Evaluating Personalized Pervasive Health Technology—But How?

Jakob E. Bardram
Technical University of Denmark

■ **IN THE CALL** for this special issue on Personalized Pervasive Health, it was stated that it "intends to provide a comprehensive view on innovative pervasive computing methods, ubiquitous technology, data-based inference algorithms, as well as evaluation studies, all related to personalized health. [. . .] All investigations must include *thorough evaluations of their approaches and methods*." (my emphasis). When reading this, I came to wonder what "thorough evaluations" actually is, in this context? Are authors supposed to show clinical evidence for the health efficacy of their technology? Or, are they supposed to show that the technology is technically sound and working? Or, that it is secure and has appropriate privacy-protection of sensitive personal data (c.f. the focus on this special issue). Or, that the technology is usable and user-friendly for the users? Or, ...?

These questions touch upon a more fundamental question, namely how researchers can evaluate novel ubiquitous computing technology in the health domain in a manner that allows them to make meaningful claims about their

utility and argue for their scientific contributions in the broader health technology domain.

Evaluation of health technology has always been difficult and subject to significant scientific disputes. From a *technological perspective*, we are mainly interested in the design of novel technology and understanding how it works under different circumstances, while gradually and iteratively improving on its technical features and capabilities. This calls for formative evaluation methods, which help us understand how the technology works and how it can be improved according to a set of design goals. In this approach, we seek to understand the technology and look "into" it, i.e., a white-box evaluation strategy.

From a *health perspective*, we are mainly interested in the efficacy of the technology in terms of clinical outcome related to screening, diagnosis, treatment, or care of patients. We care about establishing solid evidence for the clinical claims of a technology and less about how it actually works from a technical point-of-view. This calls for summative evaluation methods, which compare the outcome of using the technology to some control situation. In this approach, we seek to understand the outcome

of its use, and care less about what is inside the technology, i.e., a black-box evaluation strategy.

This tension between different methodological approaches in the technical and health sciences has been managed for many years in the more traditional medical device domain. However, with the increasing profiliation of novel technological opportunities—not least coming from the mobile and ubiquitous computing area—this tension has increased. We have seen an explosion of the use of mobile and wearable technology in health. For example, in a recent review in this magazine, we found 46 different mobile and wearable applications that have been introduced in the mental health domain over the last decade.[1] Similarly, a recent consensus report on diabetes digital app technology found, across the USA and Europe, that mobile health (mHealth) apps intended to manage diabetes health and wellness were largely unregulated and lacked any kind of evidence for their safety and efficacy.[2]

This challenge of providing appropriate evaluation methods for health technology is a topic that, fortunately, has gained increased attention lately. In this spotlight article, I will first try to outline how different attempts to address this challenge have emerged lately from both the technological as well as the health sciences. Then, I will share some insight and experience on how we have approached evaluation of personal health technology in the Copenhagen Center for Health Technology (CACHET). I hope this can be of use for others who are facing the question of how to evaluate personal health technology and guide them in the design of an appropriate evaluation strategy.

EVALUATION APPROACHES

Methodologically, the design of health technology, including the growing research into "Personalized Pervasive Health Technology,"[3] sits at the intersection of design science and health science.

On the one hand, health technologies need to be designed, developed, and refined in a design process, which often relies on technological and user-centered design methodologies. In the biomedical engineering sciences, a novel technology is often evaluated according to a performance standard. For example, a classic approach to evaluate the accuracy of a novel approach to detect hearth arrhythmia, such as atrial fibrillation, is to compare the proposed approach to labeled data, such as the PhysioNet database (e.g., the article by Peimankar and S. Puthusserypady[4]). In computer science, the design of a novel mHealth application is typically evaluated using a user-centered approach. For example, in a recent study of a recommender system for depressive patients, the Unified Theory of Acceptance and Use of Technology methodology[5] was used to access the perceived usefulness and usability of the system.[6]

On the other hand, health technologies need to be clinically verified in order to assess clinical safety, efficacy, and effectiveness. From a health-oriented perspective, a carefully designed randomized controlled trial (RCT) represents the highest level of evidence in evidence-based medicine (EBM), and is the "gold standard" for determining the effectiveness of pharmacological agents. This approach has been transferred to evaluating nonpharmacological interventions, including health technology. The Consolidated Standards of Reporting Trials (CONSORT) statement is a guideline that specifies how to report RCTs and is applied in all papers within EBM.[7] In order to accommodate specific issues related to describing an mHealth intervention, the CONSORT guidelines were extended to the Consolidated Standards of Reporting Trials of Electronic and Mobile HEalth Applications and onLine TeleHealth (CONSORT-EHEALTH).[8] This guideline suggests to expand the description of the intervention of a mHealth study to include more detailed information on the technology, including how it was funded, designed, developed, and deployed. It is also recommended to provide access to the technology and its source code, by, e.g., releasing it as an open source, or at least release a video containing a detailed walk-through of the system and its features.

There is, however, a growing awareness that these two methodological positions represent two opposite "poles" that have quite different scientific standards for the different stages in research, including how to evaluate and measure the applicability of health technology.[9] The

traditional RCT has limitations when used to evaluate health technology; for example, it does not permit iterative improvements to the design and the technology may be outdated by the time the trial is completed.[10] Moreover, since the RCT methodology is summative—i.e., measures outcomes before and after an intervention—it treats the intervention as a "black-box" and is not suited to helping researchers understand which parts of the intervention (i.e., features of the technology) are actually causing an effect. For example, even though it seems like mHealth interventions can reduce depressive symptoms,[11] it is unclear which features of these systems actually account for this effect. Is it, for example, important to provide support for self-assessment? Or in-person feedback? Or cognitive training? Such detailed information cannot be collected from a traditional RCT, since it treats the technology as an all-or-nothing intervention.

For this reason, a number of alternative evaluation methodologies have been proposed. From a health science perspective, the CONSORT guidelines have been extended to also encompass pilot and feasibility trials conducted in advance of a future definitive RCT.[12] The primary aim of a pilot or feasibility trial is to assess feasibility of conducting the future definitive RCT. Such a feasibility study can be useful to test the applicability of a novel technology before moving into an actual RCT. The study can test the feasibility of using the technology for a specific group of patients, the recruitment and on-boarding procedure, and the technical stability and scalability of the technology in real use.

It is, however, important to bear in mind that a clinical study—both a feasibility study and a RCT—is costly and lengthy. It often requires the enrollment of hundreds of patients to ensure statistically significant results, it runs over several months, and it needs a staff of skilled doctors and healthcare professionals to run them. It has been estimated that an RCT costs $41,000 per patient.[13] Therefore, while RCTs are important for evaluation of clinical effectiveness, they are best undertaken only when: i) the intervention and its delivery package are stable, ii) the intervention can be implemented with high fidelity, and iii) there is a reasonable likelihood that the overall benefits will be clinically meaningful.[14]

From a design science perspective, it has been argued that health technology should be evaluated from a more formative and qualitative approach before moving into costly clinical trials. Klasnja et al.[15] provide a set of compelling arguments for why it is important to evaluate behavior change and health technology from a human–computer interaction perspective and apply user-centered methods. This helps "open the black box" to uncover potential problems in the technology and understand the details of its use. It is, for example, a fundamental problem if the user experiences difficulties entering food items in a food tracking app, and a (small) usability problem like this might jeopardize a large and expensive clinical trial. Murray et al.[14] propose a heuristic evaluation method for digital health interventions (DHIs), which consists of 13 questions assessing whether there is a need and population for the technology, if it might be of benefit and impact, and other core aspects. This simple, question-driven approach to the evaluation of a DHI can lead to an accumulating knowledge base around an intervention in a timely and cost-efficient manner.

In summary, there are different questions that are relevant to address in different stages of the development of health technology. Figure 1 illustrates this; in the early stage of technology development, it is mostly relevant to investigate the technical feasibility of the technology, including its technical capability and accuracy. Once this is in place, it becomes relevant to assess the usefulness and usability of the technology. When the technology is more mature and stable, assessment of potential health effects can be commenced, and in the end clinical evidence can be established. When reporting an "evaluation" of a novel health technology, it is essential to be explicit about where in this space the evaluation is positioned and, subsequently, which kind of "claim" or conclusion can be drawn from such an evaluation. It is, for example, impossible to draw any conclusions about clinical efficacy or utility based on technical or usability studies.

ASSESSING CLINICAL FEASIBILITY

There is a growing need to be able to design and develop health technologies while being

Figure 1. Addressing different evaluation questions during different phases of the technology development and hence technology readiness levels. CUMACF: CACHET Unified Methodology for Assessment of Clinical Feasibility.

able to point to health benefits–in particular in the early stages of technology development and evaluation. For this purpose, technical and health scientists at the Copenhagen Center for Health Technology (CACHET) worked together to create the CACHET Unified Methodology for Assessment of Clinical Feasibility (CUMACF) methodology. The goal of CUMACF is to help researchers in the process of designing and developing health technology to run what we call "*clinical feasibility studies*," i.e., studies that help researchers understand whether the technology under design would be feasible to use in future health interventions, if implemented in clinical use. The position of CUMACF in the evaluation space is illustrated in Figure 1.

The purpose of CUMACF is twofold. First, borrowing from a design science perspective, CUMACF seeks to support an iterative design process, with frequent design and evaluation sessions involving end-users. The idea is to investigate the feasibility of the technology under design as early as possible—this saves time, effort, and money. Moreover, in contrast to a traditional RCT in which the intervention is treated as a black-box, CUMACF seeks to provide an understanding of the intervention (i.e., the technology under design) by providing insights into which parts of the technology (i.e., which features) help achieve a health outcome.

Second, borrowing from a health science perspective, CUMACF seeks to investigate health efficacy, i.e., the extent to which an intervention does more good than harm under ideal circumstances.[16] The goal is to gather early evidence on potential efficacy during design. CUMACF will not establish a high level of evidence since the study typically does not involve a control group and has insufficient statistical power. But a clinical feasibility study will help researchers understand the potential of the technology for health efficacy at an early stage and help researchers understand which other parameters, besides the technology itself, need to be (re)designed in order to obtain the desired health outcome.

A practical guide on how to use CUMACF is described in a technical report.[17] Below, we will provide an overview of the method and provide insights into its use.

METHOD

CUMACF targets the evaluation of "Personal Health Technology"[3] designed for the personal use by a patient. It is divided into three parts, assessing; (i) usage adoption, i.e., the degree to which the patient uses the technology; (ii) perceived usefulness and usability, i.e., the likelihood of successful adoption of the technology and acceptance by users; and (iii) health
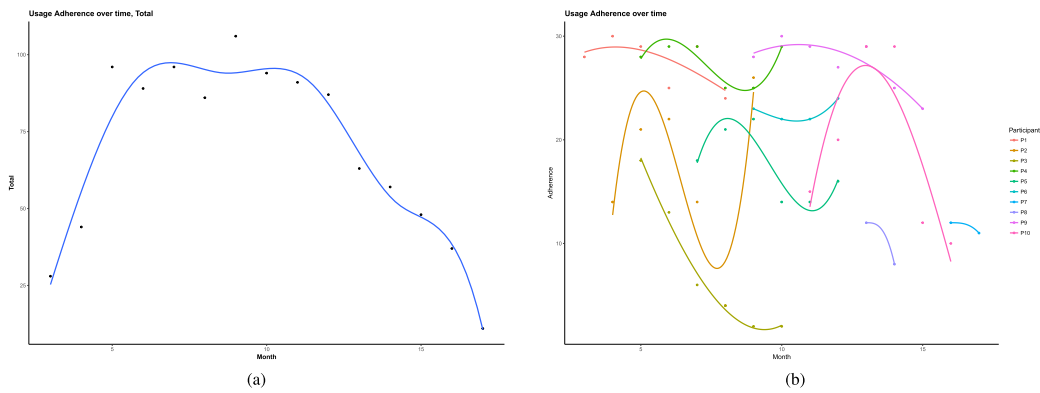
Figure 2. Example of usage adoption over the duration of a study (smoothed). (a) Total for all participants, and (b) per participant.

efficacy, i.e., the capacity for beneficial change or therapeutic effect of the intervention provided by the technology. All three things are assessed simultaneously during a clinical deployment of the technology.

Usage Adoption

A core prerequisite for assessing the feasibility of a health technology is to know whether the patient (i) uses the system in the first place, and (ii) uses the technology as instructed and prescribed. To verify this, assessment of *usage adoption** is beneficial. Usage adoption is a relative measure; it assesses to what degree the user uses the technology as compared to what is expected. For example, if a patient is asked to assess the depression level as a daily mood score, the adoption rate is the percentage of days of self-reported mood compared to the number of days the patient was enrolled in the study.

Calculating usage adoption depends on knowing the baseline, i.e., how often the user is supposed to use the system. This depends on what the user has been instructed to do (e.g., filling in a daily mood score), the duration of the study per participant and the availability of the system. Given this, adoption for a participant i can be calculated as

$$\text{adoption}_i = \frac{\text{usage}_i}{\text{duration}_i - \text{downtime}_i}.$$

If the abovementioned data are collected, detailed statistics on usage adoption can be reported; (i) overall, (ii) over time, and (iii) per participant, all of which take system availability into account. Figure 2 shows an example of how this analysis can be done.

Perceived Usefulness and Usability

The second part of CUMACF is to measure *perceived usefulness and usability*. According to research into psychometric assessment of technology acceptance, there is a strong correlation between users' perceived usefulness and usability of a system and the likelihood of future successful adoption and acceptance of the technology.[18] For example, a study showed that 80% of all activity tracking devices were abandoned after two months because "participants perceived the data collected as not useful."[19] Hence, perceived usefulness is the key for technology adoption.

CUMACF follows the Unified Theory of Acceptance and Use of Technology (UTAUT) methodology[5] combined with a few usability and behavior change questions, and is designed to assess the user's intention for *future* acceptance of the technology.

CUMACF applies a questionnaire, which is specifically designed to collect data on the likelihood of successful adoption of technology, its acceptance by the users, and their intentions to

---

*Some call this *"adherence"* to technology use, borrowing the term from clinical studies and treatment. However, we prefer not to use the term "adherence" since it carries a connotation that the technology is "prescribed" and "should" be used like a medical drug. This is often not the case with technology, which is seldom "prescribed."

	Health Professional	Patient	Automatic
Measurement (device)	medical device in clinic example: 12-led ECG	at home by patient example: 1-led ECG	automatic logging example: implanted ECG
Assessment (human)	clinical assessment example: HAM-D	self-report by patient example: PHQ	automatic assessment example: depression based on location

Figure 3. Taxonomy of health outcome measures. ECG: Electrocardiography; HAM-D: Hamilton Depression Rating Scale; PHQ: Patient Health Questionnaire.

use it for the intended health outcome. This is done according to the following five dimensions, adopted from Unified Theory of Acceptance and Use of Technology (UTAUT): (i) *health expectancy*, assessing the degree to which an individual believes that using the system will help the individual to attain *gains in health*; (ii) *effort expectancy*, assessing the degree to which an individual believes that *ease* is associated with the use of system; (iii) *social influence*, assessing the degree to which an individual perceives that *important others* believe the individual should use the system; (iv) *facilitating conditions*, assessing the degree to which an individual believes that an *organizational and technical infrastructure* exists to support use of the system; and (v) *behavioral intention*, investigating the degree to which an individual *intends to use the system*.

The CUMACF technical report[17] contains the entire questionnaire and detailed instructions on how to deploy it, and how to analyze and present statistics on perceived usefulness and usability.

Health Efficacy

Since CUMACF focuses on "feasibility," the methodology focuses on establishing *health efficacy*, i.e., involving patients who are carefully diagnosed, have significant symptoms from the disease in question, lack other serious illnesses, and are likely to follow and respond to the treatment based on the technology.[16]

Based on our experience in running several clinical feasibility studies, we have established the following general guidelines. (i) The number of participants (N) should be circa 20. (ii) Patients should be *recruited* who are carefully diagnosed and who potentially can benefit from the intervention, while at the same time are early adopters, i.e., have the skills, motivation, and ability to use the technology. (iii) Even though it needs to be adapted to the specific technology, the *duration* of intervention per patient should seldom be longer than six months. (iv) *Compensation* should be tailored to local ethics guidelines, but the technology and the infrastructure should be provided free of charge, including mobile and wearable devices. (v) The study protocol should allow for *adaptation* during the study. However, this should be restricted to adaptation, which only has a limited effect on the outcome measure of the study, and should primarily be addressing technical enhancements of nonfunctional software qualities, such as robustness, security, usability, and scalability.

Defining clinical outcome measures is clearly dependent on the health topic in question and the type of disease being addressed. For example, depressive symptoms may be measured using the Patient Health Questionnaire (PHQ) scale, whereas cardio-vascular disease symptoms are measured using an electrocardiography (ECG) device. In order to inspire researchers to come up with different outcome measures for a study, CUMACF provides a taxonomy that makes a distinction according to (i) *how* health outcome measures are obtained versus (ii) *who* measures it. This taxonomy, with some examples, is illustrated in Figure 3.

CONCLUSION

These years, we see an increasing number of mobile and wearable health technologies being designed and put into use. Most of these are designed to meet the healthcare challenges we are facing in terms of a growing demand with reduced availability of clinical resources. It is important to establish the clinical evidence of such technologies, i.e., do they actually address the clinical need they claim to? For this evaluation, the current golden standard is a RCT. However, an RCT is expensive, time consuming, and treats the technology as a black-box. For this reason, this method is less useful during the early design, development, and feasibility testing of novel technology.

There is a need to strike a balance in order to assess the clinical feasibility of a new technology before planning a full RCT. For this purpose, we have been crafting the CACHET Unified Methodology for Assessment of Clinical Feasibility (CUMACF). The overall objective of CUMACF is to provide a standardized way to assess the "feasibility" of a health technology during design and development. Such a standardized method will help to compare test results both *within* the iterative design of one specific technology as well as *between* different technologies. The former implies that a design team can assess the progression of its design across multiple iterations of the technology, whereas the latter implies that different technologies—maybe targeting the same health outcome—can be evaluated and compared in a more standardized manner. We hope that the reader finds inspiration in the CUMACF approach to evaluating clinical feasibility.

■ REFERENCES

1. J. E. Bardram and A. Matic, "A decade of ubiquitous computing research in mental health," *IEEE Pervasive Comput.*, vol. 19, no. 1, pp. 62–72, Jan.–Mar. 2020.
2. G. A. Fleming, J. R. Petrie, R. M. Bergenstal, R. W. Holl, A. L. Peters, and L. Heinemann, "Diabetes digital app technology: Benefits, challenges, and recommendations. A consensus report by the European Association for the Study of Diabetes (EASD) and the American Diabetes Association (ADA) Diabetes Technology Working Group," *Diabetes Care*, vol. 43, no. 1, pp. 250–260, 2020.
3. J. E. Bardram and M. Frost, "The personal health technology design space," *IEEE Pervasive Comput.*, vol. 15, no. 2, pp. 70–78, Apr.–Jun. 2016.
4. A. Peimankar and S. Puthusserypady, "An ensemble of deep recurrent neural networks for p-wave detection in electrocardiogram," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2019, pp. 1284–1288.
5. V. Venkatesh, M. G. Morris, G. B. Davis, and F. D. Davis, "User acceptance of information technology: Toward a unified view," *MIS Quart.*, vol. 27, pp. 425–478, 2003.
6. D. A. Rohani, A. Q. Lopategui, N. Tuxen, M. Faurholt-Jepsen, L. V. Kessing, and J. E. Bardram, "MUBS: A personalized recommender system for behavioral activation in mental health," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2020, pp. 1–13.
7. K. F. Schulz, D. G. Altman, and D. Moher, "Consort 2010 statement: Updated guidelines for reporting parallel group randomised trials," *Trials*, vol. 11, no. 1, p. 32, 2010.
8. G. Eysenbach, "CONSORT-EHEALTH: Improving and standardizing evaluation reports of web-based and mobile health interventions," *J. Med. Internet Res.*, vol. 13, no. 4, Dec. 2011, Art. no. e126.
9. A. Blandford, J. Gibbs, N. Newhouse, O. Perski, A. Singh, and E. Murray, "Seven lessons for interdisciplinary research on interactive digital health interventions," *Digit. Health*, vol. 4, Jan. 2018. [Online]. Available: https://doi.org/10.1177/2055207618770325
10. D. C. Mohr et al., "Trials of intervention principles: Evaluation methods for evolving behavioral intervention technologies," *J. Med. Internet Res.*, vol. 17, no. 7, 2015, Art. no. e166.
11. J. Firth et al., "The efficacy of smartphone-based mental health interventions for depressive symptoms: A meta-analysis of randomized controlled trials," *World Psychiatry*, vol. 16, no. 3, pp. 287–298, 2017.
12. S. M. Eldridge et al., "Trials of intervention principles: evaluation methods for evolving behavioral intervention technologies," *J. Medical Internet Res.*, vol. 17, no. 7, p. e166.
13. T. J. Moore, H. Zhang, G. Anderson, and G. C. Alexander, "Estimated costs of pivotal trials for novel therapeutic agents approved by the us food and drug administration, 2015-2016," *JAMA Internal Med.*, vol. 178, no. 11, pp. 1451–1457, Nov. 2018.
14. E. Murray et al., "Evaluating digital health interventions: Key questions and approaches," *Amer. J. Prev. Med.*, vol. 51, no. 5, p. 843–851, Nov. 2016.

15. P. Klasnja, S. Consolvo, and W. Pratt, "How to evaluate technologies for health behavior change in HCI research," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2011, pp. 3063–3072.

16. B. Haynes, "Can it work? Does it work? Is it worth it?: The testing of healthcare interventions is evolving," *Brit. Med. J.*, vol. 319, no. 7211, pp. 652–653, 1999, doi: 10.1136/bmj.319.7211.652.

17. J. E. Bardram, "CACHET unified methodology for assessment of clinical feasibility," Copenhagen Center for Health Technology, Copenhagen, Denmark, Tech. Rep., 2018. [Online]. Available: http://www.cachet.dk/research/cumacf.

18. F. D. Davis, "Perceived usefulness, perceived ease of use, and user acceptance of information technology," *MIS Quart.*, vol. 13, no. 3, pp. 319–339, Sep. 1989.

19. A. Lazar, C. Koehler, J. Tanenbaum, and D. H. Nguyen, "Why we use and abandon smart devices," in *Proc. ACM Int. Joint Conf. Pervasive Ubiquitous Comput.*, 2015, pp. 635–646.

Jakob E. Bardram is a professor in computer science at the Technical University of Denmark [www.dtu.dk] and is the director of the Copenhagen Center for Health Technology [www.cachet.dk]. His research interests include software architecture for mobile and ubiquitous computing, mobile sensing, and human-computer interaction. He received the Ph.D. degree from the University of Aarhus, Denmark. Contact him at jakba@dtu.dk.