Intro to Ethics for Digital Health
UR Health Lab

Course Introduction (00:00-07:21)

Hi, and welcome to an introduction to Digital Health Ethics. If you're watching this then you are likely a member of one of the various teams throughout the University of Rochester whose work is at the intersection of digital technologies and health care, whether that is through ISD, the UR Health Lab, CTSI, or any number of other collaborations that routinely happen here. Again, welcome.

My name is Kevin Boyd, and I am the Associate Director of Chaplaincy Services, but I am also trained as an ethicist. In this presentation we will introduce you to some basic ethical principles that are relevant to the design, development, and deployment of digital health tools. The goal is to help you cultivate your "ethical sensitivity." Think of it as an ethics-related spidey sense, that little flash that signals that you've run into an area where we need to think carefully before proceeding. We don't expect you to leave this presentation with a sense that you have all of the answers. Just an increased awareness of when you might need to ask more questions, and when some outside consultation might be beneficial.

You will encounter ethical questions at every stage of the digital health project. These projects may be for research, quality improvement or direct patient care. While no unified framework exists, we draw on principles from clinical and research ethics that may be familiar to some of you: (1) respecting persons, (2) benefiting individuals, and (3) promoting justice. Keeping these principles in mind at each stage of creation and use of a digital health tool will help you sense when more reflection is needed.

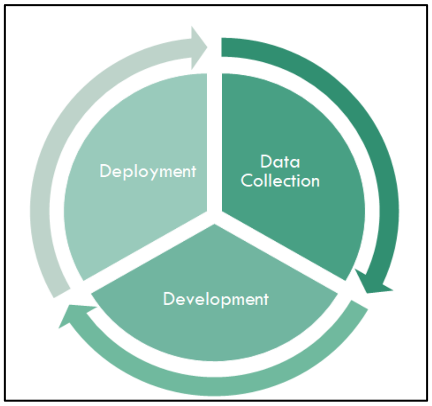

But why do any of this in the first place? What is the role of ethics in the digital health space? Why should we include this extra layer of consideration into already complex and time-consuming projects? Let me offer a few stories that illuminate the urgency of this issue.

Example 1: Racial Bias in Facial Recognition (1:56-4:09)

Joy Buolamwini was an undergrad in computer science at Georgia Tech trying to write a program that could play peekaboo. The software worked for the most part, only it consistently failed to recognize Buolamwini's face. Instead, she substituted her roommate's face and finished the project. Years later, while attending an entrepreneurship conference in Hong Kong she attended a demonstration of a "social robot", trained to respond and interact with human counterparts. Only, the robot did not engage with Buolamwini at all. Finally, years later, after finishing her Rhodes Scholarship at Oxford, Dr. Buolamwini came to the MIT Media Lab, where she worked on an augmented reality project she called "Aspire Mirror" which was designed to project uplifting visual images onto a user's face. Except the project wouldn't work on Dr. Buolamwini, unless she first put on a white mask.

So, what was going on? Well, it took a while to figure it out completely, but eventually a group of researchers looked at the data sets most widely used to train machine learning models/AI projects related to facial recognition. One of the most popular, Labeled Faces of the World, a collection of images from online news articles and image captions from UMass Amherst, was found to be more than 77% male and more than 83% white. Furthermore, since many of the images were from online news sources, they had a disproportionately high representation of quote-unquote "newsworthy" individuals-the U.S. President at the time, George W. Bush, was seen in 530 of the collection's images, more than twice the entire number of black women in the dataset.

The programs themselves lack intentionality, but they were absolutely perpetuating racial bias and discrimination because they couldn't even see people of color as people. Why? Because there weren't people of color adequately represented in the datasets used to train the machine. Biased datasets come to biased conclusions, and biased conclusions hurt people in very real ways.

Example 2: Pneumonia Risk Assessment (4:11-5:51)

Here's a clinical application. In the mid-1990s, pneumonia was the sixth leading cause of death in the US, with a mortality rate of about 10%. A group of researchers started working on a project to design an algorithm that would determine one very important, and very early medical decision with a pneumonia patient: whether to treat them as an inpatient or an outpatient. One of the researchers on that project, Richard Ambrosino, was attempting to train a rule-based model on the dataset of about 15,000 patients he was given. One day he noticed an odd rule that his system pulled from the data. The rule was, "If the patient has a history of asthma, then they are low-risk and should be treated as an outpatient."

Ambrosino is not a clinician, but he is a human being with common sense. Clearly, a patient with a significant history of asthma must be more vulnerable to pneumonia. He started talking to the physicians. He found a nuance that escaped the limits of a machine learning model. Patients with asthma were most often immediately admitted to the hospital, and often directly to the ICU for critical care. As such, those with asthma were actually less likely to die, on average, from their pneumonia than the general population – but it was precisely *because* they received high levels of care from the outset. The model, purely looking at rates of mortality, saw them as low risk. The reality is that they were low risk only because they were provided early access to inpatient care. The poorly defined outcomes would have led to a life-threateningly dangerous clinical practice if implemented.

Conclusion (5:53-7:21)

These are just a few of the complicated issues that suggest that ethics has a vital role to play in the use of technology in health care. Poorly defined outcomes can impact patient care. Poor datasets can compromise results. As the use of artificial intelligence becomes more frequent in healthcare, we run into real questions about transparency and value alignment. We have serious questions to answer about data capture, data access, and data storage. Are there secondary or unintended uses of this technology that might be harmful?

A common refrain in the technology world is, "move fast, break things." In the health care context, that sentiment can have devastating, even deadly consequences, because the "things" we are talking about are real people who we encounter in moments of extreme vulnerability. The digital health movement seeks to leverage emerging technologies to increase patient access and ultimately to improve patient outcomes. With the implementation of all new technologies, we also are confronted with new sets of questions and concerns. As we move into the subsequent modules of this presentation, we will engage in a more specific discussion of respect for persons, beneficence, and justice. In doing so we hope to elucidate some of these important questions and offer some resources for how you might address them at the University of Rochester.

Respect for Persons (07:24-17:56)

Introduction (7:24-11:17)

Hi, I'm Jon Herington, an assistant professor of bioethics at the University of Rochester. In what follows, I'll talk about a core principle in both clinical and research ethics: Respect for Persons. Fundamentally, showing respect for individuals requires acknowledging each person's capacity to make informed decisions about their own life: including whether and how they use a technology, how their data is used, and its influence over their daily care within the hospital system. Because we are working in health care, these issues are especially important because digital health projects may have a profound impact on people's health, their bodily privacy, and their autonomy. Because we work through the hospital, we are also subject to enhanced scrutiny regarding the collection, use, and storage of healthcare data.

One classic case of a failure to respect persons is the use of so-called "dark patterns", that use user interface designs to extract unnecessary information from users or nudge them towards default options. For instance, this is a login page for a popular holiday review app. In the center of the page, it prominently displays login options using social media accounts, and relegates other options to smaller, less central areas of the page. Because of the button placement, the vast majority of users will sign in with Facebook or Google, providing access to their friend lists and public photos; fewer will create a disconnected TripAdvisor account using the small text on the bottom; and fewer still will skip the data harvesting altogether using the link on the top right. This use of design tricks attempts to subvert people's reasoning about their own data privacy. Note that there may be legitimate reasons for using designs that "nudge individuals" to take default options, or that lower the cognitive burden of using an application, but doing so requires careful balancing of patient autonomy with patient benefit. If in doubt, seek advice.

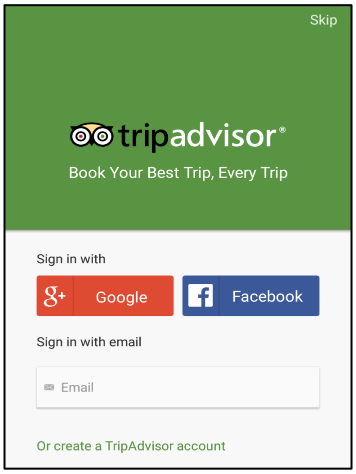

Another way in which we might disrespect persons is by failing to notify individuals of the collection of their data. For instance, so-called ambient intelligence systems have been developed to continuously collect video, thermal, audio, telemetric and wearable data from patients, healthcare workers, and visitors in wards and aged care facilities. These systems promise to help monitor patients for falls, optimize clinician workflows, and allocate resources. But collecting such sensitive data secretly and without consent undermines people's ability to control their own privacy, often during a time of intense vulnerability.

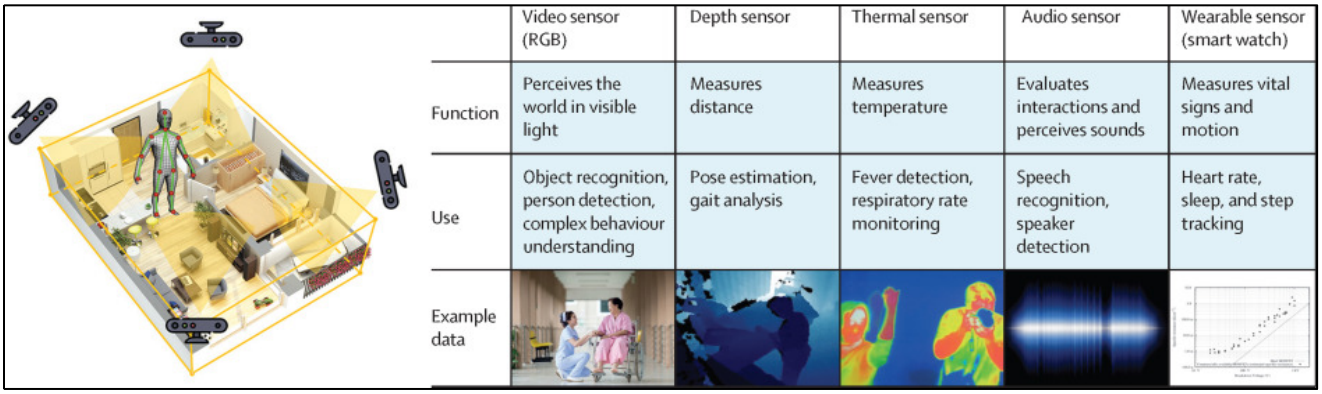

	Video sensor (RGB)	Depth sensor	Thermal sensor	Audio sensor	Wearable sensor (smart watch)
Function	Perceives the world in visible light	Measures distance	Measures temperature	Evaluates interactions and perceives sounds	Measures vital signs and motion
Use	Object recognition, person detection, complex behaviour understanding	Pose estimation, gait analysis	Fever detection, respiratory rate monitoring	Speech recognition, speaker detection	Heart rate, sleep, and step tracking
Example data					

So, what should you watch out for? Things that pose a high risk of violating an individual's autonomy in digital health projects include (1) a lack of transparency, (2) a lack of a consent process, (3) collection of identifiable or sensitive data, and (4) use of data beyond its original purpose. We consider each in turn.

Transparency (11:21-12:21)

Digital health projects should transparently disclose: (1) their existence, (2) their intended purpose, (3) the kinds of data they collect, and (4) the way that data is stored and used. These disclosures should be presented in a format that is easily digestible for people with low familiarity with standard data collection practices. Avoid walls of text that are easy to skip or ignore.

For applications that make predictions or recommendations, the performance of the application should be evaluated and disclosed to end users. Disclosures should include easily digestible information about the intended use of the technology, overall accuracy of the predictions, false positive and false negative rates.

Consent (12:23-13:52)

Digital health projects should give subjects an opportunity to make informed choices about the collection and use of their data. Crucially, this requires informing subjects about the application, asking for explicit *consent* to collect the data, and creating easy methods for subjects to withdraw their data.

Technologies that record data about people who are not the primary users of the application are particularly sensitive. This can occur when caregivers use applications to record data about patients without their explicit consent, or when home health apps inadvertently collect data on other household members. While this kind of collection often occurs anonymously, subjects should ideally be notified and provided an opportunity to refuse to participate in data collection.

Second, even where there are no direct participants- such as retrospective analyses of electronic health records - some kinds of projects should include community consultation. Any project that involves historically marginalized communities should be intentional about involving representatives of that community in the design, implementation and analysis of the research.

Privacy (13:55-15:24)

Digital health projects should respect individual privacy. This requires (1) limiting the collection of data to the features required to complete the intended task, (2) transmitting and storing data securely, and (3) limiting access to data to only those with a legitimate need to know.

Moreover, two special considerations apply to some kinds of projects. First, projects that collect video, audio, or wearables data should be carefully vetted to evaluate the risk that they will inadvertently collect or transmit the private activities of end users. Technical or policy controls should be enacted to minimize these risks, and the residual risks of video or audio capture should be outweighed by the benefits to the end users.

Second, most digital health projects should obey HIPAA privacy regulations. Additionally, most research projects must abide by the confidentiality provisions in the Common Rule on human subject protections. If you are unsure whether your project counts as research, consult the Research Subjects Review Board.

Responsible Data Re-use (15:28-16:07)

Finally, digital health projects should reuse anonymized data sensitively. While research subject and HIPAA regulations mostly allow unrestricted use of deidentified data, there are some cases where such data can be used to re-identify individuals. For instance, using anonymized transcripts of telehealth conversations between patients and mental health counselors may raise issues regarding dignity and privacy, even if the data is technically de-identified.

Conclusion (16:10-17:01)

Ultimately, respecting the personhood of our patients requires thinking carefully about whether our technological solutions treat them with dignity. This typically involves providing users with more control over the types of data they make available to researchers. This is traditionally accomplished by complex privacy notices with an opportunity to decline participation, but should ideally involve giving individuals more fine-grained control over the kinds of data they provide and what purposes they allow it to be used for. Ultimately, if we suspect that our projects are not transparent, or are making it more difficult for ordinary users to direct their own care, then we should stop what we are doing and seek advice.

Summary (17:04-17:55)

To review our section on Respect for Persons, let me offer three brief summary statements:

1. First, respecting persons requires transparency. It requires notifying subjects about the existence of digital health projects, the kinds of data we collect, and the value judgments we made along the way.
2. Second, respecting persons requires seeking their consent. Documentation should be easy to access and understand, and subjects should be able to exercise control over their participation and their data.
3. Finally, respecting persons requires protecting their privacy, limiting the data we collect to only those things which are necessary, and storing it securely and safely.

Beneficence (18:00-26:39)

Introduction (18:00-21:42)

Digital health is supposed to facilitate healthcare that is higher quality, more accessible, more efficient and more reliable. In this module, we discuss the ethical principle of Beneficence - by which I mean benefiting patients and society. This principle underpins most of the quality assurance and performance testing that we do. As applied to health technologies, beneficence requires ensuring that the use of our technologies (1) does not cause harm to patients or their families, (2) that it improves their care or experience, and (3) that the performance of the technology is robust across different contexts. As we shall see, there may be tradeoffs between these different goals, and thus, ensuring we discharge our duties to patients can be more difficult than it appears.

Let’s start with a couple of examples of digital health projects gone wrong, before moving to some general considerations during data collection, development, and deployment.

Epic Systems, a nationwide provider of electronic health record systems, developed a module that purported to warn clinicians if a patient was likely to develop sepsis during hospitalization. The system was based on about 400,000 patient encounters that occurred between 2013 and 2015 across three different health systems. Epic’s internal validation of the system at the three different hospitals suggested good performance, better than clinicians in most cases (with an area under the curve of between 0.73 and 0.86). Epic’s system was integrated into their Electronic Health Record system and was deployed at other 170 hospitals. In 2021, however, an external validation of the model in a large cohort of Michigan hospitals showed very poor performance- identifying only a third of patients with sepsis, despite generating alerts for every fifth patient who entered the emergency room. The high rate of false positives and false negatives generated a large amount of alarm fatigue amongst clinicians and likely diminished the quality of patient care. Epic ultimately recommended that users of its system retrain the model on a hospital's own data before it was clinically deployed in order to improve the accuracy of the model.

A second example involves the usability of applications. For instance, a 2016 study of the usability of a web-portal and SMS-based diabetes management system, found that two out of five key tasks – for instance, correcting a glucose measurement or exporting a glucose value – could not be completed by 30-40% of patients. Other studies have shown that usability is a key friction point for digital health projects that contributes to low usage and user drop-off, undermining self-management of chronic conditions.

So, what things help us avoid harm and benefit our patients?

User-Centered Goals and Design (21:45-22:31)

Applications which are developed with the input of their end users are more likely to serve their needs. Users also have different understandings of, and access to, technology. User-centered design tackles these problems by consulting with users to identify their specific needs and use contexts. It is particularly focused on ensuring the accessibility of the digital health application to all people and engaging in usability testing before deployment. This not only respects our users, but also ensures that unknown harms are avoided, and benefits are fully realized.

Data Quality (22:33-23:41)

We are all familiar with the slogan- "garbage in, garbage out"- and while there may have been big strides in the quality of data collection over recent years, applications are still being built on top of incomplete, convenient, or biased datasets. In general, developers should ensure data quality by: (1) understanding the source of their data (who collected it, for what purpose, from which location and time period), (2) involving subject matter experts or community members to interpret data variables or structures, and (3) consider collecting custom datasets where there are questions about the accuracy or structure of the existing data. Avoiding the use of convenience datasets is an important way in which we will generate high quality, efficient, and robust digital health projects that minimize the risk of harm and maximize benefits.

Performance Validation (23:44-25:11)

Digital health projects should have their performance validated more rigorously depending on the level of risk their deployment poses to patient health and wellbeing. For simple proof-of-concept models and applications, this may involve a simple test-training split of existing data. For those digital health projects which involve behavior management of low-risk conditions in the community, simple cohort testing and qualitative assessment may be sufficient. For applications with diagnostic, therapeutic or resource allocation implications for inpatients, it may be appropriate to consider a full clinical trial before widespread deployment. Likewise, applications which distribute resources amongst community members, or have implications for individuals' access to healthcare, should warrant a very high level of scrutiny. Finally, performance testing should be robust- evaluating performance and usability for different platforms, clinical contexts, and users. This was seen in the Epic sepsis predictor, where its performance degraded in lower-resource hospitals. Performance in one context does not ensure the same performance in a different context.

Conclusion (25:13-25:38)

Preventing harm and maximizing benefits to our patients requires careful consideration throughout the design, data collection, testing and deployment phases. If an app is being developed without input from end users or with convenience datasets, or deployed without robust and verifiable testing, then we should stop what we are doing and seek advice.

Summary (25:41-26:39)

To review our section on Beneficence, let me offer three brief summary statements:

1. First, benefiting our patients requires listening to them. We should ensure that we have asked end users what their needs are and allowed them to review prototypes.
2. Second, benefiting our patients requires ensuring our data is high-quality. We should have a good understanding of where, when and why data was collected. Data used to train machine learning models should be appropriately representative of the subjects who will be using the final application.
3. Finally, benefiting our patients requires carefully validating the performance of our digital health projects. Ideally, we will validate each project's performance using real-world trials, before it is deployed as part of patient care.

Justice (26:44-38:18)

Introduction (26:44-27:18)

Hi, this is Kevin, back to wrap up our final section: Justice. Justice is a huge concept and can mean a number of different things to different people. For digital health projects, justice ensures that 1) people are treated fairly and with equity, 2) that we are inclusive, and 3) that we have done as much work as we can to think through the unintended potential uses and abuses of the digital health tools that we are developing.

Equity (27:19-30:06)

We are intentional in our use of the word equity here and distinguish that from the idea of equality. Equality would mean that each individual or group of people is given the same resource or opportunity. Equity realizes that each person has different circumstances and allocates the resources and opportunities needed to provide equitable access and opportunity. What does that mean in the real world, or in a health care context? Equality would simply require disseminating important information—for instance, a post-visit summary from a health care clinic—in the same language (usually English), or making sure that the information on your hospital website is presented in a consistent English-language format. However, equity would require you to account for the non-native speakers and readers in your patient population and provide important health information in multiple languages. Equity would require that your website provide various accessibility tools to ensure that users with varied physical, cognitive, and linguistic differences can have access to necessary information.

Let's take scheduling as another example. If you've ever tried to call a physician's office to schedule an appointment, you know that it can be a frustrating and time-consuming process. Why not streamline that into an online system where patients could select their own physicians along with open clinic appointments that meet their scheduling needs, like the OpenTable app, but with health care instead of a restaurant reservation. From the perspective of simple equality, it could make sense: everyone has equal opportunity to select the appointment that they want from the convenience of their own computer and at any time during the day. From an equity standpoint, though, several issues arise. First, not everyone has reliable access to the internet. There are still many areas in our country without dependable service. Secondly, not everyone has a mobile phone, tablet, or computer that can access the internet. Individuals with lower socio-economic status, often the same people who already encounter barriers to accessing health care, could be shut out altogether. Finally, not all people use the internet in the same way. There may be cultural or demographic factors that impact how one uses the internet, and they could inadvertently create adverse health impacts if we were to shift to this entirely online mode of scheduling. An equitable version of this model would likely require retaining some ability to schedule in person or via phone, as well as ensuring the online interface is equipped with accessibility software.

Inclusion (30:07-32:32)

The long history of medicine is laced with paternalism, by which I expressly mean an instance when a healthcare provider or the larger health system assume they know what will be in their patients' best interests and make decisions for those patients without their consent. Too often, the patients themselves, or the communities who are the targets of larger public health initiatives, are entirely left out of these projects, or are included in only marginal ways. A commitment to inclusion means that the project invites input from stakeholders at every stage of its development, including the initial planning and design phases. This should be true across the board but especially so in cases where the digital health tool is engaging with vulnerable or historically marginalized groups.

Here are a few questions to ask as you build your solution. Were individuals from these groups included in the initial phases of problem identification? Were their experiences included in design and data capture? In other words, is the digital health project truly collaborative in nature with the people that it intends to help, and is it actually helping them?

A specialty care clinic with a focus on under- or uninsured patients might be able to meet a significant need in a large urban area, but if it is far away from the city's public transit routes, it isn't going to do much good. A mental health digital application might contain wonderful resources, but if they are all presented in ultra-high-definition videos that require high-speed data plans to stream, they aren't likely to benefit all of their intended audience. In both instances, intended users of the service could quickly identify the problem, while planners and developers outside of that lived experience might not be able to see it on their own.

Inclusion asks us to make space in the project to respectfully consider a variety of lived experiences and to include them when we collect and analyze data, and when we design, develop, and deploy digital health projects. Finally, as we have mentioned throughout this presentation, the idea of inclusion also requires that we do everything possible to make sure that the digital tools that we design are accessible to those who may have physical, cognitive, or linguistic differences.

Substantive fairness and secondary or unintended consequences (32:34-35:52)

When we speak of substantive fairness in digital health tools, we mean the duty to consider the potential for bias at every stage of a project's lifecycle. Recall the example from our introduction: facial recognition technology trained primarily on white, male faces. That dataset of mainly white and mainly male faces was clearly biased, and it clearly led to biased outcomes. This is a complicated problem. Training datasets are large, and it is not always easy to find out what data was used in their training process. With larger AI projects, the training set might be of an impossible size; large language models from Open AI and Google are scraping the entire internet as their training data. This can raise significant questions about the utility of the data, not only in terms of sheer size but also in relation to its geographic variation and collection methods.

We should also be diligent about finding out all that we can about the datasets that we use to train machine learning and AI models. If they are large, publicly available sets, then the data used for training should also be made available. If there are internal data sets from our own institutions or research, then we should investigate the data that was captured, how it was captured, and how it was interpreted. Although we often forget this, data is not recorded or analyzed in a vacuum, and the same biases that exist to create health disparities in the first place around race, gender, and socioeconomic status can easily be inadvertently reproduced in how we go about collecting our data. We should seek to eliminate as much bias as possible in the tools we develop.

Finally, to put it simply, justice means that we need to ask, "Is this something that is going to make a positive difference," "does this do good," and in some cases, "does this do more good than harm?" Justice reminds us to ask the question, "Could this hurt people, and if so, how?" For example, a recent AI bot from a large foreign supermarket chain was intended to help people reduce their food waste by generating recipe ideas from leftover food items. Basically, you could enter whatever things are in your pantry, or about to go bad, and this bot would give you a recipe idea. It sounds good, and it is surely an environmentally aware and sustainable idea—on the whole, a good. However, mischievous users quickly demonstrated that this bot would also offer "recipes" for things like deadly chlorine gas when the right combination of ingredients were input. The bot entitled that particular recipe, "Aromatic Water Mix." This was clearly not the intended use of this technology, but at the same time, it was entirely predictable given what we know about human nature and our propensity to immediately attempt to break or pervert any new technology.

Conclusion (35:53-37:20)

No one could anticipate every potential misuse of digital health technologies that they create. However, good ethics does require us to think seriously about the potential misuses of the things that we create, and to put some guardrails in place to prevent or minimize that harm. This is an incredibly crucial point, which is why I am ending with it. The capacities of large language model artificial intelligence systems have increased in exponential ways. The rate of technological progression always outpaces thoughtful reflection on the appropriate use of that technology. With artificial intelligence in particular, we are talking about a tool that has potential impacts across nearly every aspect of our lives. We need to be thoughtful, especially at these initial stages of integrating this technology into our lives and practice. Ethicists across the board have sounded caution, if not alarm, at the rate of advance, but as workers in the field of health care, we need to be extra careful in our use of these developing technologies. To be clear, we are not against technology in the least; we are simply advocating for thoughtful consideration of how the tools that we develop might be used and misused in the real world. We are responsible for them. This is an area where outside consultation, an extra set of eyes, can be especially valuable.

Summary (37:23-38:17)

To review our final section, let me offer three brief summary statements:

1. Justice takes into account different circumstances and contexts; it promotes a culture of equity in the attempt to create fair opportunity and access to all.
2. Justice is inclusive. Justice demands we seek and incorporate the feedback and experience of all users and that they be included in the design, development, and deployment of digital health tools.
3. Finally, justice requires us to think seriously about the potential unintended uses and abuses of the tools we create. We are responsible not only for the good they do, but also for the harm that they can produce. We need to think carefully not only about how we intend for them to be used, but how they could be used in malicious ways.

Course Conclusion (38:21-38:42)

This brings us to the end of our presentation. We know that it was a lot to cover, and even so, we also know that we've only touched on some of the core ethical issues. We don't expect anyone to feel like a content expert, but we do hope that we have cultivated a sensitivity to the role of ethics in the digital health space. Thank you.

Contact URMC Digital Health Ethics Consultants

Rev. Kevin M. Boyd
Associate Director, Chaplaincy Services
kevin_boyd@urmc.rochester.edu

UR Health Lab

Jonathan Herington
Assistant Professor, Department of Philosophy
Assistant Professor, Department of Health Humanities and Bioethics
jhering2@ur.rochester.edu

References + Recommended Reading

Introduction

- Digital Health (Yasmyne Ronquillo, Arlen Meyers, Scott J. Korvek; *National Library of Medicine*, 2023)
- From Data Collection to Model Deployment: A Comprehensive Guide to Building Successful ML/AI Models for Data Scientists and Engineers (3D Statistical Learning; *LinkedIn*; 2023)
- Study finds gender and skin-type bias in commercial artificial-intelligence systems (Larry Hardesty; *MIT News*; 2018)
- An evaluation of machine-learning methods for predicting pneumonia mortality (Gregory F. Cooper, Constantin F. Aliferis, Richard Ambrosino et al.; *Artificial Intelligence in Medicine*; 1997)

Respect for Persons

- The Belmont Report: Ethical Principles and Guidelines for the Protection of Human Subjects of Research (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research; *US Department of Health & Human Services*; 1979)
- Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns (Christoph Bösch, Benjamin Erb, Frank Kargl et al.; *Proceedings on Privacy Enhancing Technologies*; 2016)
- Ethical issues in using ambient intelligence in health-care settings (Nicole Martinez-Martin, Zelun Luo, Amit Kaushal et al.; *Lancet Digital Health*, 2021)

Beneficence

- Beneficence and the professional's moral imperative (Frank Stuart Kinsinger; *Journal of Chiropractic Humanities*; 2010)
- External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients (Andrew Wong, Erkin Otles, John P. Donnelly et al.; *JAMA Internal Medicine*; 2021)
- Factors Associated With Variability in the Performance of a Proprietary Sepsis Prediction Model Across 9 Networked Hospitals in the US (Patrick G. Lyons, Mackenzie R. Hofford, Sean C. Yu et al.; *JAMA Internal Medicine*; 2023)
- Quantifying usability: an evaluation of a diabetes mHealth system on effectiveness, efficiency, and satisfaction metrics with associated user characteristics (Mattias Georgsson, Nancy Staggers; *Journal of the American Medical Informatics Association*; 2016)
- Lack of Adoption of a Mobile App to Support Patient Self-Management of Diabetes and Hypertension in a Federally Qualified Health Center: Interview Analysis of Staff and Patients in a Failed Randomized Trial (Kathleen Thies, Daren Anderson, Benjamin Cramer; *JMIR Human Factors*; 2017)

Justice

- Introduction: What is Health Justice? (ed. Lindsay F. Wiley, Ruqaiijah Yearby, Brietta R. Clark et al.; *The Journal of Law, Medicine & Ethics*; 2022)
- Supermarket AI Gives Horrifying Recipes For Poison Sandwiches And Deadly Chlorine Gas (Matt Novak; *Forbes*; 2023)
- Study finds gender and skin-type bias in commercial artificial-intelligence systems (Larry Hardesty; *MIT News*; 2018)

Summary: Ethics Concerns in Digital Health Projects

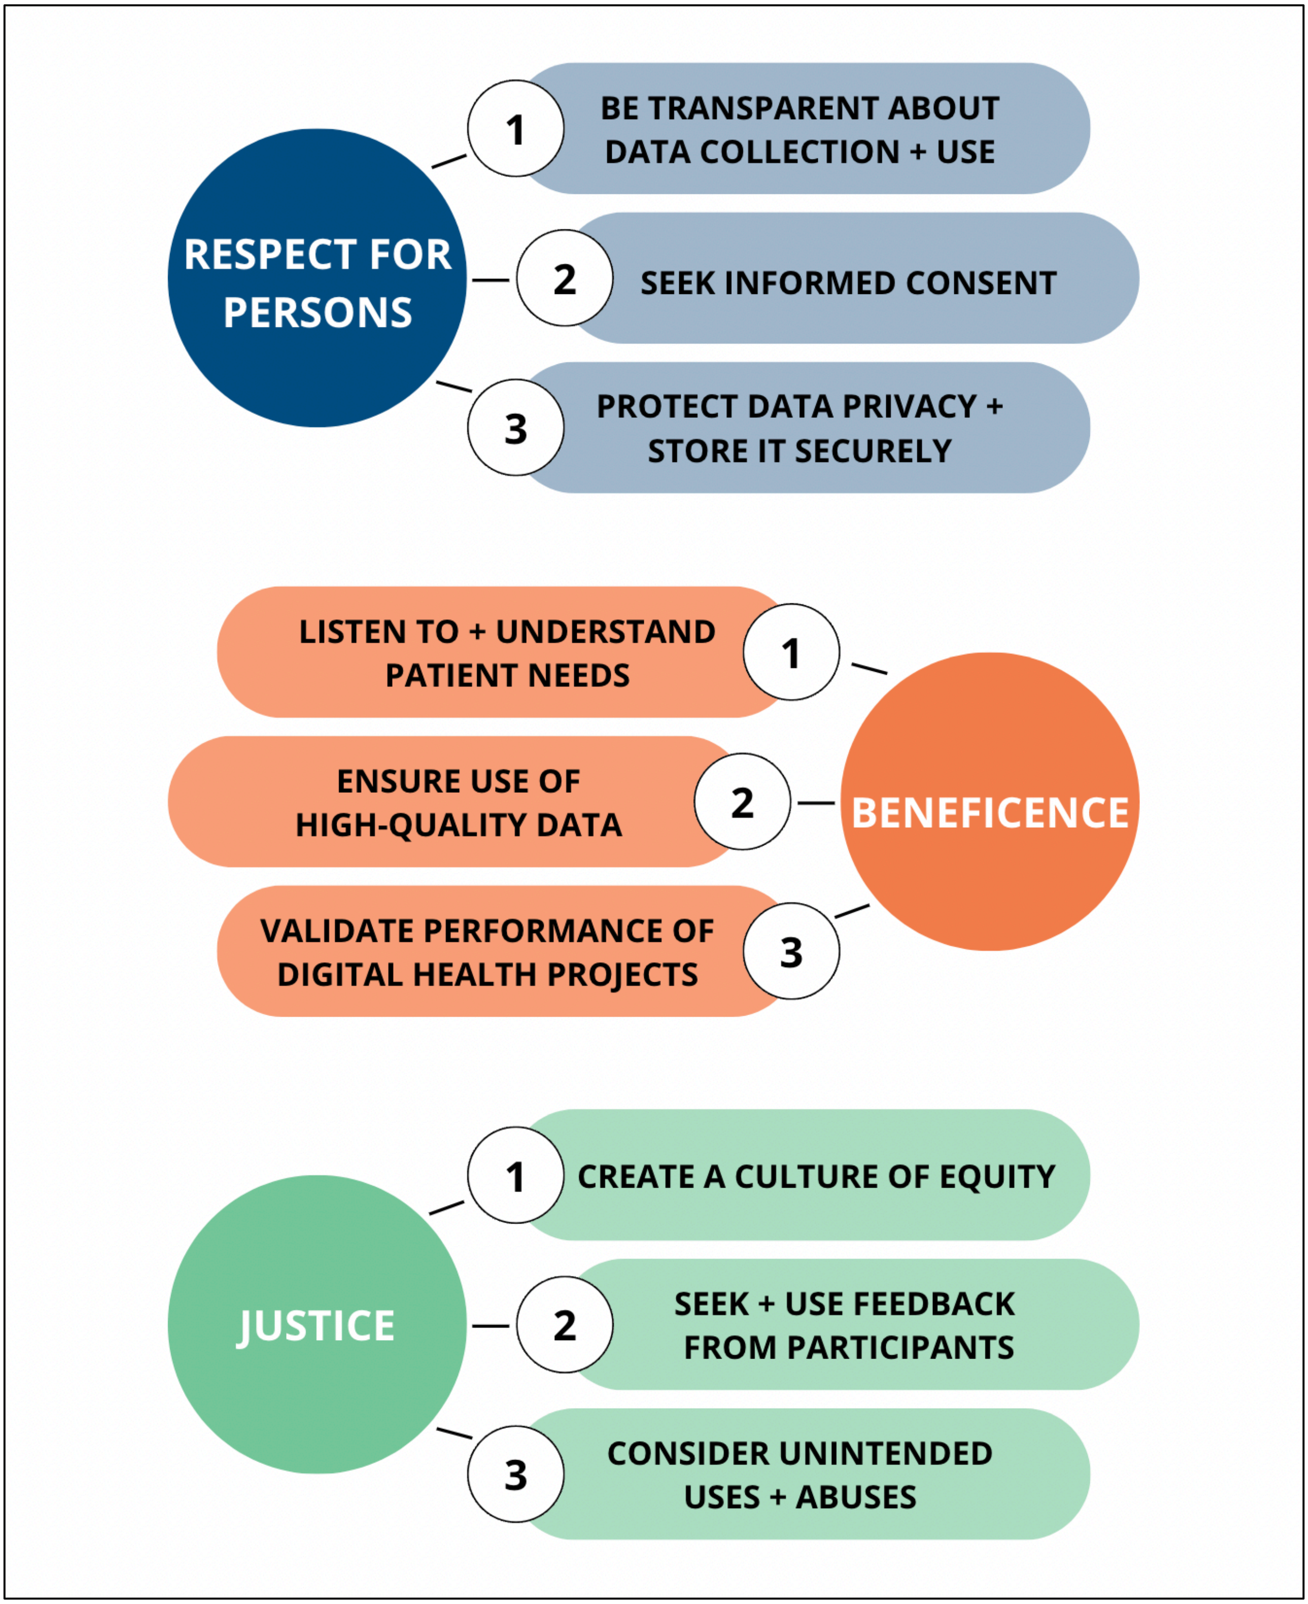

RESPECT FOR PERSONS

1. BE TRANSPARENT ABOUT DATA COLLECTION + USE
2. SEEK INFORMED CONSENT
3. PROTECT DATA PRIVACY + STORE IT SECURELY

BENEFICENCE

1. LISTEN TO + UNDERSTAND PATIENT NEEDS
2. ENSURE USE OF HIGH-QUALITY DATA
3. VALIDATE PERFORMANCE OF DIGITAL HEALTH PROJECTS

JUSTICE

1. CREATE A CULTURE OF EQUITY
2. SEEK + USE FEEDBACK FROM PARTICIPANTS
3. CONSIDER UNINTENDED USES + ABUSES