. M. Verbeek, L. V. van de Poll-Franse, S. C. Pauws, and E. J. Krahmer. Need for numbers: Assessing cancer survivors' needs for personalized and generic statistical information. *BMC Medical Informatics and Decision Making*, 22(1):1–14, 2022.

62 R. D. Vromans, S. C. Pauws, N. Bol, L. V. van de Poll-Franse, and E. J. Krahmer. Communicating tailored risk information of cancer treatment side effects: Only words or also numbers? *BMC Medical Informatics and Decision Making*, 20:1–12, 2020.

63 R. D. Vromans, S. C. Pauws, L. V. van de Poll-Franse, and E. J. Krahmer. Effects of comparative information when communicating personalized risks of treatment outcomes: An experimental study. *Journal of Risk Research*, 26(3):324–343, 2022.

64 R. D. Vromans, C. N. Tillier, S. C. Pauws, H. G. van der Poel, L. V. van de Poll-Franse, and E. J. Krahmer. Communication, perception, and use of personalized side-effect risks in prostate cancer treatment decision-making: An observational and interview study. *Patient Education and Counseling*, 105(8):2731–2739, 2022.

65 R. D. Vromans, C. C. M. van de Ven, S. J. W. Willems, E. J. Krahmer, and M. Swerts. It is, uh, very likely? the impact of prosodic uncertainty cues on the perception and interpretation of spoken verbal probability phrases. *Risk Analysis*, pages 1–20, 2024.

66 R. D. Vromans, M. C. van Eenbergen, G. Geleijnse, S. C. Pauws, L. V. van de Poll-Franse, and E. J. Krahmer. Exploring cancer survivor needs and preferences for communicating personalized cancer statistics from registry data: Qualitative multimethod study. *JMIR Cancer*, 7(4):e25659, 2021.

67 R. D. Vromans, M. C. van Eenbergen, S. C. Pauws, G. Geleijnse, H. G. van der Poel, L. V. van de Poll-Franse, and E. J. Krahmer. Communicative aspects of decision aids for localized prostate cancer treatment–a systematic review. *Urologic Oncology: Seminars and Original Investigations*, 37(7):409–429, 2019.

68 G. C. Wishart, E. M. Azzato, D. C. Greenberg, J. Rashbass, O. Kearins, G. Lawrence, C. Caldas, and P. D. P. Pharoah. Predict: a new uk prognostic model that predicts survival following surgery for invasive breast cancer. *Breast Cancer Research*, 12(1):1–10, 2010.

69 BigScience Workshop:, T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilić, D. Hesslow, ..., and P. Villegas. Bloom: A 176b-parameter open-access multilingual language model, 2022. arXiv preprint. `arXiv:2211.05100`.

70 M. Zeguers, H. C. de Haes, L. C. Zandbelt, C. L. Ter Hoeven, S. J. Franssen, D. D. Geijsen, C. C. Koning, and E. M. Smets. The information needs of new radiotherapy patients: how to measure? do they want to know everything? and if not, why? *International Journal of Radiation Oncology, Biology, Physics*, 82(1):418–424, 2012.

71 B. J. Zikmund-Fisher. The right tool is what they need, not what we have: a taxonomy of appropriate levels of precision in patient risk communication. *Medical Care Research and Review*, 70(1), 2013.

72 B. J. Zikmund-Fisher. Helping people know whether measurements have good or bad implications: increasing the evaluability of health and science data communications. *Policy Insights from the Behavioral and Brain Sciences*, 6(1):29–37, 2019.