Health Personalisation: From Wellbeing to Medicine

SHLOMO BERKOVSKY and ENRICO COIERA
Centre for Health Informatics, Macquarie University

The current agenda in health personalisation research mainly revolves around supporting lifestyle and wellbeing. Personalised recommendations for patients and consumers have been explored for areas like physical activity, food intake, mental support, and health information consumption. Strikingly little attention has been paid to personalised medical applications supporting clinical users. In this paper, we turn the spotlight on such medical use cases and the advantages personalised decision-support can bring. We discuss the differences between patient- and clinician-centric personalisation and highlight touch points, where personalised support might improve clinicians' decision-making.

DOI: 10.1145/3473044.3473048 http://doi.acm.org/10.1145/3473044.3473048

1. INTRODUCTION

Personalised technologies and recommender systems have evolved over the last two decades from a niche feature for Web browsing and eCommerce to a popular service found in virtually every online application. Various personalised solutions are currently available for information access, eLearning, entertainment, shopping, tourism, road navigation, personal assistance, and many more [Berkovsky and Freyne 2015].

While personalisation has received attention in the medical domain, its focus has been relatively narrow. Typical personalised services support the wellbeing of their users by recommending recipes and meal plans, suggesting exercises and workouts, promoting physical activity and exergaming, assisting in health information access, providing personalised social support, and so on [Kocaballi et al. 2019]. A common thread among these applications is the support for maintaining a healthy lifestyle that they offer to lay users with a limited medical knowledge.

In contrast, very little personalisation research has targeted clinicians, i.e., doctors, nurses, or allied health professionals, as the recipients of the service. As an illustrative but representative example, consider the word cloud of titles of recently published papers on health recommender systems shown in Figure 1. The cloud is mainly composed of terms directly associated with lifestyle and wellbeing, with the sole exception being the generic term 'medical'.

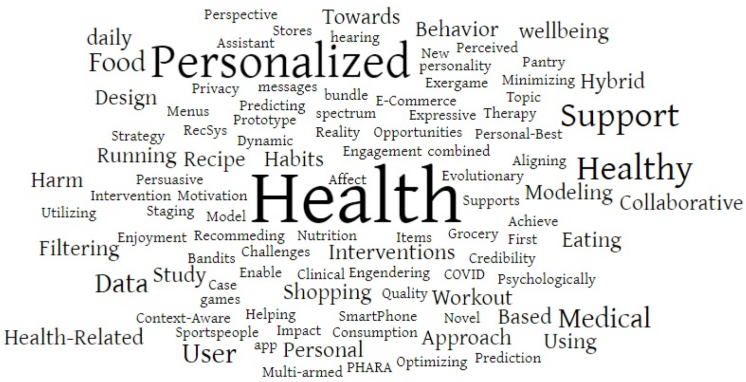

Fig. 1. Word cloud of titles of papers published in ACM RecSys and the co-located HealthRecSys workshop. Words 'recommend', 'recommender', and 'system' were omitted from the cloud.

In this paper, we examine the notion of *clinician-centric personalisation*, the barriers it is facing, and promising veins of future work in this space. We initially discuss the differences between patient- and clinician-centric personalisation and then highlight where personalisation can fit existing healthcare decisions [Coiera 2012]. We use this paper as a call for collaborative action involving both the personalisation and health informatics research communities.

2. CLINICIAN-CENTRIC PERSONALISATION

To analyse the differences between clinician-centric and general-purpose personalisation, we focus on three facets: medicine as the application domain, unique aspects of personalisation in medicine, and clinicians as the recipients of the personalised service. They primarily converge to a different personalisation use-case: while general-purpose personalisation often implies a lean back setting (e.g., entertainment recommendations or news consumption), clinician-centric personalisation is delivered to a skilled professional within an established workflow (e.g., seeing patients or operating in a hospital). Thus, the service needs to fit the existing setting, reliably support the clinicians, and to be as transparent and intuitive as possible.

2.1 The Medical Domain

Despite the wide range of domains, in which personalisation has been exploited, medical applications are still uncommon. One of the reasons for this is the *scarcity of available data* for personalisation researchers to harness [Liu et al. 2017]. While typical data used to personalise in other domains – be it Web browsing logs, item preferences, or user-generated content – are abundant, personal health records facilitating personalisation are sensitive and hard to obtain. Collection and sharing of such data are typically restricted by privacy and ethical regulations. Even if shared, the data are often anonymised and de-identified, which hinders linkage of datasets and construction of rich patient models.

Another challenge hindering personalised solutions in the medical domain is the *substan-*

tial risks associated with the provision of wrong or inappropriate advice. Whilst a bad movie or shopping recommendation may lead to unnecessary consumption and potentially loss of time and money, these are minor consequences compared to a disease misdiagnosis, wrong dosage of a drug, or avoidable side effect of an incorrect prescription [Bright et al. 2012]. These increased human and economic costs of medical errors, combined with the scarcity of available data, pose significant barriers for developers of personalised solutions.

In relation to errors, it is essential to also highlight the different *error metrics* used in the medical domain. Many general-purpose personalised services rely on classification metrics that assign a similar importance to false negative and false positive errors, e.g., F1 – an equal-weight harmonic mean of precision and recall. That said, in medicine a false negative would cause an under-diagnosed patient not being treated, whereas a false positive would lead to an over-diagnosed case and unnecessary treatment. The outcomes of these two are substantially different and the performance metrics should account for this error asymmetry accordingly.

2.2 Personalisation in Medicine

The likely use-case of medical personalised solutions is provision of intelligent support to clinicians, in terms of information access, decision support, task prioritisation, data visualisation and annotation, information search, evidence analysis, and more. To demonstrate practical value, these services will need to be first and foremost exploited and taken up by the clinicians. The heightened risk and responsibility that come with the personalised services make it imperative to *build and sustain the trust* of clinicians [Bussone et al. 2015].

A key factor in instilling trust in personalised support is *explainability* [Holzinger et al. 2019]. Independently of the technical quality of decision support, clinicians may want to explore the reasoning behind it. As different modalities of data are involved in decision making processes, such explainability can take various forms, e.g., summarisation and entity extraction in clinical text, segmentation and annotation of medical images, and sequencing and selection of multi-dimensional 'omics' data. An inability to explain reasoning in a way that instils clinician's trust and confidence may limit the application of non-interpretable machine learning methods and deep learning architectures.

Similarly, any personalised service will need to fit within the clinical workflow, and such workflows can vary between notionally similar clinicians. As clinicians are likely to be receiving recommendations while engaging with patients, their ability to provide sufficient attention and corroborate machine-provided information will be restricted. For instance, they may be unable to read scientific papers, guideline recommendations, or search for clinical evidence. Nor they will be able to consider a long list of alternatives. Hence, any clinical decision support system will need to generate *concise and factual output*, only including a small number of self-contained and prioritised options, supported by brief explanations or easy-to-grasp visualisations.

2.3 Personalising for Clinicians

Clinicians are *knowledgeable recipients*, with strong domain expertise, and often many years of training and practical experience. This translates into different preferences for the type and frequency of the support required, as well as the specific task, for which the support is needed. For example, more junior clinicians may favour advice-oriented support, whereas experts who are clear on the options available are likely to want the key data needed to allow them to make a decision [Sintchenko et al. 2005]. This is different from general-purpose personalisation, where the users may only be vaguely familiar with the domain and the specific options they are offered.

In addition, clinicians can be considered as rather *risk-averse recipients*. Consider a movie recommendation scenario, where features like novelty and serendipity of the recommendations are perceived advantageous for users. The same features are likely to be disadvantages, if present in a medical support service, as such a support is likely to require additional cognitive effort and may undermine the clinician's willingness to use the service. As such, the service would need to strike the balance between raising potentially overlooked alternatives, and providing consistent and comprehensible support to clinicians [Anckler et al. 2017].

Lastly, it is important to *differentiate between the recipient and the subject* of the personalised service in medicine. Normally, the two would refer to the same person; e.g., sightseeing recommendations will be tailored to a traveller's profile and interests, and they will naturally be delivered to the same traveller. On the contrary, personalised support in medicine will typically be tailored to the patient data (such as present symptoms, medical history, examinations, co-morbidities, and so on), although the generated personalisation will rather be delivered to the clinician, who is the recipient of the service.

3. CLINICAL DECISION-SUPPORT

Considering the above characteristics of the application domain, type of service, and recipients, we argue that the service can be classified as *personalised decision-support*. In essence, the envisaged use-case can be described as a collaborative human-machine team, where the human clinician is the decision-maker (although often acting in consultation with the patient) and the machine provides a "second opinion" decision-support to the human. In this section, we discuss how such a decision-support fits common healthcare processes and what computational personalisation challenges would the provision of such a decision-support entail.

3.1 Fit to Healthcare Decisions

There are three major classes of clinical decisions [Szolovits 2019].The first is *diagnostics*, which deals with identifying the medical condition of a patient or the reasons for the patient's complaint. Once a diagnosis is established, the following components are prognostics and therapeutics. *Prognostics* is about predicting the patient's trajectory and likely progression of the disease, whereas *therapeutics* is about determining the most appropriate treatment for the patient. It should be highlighted that the latter two are strongly interrelated, as the prognosis may influence the treatment and, likewise, treatment may affect the patient's prognosis.

Which of these three tasks might benefit from a personalised decision-support for clinicians? We believe that all of them. The decision regarding the patient's diagnosis is evidently personalised, taking into consideration the patient's demographic information, medical history, current symptoms, examination and pathology results, other co-morbidities, and so on. Prognostic decisions will mainly be influenced by the diagnosis, but also by the ongoing treatment and the disease progression, both specific to the patient and their condition. Correspondingly, treatment decisions will be affected by the diagnosis, prognosis, and observed progression of the disease, all also being tailored to the patient [Montani and Striani 2019].

In practice, we observe personalised decision-support solutions primarily in the diagnostic component. Specifically, we note a wide range of imaging software tools supporting clinicians in interpreting medical images, e.g., radiography, ultrasound, tomography, magnetic resonance imaging, and more. Another prominent example is primary care, where various decision-support and predictive modelling tools are used by general practitioners and family doctors. Lastly, another promising application of personalised decision-support is mental health, where online screening tools can help clinical psychologists prioritise diagnostic interviews.

Some personalised tools for therapeutic purposes have also been developed. For example, there exist personalised tools capable of predicting individualised response of a patient to a treatment (or drug). The work on personalised prognostics has primarily focussed on statistical methods for modelling lifetime risks and on personalised predictive models of patient recovery after surgical interventions.

3.2 Personalisation Challenges

Although some progress has been observed, the technical questions related to materialising and translating new personalised decision-support solutions for clinicians are non-negligible. To this end, we highlight research advancements in personalised technologies, which could be adopted also for the medical domain.

The task of creating and consolidating patient models can leverage the progress in the *user modelling* research [Cirillo and Valencia 2019]. There are many questions around the representation and storage of heterogeneous, multi-faceted, and multi-modal patient data, which could be informed by works on user model representation and knowledge transfer. In similar to user modelling applications, clinician-facing personalisation can benefit from various sensors and wearable devices capturing patient data. Next questions deal with the aggregation of such a diverse data and ability to assemble updated, transparent, and unbiased patient models, which have received an increasing attention in user modelling research. This also entails the question of selecting the data, most valuable and informative for the provision of personalised services.

Once a reliable patient model is formed, the next challenge deals with the *algorithmic methods* for generation of personalised services [Szolovits 2019; Montani and Striani 2019]. While many methods exploited by the personalisation community can be applied, we highlight the potential of content-based methods, owing to the availability of rich domain knowledge and the scarcity of large-scale datasets. The open questions refer to the constraints of the medical domain, such as the need for accuracy in decision-support and the explainability constraint, which may eliminate uninterpretable models. Another question touches upon the generalisation of the methods and ability to transfer knowledge from one sub-discipline (or condition) to another, in similar to cross-domain personalisation. Lastly, we note the potential of contextual personalisation methods that can consider patient's lifestyle and co-morbidities as contextual factors.

The next group of questions deals with *usability and user interaction* with personalised clinical decision-support [Rundo et al. 2020]. As flagged earlier, the explainability and interpretability of decision-support are critical for the uptake of the service by clinicians. This brings to the fore recent advancements in user interaction with personalised systems, intelligent tutoring systems, personalised text summarisation, and visualisation of high-dimensional data. At the same time, we note the need from building and sustaining clinician's trust, as a means to maximise the uptake of the decision-support. This potentially entails methods for adaptive resource allocation, to guarantee that the deployed methods can generate accurate and reliable decision-support, despite the limitations of the data and the requirement to produce near real-time output.

Last but not the least, we emphasise practical questions related to the provision of personalised services, which have been studied in personalisation research. One of them deals with the linkage of medical data, often stored by different clinicians and health systems, represented in various formats, and having different modalities. While linking these poses a serious challenge, the ability to assemble accurate and up-to-date patient models can boost the quality of personalised decision-support. The other challenge surrounds the need to balance the personalisation goals with ethical and privacy constraints. As noted, medical data often come with strict access limitations, privacy legislation, and ethics approvals, encumbering the access to and sharing of such data. Thus, privacy-preserving personalisation methods can offer a sound alternative satisfying both the personalisation goals and privacy constraints [Friedman et al. 2016].

4. SUMMARY

This paper highlighted the importance of personalised decision-support services in the medical domain. We initially discussed the aspects specific to the provision of clinician-centric personalised decision-support, and then turned to characterising such a service, both in terms of its fit to existing healthcare workflows and compatibility with recent personalisation research advancements.

Having identified the commonalities and potential touch points, we emphasise the existing gaps between the two research communities. Indeed, very few personalisation works target clinicians as the recipients of the personalised service and, likewise, very few precision health and health informatics works adopt methods and solutions developed by the personalisation community. We would like to use this article as a call for collaborative action. We argue that there is a space and need for the two communities to join forces and collaborate closer. This has the potential to pave the way to innovative applications of core personalisation methods for provision of clinician-centric personalised decision-support services.

REFERENCES

ANCKLER, J. S., EDWARDS, A., NOSAL, S., HAUSER, D., MAUER, E., AND KAUSHAL, R. 2017. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. *BMC Medical Informatics and Decision Making 17,* 1, 1–9.

BERKOVSKY, S. AND FREYNE, J. 2015. Web personalization and recommender systems. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2307–2308.

BRIGHT, T. J., WONG, A., DHURJATI, R., BRISTOW, E., BASTIAN, L., COEYTAUX, R. R., SAMSA, G., HASSELBLAD, V., WILLIAMS, J. W., MUSTY, M. D., AND WING, L. 2012. Effect of clinical decision-support systems: a systematic review. *Annals of internal medicine 157,* 1, 29–43.

BUSSONE, A., STUMPF, S., AND O'SULLIVAN, D. 2015. The role of explanations on trust and reliance in clinical decision support systems. In *Proceedings of the International Conference on Healthcare Informatics*. 160–169.

CIRILLO, D. AND VALENCIA, A. 2019. Big data analytics for personalized medicine. *Current opinion in biotechnology 58*, 161–167.

COIERA, E. 2012. The true meaning of personalized medicine. *Yearbook of Medical Informatics 7*, 4–6.

FRIEDMAN, A., BERKOVSKY, S., AND KAAFAR, M. A. 2016. A differential privacy framework for matrix factorization recommender systems. *User Modeling and User-Adapted Interaction 26,* 5, 425–458.

HOLZINGER, A., LANGS, G., DENKS, H., ZATLOUKAL, K., AND MÜLLER, H. 2019. Causability and explainability of artificial intelligence in medicine. *Data Mining and Knowledge Discovery 9,* 4, 1312.

KOCABALLI, A. B., BERKOVSKY, S., QUIROZ, J. C., LARANJO, L., TONG, H. L., REZAZADEGAN, D., BRIATORE, A., AND COIERA, E. 2019. The personalization of conversational agents in health care: systematic review. *Journal of medical Internet research 21,* 11, 15360.

LIU, Z., REXACHS, D., EPELDE, F., AND LUQUE, E. 2017. A simulation and optimization method for calibrating agent-based emergency department models under data scarcity. *Computers & Industrial Engineering 103*, 300–309.

MONTANI, S. AND STRIANI, M. 2019. Artificial intelligence in clinical decision support: a focused literature survey. *Yearbook of Medical Informatics 28,* 1, 120.

RUNDO, L., PIRRONE, R., VITABILE, S., SALA, E., AND GAMBINO, O. 2020. Recent advances of HCI in decision-making tasks for optimized clinical workflows and precision medicine. *Journal of biomedical informatics*, 103479.

SINTCHENKO, V., IREDELL, J. R., GILBERT, G. L., AND COIERA, E. 2005. Handheld computer-based decision support reduces patient length of stay and antibiotic prescribing in critical care. *Journal of the American Medical Informatics Association 12,* 4, 398–402.

SZOLOVITS, P. 2019. *Artificial intelligence in medicine*. Routledge.

Shlomo Berkovsky is the leader of the Precision Health research stream at the Centre for Health Informatics, Macquarie University. The stream focusses on the use of machine learning methods to develop patient models and personalised predictions of diagnosis and care. His original areas of expertise include user modelling, online personalisation, and persuasive technologies.

Enrico Coiera is the Foundation Professor in Medical Informatics at Macquarie University and Director of the Centre for Health Informatics he co-founded in 2000. With a research background in industry and academia, he has an international reputation for work on decision support and communication processes in biomedicine. His textbook, *Guide to Health Informatics* is in its third edition, widely used internationally, and translated into several languages.