Design Principles for (X)AI-based Patient Education Systems

Nicolas Pfeuffer [1]

Abstract: Recently, the management of chronic diseases has advanced to a prime topic for Information Systems (IS) research and practice. With increasing capability of Information Technology, patients are empowered to engage in self-management of chronic diseases connected to promises of health benefits for the individual as well as an unburdening of clinics and economic advantages for health care systems. Nevertheless, patients must be adequately educated about risks, screening and examination options to make patient self-management effective, sustainable and profitable. In this regard, Explainable Artificial Intelligence ((X)AI)-based Patient Education Systems (PES) may be an opportunity to provide patient education in an interactive, intelligible and intelligent manner. By establishing Design Principles (DP) for the engineering of effective (X)AI-based PES, instantiating them in a system prototype and evaluating the DP with the help of general practitioners, this paper contributes to the body of knowledge in designing health IS.

Keywords: Patient Education, Patient Adherence, Explainable Artificial Intelligence, Design Science

1 Introduction

The management of chronic diseases with Health Information Technology (HIT) is increasingly gaining attention in IS research [e.g., BCK20]. Among research on effective use of electronic health records (EHR), mobile and social HIT, the self management of diseases is of increasing importance, alleviating pressure on health care providers and improving the life of chronically ill patients on the go [e.g., JC20]. Research suggests that self management and especially awareness of health factors critical for chronic disease management may improve life satisfaction and health outcomes [e.g., JC20].

In self management, the education of patients and the establishment of health-literacy (e.g., the knowledge about risk factors or treatment options about diseases) takes up an important role [e.g., BCK20, p.186]. Educated patients have a decreased risk of taking on health harming behaviors compared to others [e.g., JC20, p.467]. Patient education also increases the potential of patient satisfaction with the treatment and adherence in treatment [e.g., Ho10, p.277]. The economic relevance of patient education for health care systems and individuals is thus indisputable, which is also underlined by research examining the economic impact of patient education [e.g., St18].

[1] Goethe Universität Frankfurt, Information Systems and Information Management,
Theodor-W.-Adorno-Platz 4, 60323 Frankfurt am Main, Germany, pfeuffer@wiwi.uni-frankfurt.de

cbe doi:18.18420/wifo2021

While the application of artificial intelligence (AI) technologies is progressively transforming health care and research on HIT [e.g., BCK20, p. 191] regarding the treatment, prevention and diagnosis of diseases, little technological approaches and research has been conducted in augmenting the process of patient education with AI. Although especially approaches involving explainable AI (XAI) have shown that users are able to assemble knowledge and understand the underlying AI model [e.g., Ku15, RSG16], existing research has not yet grasped the opportunity of designing such systems for patient education. Such systems would not only open up opportunities in distributed, ubiquitous education of patients. It could also alleviate the work of health personnel engaging in patient education and could lead to increased patient adherence, which is considered to be a main influencing factor in treatment and preventing patient readmission [e.g., Xi21].

This research paper aims to remedy this circumstance by providing Design Principles (DP) for designing (X)AI-based Patient Education Systems (PES). This paper proceeds as follows: First, we will illuminate background information about the topic of patient education and IT-based Self Monitoring and Management of diseases (ITSM). Next, we will take a look at our Design Science methodology, and explain the steps that lead us to the emergence of the DP and a prototype. Consequently, we will show the results of the evaluation of the prototype and DP and derive concluding thoughts for IS research and the future of AI-based PES.

2 Background and Related Research

2.1 Patient Education as a Pillar of Health Care

In the past few decades, the active role of the patient has revealed to be essential in the effective treatment of diseases [e.g., Ho10]. Patient adherence with medical treatment, such as screenings, is thus important, yet a substantial share of patients is considered to be non-adherent [e.g., Di04]. In this regard, patient education does not only lead to better understanding for preventive care of diseases, but also to a reduction of emotional distress and an increase in the conformity with examination procedures [Ho10, p.277]. Studies on preventing heart failure, for example, highlight the various purposes and importance of patient education [St05]. Miller [Mi16] found that patient education and subsequent health-literacy has significant positive effects on patient adherence and adherence outcomes and suggests to explore further measures to improve health literacy and patient adherence.

2.2 IT-Supported Patient Education, Adherence and AI

In the past, IT has been used as a tool to promote patient education for various purposes, for example teaching about self-management techniques, knowledge about diseases as well as education about treatment options [e.g., JC20, p.458]. As Jiang et al. [e.g., JC20, p.458] found, the majority of these IT tools is either web-application or mobile app based. Jiang et al. list three types of outcomes of patient education with ITSM tools [e.g., JC20, p.474]: self-understanding, self-efficacy and health literacy. Although several research projects appear to support the goal of health literacy, and several projects emphasize the importance of supporting patient adherence with ITSM tools [e.g., Ha15], there appears to be little research and knowledge on designing effective tools connecting the two areas.

Furthermore, although AI-based tools to support the management of chronic diseases are increasingly being illuminated [e.g., BCK20], little research has explored the merits of using AI-based tools for ITSM and patient education [e.g., JC20]. In this regard, self-monitoring analytics systems help patients track immediate health relevant metrics and their trajectories, such as heart rate during running or daily nutrient intake [e.g., JC20, p.474]. In contrast to the already advanced self-monitoring analytics systems [e.g., JC20, p.474], AI-based systems, especially those equipped with XAI, could help patients better understand the importance [e.g., Ku15, RSG16] of health-related factors directly contributing to a higher disease risk, thus increasing health literacy with potentials to more effective ITSM. While previous generations of HIT already recognized the relevance of transparent algorithmic decision making (e.g., rule based systems, such as MYCIN [BS84]), AI-based PES powered by ML provide high potentials through their statistical approach in contrast to the often rigid rule-based approach. Although research suggests the value of AI and Machine Learning (ML) in improving patient education [e.g., Sa18] the idea of using it to educate patients to promote patient adherence seems to have been elusive until now.

3 Methodology

To tackle the problem of missing design knowledge on AI-based PES adequately, we conduct a design science research (DSR) project based on the guidelines of Kuechler and Vaishnavi [KV08] (see Fig. 1) and thus also follow the design guidelines of Baskerville et al. [BKS11] for a level 1 research design. In doing so, we successively perform the usual phases of DSR: *Awareness of the Problem, Suggestion, Development, Evaluation* and *Conclusion*. We first raise awareness about the problem of missing design knowledge on AI-based PES based on reviewing extant literature. Next, we suggest a range of DP that are derived from existing literature on XAI, interactive ML (IML) and patient education. In the development phase, we instantiate a technical prototype of an XAI-based PES that incorporates the proposed DP. In doing so, we show that our DP and a prototype system can be implemented by IS practitioners, thus obeying the guidelines for reusability of DP [IRH21].

Design Phase	Problem Awareness	Suggestion	Development	Evaluation	Conclusion
Research Activity	Analysis of Knowledge Base on ML and XAI for Patient Education	Deriving Design Knowledge from Literature on XAI, IML and Patient Education	Instantiation of DP in a Technical Prototype	Evaluation of the Prototype in Semi-Structured Expert Interaction Interviews	Summarizing Knowledge and Drawing it Back to the Design Problem
Research Results	Little Extant Research on ML for Patient Education; No Design Knowledge for XAI-based PES	Presenting Design Principles (DP) derived from Literature	Technical Prototype in the form of an Interactive Heart Disease Risk Companion System with XAI and IML	Evaluated DP and Prototype; Implications for Future Research and Practice	Description of Contribution to the Design Knowledge in Information Systems Research

Fig. 1: Research Process adapted from Kuechler and Vaishnavi 2008 [KV08]

This system and especially the DP are then evaluated in semi-structured expert interaction interviews with general practitioners, who are pivotal in generating feedback since they are potential system managers and power users for such systems. In terms of evaluation strategies according to Venable et al. [VPB16], we opted for the technical risk and efficacy strategy for several reasons. First (1) the system is a technical prototype which still requires formal evaluation, before running the risk of influencing patient decision-making. Second (2) the data and model, although they have been used many times before in research, may need to be updated to current clinical standards, tested and evaluated against modern data as well before letting patients use them. Lastly (3), due to the Covid-19 pandemic, observing interactions with patients in e.g., private practice settings would lead to unrecommended contact and may pose unnecessary health risks to participants. In the interviews, the capabilities of the system are demonstrated to the general practitioners, they may interact with it and finally evaluate the DP. This feedback is pivotal in evaluating and improving the prototype and the connected DP. Lastly, the gained knowledge is assembled and promoted to the IS community.

4 Lack of Design Knowledge on AI-based PES: A Design Problem

In the review of extant work on designing AI-based PES including meta-reviews on ITSM [e.g., JC20, BCK20] and patient education, we find a lack of according systems. Based on the little research we are able to find, there is high potential in enriching the design knowledge in this area. Among the little research that can be found, Yu et al. [Yu20] provide a web-based, easily accessible self-diagnosis system. The intention behind the system is that patients may easily self-diagnose with available diagnostic test data while also gaining insight into factors relevant to system predictions. The authors use CART algorithms and visualizations that transparently depict the algorithms overall model and local model for their own prediction. Problematic in this case, however, is that patients may rarely conduct all the required tests and measurements themselves - in this case, the system imputes their data, such that it can still generate a prediction. Such algorithmic behavior may not only reduce the usefulness and accuracy for individual cases, but may also harm the patient by providing either overly optimistic predictions or predictions that are far from the ground truth. A second problematic issue of systems similar to Yu et al.'s [Yu20] are that patients are not able to compare their results with other individual patients, although research suggests the usefulness of analyzing other patient data in increasing health-literacy [Ba08] and motivation in seeking medical treatment [FM08].

In general, the paucity of existing approaches to AI-based PES to promote patient adherence must be regarded critical. Not only does this circumstance leave practitioners at a loss as to how to design such systems and endangers the emergence of secure and reliable guidelines. It also leads to missed opportunities in supporting the main actors in patient education (i.e. general practitioners and nursing personnel) and alleviating the economic and social burden of non-adherence on the health care system [e.g., KAL18].

5 Suggesting Design Principles for (X)AI-based PES

This section is dedicated to the elaboration of our DP. In formulating the DP, we take care to do so according to the schema presented by Gregor et al. [GCS20] to clearly communicate prescriptive design knowledge. As the primary goal of AI-based PES is to provide knowledge about diseases, treatment, screening or examination options, not only must the respective dataset for which the PES is to be engineered be explainable, but the chosen model should be intelligible. This means, for example, an AI-based PES for educating about screening for Cardiovascular Disease Risk, must be interpretable not only by physicians, but also by patients. Because intelligibility is an important topic for the engineering of humane AI, leading scholars suggest the usage of intelligible models instead of black-box models as well [Ru19]. Arguably, intelligible models, such as Generalizable Additive Models (GAM) have shown to perform favourably with medical data in the past [Ca15]. We thus propose:

DP1: For developers to enable patients to interpret the model when using AI-based PES, developers should be using intelligible models whenever possible, to promote the chances of increasing health literacy-induced adherence.

In accordance with DP1, we also believe that patients should have the option to view not only global explanations, i.e. the overall important features for the model, but also local explanations, i.e., instance-based explanations, which have been shown to lead to greater satisfaction and understanding of models in the past [e.g., Ku15, RSG16]. Thus, we propose:

DP2: For developers to enable patients to interpret the overall decision processes and factors of the model when using AI-based PES, developers should provide global explanations, to enable an increase in health literacy.

DP3: For developers to enable patients to form a better understanding of the matter at hand with greater satisfaction of the system when using AI-based PES, developers should provide local, instance-based explanations, to enable an increase in health literacy.

Furthermore, in accordance to the lacks of systems such as the one provided by Yu et al. [Yu20], we believe that the model itself must be able to adjust to the situation of the patient and thus, the user must be able to interactively change the models' features in use. This can be achieved by introducing elements from IML that allow the user to exclude and include features [e.g., Ku15]. This leads to the following DP:

DP4: For developers to empower patients to tailor the PES to their situation, developers should provide IML elements to alter model specifications that are easy to use for patients, thus increasing the chances for increased health-literacy.

Counterfactual or contrastive explanations have been found to exceedingly inform human understanding not only in interactions with others, but also with algorithms [e.g., Mi19, p.16-17]. Furthermore, regarding patient health-literacy, patients have been observed to learn from observing other patients EHR and treatment data [FM08] and therefore seek treatment from appropriate physicians [Ba08]. We thus propose the following DP:

DP5: For developers to empower patients to form contrastive understanding, developers should provide interactive elements to visualize patient data that provides contrastive explanatory instances, thus increasing the chances for increased health-literacy.

Lastly, we build on the ideas of Spinner et al. [Sp20, p.1072] and propose to include *meta-explanations* for the transparency features. Such meta-explanations may help the users to better digest and understand the information they see and also give them the opportunity to inform themselves further, if they need or want to. Thus, we believe that meta-explanations may be conducive to increasing health-literacy and the associated patient adherence.

DP6: For developers to empower patients to better understand the transparency features of the AI-based PES, developers should provide meta-explanations for each explanatory indicator of the transparency features, thus increasing the chances for increased health-literacy.

6 Instantiation of an AI-based PES for General Practitioners

To provide technical proof of the usefulness of our DP and that they are realisable, pertaining to the demands of Iivari et al. [IRH21], we instantiate them in an artifact called "Patient Information Companion" (to be seen in Fig. 2 and Fig. 3). Our Patient Information Companion is built with the purpose of being used to educate patients on health indicators found in examinations. Such a companion system is intended to educate patients on a range of topics, including "Heart Disease / Cardiovascular Disease" or "Lung Disease" to increase their awareness of the importance of preemptive screenings and adherence with examination procedures.

For our prototypical instantiation, we choose to implement the "Heart Companion" as a web-application. As a backend framework for web-applications, the open source Python library "Anvil"[2] was used. As a dataset to train our AI component, we select the well known and publicly available "UCI Cleveland Heart Examination" Repository [DG17], [De17]. The Cleveland Heart Examination dataset (created by Robert Detrano, MD, PhD at VA Medical centre, the Cleveland Clinic Foundation [e.g., De89]) has been used in numerous works on ML in the past and in recent years [e.g., Br97, Na13]. For today's standards on Big Data, the dataset is quite small, containing 297 full records on patient

[2] Github Repository: https://github.com/anvil-works/anvil-runtime; Website: https://anvil.works/, Last checked 11.10.2021

examination. Nevertheless, it is rich in informative attributes on various patient examination procedures and as such, ML algorithms in the past have consistently achieved adequate performance in prediction tasks [e.g., Na13].

In this paper, we use it to instantiate our DP 1-5. DP 1 is instantiated by preparing and using the data to train a binary Explainable Boosting Machine (EBM)[3] [No19]. EBMs are a variant of GAM that are intelligible and have shown to perform favourably with medical data in the past [Ca15]. As a consequence of its intelligibility, DP 2 and 3 could be generated from the intelligible, additive explanatory factors produced by the EBM. We instantiate DP 4 in providing an interface to exclude certain factors (i.e., *Fluoroscopy* or *Nuclear Stress Test*, if the user is afraid of doing them or if the physician does not conduct such tests).

By using this interactive interface, the user gets to know about changes in the explanatory factors (i.e., disease indicators / exercise indicators) when excluding certain examination criteria. Through a performance label, which we instantiate with an AUC-ROC score, the user is informed about changes in performance of the AI (i.e., decrease in performance if one or several factors are excluded). In interactively changing the model, of course, the individual explanations are adjusted as well. DP 5 is instantiated through an interactive interface that allows to select from a range of patients to calculate and plot the explanations appropriately.

Lastly, we instantiate DP 6 in the form of a list of indicators that are included in the AI-based PES and shown as indicators in the XAI graphs. The list itself is interactive, includes a short informative text per list item/indicator, as well as a link to the source of that information. This way, the patient/user is neither forced nor overwhelmed to look at large amounts of information and still is enabled and encouraged to inform themselves further.

[3] The binary classifier is built to distinguish *high-risk* and *low risk* for cardiovascular disease. In 10-fold cross-validation of the full model (using all features), mean ROC-AUC was 91.62 %, min ROC-AUC was 85.27 % with a standard deviation of 0.03. The full model used in the prototype had a ROC-AUC score of 86.81 %.

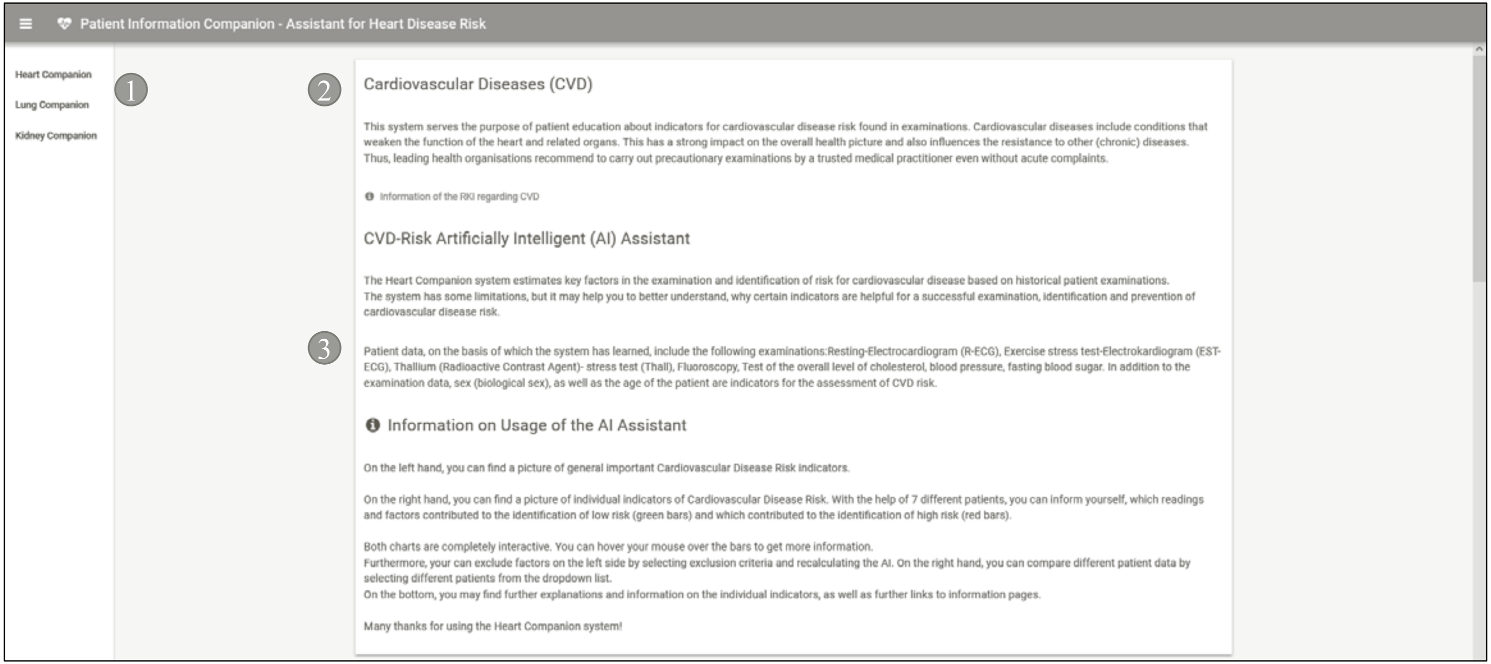

Fig. 2: Introduction View of the Patient Information Companion, showing (1) an overview of potential functions, (2) an introductory explanation on the relevance of educating oneself and (3) instructions on system usage.

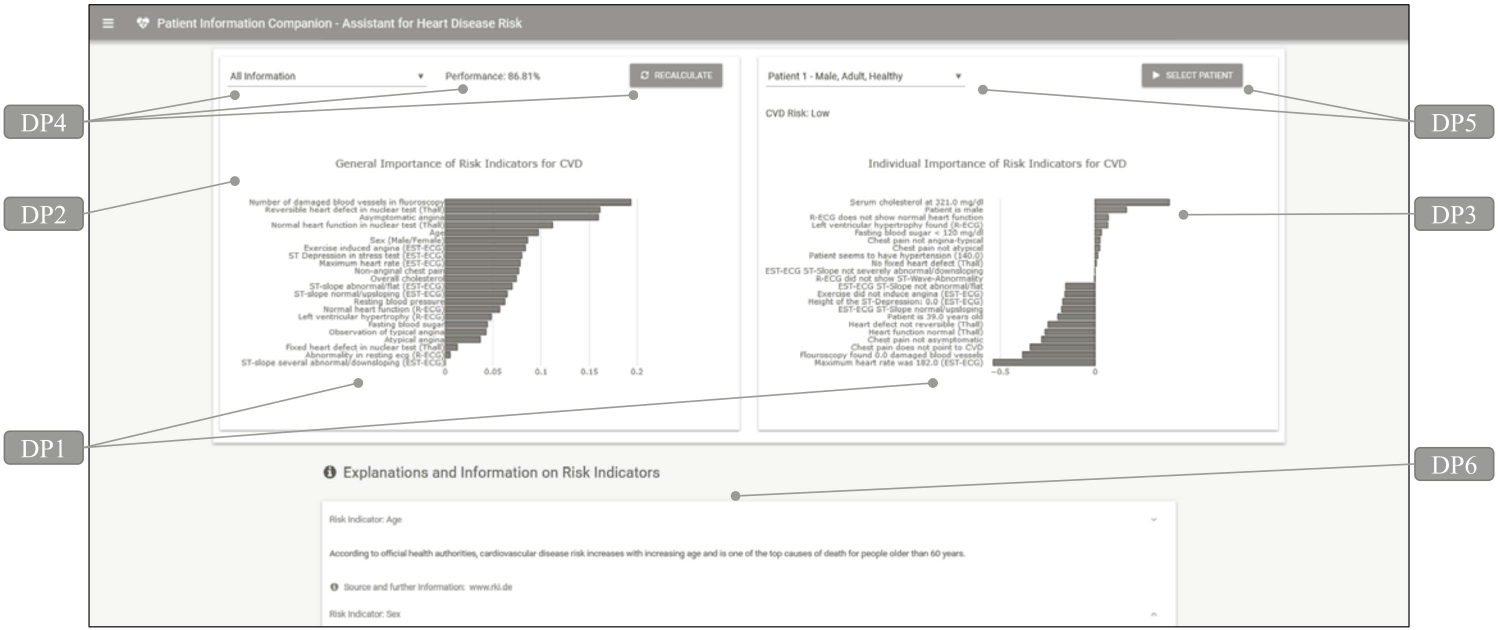

Fig. 3: Instantiation of the DP 1-6 in the Patient Information Companion. Particularly noteworthy are the importance bars on the right XAI plot (based on InterpretML [No19]), indicating beneficial readings (green, contributing to low risk) and harmful readings (red, contributing to high risk).

7 Evaluation

As mentioned in the methodology section, we opt for the technical risk and efficacy strategy [VPB16] for evaluation and engaged in a series of interaction interviews with general practitioners (GP). Three interaction interviews were conducted with experienced GP. The interviews were structured with the help of a semi-structured questionnaire. First, general questions on characteristics of the participants and their experience with AI in health care were asked (see Tab. 1).

Participant	Age	Professional Experience in Years	Experience with AI in Health Care
GP 1	37	10	No Awareness of Previous Experience
GP 2	65	39	Automated Diagnostics
GP 3	60	35	Automated Diagnostics

Tab. 1: General Practitioners in the Study

Next, the system functionalities were demonstrated and the physicians were instructed on how to use the system. Subsequently, the physicians engaged in interactions with the AI-based PES and explored the different functionalities. After they had gathered a good impression of the system, all six DP were evaluated by confronting and questioning the physicians on the respective functionalities. Lastly, two follow-up questions on the potential of such systems in patient education and increasing patient adherence concluded the interviews.

The expert interviews lasted on average 38 minutes and were accompanied by taking field notes and audio recordings on the evaluations to caption the statements and thoughts of the physicians on the DP and the prototype. As we can see from Tab. 2, all three physicians approved of all six DP.

Unit of Analysis	Brief Description	Verdict		
		GP1	GP2	GP3
DP 1	Intelligible Algorithm	✓	✓	✓
DP 2	Global Transparency	✓	✓	✓
DP 3	Local/Individual Transparency	✓	✓	✓
DP 4	Interactive ML	✓	✓	✓
DP 5	Contrastive Transparency	✓	✓	✓
DP 6	Meta-Explanations on XAI Indicators	✓	✓	✓
System Prototype	PES Prototype	✓	✓	✓

Tab. 2: Evaluation of the Design Principles and PES Prototype with General Practitioners

Especially DP 4 and the interactive ML was highlighted as being extremely helpful, since the possibility to include and exclude certain exercises in conjunction with the changing AUC-ROC metric would give the potential to increase adherence to certain exercises. DP 5 and the contrastive explanations was highlighted as well, giving the patient much better potential of increasing health-literacy and potential to increasing adherence. In general, the idea of such a system was considered favourably while the physicians also provided some additional feedback on how they imagined such systems could be brought into practice and provide the most benefit.

First, GP1 considered the self-prediction of patients with such systems as critical, giving potential to emergent psychological pressure. Instead, GP1 approved the usage for selected patients and rather from a system that does not enable the comparison of the user to other patients. Furthermore, if the user was enabled to learn from the data of other patients without requiring to self-diagnose, then GP1 would approve of such a system as well. GP2 argued that they would see AI-based PES "as something that you offer in the waiting area" and probably not on the website of the practice, because it could be too complex for elderly users. GP3 similarly stated that such a tool "may be wanted by younger generations" and that it would be more easily usable by them. In contrast, GP1 and GP3 could imagine such a system on the website as well as in their practice. Still, as GP3 noted, for older generations, the system functionality would require more thorough explanation. GP2 and GP3 similarly highlighted the role of GP assistants in using and explaining the tool to engage with select patients to increase their adherence.

In summary, from the domain expert perspective, AI-based PES should be tailored to the specifics of individual patient needs (e.g., the psychological stance, option to self-diagnose, elderly user who requires assistance by GP assistants, prior adherence behavior) to unleash their true potential. Finally, in their general assessment of the system's benefits, GP2 praised the potential of AI-based PES for increasing work efficiency, facilitating work and increasing patient adherence at the same time. Likewise, GP3 pondered "we have brochures [...] a TV that is always running and which provides information, but that is only general information. [With this system] here, you can inform patients specifically about their disease and risk indicators [...] And honestly, I think this is the future [...]."

8 Conclusion

In this paper, we have established initial DP for designing (X)AI-based PES and conducted a first evaluation with the help of general practitioners. In addition, this paper shows that a prototype including these DP is not only regarded as highly useful by physicians but is realizible by IS practitioners, since it was built with open access libraries in Python (also including HTML and CSS elements through the web app) and thus also conforms with the demands of Iivari et al. [IRH21]. The established DP appeared to be highly useful in the eyes of the practitioners, leading us to conclude that they may be sound guidelines in designing (X)AI-based PES. Establishing these DP can be considered the first step in exploring the usefulness of (X)AI-based PES in patient education. Thus laboratory, natural and field experiments are needed, to further evaluate the usefulness of such systems in practice and to generate a deep theoretical understanding of their impact and use. We are confident that we laid the foundation for fruitful future research work in contributing our DP for (X)AI-based PES to the knowledge base of IS design.

Bibliography

[Ba08] Barker, K. K.: Electronic Support Groups, Patient-Consumers, and Medicalization: The Case of Contested Illness. Journal of Health and Social Behavior 49/1, pp. 20–36, 2008.

[BCK20] Bardhan, I.; Chen, H.; Karahanna, E.: Connecting Systems, Data, and People: A Multidisciplinary Research Roadmap for Chronic Disease Management. MIS Quarterly 44/1, pp. 185–200, 2020.

[BKS11] Baskerville, R.; Kaul, M.; Storey, V.: Unpacking the Duality of Design Science. In: Proceedings of the Thirty-Second International Conference on Information Systems (ICIS) 2011, pp. 4168–4182, 2011.

[Br97] Bradley, A. P.: The Use of the Area Under the ROC Curve in the Evaluation of Machine Learning Algorithms. Pattern Recognition 30/7, pp. 1145–1159, 1997.

[BS84] Buchanan, B. G.; Shortliffe, E. H.: Rule-Based Expert Systems: The Mycin Experiments of the Stanford Heuristic Programming Project. CUMINCAD, Boston, 1984.

[Ca15] Caruana, R. et al.: Intelligible Models for Healthcare: Predicting Pneumonia Risk and Hospital 30-Day Readmission. In (Cao, L. et al. Eds.): Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1721–1730, 2015.

[De17] Detrano, R.: UCI Machine Learning Repository|Cleveland Database Heart Disease Data Set, 2017, url: https://archive.ics.uci.edu/ml/datasets/heart+disease.

[De89] Detrano, R. et al.: International Application of a New Probability Algorithm for the Diagnosis of Coronary Artery Disease. The American Journal of Cardiology 64/5, pp. 304–310, 1989.

[DG17] Dua, D.; Graff, C.: UCI Machine Learning Repository, 2017, url: https://archive.ics.uci.edu/ml.

[Di04] DiMatteo, M. R.: Social Support and Patient Adherence to Medical Treatment: A Meta-Analysis. Health Psychology 23/2, pp. 207–218, 2004.

[FM08] Frost, J.; Massagli, M.: Social Uses of Personal Health Information Within Patientslikeme, an Online Patient Community: What Can Happen When Patients Have Access to One Another's Data. Journal of Medical Internet Research 10/3, e15, 2008.

[GCS20] Gregor, S.; Chandra Kruse, L.; Seidel, S.: Research Perspectives: The Anatomy of a Design Principle. Journal of the Association for Information Systems 21/6, pp. 1622–1652, 2020.

[Ha15] Hamine, S. et al.: Impact of mHealth Chronic Disease Management on Treatment Adherence and Patient Outcomes: A Systematic Review. Journal of Medical Internet Research 17/2, pp. 1-15 (e52), 2015.

[Ho10] Hoving, C. et al.: A History of Patient Education by Health Professionals in Europe and North America: From Authority to Shared Decision Making Education. Patient Education and Counseling 78/3, pp. 275–281, 2010.

[IRH21] Iivari, J.; Rotvit Perlt Hansen, M.; Haj-Bolouri, A.: A Proposal for Minimum Reusability Evaluation of Design Principles. European Journal of Information Systems 30/3, pp. 286–303, 2021.

[JC20] Jiang, J.; Cameron, A.-F.: IT-Enabled Self-Monitoring for Chronic Disease Self-Management: An Interdisciplinary Review. MIS Quarterly 44/1, pp. 451–508, 2020.

[KAL18] Kvarnström, K.; Airaksinen, M.; Liira, H.: Barriers and Facilitators to Medication Adherence: A Qualitative Study With General Practitioners. BMJ open 8/1, 1-8 (e015332), 2018.

[Ku15] Kulesza, T. et al.: Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In (Brdiczka, O. et al. Eds.): IUI'15. Proceedings of the 20th ACM International Conference on Intelligent User Interfaces. Association for Computing Machinery, New York, NY, USA, pp. 126–137, 2015.

[KV08] Kuechler, B.; Vaishnavi, V.: On Theory Development in Design Science Research: Anatomy of a Research Project. European Journal of Information Systems 17/5, pp. 489–504, 2008.

[Mi16] Miller, T. A.: Health Literacy and Adherence to Medical Treatment in Chronic and Acute Illness: A Meta-Analysis. Patient Education and Counseling 99/7, pp. 1079–1086, 2016.

[Mi19] Miller, T.: Explanation in Artificial Intelligence: Insights From the Social Sciences. Artificial Intelligence 267, pp. 1–38, 2019.

[Na13] Nahar, J. et al.: Association Rule Mining to Detect Factors Which Contribute to Heart Disease in Males and Females. Expert Systems with Applications 40/4, pp. 1086–1093, 2013.

[No19] Nori, H. et al.: InterpretML: A Unified Framework for Machine Learning Interpretability. arXiv preprint arXiv:1909.09223, 2019.

[RSG16] Ribeiro, M. T.; Singh, S.; Guestrin, C.: "Why Should I Trust You?" Explaining the Predictions of Any Classifier. In (Balaji Krishnapuram et al. Eds.): Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. Association for Computing Machinery, New York, New York, USA, pp. 1135–1144, 2016.

[Ru19] Rudin, C.: Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. Nature Machine Intelligence 1/5, pp. 206–215, 2019.

[Sa18] Saunders, C. H. et al.: Bring on the Machines: Could Machine Learning Improve the Quality of Patient Education Materials? A Systematic Search and Rapid Review. JCO Clinical Cancer Informatics 2, pp. 1–16, 2018.

[Sp20] Spinner, T. et al.: explAIner: A Visual Analytics Framework for Interactive and Explainable Machine Learning. IEEE Transactions on Visualization and Computer Graphics 26/1, pp. 1064–1074, 2020.

[St05] Strömberg, A.: The Crucial Role of Patient Education in Heart Failure. European Journal of Heart Failure 7/3, pp. 363–369, 2005.

[St18] Stenberg, U. et al.: Health Economic Evaluations of Patient Education Interventions a Scoping Review of the Literature. Patient Education and Counseling 101/6, pp. 1006–1035, 2018.

[VPB16] Venable, J.; Pries-Heje, J.; Baskerville, R.: FEDS: A Framework for Evaluation in Design Science Research. European Journal of Information Systems 25/1, pp. 77–89, 2016.

[Xi21] Xie, J. et al.: Understanding Medication Nonadherence From Social Media: A Sentiment-Enriched Deep Learning Approach. MIS Quarterly Forthcoming, Preprint available at SSRN 3091923: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3091923, 2021.

[Yu20] Yu, C.-S. et al.: Development of an Online Health Care Assessment for Preventive Medicine: A Machine Learning Approach. Journal of Medical Internet Research 22/6, pp. 1-18 (e18585), 2020.