Check for updates

RESEARCH ARTICLE

The Artificial Intelligence in Public Health Toolkit: A novel resource for stakeholder engagement [version 1; peer review: awaiting peer review]

Alison J. Wright [1,2]*, Eva Jermutus [3]*, Ella Howes [2,4], Alison O'Mara-Eves [5], Clement Veall [2,6], Jennie Ives[7], Susan Michie [2]

[1]Institute of Pharmaceutical Science, King's College London, London, England, UK
[2]Centre for Behaviour Change, University College London, London, England, UK
[3]Social Science Research Unit, UCL Social Research Institute, University College London, London, England, UK
[4]Leeds Institute of Clinical Trials Research, University of Leeds, Leeds, England, UK
[5]EPPI-Centre, University College London, London, England, UK
[6]Department of Clinical and Movement Neurosciences, UCL Queen Square Institute of Neurology, University College London, London, England, UK
[7]Jennie Ives Freelance, Warwick, UK

* Equal contributors

Open Peer Review

Approval Status *AWAITING PEER REVIEW*

Any reports and responses or comments on the article can be found at the end of the article.

Abstract

Background

Artificial intelligence (AI) has considerable potential to enhance public health. People using AI systems for public health decisions, or who are affected by such decisions, may need to understand how these systems work, or articulate how much they want decision-makers to trust the system. This public engagement project, part of the Human Behaviour-Change Project, aimed to a) explore people's views regarding trust in, and use of, AI for public health decisions and, based on that, b) create a toolkit of resources to facilitate people critically questioning the use of an AI system.

Methods

Six online, public engagement workshops were conducted in England in 2021 to inform the content and design of the toolkit. Twenty-four people including members of the public, public health professionals, and researchers worked with a graphic designer to create the toolkit.

Results

The resulting 'AI in Public Health Toolkit' contains resources to enable people to evaluate AI systems and provides a roadmap for the decision process, a set of suggested questions to ask about an AI system, a guide to features of good answers and a 'personal views tool' prompting reflection on the answers received. Participants suggested that public health decision-makers should use the Toolkit to consult people representative of those affected by the decision to recommend whether an AI system should be used in that instance.

Conclusions

The 'AI in Public Health Toolkit' has the potential to facilitate public engagement in the use of AI in public health. The Toolkit gives those developing AI-driven systems a sense of the public's queries regarding such systems. The resources in the Toolkit can also facilitate conversations about broader AI applications to healthcare and public services.

Plain language summary
Artificial intelligence (AI) could help improve public health, but there are worries about its use. Public health decision-makers using AI systems need to be transparent and accountable. People who are affected by the decisions may want to learn more about the AI system and express how much they want decision-makers to trust it. This project aimed to explore people's views on using AI in public health and create a toolkit to enable people to question and understand AI systems.

Members of the public, public health workers, researchers and a graphic designer, all from England, participated in six online workshops in 2021. During the workshops, they discussed what content the toolkit should include and how it should look. Together, they created the "AI in Public Health Toolkit," to help people ask relevant questions about the pros and cons of AI systems and reflect on what they find out.

Participants proposed that public health leaders use the Toolkit to consult people affected by a public health decision about whether an AI system should be used. The Toolkit also illustrates people's key concerns about AI systems and can encourage discussions about AI's use in healthcare and public services.

Keywords
artificial intelligence, public engagement, trust, public health, decision making, toolkit

HB
CP

This article is included in the Human Behaviour-

Change Project collection.

Corresponding authors: Alison J. Wright (alison.j.wright@kcl.ac.uk), Susan Michie (s.michie@ucl.ac.uk)

Author roles: Wright AJ: Conceptualization, Data Curation, Formal Analysis, Investigation, Methodology, Project Administration, Resources, Supervision, Visualization, Writing – Original Draft Preparation, Writing – Review & Editing; **Jermutus E**: Conceptualization, Data Curation, Formal Analysis, Investigation, Methodology, Project Administration, Resources, Visualization, Writing – Original Draft Preparation, Writing – Review & Editing; **Howes E**: Conceptualization, Data Curation, Formal Analysis, Investigation, Methodology, Project Administration, Resources, Visualization, Writing – Review & Editing; **O'Mara-Eves A**: Conceptualization, Data Curation, Formal Analysis, Investigation, Methodology, Project Administration, Resources, Visualization, Writing – Review & Editing; **Veall C**: Conceptualization, Data Curation, Formal Analysis, Investigation, Methodology, Project Administration, Resources, Visualization, Writing – Review & Editing; **Ives J**: Conceptualization, Investigation, Methodology, Resources, Writing – Review & Editing; **Michie S**: Conceptualization, Funding Acquisition, Supervision, Writing – Review & Editing

How to cite this article: Wright AJ, Jermutus E, Howes E *et al.* **The Artificial Intelligence in Public Health Toolkit: A novel resource for stakeholder engagement [version 1; peer review: awaiting peer review]** Wellcome Open Research 2024, **9**:420 https://doi.org/10.12688/wellcomeopenres.22480.1

First published: 29 Jul 2024, **9**:420 https://doi.org/10.12688/wellcomeopenres.22480.1

Introduction

Artificial intelligence (AI), defined as a system's ability to "correctly interpret external data, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation" (Kaplan & Haenlein, 2019; p.17) has great potential to improve health outcomes. Considerable attention has been paid to applications of AI in clinical medicine (e.g. Matheny *et al.*, 2020; Yu *et al.*, 2018). There is also growing work on the uses of AI and other machine learning to support public health tasks (Fisher & Rosella, 2022; Mhasawade *et al.*, 2021). AI combined with rich new data sources, such as social media, apps and wearable or environmental sensors, has enabled new approaches to public health surveillance (e.g. Serban *et al.*, 2019). In addition, AI can be used in predicting future health outcomes and identifying at risk populations who may most benefit from public health interventions. For example, machine learning approaches have been used to identify housing units more likely to be unsafe, enabling housing inspectors to better target their efforts (Robb *et al.*, 2022). AI could also aid the design of public health interventions by integrating evidence more efficiently than humans can, predicting which interventions are most likely to work for a given population or in a given setting (Michie *et al.*, 2017).

Despite the potential benefits of applying AI systems to public health, several potential issues have been raised. The first concerns the adequacy of the performance of the AI systems for the task at hand. What constitutes "adequate performance" will depend on the task and the consequences of poor performance. While people may not mind receiving AI-based movie recommendations from an online streaming service that do not reflect their actual taste in movies, they are likely to be more concerned about inaccurate performance of health-related AI applications.

The second issue concerns the risk that applications of AI in public health could exacerbate existing health inequalities. Importantly, while an AI system's performance may be reasonable on average, this may disguise poorer performance for members of minority groups. Poorer performance for minority groups can result from using training data that is not representative of the population or the context in which the AI system is being applied, or where important confounding variables are not included (Chen *et al.*, 2021). Alternatively, if training data reflects societal biases, some AI training approaches can learn and even amplify these biases (Fisher & Rosella, 2022).

Third, an AI system can raise further ethical concerns if the training data required includes variables that spark privacy concerns. Individuals may object to having their sensitive data collected for one purpose (e.g. providing the individual with healthcare) used for another (e.g. developing an AI application) that does not benefit them directly. Such concerns may be exacerbated when data is passed to a third-party organisation or without the data subjects' knowledge or consent (e.g. Hern, 2017). Concerns may also relate to the risk of sensitive personal information included in training data being exposed if a data breach occurs.

Fourth, AI systems can have poor explainability or interpretability. AI algorithms can contain many variables modelled in complex ways, making it virtually impossible for humans to understand how an AI system arrives at its output. Lack of interpretability can adversely affect trust in an AI system (Fisher & Rosella, 2022) and may exacerbate concerns about algorithmic bias, due to the difficulty of verifying what factors are influencing the system's outputs.

Finally, going beyond issues with an AI system itself, there are concerns about how a system might be deployed as part of a public health decision-making process (Smith *et al.*, 2020). In particular, there may be no mechanisms that enable people affected by the use of the AI system to have input into - and challenge - the decision to use the system.

The Human Behaviour-Change Project (HBCP), a collaboration between behavioural, social, and computer scientists, included a public engagement project to explore potential users' trust in the use of AI to support public health decisions (Michie *et al.*, 2020). The aim of the HBCP was to develop a "knowledge system" that uses AI to make predictions about the outcomes of behaviour change interventions in a given context, based on evidence from published randomised controlled trials. The system has been developed for the initial use cases of smoking cessation and physical activity, given their importance to health and wellbeing (Michie *et al.*, 2017). It is envisaged that the system will eventually be applicable to other health-related behaviours, and indeed, any behaviours.

The HBCP knowledge system has a long-term vision of being used in public health decision-making to predict, for example, what behaviour change services or support would be most effective for residents of a certain area. Instead of relying on time consuming or incomplete approaches to synthesise evidence, the aspiration is that decision-makers can instead request the knowledge system to predict the outcomes of different interventions applied to the relevant population in the relevant setting. This could then feed into decisions about what services to fund or commission. However, a key question concerns the extent to which public health decision-makers be accountable to the stakeholders affected by the use of AI systems, and whether key groups of stakeholders will and consider the AI system's use acceptable.

This public engagement project seeks to address these questions and to provide resources to assist those using AI systems to help decision-making. The project uses the type of knowledge system being developed by the HBCP as an exemplar for exploring and facilitating appropriate trust in AI applied to public health decisions more broadly. Appropriate

trust here means calibrating trust in line with an AI system's capabilities as well as one's own priorities (Schlicker & Langer, 2021). Appropriate trust is considered central to encouraging acceptance of AI that is well-founded and beneficial. The project brought together health-interested publics, scientists, public health practitioners and a designer to engage with key aspects of AI systems as applied to public health. In doing so, the aim was to produce a toolkit of resources to facilitate people's ability to critically question the nature of an AI system, thereby decide how much they want public health decision-makers to trust the system.

This paper describes the workshop process by which the toolkit was developed, presents the resulting toolkit, and outlines how the toolkit can be used. A separate forthcoming paper will present more detail and evaluation of the public engagement approach used in this project – here we are seeking to introduce the Toolkit as a resource for potential users, explain how it was created and what it could be useful for.

Methods
Bringing together the team of participants
Due to the coronavirus pandemic, initial plans for face-to-face workshops had to be replaced by online workshops. Members of the public with an interest in health interventions were invited from two sources. First, we approached people who were current or previous users of smoking cessation services offered by Everyone Health, a company that provides community health and wellbeing services, including smoking cessation, weight management and physical activity services. In England and Wales, responsibility for health improvement and some health protection elements of public health rests with local authorities. Everyone Health is commissioned by a number of local authorities to provide behaviour change services for their residents for free at the point of delivery. We also invited people involved in delivering behaviour change services whose jobs could be affected by use of AI for public health decision making, namely Everyone Health staff members, to participate in the workshop.

Second, we advertised the project through the UCL Co-Production Collective, to its list of people who had expressed interest in being involved in health-related public engagement projects. These three groups were joined by a group of scientists from the HBCP team, with backgrounds in behavioural science, health psychology, evidence synthesis and trust in artificial intelligence, and a graphic designer with experience in designing materials in collaboration with public engagement groups.

Workshops
Six 2-hour online workshops were held via Zoom in the evenings between September and October 2021. The HBCP team started with an initial plan for the workshops and the structure; the content of the workshops was refined each week based on the discussions at previous workshops and feedback from

participants. A separate forthcoming paper will present more detail and an evaluation of the approach taken to public engagement in this project.

The workshops were facilitated by the HBCP research team and a graphic designer. An outline of workshop content is provided in Table 1. The toolkit was revised based on feedback received in Workshop 6 and this version sent to participants via email, for another round of feedback. This feedback was then incorporated into the final version.

Results
A process for recommending whether an AI system should be used for public health decision-making
The group suggested a process for public health care decisions that could be made using AI system outputs, resulting in a Toolkit containing resources for asking questions to learn about the AI system; evaluating the quality of answers to these questions; and reflecting on one's perspective on the AI system in light of the acquired knowledge.

These were organised into a "roadmap" of five steps (see Table 2). At the core of the process is the idea that decision-makers (e.g., a public health policy-maker) need to consult a panel of people, representative of those affected by these decisions, to determine whether an AI system should contribute to making such decisions.

In such a consultation, the panel would be given information about the public health decision to be made, such as why it needed to be made and the possible outcomes. Panel members would then be given the opportunity to find out more about AI, ask questions about the AI system and reflect on how they felt about the answers received. Following this process, panel members would share their views about whether the AI system's recommendation or other output should be used as part of the decision-making process. This could include taking a vote on whether the AI system's outputs should be used as part of the decision.

The AI in Public Health Decisions Toolkit components
The Toolkit has four components, all available as pdfs for download at the AI in Public Health Decisions Toolkit website (https://www.artificialintelligenceinpublichealth.org/home). Much of the content, the choice of colour scheme, images, layout and fonts were based on views of the public engagement workshop participants. The components are:

1. The Roadmap - "Introducing the five-step process":
A summary of the five-step process. It introduces the concepts of public health and AI systems applied to public health decisions and notes that AI systems have pros and cons, hence the need for this resource.

2. Suggested questions – "Not sure what to ask"
Lists questions that people might want to ask about an AI system, grouped into eight topic areas, based on questions

Table 1. Summary of workshop content.

Workshop	Topics	Exemplary activities of workshop
1	Introductions to team members and to the project Discussed the influence decision-makers have on the services available in a certain area	Ice breaker: abstract image for members to comment on to illustrate that there is a range of views which are all valid Introduction to concept of behaviour change intervention Discussing members personal experience of changing a behaviour such as smoking via Jamboard by questions such as 'Did you use any sort of organised programme or product to help?'
2	Understanding AI – how it works and potential problems	Using Mentimeter to gauge what people associate with term 'intelligence' and which services using AI they might already use (e.g. Netflix, GoogleMaps, Amazon, Uber…) Introduction to how AI works Introduction to issues with AI (bias, accuracy, training data,..)
3	Benefits of AI and weighing them against potential problems Considering what trust is, what influences trust and what people wanted to know to decide whether to trust an AI system	Recap of AI-related issues from previous week Introduction to benefits of AI Introduction to trust including a series of Jamboard questions to gauge what people associate with trust; what they would want to know about AI before deciding to trust it etc.
4	Introduction to graphic design – considering how we want the resources to look and what we want the audience to feel	Recap of previous week Introduction to graphic design and its principles through members commenting on materials prepared by graphic designer Jamboard activities used to explore what group wants audience to feel/think and do when looking at resources in the toolkit.
5	Bringing together design and content, reviewing draft designs	Recap aims of project and what we want target audience to think/ feel/ do Discussion and feedback on drafts of design and content of the resources
6	Reviewing the updated designs	Recap of resources in toolkit Review and discuss updated designs of resources Reflection on experience of project through Jamboard questions such as 'What has been the most interesting thing about these workshops for you?'

Table 2. Overview of five-step process outlined in the Roadmap resource and how other resources map onto those steps.

Step	Description	Resource
1	Finding out about the decision being made and the people affected by it	
2	Optional step to learn more about AI in general by watching a video	Link to external video by KI-Campus (2020)
3	Asking questions to find out more regarding the AI system being applied to the current decision	"Not sure what to ask" to help choose questions "Not sure what's a good answer" to see what to look for in the answers
4	Considering one's opinion of the AI system, based on the answers received to your questions	"Your personal views"
5	Providing feedback to public health practitioners who are considering using the AI system	

generated by the workshop participants. Readers are invited to select which questions they would like answered. The resource includes:

- A "question bank" where the eight groups of questions are presented in speech bubbles, across two pages (see Table 3 for the full list of questions).

- A mind map with all eight groups of questions presented on one page

  o The mind map also includes a blank mind map diagram on the second page, for people who wish to add their own questions

The two formats reflect feedback from workshop participants, some of whom felt that a mind map was a very helpful way to organise questions on a complex topic, while others found the presentation of all eight groups of questions on one page overwhelming and preferred the question bank design. It was agreed that both options should be retained in the Toolkit to maximise accessibility. Those using the resource to conduct a public consultation around the use of an AI system can choose one or both options to share with participants.

3. Guide to features of good answers – "Not sure what's a good answer?":
The guide is arranged into the same eight topic areas as the suggested questions. We have avoided suggesting specific

cut-offs for properties of the system such as accuracy of predictions. "Acceptable" accuracy may vary according to the problem the AI system is being used for and the consequences of the decision. Participants emphasised that they did not want to assess detailed technical data about the performance of the AI system themselves, but instead wanted to know what someone with suitable expertise would make of the system. For example, they wanted to know whether someone with suitable training would be able to explain how the AI system worked, rather than assessing the explainability of the AI algorithms for themselves.

4. Personal views tool – "Your personal views":
The tool gives people a chance to reflect on how they feel about the answers they have received to their questions in a structured fashion. It was initially inspired by values clarification tools in the healthcare decision aids literature (see e.g. Witteman *et al.*, 2021) and then refined based on input from workshop participants. People are asked to respond to 11 statements (Table 4) regarding how satisfied they are with what they have found out about the AI system being applied to the specific public health decision. For each statement, people can select strongly disagree, disagree, unsure, agree or strongly agree.

Responses recorded on the personal views tool can form the basis for discussions with public health leaders about whether an AI system's outputs should be used as part of a particular public health decision. If the public engagement process

Table 3. Overview of the eight topic areas and their suggested questions.

Topic area	Questions
What's the decision about?	• What is the AI system being used for? • Does the person using the AI system have enough time and training to use it correctly?
The AI system's background	• Who developed the AI system? • Why did they want to make the AI system?
What data was used to train the AI system?	• Was a broad range of data used? • Was the data relevant to the people affected by this decision? • Is people's personal information being used? If so, how will it be kept safe?
How well does the AI system work?	• How reliable and accurate is the system? • What's the evidence for this?
Do other people recommend the AI system?	• Is the AI system recommended by previous users? • Has the AI system been approved by an independent expert body?
How fair and unbiased is the AI system?	• Does the system work equally well for all groups of people? • How well does the system take diversity into account? • Has anything unfair ever happened, due to someone using this AI system to help make decisions?
Is there a clear explanation for how the AI system works?	• How much do we know about how the AI system comes up with its results?
Human support:	• Is there a human who can answer questions about the AI system, if it's giving unclear or strange results?

Table 4. Statements from the "your personal views" tool.

Statement
I feel I understand the decision that the AI system is being used to help with
I'm satisfied with what I found out about who is using the AI system
I'm satisfied with what I've found out about who developed the AI system and why they made it
I'm satisfied with what I found out about whether the AI system was trained on data relevant to the decision and the people the decision affects
I'm satisfied with what I found out about whether the system uses people's personal data
I'm satisfied with what I found out about how well the AI system works
I'm satisfied with what I've found out about whether other people recommend the AI system
I'm satisfied with what I found out about whether the system has been approved by an expert group
I'm satisfied with what I've found out about whether the AI system is fair and unbiased
I'm satisfied with what I've found out about whether there's an explanation of how the AI system works
I'm satisfied with whether there's someone who can give advice about the AI system if needed

includes taking a vote on whether an AI system should be used as part of a particular public health decision, the personal views tool may help people decide how to vote.

Discussion

The AI in Public Health Decisions Toolkit provides a process that enables citizens to share their views about whether a given AI system's recommendations should be used as part of a public health decision. The Toolkit is generic, designed to work across a variety of AI systems and types of public health decisions. Users can adapt it based on the particular AI system. For example, questions around the use of personal data may not be relevant for all systems. The Toolkit and the associated process provide public health leaders at national or local levels with a way of providing accountability to citizens affected by public health decisions made with AI. It also enables the public to critically question the use of AI systems as well as supporting individuals in deciding how much to trust a given system.

The Toolkit questions provide the opportunity to anticipate what the public might want to know about an AI system that is being used for public health decision-making. For people affected by the use of AI systems, the question bank offers a guide for querying the system and a starting point for critically engaging with the technology.

For organisations that develop AI systems, the Toolkit provides questions that members of the public want answered if an AI system is to be applied in a public health context. This can guide what information about the AI system should be provided to public health decision-makers who are considering using the system. It also gives AI system developers an opportunity to increase transparency in areas of most concern to members of the public. It may also be useful for developers to consider these questions as a thought exercise while developing the system, so that they can foresee potential objections and take steps to mitigate them.

Although the impetus for creating the Toolkit was the type of AI-based knowledge system the HBCP aimed to develop, the Toolkit is relevant to a wide variety of AI systems applied to public health, healthcare and public services. The Toolkit component, "What's a good answer?", can be useful even in contexts where the decision has already been made to use an AI system. The Toolkit provides an indication of what people's concerns and information needs might be and supports public health, healthcare and public service leaders in addressing them.

Strengths & limitations

The Toolkit was developed as result of a public engagement process involving a diverse group of people. We enabled a broad range of people to contribute by organising the online workshops in evenings, based on a poll of what time potential participants would find easiest to attend. We also provided compensation in recognition for participants' time in joining the workshops and provided data vouchers to cover the costs of joining the workshops online. Nevertheless, we inevitably will not have reached the full range of people potentially affected by the use of AI in public health. In particular, our recruitment materials described the project as being about a "computer system," which may have deterred individuals who were particularly unconfident with, uninterested in, or mistrustful of, digital technology.

The Toolkit design attempts to address the needs of neurodivergent people, based on feedback from neurodivergent workshop participants. For example, the text has high readability, we used dyslexia-friendly fonts and selected design elements based on what neurodivergent participants said they found most helpful for processing information. The Toolkit is

applicable to a variety of AI applications; can be adapted to suit different domains and is available for others to use at https://www.artificialintelligenceinpublichealth.org/home.

Following the five-step process within the Toolkit requires time and investment from public health or other healthcare or public service leaders. In particular, leaders will need to articulate the nature of the decision and why they are considering using an AI system, and to have access to suitable technical information to be able to answer people's questions. However, investing this time could avoid later objections to the AI system's use due to lack of transparency and avoid negative media coverage of the use of the AI system. AI system developers will have a role to play in providing technical information to enable leaders to answer the public's questions about the AI system. They will need to provide information in everyday, highly accessible language, avoiding technical and specialist terminology. While developers may have concerns about potentially sharing commercially sensitive information about the AI system, they do not have to provide the technical detail about how exactly an AI system works to provide the information that the public are likely to want about an AI system.

Conclusions

AI systems have the potential to improve public health through a range of avenues. However, no AI system is perfect and those affected by the use of AI systems can have a range of concerns. Applying the 'AI in Public Health Toolkit' provides a way for people considering the use of AI for a particular public health decision to provide accountability to those who will be affected by the AI system's use. The Toolkit can also inform conversations about applications of AI to decisions in public health and healthcare.

Data availability

No data are associated with the article.

References

Chen IY, Pierson E, Rose S, *et al.*: **Ethical Machine Learning in healthcare.** *Annu Rev Biomed Data Sci.* 2021; **4**: 123–144.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Fisher S, Rosella LC: **Priorities for successful use of Artificial Intelligence by public health organizations: a literature review.** *BMC Public Health.* 2022; **22**(1): 2146.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Hern A: **Royal free breached UK data law in 1.6 m patient deal with google's deepmind.** *The Guardian.* 2017; **3**.
Reference Source

Kaplan A, Haenlein M: **Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of Artificial Intelligence.** *Bus Horiz.* 2019; **62**(1): 15–25.
Publisher Full Text

KI-Campus: **Artificial Intelligence explained in 2 minutes: What exactly is AI?** [Video]. 2020.
Reference Source

Matheny ME, Whicher D, Israni ST: **Artificial Intelligence in health care: a report from the national academy of medicine.** *JAMA.* 2020; **323**(6): 509–510.
PubMed Abstract | **Publisher Full Text**

Mhasawade V, Zhao Y, Chunara R: **Machine learning and algorithmic fairness in public and population health.** *Nat Mach Intell.* 2021; **3**: 659–666.
Publisher Full Text

Michie S, Thomas J, Johnston M, *et al.*: **The Human Behaviour-Change Project: harnessing the power of Artificial Intelligence and machine learning for evidence synthesis and interpretation.** *Implement Sci.* 2017; **12**(1): 121.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Michie S, Thomas J, Mac Aonghusa P, *et al.*: **The Human Behaviour-Change Project: an Artificial Intelligence system to answer questions about changing behaviour [version 1; peer review: not peer reviewed].** *Wellcome Open Res.* 2020; **5**: 122.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Robb K, Amigo ND, Marcoux A, *et al.*: **Using integrated city data and machine learning to identify and intervene early on housing-related public health problems.** *J Public Health Manag Pract.* 2022; **28**(2): E497–E505.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Schlicker N, Langer M: **Towards warranted trust: a model on the relation between actual and perceived system trustworthiness.** In: *Proceedings of Mensch und Computer 2021.* 2021; 325–329.
Publisher Full Text

Serban O, Thapen N, Maginnis B, *et al.*: **Real-time processing of social media with SENTINEL: a syndromic surveillance system incorporating deep learning for health classification.** *Inf Process Manag.* 2019; **56**(3): 1166–1184.
Publisher Full Text

Smith MJ, Axler R, Bean S, *et al.*: **Four equity considerations for the use of Artificial Intelligence in public health.** *Bull World Health Organ.* 2020; **98**(4): 290–292.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Witteman HO, Ndjaboue R, Vaisson G, *et al.*: **Clarifying values: an updated and expanded systematic review and meta-analysis.** *Med Decis Making.* 2021; **41**(7): 801–820.
PubMed Abstract | **Publisher Full Text** | **Free Full Text**

Yu KH, Beam AL, Kohane IS: **Artificial Intelligence in healthcare.** *Nat Biomed Eng.* 2018; **2**(10): 719–731.
PubMed Abstract | **Publisher Full Text**