

WHITE PAPER | AUGUST 2022

Conducting an Effective IRB Review of Artificial Intelligence Human Subjects Research (AI HSR)

Tamiko Eto, MS CIP

Technology in Human Subjects Research (HSR)

<https://TechInHSR.com>

Technology in HSR.com

Table of Contents

03	Executive Summary
04	Scope of This Document & Definitions
05	AI Human Subjects Research (AI HSR)
06	Applying Federal Definitions
08	Defining "Clinical Investigations" and "Research"
09	Defining "Human Subjects"
11	Current Regulatory Framework and IRB Function
12	Outstanding Challenges & Future Recommendations
15	Limitations of These Recommendations & Conclusion
16	Addendum A: AI HSR Decision Tree
17	Addendum B: Exempt Category Decision Tree
18	Works Cited / Resources
20	Author Bibliographies, Copyright & Disclaimer
21	Addendum C: AI HSR IRB Reviewer Checklist

NOTE: The Secretary's Advisory Committee on Human Research Protections (SACHRP) is a federal advisory committee that provides advice and recommendations to the Secretary of HHS on issues involving the protection of human subjects in research.

The committee meets three times a year. Meetings are open to the public. In 2021, we were asked to present on making determination of "human subjects" involvement in Artificial Intelligence research projects, and ethical considerations when conducting such review. Additional SACHRP meetings relating to AI HSR can be found on their website here:

<https://www.hhs.gov/ohrp/sachrp-committee/meetings/index.html>



Executive Summary

Institutional Review Boards (IRBs) are formally designated independent groups charged with the review and ethical oversight of research involving human subjects. The IRB is composed of knowledgeable experts in various fields to provide guidance to researchers to minimize risks and maximize benefits for research participants. Moreover, the IRB is in place to protect the rights and welfare of human subjects in research projects. IRBs inform their decisions based on the principles of the Belmont Report, and established regulations and policies from the Code of Federal Regulations and Food and Drug Administration (FDA) (if applicable).

IRB oversight has been required for human subject research dating back to 1974¹; however, the terms research and human subjects are often misunderstood and inconsistently applied today. Federal guidelines were altered in 2018 to define human subjects to include, “information about [not just physical interventions and interactions with] a living individual”. Artificial intelligence and machine learning (AI and ML) research involving human data challenges the federal human subjects guidelines stemming from the difficulty in defining “about whom” the data is being collected.

This White Paper is intended to be used as a basis for further discussion. We seek feedback on it to inform future iterations of the recommendations it contains. Our aim is to help IRBs build their capacity as regulatory bodies responsible for protecting human subjects in research. We provide recommendations on how AI HSR can be reviewed and adequately overseen within the current regulatory framework until a more thorough regulatory framework can be developed. We also include a decision tree for human subjects and exempt category four (4) (secondary use) determinations, based off the Office of Human Research Protections (OHRP) current guidance².

For IRB professionals the questions arise in two realms: Is the activity “human subjects research” and, if yes, does it meet Exempt criteria?

Scope of This Document

The scope of this paper is limited to Human Subjects Research (HSR) that is regulated by the U.S. department of Health and Human Services (DHHS) under the Office of Human Research Protections (OHRP) and Food and Drug Administration (FDA). The scope of this white paper is intentionally limited to these regulated projects, as OHRP and FDA are currently the only regulatory bodies with established federal definitions and guidance that Institutional Review Board (IRB) use in their ethical and regulatory oversight responsibilities of *human subjects research*.

This document does not go into detail about FDA regulated devices. The FDA definitions have not been harmonized with the Revised Common Rule.

What is Artificial Intelligence and Machine Learning

Artificial Intelligence (AI):

In alignment with technological advances, AI has evolved. To provide clarity, Nilsson's (2010)³ definition on AI will be used throughout the document: AI definition as: “an activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment.”

Machine learning (ML):

ML “is a paradigm that enables systems to automatically improve their performance at a task by observing relevant data” (Stone et al. 2016). ML techniques enable systems to “learn” from the data and perform tasks without being explicitly programmed. Supervised and unsupervised learning are further subcategories of ML. In supervised learning, the system learns from examples of labeled input-output pairs. In comparison, unsupervised learning allows the system to detect patterns without being provided any pre-labeled examples³.

Predictive Model:

Predictive modelling is a statistical technique that uses machine learning and data mining to train a model to predict likely outcomes, usually with retrospective data. A predictive model must be validated and iterated regularly to accommodate shift (changes in the underlying data).

Automated Decision-Making Systems (ADS)

The process of making a decision by automated means without any human involvement. These decisions can be based on factual data, as well as on digitally created profiles or inferred data²⁸. Automated Decision System includes any technology that either assists or replaces the judgement of human decision-maker. These systems draw from fields like statistics, linguistics, and computer science, and use techniques such as rules-based systems, regression, predictive analytics, machine learning, deep learning, and neural nets²⁹.

Note: The difference between an algorithm and a predictive model is that an algorithm is the instructions or code, whereas the model is the output or “learning” that was a result of those instructions.

AI Human Subjects Research (AI HSR):
Research conducted to develop AI tools
involving human subjects.
(Canca & Eto, 2020)

Note: AI does not have to run on big data or use machine learning (ML) techniques.



The first step in an AI HSR review is to break down the term *Artificial Intelligence Human Subjects Research*.

There are two federally established built-in definitions that are vital to understanding what constitutes AI HSR.

The federal definitions that aid in making AI HSR determinations are:



Research



Human Subjects

These terms, and their relationship to AI, are expanded upon below.



In every regulatory system there must be a jurisdictional line informing what is regulated and what is not. The Common Rule only requires IRB oversight of research that is supported or conducted by the Department of Health and Human Services (HHS). Therefore, IRBs want to rule out research projects that do not fall under their jurisdiction.

In the initial research application review, the IRB will determine what level of review, if any, is required. This determination is called many things, but for the purpose of this paper we will call it a Human Subjects Research (HSR) Determination. They do this by first eliminating those that do not meet the federal definition of “research”.

If the project is considered “research” per the federal definition, then they eliminate those that do not constitute “human subjects”, per that federal definition. This is a simple 4-step process:

- STEP 1:** determine if the study is a clinical investigation (per the FDA definition)
- STEP 2:** define *research*
- STEP 3:** identify if the data is *human focused* by understanding the AI’s role in the project
- STEP 4:** define *human subjects*. *Human Subjects* and *Research* have already been defined by the Code of Federal Regulations



HUMAN SUBJECTS

A living individual *about whom* an investigator (whether professional or student) conducting research:

Obtains information or biospecimens *through intervention or interaction* with the individual, and uses, studies, or analyzes the information or biospecimens

or

Obtains, uses, studies, analyzes, or **generates identifiable** information or identifiable biospecimens.

RESEARCH

A systematic investigation including research development, testing, and evaluation designed to develop or *contribute to generalizable knowledge*.


GENERALIZABLE KNOWLEDGE

Information where the intended use of the research findings can be applied to situations or populations beyond the current project.

Federal Definitions

The user then applies those definitions to each specific project activity; each “activity” being how information is collected (e.g., via medical record abstraction, interviews, interventions, etc.). If there is more than one activity, each activity counts.

For example, if a study involves venepuncture and interviews, we wouldn’t refer to the study as “interview only” or “blood draw only”. When reviewers don’t use this approach and make determinations subjectively or out of memory, determinations are made inconsistently.



NOTE:

Under the Revised Common Rule, the definition of "human subjects" was revised to clarify the regulatory intentions. "Data" was revised to read "information or biospecimens". Additionally, the language: "using, studying, or analyzing individuals' information or biospecimens or generating identifiable private information or identifiable biospecimens" was added to clarify what "obtain" meant. Most importantly, under the Revised Common Rule, the revised definition of "human subjects" **includes a provision requiring IRBs to assess whether there are analytic technologies that should be considered by investigators to generate identifiable private information**⁶.

Because AI/ML research is data-driven and may or may not have human intervention or interaction components, we have modified the OHRP decision tree to be relevant to AI research projects. Further, with the Revised Common Rule’s updates to the definition of “human subjects” to clarify “data” was intended to mean “information”, for the purpose of this paper, we use “data” and “information” interchangeably.

The AI-modified OHRP HSR Decision Tree (Addendum A) breaks down key terms of the federal definitions of *research* and *human subjects*, and introduces a “human focused data” aspect to further address the question, “about whom” is the data being collected (a key term in the federal definition).

Using the AI-modified decision tree, in concert with the **AI HSR IRB Reviewer Checklist (Addendum C)**, is especially useful when making HSR determinations for AI projects because it requires the reviewer to fully understand the role of AI in each project. This is key to making an accurate determination.

Investigational Technology and Clinical Investigations

Unlike typical data-only research projects, AI/ML is oftentimes introduced through software and devices. Therefore, before we can begin the HSR determination, we first need to determine if FDA regulations apply (i.e., 21 CFR §50, 56, 812, or 820), and whether the project is a *clinical investigation*. Even if the project is ***not designed to contribute to generalizable knowledge, clinical investigations*** require IRB review under 21 CFR 50 and 56.

Further, if the technology is investigational, it may also be bound to 21 CFR §812 and §820.

If the project is not a clinical investigation, we continue through the AI HSR Decision Tree by applying the federal definitions of “research” and “human subjects”. This requires two definitions and decisions:

(1) Is the activity “Research”? *and* (2) Does the activity involve “Human Subjects”?

Is the Activity "Research"?

This is often the easier of the definitions. Is the activity a “systematic investigation” and is the intent to develop “generalizable knowledge”?

Not all AI projects fit the federal definition of *research*. “Research development, testing, and evaluation” is included in the federal definition. As such, algorithm development, validation, and evaluation often fall within the federal definition of *research* ***depending on the AI’s role in the study***. When these projects are designed to contribute to generalizable knowledge, the study is considered research per the federal definition.

Generalizable Knowledge

“ Research conducted with the intention of drawing conclusions that have some general applicability; uses a commonly accepted scientific method; or research findings intended to be used to apply to situations and populations beyond the current project [7] ”

Using the above established definitions in this step-wise process, in conjunction with the AI HSR IRB Reviewer Checklist (which will help identify the technology’s role in the project), IRBs can determine if the project is *research* (HSR) per the federal definitions.

Defining "Human Subjects" in AI "Research"

Machine learning development (algorithm or model development and validation) is generally considered a "systematic investigation, carried out according to a plan, designed to develop or contribute to generalizable knowledge" when that AI/ML is used to develop software or tools that will be used to make decisions that affect humans.

We conclude that AI/ML in that context is likely "research" per federal definition. The larger question remains: does the project constitute "human subject research"? In other words, at what point does AI/ML "involve human subjects" (and therefore require IRB oversight)?

AI "human subjects" determinations depend on (1) the *identifiability or re-identifiability* of the information and (2) the **purpose/role of AI** in the project; the latter addresses the "about whom" of the federal definition of "human subjects". If the information is (a) identifiable or generates identifiable information, and (b) "human-focused", it is *human subjects*, and therefore satisfies the regulatory definition of "*human subject research*".

NOTE: AI/ML has been known to make re-identification easier, especially for patients in healthcare settings¹. Further, deidentified data may become reidentifiable through data triangulation from other datasets². IRBs should consider these factors as they review data sources and future use proposals.

“ Many assume that 'anonymized data cannot be used to reidentify the subject of the data. Unfortunately, as data sets proliferate, the ability to combine multiple datasets may defeat the deidentification strategy

Price, 2019

”

Assessing "About Whom" the Data is Collected

Once the data has been determined identifiable or generating identifiable information, we move to the next step, which is to identify "about whom" the information is collected. To do this, we first define the role AI has in the project. **Only by understanding the role of the AI can we determine if the AI is *Human-Focused* or *Not Human-Focused*.** This can be addressed by asking if the technology is developed to model human thought, understand or treat a human condition, empower machines to act on their own, or to perform functions similar to human intelligence, such as the ability to perceive, learn, reason, and act? If so, the AI is Human-Focused.

If the research objectives are collecting data that might involve human data (even if it is identifiable) but not intended to model human behavior or understand it, this would be considered ***Not Human-Focused***. Similarly, if the AI is solely being used as a tool, for example, when AI is used as a form of data management, text record mining, or record abstraction, even if it is collecting identifiable human data, we would also consider this ***Not Human-Focused***.

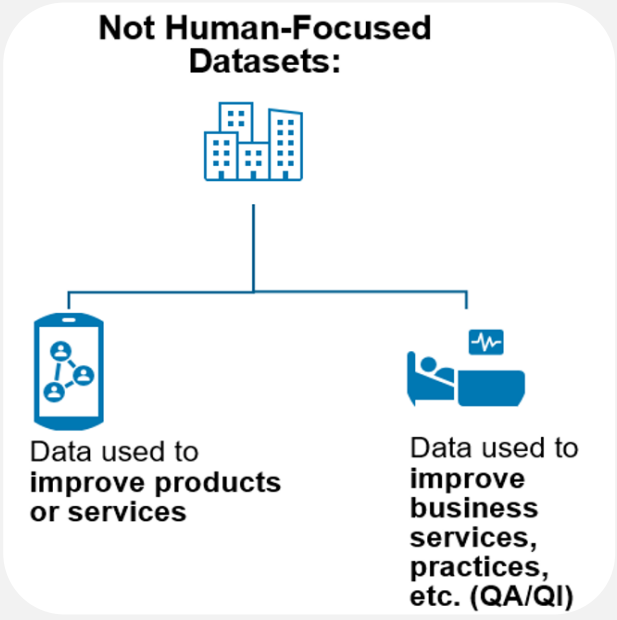
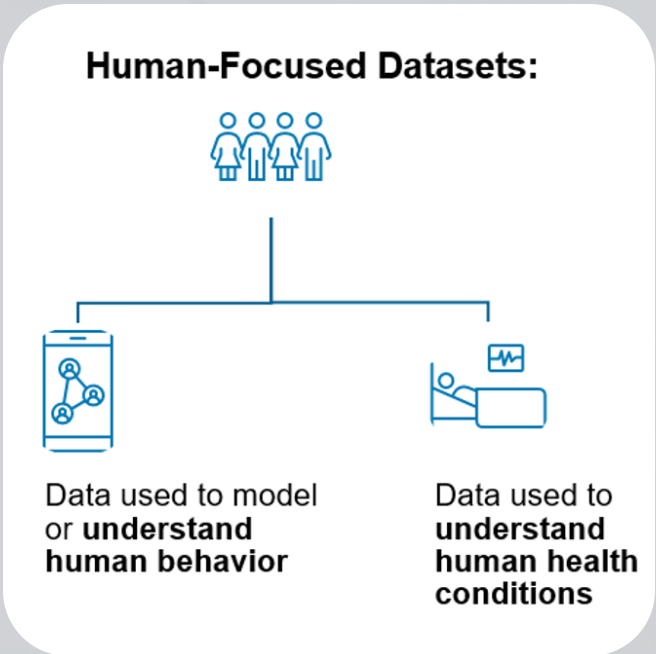
Assessing the Purpose / Role of AI

Locating the AI's role in the research application may not always be mentioned up front. Some reviewers may misinterpret the technology's role to be basic software development, an analytical or statistical technique, or form of data management, as a result of how the technology is introduced in the application. It is helpful to keep an eye out for common AI terminology like "machine learning", "natural language", "neural networks", "deep learning", or "predictive modelling".

To make an AI HSR determination, we must first define the role of AI in that project because not all AI projects are *Human Subjects Research* projects. Similarly, not all AI projects rely on big data or involve machine learning.

When Data is Human-Focused:

AI HSR projects are different from the typical data analysis project in that they are human-centric (Human-Focused); the technology is developed to understand or treat a human condition, model human thought, empower machines to act on their own or perform functions similar to human intelligence, such as the ability to perceive, learn, reason, and act¹⁰. These projects are heavily dependent on human-focused data and continuous iteration. Without relevant human data that accurately reflects the human population in which the technology is developed to serve, the technology will fail. Therefore, identifying the difference between *Human-Focused* and *Not Human-Focused* data is key to making AI Human Subjects Research determinations.



When Data is Not Human-Focused

When the focus or role of the data is solely to improve a platform, product, or service—then the project is likely *Not Human Subjects Research (NHSR)*. Similarly, when the AI research is not meant to help us understand humans, human behaviour, or human conditions (and is not intended to model human behaviour or treat human conditions), the project would not generally be considered AI HSR.

Note: *Not Human-Focused* projects generally focus more on products, institutional output, and processes. Similar to the way the Common Rule Exempt determinations work, if a project has both Human-Focused and Not Human-Focused aspects, it is treated as a Human-Focused (AI HSR) project.

In summary, AI HSR determinations largely depend on (1) identifiability of the data, and (2) the AI’s purpose/role in the project. Assuming (a) the researcher uses/analyses identifiable data, **or generates identifiable data, and (b) the data is Human-Focused**, the project is likely *AI HSR*. If the project is HSR, it may or may not be eligible for one of several exempt determinations. We have provided an *Exempt Determination Decision Tree in Addendum B*.

Some recommendation engines (e.g., recommending books you might like) constitute HSR, under the current regulatory framework.

The **recommendation engine** in Example 1, to the right, is *not* HSR. Even though the company is collecting identifiable information and interacting with the subjects via “robots”, it does not constitute HSR because:

- a. neither the company or the project is regulated by DHHS or the FDA, and
- b. the project is not intended to be generalized outside of this specific company's specific product.

Similarly, some **vision-based AI** projects are HSR. However, in the case of Example 2, (a) the data is not “about a human”, and (b) the researchers are not trying to understand or model human behavior.

1. Recommendation Engine

Subscription-based streaming entertainment business uses AI/ML. Recommendation engines are trained to recommend content (movies, etc.) to individual users.

The company collects and uses/analyzes human data and interacts with them (via the technology).

2. Vision-based AI Grading

A teacher takes a picture of a student's test/homework.

A trained machine will grade it.

The machine is capturing and analyzing data.

IMPORTANT NOTE:

While these examples are likely *not AI HSR*, there may be other applicable regulations, policies, ethics, and privacy issues to consider.

Current Regulatory Framework and IRB Function

Most established IRBs embed ethical principles into their IRB application in order to ensure the necessary ethical principles are incorporated into their review processes. For example, the ethical principle of Respect for Persons requires that all subjects (unless waived by the IRB) must engage in a Fully Informed Consent process and agree to participate in research.

The IRB application typically has a section that spells out the proposed process (or requires a strong justification for a waiver of informed consent). IRB questions are generally iterations of ethical principles addressed in the Nuremberg Code¹¹, Declaration of Helsinki¹², Belmont Report¹³, and the 21st Century Cures Act¹⁴.

Details on how IRBs incorporate these ethical principles into their IRB applications is beyond the scope of this paper but can be found in a recent CITI program training module called Artificial Intelligence (AI) and Ethics in Human Subjects Research.

As we have seen from the most recently revised 2018 Common Rule¹⁵, regulations and principles governing human subjects research will continue to evolve in sophistication and complexity, especially as we evolve technologically¹⁶.

Core ethical values in the Belmont Report¹³ (Respect for Persons, Beneficence, and Justice) are the current backbone of IRB review and guide the design and review of ethical research, and it is these principles that serve as the skeleton of the *AI HSR IRB Reviewer Checklist*.

Ethical principles embedded in the regulations guide IRBs and investigators in the design and responsible conduct of research. The combination of these principles with other federal regulations (Common Rule and FDA), as well as institutional policies, create a structure that accommodate our desire as a society to expand our scientific knowledge, while at the same time enable us to, and ensure that, we protect and show compassion for the people who volunteer for our research studies.

Some states have implemented laws against Automated Decision-Making Systems³⁰. IRBs need to be mindful of state laws in this regard.

Unfortunately, regulations alone do not guarantee that a study will be ethical. Additionally, unique, and rapidly evolving AI calls for revisiting our current ethical and regulatory framework, as principles and guidelines do not enforce but only recommend.

Outstanding Challenges



There are three outstanding challenges in the current IRB regulatory framework : **1) Risk Assessments**, **2) Expert Reviewers**, and **3) Siloed Thinking in the Review of Technology**.

Challenge #1: Risk Assessments:

It is often assumed that IRB exemptions were written for situations that are very low risk. Although many institutions have encouraged consideration of risk in the exemption consideration, it is not a requirement.

If a study's only involvement of human subjects (including just their data) fits in one or more of the established established exempt categories, the study is considered exempt from the requirements of the Common Rule¹⁷.

The regulatory assumption is that so long as adequate security measures are in place, secondary data-only projects involve minimal harm to the subjects, compared to the potential benefits research results may provide to individuals or society¹⁸. This assumption is highly problematic. While we may understand data, we are far from understanding the true benefits and risks of this rapidly evolving novel technology and Big Data that usually accompanies it. Furthermore, current regulatory guidelines for Expedited in Full-Board studies suggest that:

"In evaluating risks and benefits, the IRB should consider only those risks and benefits that may result from the research... The IRB should not consider possible long-range effects of applying knowledge gained in the research" [19]

Many reviewers misunderstand this to mean IRBs are prohibited from considering long-term /downstream consequences. These subjective interpretations and unknown risks are one reason why an official AI research policy has been delayed. Nevertheless, while we "wait and see"²⁰ how this technology evolves and what kind of policy should be developed, because the risk is so high, many institutions have taken it upon themselves to institute their own policies in their review of such novel technology. We hope that this document aids those institutions in conducting their review of AI HSR.

Challenge #2: Expert Reviewers:

The Code of Federal Regulations § 46.107²¹ describes IRB membership requirements:

"Each IRB shall have at least five members, with varying backgrounds to promote complete and adequate review of research activities commonly conducted by the institution. The IRB shall be sufficiently qualified through the experience and expertise of its members (professional competence)... The IRB shall therefore include persons knowledgeable in these areas."

If an institution has a considerable amount of AI HSR, compliance with this requirement would mean their IRB includes an expert in this field (for example, AI and privacy experts)²². For larger institutions with vast resources, this may not be a big ask, but can be a tremendous challenge, if not impossible, for smaller institutions with less resources.

Challenge #3: Siloed Thinking in the Review of Technology:

The field of research compliance is unfortunately siloed. AI HSR particularly, directly affects multiple departments. While Privacy, IT, Contracting, Biosafety, Radiation Safety, and IRB operations are all interconnected, they tend to operate within their own bubbles.

We must get out of siloed thinking. For example, many IRBs are uncertain where their review begins and where another takes over. IRBs may believe that the responsibility of reviewing technology is solely upon their IT department, leaving significant gaps in regulatory oversight. IT support services, likewise, may believe that the IRB is confirming that all regulatory requirements (such as Part 11, etc.) are met or assume someone else has taken care of that piece. As a result, the review is neglected, and regulatory bodies are unable to adequately mitigate privacy risks (principle of justice) and inform participants of what they are getting into (principle of respect for persons).

Future Recommendations

Recommendation**#1: Focus on the Data**

AI HSR ethical review and adequate research compliance oversight is dependent on clearly defining the role of AI in each research project and focusing on the data. Because AI/ML largely depends on a model, it is tempting to focus only on the complexities of the technology. Successful AI HSR is entirely dependent on the data.

From an IRB perspective, data is more important than the model. The data in AI HSR is the primary, core, and permanent asset of the project²³. The IRBs focus should be weighted more heavily on the data used to train the model, as opposed to the algorithm or model itself (**Note:** if the model is used in a clinical investigation or FDA-regulated, standard regulatory requirements should still be addressed).

This approach works in the IRBs favor as IRBs are more well suited to address data concerns than technology (though, as mentioned earlier, the technology may require additional ancillary reviews (e.g., FDA), or risk assessment by the IT and Contracts department).

Using an AI HSR IRB Reviewer checklist would help to ensure all important aspects are addressed during IRB review. Additionally, we recommend that institutions require adequate board member training, and ensuring they meet their regulatory obligations of adding an AI, Privacy, and data experts to the review board.

**Recommendation #2: Keep it Local**

If Possible, keep oversight local. If not possible, plan ahead. Ancillary AI Ethics Committees and commercial IRBs are innovative and helpful for institutions that can afford them. Not all institutions have the resources to develop or employ them.

While our current regulatory framework is in need of updates to accommodate novel technology, we have adequate tools and protections at our disposal in overseeing AI HSR under the current regulatory framework. Therefore, keeping the review with the local IRB is an especially beneficial approach to those institutions/projects with limited resources. While it may be ideal for institutions with no IRB to outsource their reviews, for institutions with a home IRB, below are some possible roadblocks when considering outsourcing AI HSR oversight.

(1) Cost:

The study team may need to plan for additional funding if the review is not free.

Small local IRBs may not have the necessary competence and will need to find – and pay for – a consultant. Because most in-house reviews are either free or more affordable, long term projects that foresee additional reviews (e.g., modifications or annual renewals) could end up saving the study team time and money by investing in internal resources.

Further, outsourced IRBs would still require a reliance process, which puts a partial review requirement on the local IRB anyway.

(2) Duplication of Effort:

An AI Research Review Committee (AIRC)²⁴ typically acts as an ancillary review to IRB review. Furthermore, reliance agreements also require partial local IRB oversight for ancillary review and local considerations.

Many of the issues reviewed by an outsourced IRB would parallel local IRB review, especially for covered entities that maintain Privacy oversight, which may lead to duplication of effort, time, and money.

(3) No Regulatory Teeth:

Currently, only some projects require adherence to the federal regulations and those are only enforceable by an IRB (and the FDA). The rest are voluntary.

If an AIRC²⁴ (or any AI ancillary review) has recommended changes to the protocol, the committee may lack regulatory “teeth” unless adherence is tied to funding^{25, 26}. This means that the researchers will not be required or inclined to comply with the Committee’s “suggestions”. Additionally, due to siloed thinking, ancillary reviews add yet another department for IRBs to communicate with.

Consequently, AIRC suggestions may or may not make their way to the IRB unless there are established procedures that keeps the two committees “talking to each other”.

This challenge can be mitigated by establishing a required line of communication in the IRB submission process (for example, a congruence check on what was reviewed by the Committee and what shall be approved by the IRB).

(4) Sustainability:

If an institution invests in establishing an internal AI HSR Committee, as opposed to building onto an established IRB process, this ancillary committee needs a sustainable administrative process.

Establishing this will take a significant amount of time, money, and institutional resources. The committee will also need to consider issues such as:

- Committee membership
- Frequency of Committee meetings and formality
- Record retention
- Classified or confidential projects
- Confirming and mitigating conflicts of interest

Recommendation #3: Embrace it.

AI is not going to go away. It will continue to evolve at a rapid pace, and at some point, we will need to get on board. Get ahead of the game and:

- Incorporate AI/ML into IRB **Applications**
- **Embed relevant responsible/tech-ethical** issues into these forms as well.
- Require institution-wide education for IRBs, IT support, and the research personnel on how the technology works, and how to apply the regulatory protections and ethical considerations for this novel technology.

There are many valuable training modules available online²⁷.

Limitations of These Recommendations

These recommendations are utilizing the current Revised Common Rule federal definitions of "Research" and "Human Subjects" for a limited type of AI HSR. Defining AI HSR largely depends on the definition of "generalizable knowledge" which has not been federally defined (specifically in regard to algorithm development and model validation), and distinguishing between "identifiable information" and "identifiable private information". These ambiguities are approached inconsistently. As discussed, the recommendations are limited to how current AI HSR can fit within the current regulatory framework, though it is an imperfect fit. If we are to see any change, we must make significant policy changes. These recommendations are for institutions that would like to incorporate ethical and responsible AI oversight into their regulatory obligations, without stepping outside of the current regulatory framework.

Conclusions

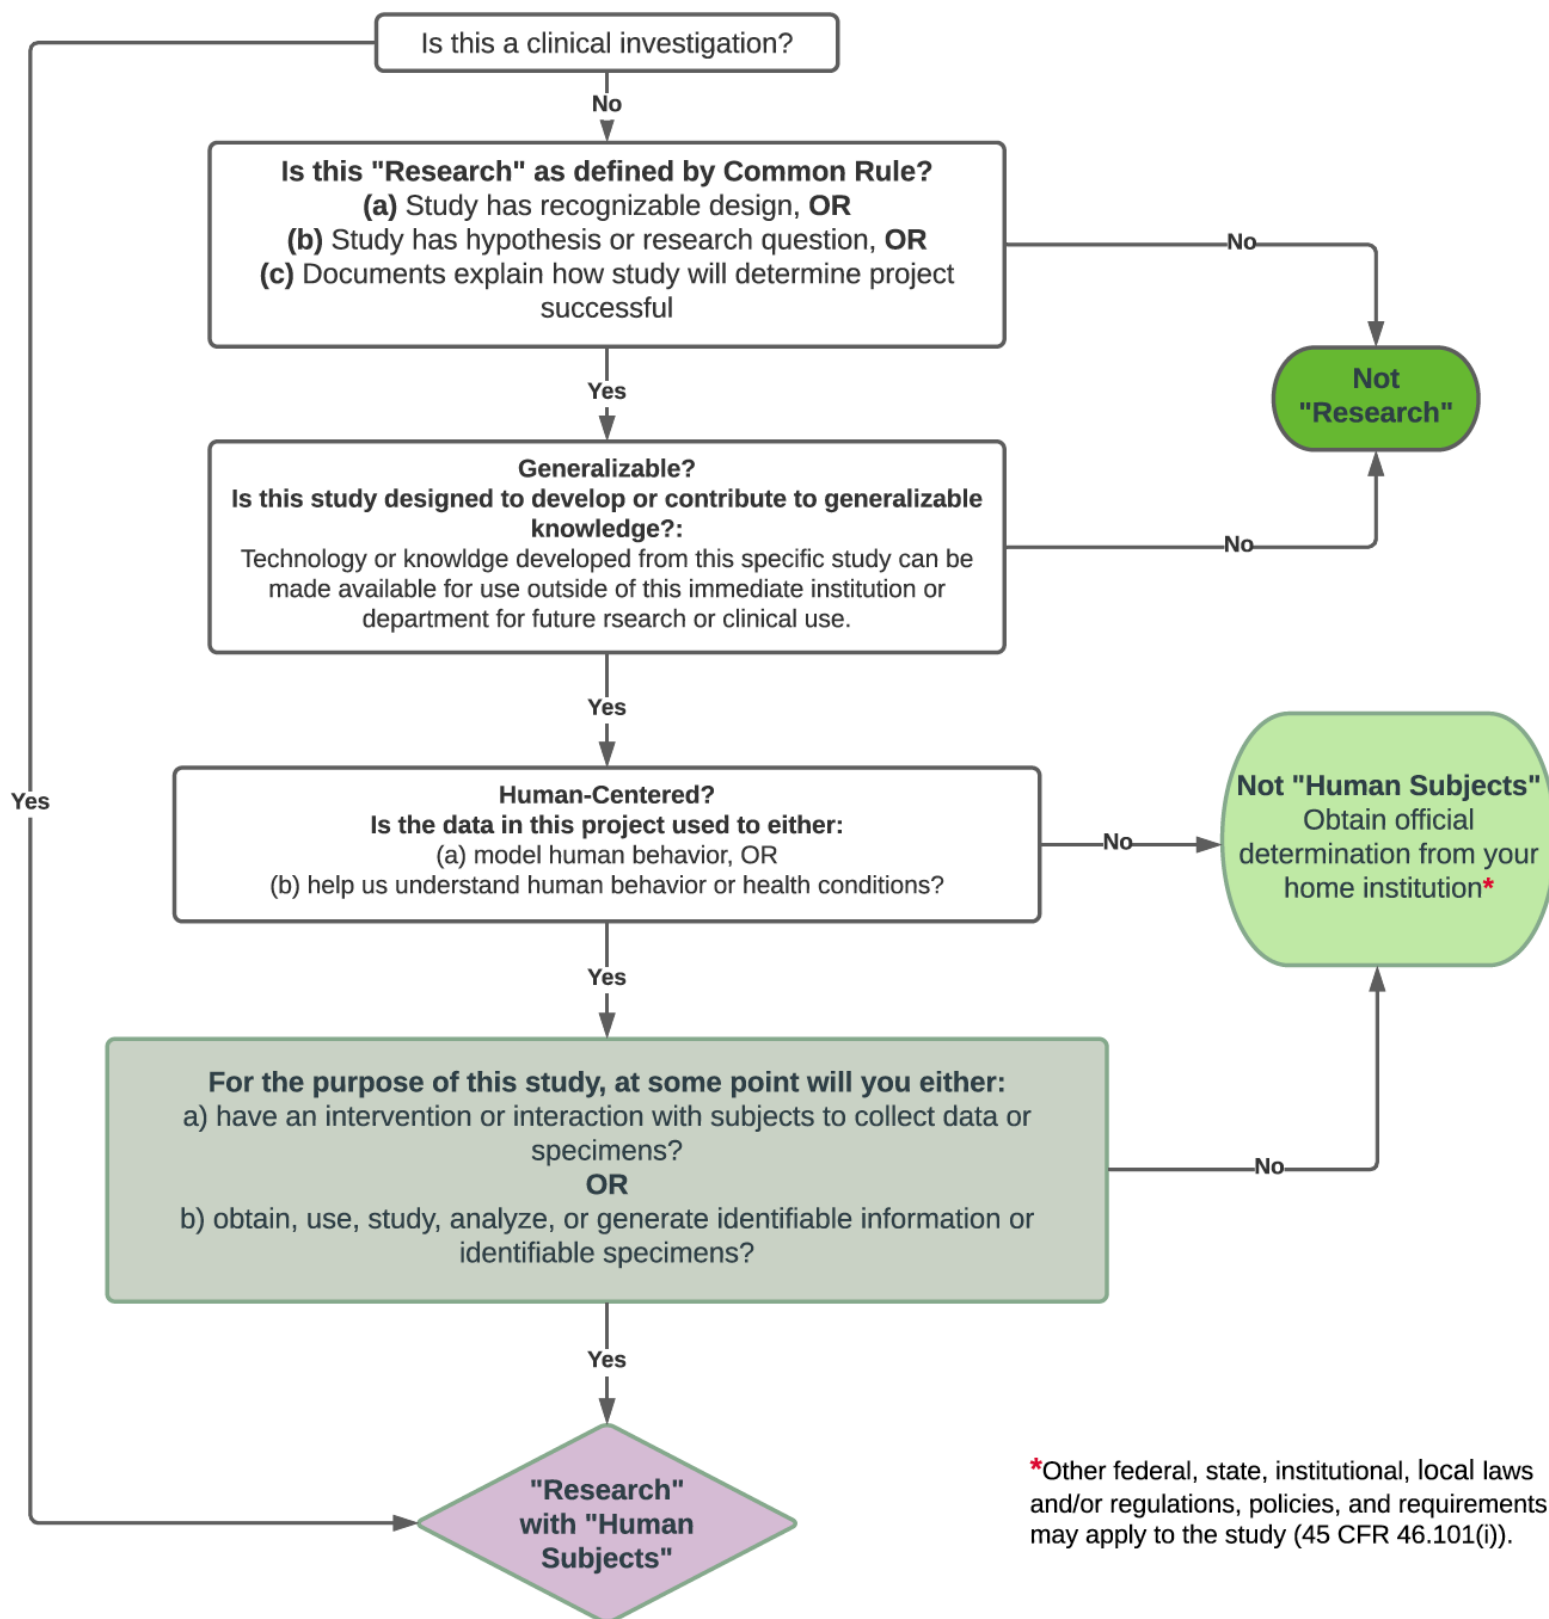
IRBs have been treading lightly when it comes to the oversight of AI HSR. This may be due to insufficient understanding of when AI research involves *human subjects*. It may also be in fear of committing scope creep. This inaction puts both research participants, society, and the institution at risk. It is going to be some time before the U.S. adopts Human Subjects Research Protection Guidance on AI HSR. Nevertheless, while we await a regulatory framework, the disparities and harms that this unregulated technology has introduced is growing exponentially. Rather than "wait and see", we need to put our focus on how the current framework applies now. It is imperative that we familiarize ourselves with AI in a Human Subjects Research context.

In response to a growing need of AI ethical and regulatory oversight, there are a number of commercial and not-for-profit AI Ethics Committees making their services available. Some institutions have tried to fit AI ethics review into an ancillary (external to IRB) review process. These ancillary AI ethics committees either take on the look and feel of a scientific review committee or treat the process like an IBC or SCRO committee.

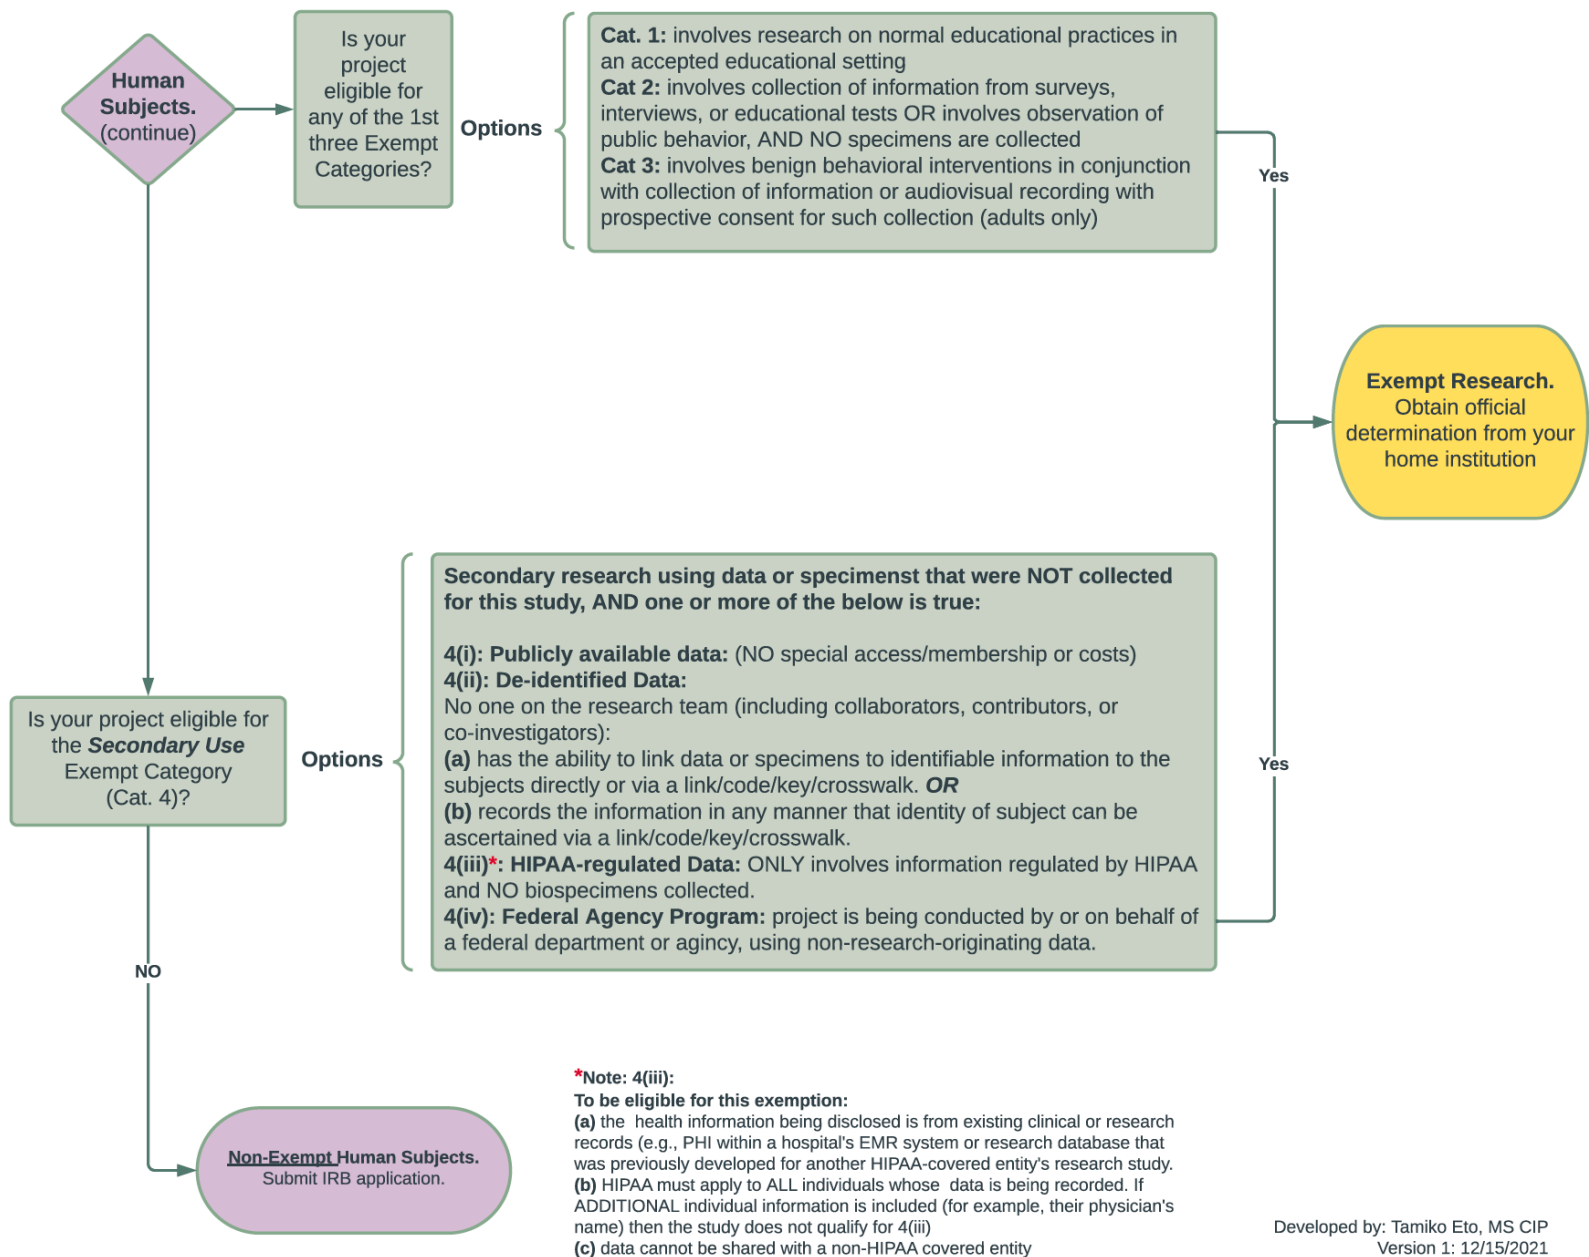
Admittedly, the current regulatory framework has limitations, regardless of if it is AI HSR or any other type of HSR. Simply moving AI HSR oversight to an ancillary committee is not an efficient solution for researchers who will still have to navigate their way through the IRB for these same projects in addition to extra bureaucratic hoops.

Ancillary AI HSR committees may delay the process to approval and disincentivize compliance. Rather than build a new AI HSR IRB or ancillary review committee, we need to provide and require the AI HSR education/training of IRB administration and remind the IRB of its duty to ensure relevant experts sit on the Board when reviewing specific research. I argue that IRBs can fit AI HSR oversight within their current IRB regulatory framework in many significant and meaningful ways using the tools outlined in this White Paper, and without committing scope creep.

**Artificial Intelligence Human Subjects Research
(AI HSR) Determination Decision Tree**
(to be used for AI/ML HSR Determinations)



*Other federal, state, institutional, local laws and/or regulations, policies, and requirements may apply to the study (45 CFR 46.101(i)).



Developed by: Tamiko Eto, MS CIP
Version 1: 12/15/2021

1. Protection of Human Subjects, 39 Fed. Reg. 105, 18914-18920 (1974) (to be codified at 45 C.F.R. [[section]] 46).
2. United States Department of Health and Human Services (DHHS). "Human Subject Regulations Decision Charts". <https://www.hhs.gov/ohrp/regulations-and-policy/decision-charts/index.html>. Accessed 30 December 2021.
3. Nilsson, Nils J. 2010. *The Quest for Artificial Intelligence: A History of Ideas and Achievements*. Cambridge, UK: Cambridge University Press.
4. Stone, Peter, Rodney Brooks, Erik Brynjolfsson, Ryan Calo, Oren Etzioni, Greg Hager, Julia Hirschberg, Shivaram Kalyanakrishnan, Ece Kamar, Sarit Kraus, Kevin Leyton-Brown, David Parkes, William Press, AnnaLee Saxenian, Julie Shah, Milind Tambe, and Astro Teller. 2016. One Hundred Year Study on Artificial Intelligence: Report of the 2015-2016 Study Panel, Stanford University, Stanford, CA. Accessed June 6, 2022.
5. Canca, Cansu; Eto, Tamiko; and Leong, Brenda. 2020. "Artificial Intelligence (AI) and Ethics in Human Subjects Research." *CITI Program* <http://www.citiprogram.org>. Accessed 30 December 2021.
6. Office of Human Research Protections (OHRP) "Revised Common Rule Q&As" <https://www.hhs.gov/ohrp/education-and-outreach/revised-common-rule/revised-common-rule-q-and-a/index.html>
7. United States Department of Health and Human Services (DHHS) Office of Research Integrity (ORI). "ORI Introduction to RCR: Chapter 3. The Protection of Human Subjects" <https://ori.hhs.gov/content/chapter-3-The-Protection-of-Human-Subjects-Definitions>. Accessed 30 December 2021
8. Ross, Patrick. "AI Use in Healthcare: Overview of Initial Steps to Develop AI Regulations/Guidances: Security and Safety Issues to Consider." *Health 2019* (2019); Sara Gerke, Timo Minssen, Glenn Cohen, Chapter 12 - Ethical and legal challenges of artificial intelligence-driven healthcare, *Artificial Intelligence in Healthcare*, Academic Press, 2020, Pages 295-336
9. Price, W.N., Cohen, I.G. Privacy in the age of medical big data. *Nat Med* 25, 37–43 (2019)
10. Ribeiro, Jair. A.I. in 2020: A Year writing about Artificial Intelligence (An Year of AI). Self-published, Amazon.com, 2021.
11. National Institutes of Health Office of NIH History and Stetten Museum. <https://history.nih.gov/display/history/Nuremberg+Code>. Accessed 30 December 2021.
12. World Medical Association. "WMA Declaration of Helsinki – Ethical Principles for Medical Research Involving Human Subjects". (Jul 2018). <https://www.wma.net/policies-post/wma-declaration-of-helsinki-ethical-principles-for-medical-research-involving-human-subjects/>. Accessed 30 December 2021.
13. United States Department of Health and Human Services (DHHS). "The Belmont Report: Ethical Principles and Guidelines for the Protections of Human Subjects of Research". <https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/index.html>. Accessed 30 December 2021.
14. US Food and Drug Administration (FDA). "21st Century Cures Act". <https://www.fda.gov/regulatory-information/selected-amendments-fdc-act/21st-century-cures-act>. Accessed 30 December 2021.
15. United States Department of Health and Human Services (DHHS). "Federal Policy for the Protection of Human Subjects ('Common Rule')". <https://www.hhs.gov/ohrp/regulations-and-policy/regulations/common-rule/index.html>. Accessed 30 December 2021.
16. Eto, Tamiko. "Ethics and Clinical Research". EtoConsulting (Sept. 2018). <https://etohconsulting.com/2018/09/15/ethics-and-clinical-research/>. Accessed 30 December 2021.
17. United States Department of Health and Human Services (DHHS). "<https://www.hhs.gov/ohrp/regulations-and-policy/regulations/45-cfr-46/common-rule-subpart-a-46104/index.html>
18. United States Department of Health and Human Services (DHHS). "<https://www.hhs.gov/ohrp/sachrp-committee/recommendations/2015-april-24-attachment-a/index.html>
19. United States Department of Health and Human Services (DHHS). "Approval of Research with Conditions: OHRP Guidance (2010)" <https://www.hhs.gov/ohrp/regulations-and-policy/guidance/guidance-on-irb-approval-of-research-with-conditions-2010/index.html>. Accessed 30 December 2021.
20. Tang, Hazel. "Reasons behind the "wait and see" approach in AI regulations". AI in Medicine Clinician Series (Oct. 2020). <https://ai-med.io/clinicians/reasons-behind-the-wait-and-see-approach-in-ai-regulations/>. Accessed 30 December 2021.
21. United States Department of Health and Human Services (DHHS). "Basic HHS Policy for Protection of Human Research Subjects". <https://www.ecfr.gov/on/2018-07-19/title-45/subtitle-A/subchapter-A/part-46>. Accessed 30 December 2021.

22. Eto, Tamiko; Canca, Cansu; and Leong, Brenda. 2020. "Artificial Intelligence, Regulatory Issues, and Other Challenges." *CITI Program*
23. McComb, Dave. "The Data-Centric Revolution: Data-Centric vs. Data-Driven". The Data Administration Newsletter (Oct. 2016) <https://tdan.com/the-data-centric-revolution-data-centric-vs-data-driven/20288>. Accessed 30 December 2021.
24. Jordan, Sara. "Designing an Artificial Intelligence Research Review Committee". Future Privacy Forum. 2019.
25. Michael S. Bernstein, Margaret Levi, David Magnus, Betsy Rajala, Debra Satz, Charla Waeiss. "ESR: Ethics and Society Review of Artificial Intelligence Research". Cornell University 2021. <https://arxiv.org/abs/2106.11521>. Accessed 30 December 2021.
26. Jensen, Beth. "A New Approach To Mitigating AI's Negative Impact". Human Centered Artificial Intelligence Center at Stanford University (Jun 2021). <https://hai.stanford.edu/news/new-approach-mitigating-ais-negative-impact>. Accessed 30 December 2021.
27. The Collaborative Institutional Training Initiative (CITI Program). <https://about.citiprogram.org>. Accessed 30 December 2021.
28. Information Commissioner's Office (ICO). Guide to General Data Protection Regulation. June 2018. <https://ico.org.uk/> Accessed 3 July 2022.
29. Government of Canada. Directive on Automated Decision Making. April 2021. <https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592> Accessed 3 July 2022.
30. Gesser, A, et. al. New Automated Decision-Making Laws: Four Tips for Compliance. Program on Compliance and Enforcement. New York University School of Law. June 29 2022. https://wp.nyu.edu/compliance_enforcement/ Accessed 3 July 2022.

Author Biography

Tamiko Eto - Ms. Eto has over 17 years of experience in the field of human subjects research protections and manages IRB review and research compliance, including the facilitation of technology risk assessment and data sharing contracts, at Kaiser Permanente's Division of Research. The research portfolio there comprises over 7 million members with a large profile of AI-related research projects including FDA-regulated Software as a Medical Device (SaMD). Concurrently, Ms. Eto also serves on the AI Ethics Advisory Board for the Institution of Experiential AI (EAI) at Northeastern University. Prior to her work at the Division of Research, Tamiko served as Acting Director at Stanford Research Institute's (SRI) Office of Research Integrity and Chair of SRI IRB, where she performed scientific reviews, policy interpretations and the development of AI-related projects. She has now leveraged her experience to implement regulatory policies to health care research projects that delve into AI research. She works closely with AI researchers and institutional /regulatory bodies in addressing ethical and regulatory challenges related to AI. To facilitate researchers and IRB professionals across the US she has developed tools and checklists for IRBs to use in their review of AI research. Developing these tools, she also actively collaborates on research to be at the forefront of developing an ethical and regulatory framework for research involving human subjects. She is now currently pursuing her second Master's Degree at Stanford University in Ethics and Policy.

Edited By:

Erica Heath. Ms. Heath founded Independent Review Consulting (IRC) in 1984 and was president until it became a part of Ethical and Independent Review Services, LLC. where she is now regulatory director. She has spent her career working in the field of protection of human research subjects in both clinical and non-clinical research. Within this field she is known for her expertise in IRB administration and review of medical devices.

Erica earned her BA from San Jose State University and her MBA from Golden Gate University. She served on the Accreditation Council for the Association for Accreditation of Human Research Protection Programs (AAHRPP) for 7 years and is currently on the Council of Certified IRB Professionals.

About TechInHSR

Research is about making the world a better place, making our lives easier, and helping us understand the world we live in. Sometimes regulations can seem intimidating, difficult to understand and adhere to, leading to delays in initiating research. While some regulations and policies may seem bureaucratic, human subject research protections are established to protect not only the people that volunteer, but the institutions, the researchers, and even society. TechInHSR (formerly EtoConsulting) has attempted to translate those regulations to simplify compliance and provide independent feedback on ethical and regulatory compliance when doing research. Our mission is to stay on top of, and make sense of human subjects research regulations, policies, and laws at all levels.

Copyright

This White Paper contains a variety of copyright material. Some of this is the intellectual property of the author. Artificial Intelligence Human Subjects Research IRB Reviewer Checklist (with AI HSR and Exempt Decision Tree) © 2021 by Tamiko Eto is licensed under CC BY-NC-SA 4.0. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-sa/4.0/>

Disclaimer

Whilst the authors and TechInHSR has attempted to ensure the information in this White Paper is as accurate as possible, the information is for personal and educational use only, and is provided in good faith without any express or implied warranty. There is no guarantee given to the accuracy or currency of information contained in this White Paper. Neither the author, nor TechInHSR, accept responsibility for any loss or damage occasioned by use of the information contained in this White Paper. The content is based on personal interpretations of the regulations. This does not, in any way, serve to be a legal resource or alternative to IRB or Ethics Committee oversight. This White Paper is not designed to provide legal advice or legal guidance. Please consult with your organization's attorneys if you have questions or concerns about the relevant laws and regulations discussed in this White Paper. The opinions in this paper are entirely those of the author's and do not reflect the institutions of which the author is employed.