Summer 8-19-2021

The Challenges of Applying Computational Legal Analysis to mHealth Security and Privacy Regulations

Brian Tung

Washington University in St. Louis

School of Engineering and Applied Science

Department of Computer Science and Engineering

Thesis Examination Committee:
Ning Zhang, Chair
Stephen Cole
Jonathan Shidal

The Challenges of Applying Computational Legal Analysis to mHealth Security and Privacy Regulations

by

Brian Beyuan Tung

A thesis presented to the McKelvey School of Engineering
of Washington University in partial fulfillment
of the requirements for the degree of

Master of Science

August 2021
Saint Louis, Missouri

Contents

List of Tables

List of Figures

Acknowledgments

This thesis would not be possible without the support of many others. First, I would like to thank my research advisor, Professor Ning Zhang. You have been a true friend and have gone above and beyond the call of duty as an advisor. Your determination to build a better world inspires me to do the same.

I would also like to thank my lab mate Zhiyuan Yu for his conversations, research assistance, and support. I thank Professor Steven Cole for many years of sincere instruction in computer systems, security, and privacy. You are a phenomenal teacher and also go above and beyond what is required of you. I would also like to thank Professor Hila Ben Abraham for her thoughtful instruction in cybersecurity. I thank Professor Jonathan Shidal for his dedication to students and for teaching me the fundamentals of systems programming.

In the social sciences, I thank my undergraduate research advisor Professor John Kuk. Thank you for exploring the social sciences with me and for indulging my intellectual curiosity. Your support and friendship have made me the scholar I am today. I am also thankful for Professor Jake Rosenfeld for his writing instruction, and for Professor David Cunningham for his fun introduction to social theory. I would also like to thank Professor Grizelda McClelland and Dean Mark Smith for their sage academic advice.

Finally, I would like to thank my friends and family. You have made me who I am today, and I am eternally grateful for your love and support.

Brian Beyuan Tung

Washington University in Saint Louis
August 2021

Dedicated to my parents.

ABSTRACT OF THE THESIS

The Challenges of Applying Computational Legal Analysis
to mHealth Security and Privacy Regulations

by

Brian Beyuan Tung

Master of Science in Computer Science

Washington University in St. Louis, August 2021

Research Advisor: Professor Ning Zhang

As our world has grown in complexity, so have our laws. By one measure, the United States Code has grown over 30x as long since 1935, and the 186,000-page Code of Federal Regulations has grown almost 10x in length since 1938. Our growing legal system is too complicated; it's impossible for people to know all the laws that apply to them. However, people are still subject to the law, even if they are unfamiliar with it. Therein lies the need for computational legal analysis. Tools of computation (e.g., data visualization, algorithms, and artificial intelligence) have the potential to transform civic education, legal procedure, and society at large. In this thesis, I try to apply computation to a small part of the law, federal regulations and recommendations about mHealth security and privacy, and explain why I believe legal computation can lead to more equitable health outcomes. I conclude that the open-ended nature of the law, the FDA's reliance upon self-certification, and complicated legal dependencies make it difficult to apply computation to mHealth security and privacy regulation. However, the use of non-legal, authoritative guidance documents that explain regulatory intent may aid in the development of computational legal auditors.

Chapter 1

Introduction

It is difficult to overstate the importance of the law in our society. The law affects every aspect of our lives, and as much as technologists like to deny it, our current technological environment is a direct product of our legal system. Our laws, and the government described therein, encode the rules of society we all must follow. Our current technologies depend on their legal protections for their existence. Patents, trade secrets, copyright, national security, civil rights, international commerce, data sharing, and free speech are just a few legal areas that are of interest to technologists. Agencies that enforce how the law is applied are also important; for example, the continued success of free software licenses like the GPL are entirely dependent on legal systems that respect, and have the power to enforce, the GPL's copy-left conditions.

As society has gotten more complicated, however, the law has also ballooned in complexity. In 1935, the U.S. Code was roughly 2,500 pages, but in 2020 the U.S. Code was over 30x as long at roughly 80,000 pages [38, 125].[1] The Code of Federal Regulations (C.F.R.), a compilation of the rules posited by executive departments and agencies, likewise continues to grow. In 1938, its first release, the C.F.R. was around 20,000 pages. In 2019, the C.F.R. was a staggering 186,000 pages [12]. Furthermore, new judicial decisions come out regularly, which when combined with all previous existing judicial decisions, comprise U.S. case law.

[1] The measurement of the law is an active research area. I readily acknowledge that measuring page-length is an imperfect way to measure the length of the law, and I discuss this measurement issue later in this thesis. Many measurement issues remain active research questions; for example, we don't know how many laws there are, how to measure how long the law is, or how to measure legal complexity [11, 107, 33, 79, 32, 100, 112, 160].

Our federalist system, where governing authority is shared among the federal, state, county, city, and even neighborhood levels, further complicates legal issues.

The law seems too complicated! It's impossible for any person to know all the laws that apply to her. And yet in the event of a transgression "ignorance of the law excuses not" or *ignorantia juris non excusat*. The law still applies, even if the transgressor is unaware of its existence. With the ballooning complexity of our government and legal systems, our society desperately needs innovation in civic education. It's our moral duty to help people understand the laws that govern our lives.

In light of this problem, my thesis asks the broad question: can computation help people understand the law? Specifically, I investigate the two questions: (1) can we teach computers to understand the law, and (2) can we automate checking for legal compliance.

1.1 Can We Teach Computers to Understand the Law?

As I previously mentioned, the law has gotten too complicated. In 1935, the U.S. Code was roughly 2,500 pages, but in 2020 it is over 30x as long, at roughly 80,000 pages [38, 125]. The Code of Federal Regulations (C.F.R.), a compilation of the rules posited by executive departments and agencies, likewise continues to grow. In 1938 the C.F.R. was around 20,000 pages, but in 2019 it was a staggering 186,000 pages [12]. The U.S. Code, the compilation and organization of legislative laws, is separate from regulations from executive agencies, compiled in the Code of Federal Regulations, which is separate from the vast repository of case law that also comprises U.S. law. Complicating things further, the law is constantly changing. Yet, people are still bound by the law, whether or not they understand it.

In response to this problem, my first question asks: is there a way to represent the law in such a way that computers can understand it? If so, we'd be able to harness the power of computation in our scholarship, practice, and civic education.

1.2 Can We Automate Checking for Legal Compliance?

If we can combine the art of legal analysis with the power of computation, tools like data visualization, AI, UI/UX, human-computer interaction, and algorithms may be applied to the legal field as well. While simply synthesizing the law into a format understandable by computers is important, the next part of my thesis investigates the possibility of automatically interpreting the law as applied to a specific situation. In other words, the second part of my thesis investigates the ability of a layperson to ask questions like: "does my app violate HIPAA regulations", or "can my doctor share my medical history with a company without notifying me". These personalized questions more closely represent the kinds of legal questions that are of interest to laypeople.

1.3 How this Thesis is Organized

This thesis begins by giving some key background on the U.S. government and legal system. Then, because trying to teach computers to understand all areas of the law is too broad, I narrow the scope of my investigation to pursue a case-study in mobile health (mHealth) security and privacy regulation. This case study is motivated by the belief that by teaching computers to automatically audit mHealth security and privacy regulations, we can ultimately encourage equitable mHealth adoption and reduce inequalities in healthcare. To make this argument, I present a theoretical model of technology adoption, directly tied to one's trust in technology.

I then investigate if it is possible to automate compliance checking given existing mHealth security and privacy regulations. I conclude that it would currently be very difficult to do so and discuss key regulatory characteristics that make it difficult for computers to understand and audit legal regulations. However, I also find that the use of non-legal, authoritative guidance documents that explain regulatory intent may help in the creation of computational legal auditors. I then discuss other work related to computational legal analysis and finish by giving some recommendations for future study.

Chapter 2

Background: U.S. Government & Legal System

Before we can study how to make the law interpretable to computers, we must first be familiar with the basics of the U.S. government and legal system.

2.1 Common Law vs. Civil Law

One of the most important distinctions among legal systems are whether they are predominantly influenced by common law or civil law. While most legal systems (including the U.S.'s) are somewhat of a blend of common, civil, or other types of law, the United States is often seen as a classic example of common law. Other examples of predominantly common law states include Canada, the United Kingdom, Ireland, Australia, New Zealand, and India. Predominantly civil law systems include many other states like a large majority of Europe and South America, Russia, China, Japan, and South Korea [126, 123].

The most important difference between common and civil law systems is the different roles of the judicial branch and the legislature. In civil law systems, generally only legislative action can "update the law". Therefore, judicial decisions generally only apply to the involved parties. A judge facing a similar problem, with different parties, does not need to follow a similar, previous ruling. However, the distinguishing factor of common law is *stare decisis* – the doctrine of the legal precedent. In common law systems, once a judge rules on a

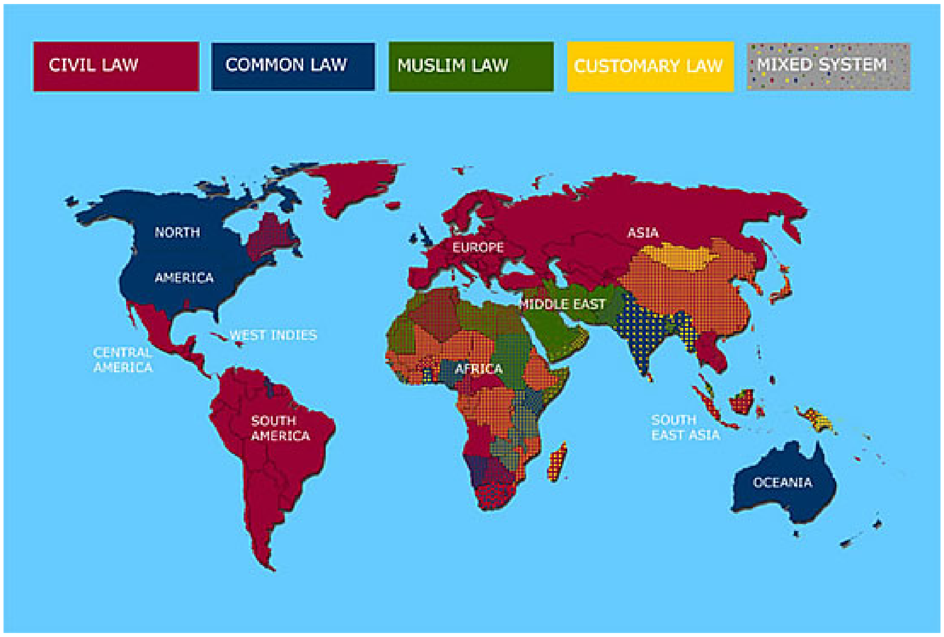

Figure 2.1: Common vs. Civil Law Adoption.
Figure from JuriGlobe at the University of Ottawa [123]

legal problem, all adequately similar future legal problems are expected to be resolved in the same way. Only a more powerful court can overrule a lower court's legal precedent (e.g., the Supreme Court of the United States can overrule legal precedents made by the U.S. District Court for the Southern District of New York) [17, 7, 96].

Case law, the aggregation of judicial decisions and their precedents, is vital to the governance of any common law jurisdiction. The large amount of previous and current judicial opinions, however, can make it difficult to interpret the law in common law systems [144, 98].

As an example of the importance of case law in the governance of American society, I list a few historically important judicial cases: *Marbury v. Madison* (1803), the judicial branch can strike down laws they find unconstitutional; *Dred Scott v. Sandford* (1857), slavery is legal, and slaves are not U.S. citizens; *Plessy v. Ferguson* (1896), racial segregation is legal; *United States v. Wong Kim Ark* (1898), children born in the U.S. have birthright citizenship; *Brown v. Board of Education* (1954), racial segregation in public schools is illegal; *Roe v. Wade* (1973), abortions are generally legal within the first two trimesters; *U.S. v. Windsor* (2013), the federal government must recognize same-sex marriage, overturns part of the 1996 Defense Against Marriage Act; and *Obergefell v. Hodges* (2015), legalizes same-sex marriage in all states [113, 34, 166, 6, 157, 145, 137]. Note how judicial opinions can change over time and how future opinions can directly overturn previously established precedents.

Judicial decisions have also had an important impact on technology; a few examples include: *Bernstein v. United States* (1996), software code is protected free speech, and encryption software can be freely shared; *Reno v. ACLU* (1997), strengthens “free speech” protections of web content; *Zeran v. America Online, Inc.* (1997), strengths CDA 230, no provider of a computer service is held liable for how users use the service; *Perfect 10, Inc. v. Amazon.com, Inc.* (2007), image search engines do not violate copyright law; and *Carpenter v. United States* (2019), police must get a warrant to access historical cell site location data [78, 136, 106, 114, 105, 154].

These judicial decisions are just a few of the critical opinions that form U.S. law. Those familiar with U.S. history and recent events ought to know the importance of the judicial branch in governing American society. The importance of case law in the U.S. legal system must be underscored, especially for those who are unfamiliar with common law societies.

2.2 Federalism

Another distinctive aspect of the United States is federalism. Federalism is the practice of sharing power between a national (i.e., federal) government and local governments. These local governments (e.g., each individual state in the U.S.) often have constitutionally protected sovereignty and powers. The United States is often cited as a canonical example of federalism and one of the first countries to experiment with federalism on a national scale [23, 90].

Federalism is often compared to the unitary state, a system in which the central government exerts more control over local governments. These local governments often do not have constitutionally protected sovereignty, and central governments, if they so choose, may abolish or override these local governments. France is often thought of as a canonical example of a unitary state [120, 23].

Federalism is also compared to confederalism, a system in which local governments trump more general, centralized governments. The first government of the United States, under the Articles of Confederation, had many characteristics of confederalism [29, 119]. The European Union is another example of confederalism, as it balances both federalism with confederalism in a system sometimes called a "federal union of states" [104, 42, 84].

While the United States is often a canonical example of federalism, and France a canonical example of a unitary state, most countries share federalist and unitary characteristics. Nonetheless, like the difference between common vs. civil law states, countries can often be categorized as primarily federalist or unitary. States that are primarily federalist include: the United States, Canada, Mexico, Germany, Spain, Brazil, Argentina, South Africa, India, Pakistan, Russia, the United Arab Emirates, and Australia. Unitary states primarily comprise of the rest of the world and a majority of the world's population.

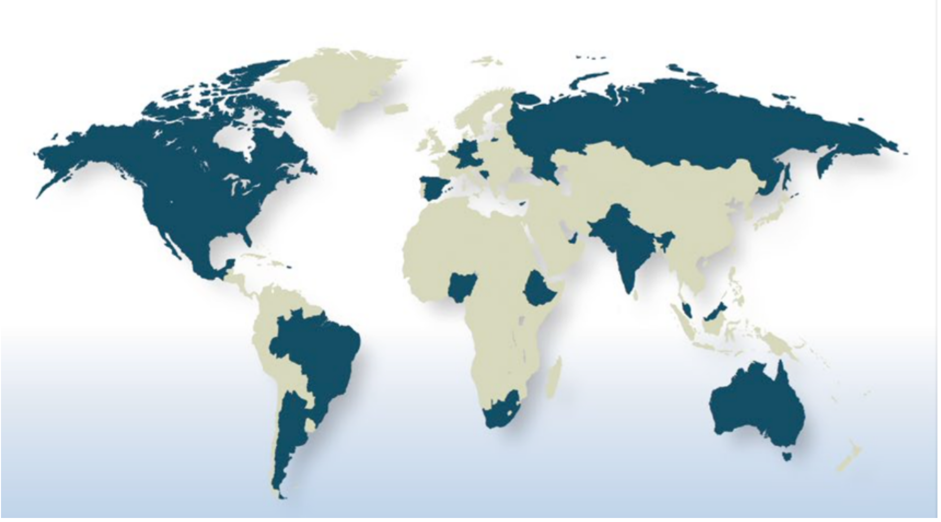

Figure 2.2: Federalist vs. Unitary States.
Figure from The Forum of Federations Handbook of Federal Countries (2020) [86]

The United States, with its strong system of federalism, complicates our legal analysis. Because it is not always clear what powers are delegated to the federal government and what powers are delegated to the states, federalism remains one of the most basic, unresolved, contentious areas of debate and research in America. Furthermore, some of the most controversial topics in America today (like abortion, gun control, voting, and COVID restrictions) revolve around federalism.

If our goal is to computationally analyze the current state of the law, we must not only be able to analyze federal, state, county, and city laws, but also be able to interpret when federal laws should trump local laws, and vice versa. Furthermore, each state has its own body of case law (recall the section: *Common Law vs. Civil Law*), somewhat independent from federal case law. For each given legal situation, it can be difficult to determine the relationship between relevant federal and state case law.

The analysis of legal questions reminds me of Moravec's paradox, summed up in this tidy xkcd cartoon. Like in computer science, some legal questions are fairly simple. However, other seemingly simple legal questions may be deceptively difficult. In the United States, a canonical common law and federalist society, legal questions may be especially difficult to resolve.

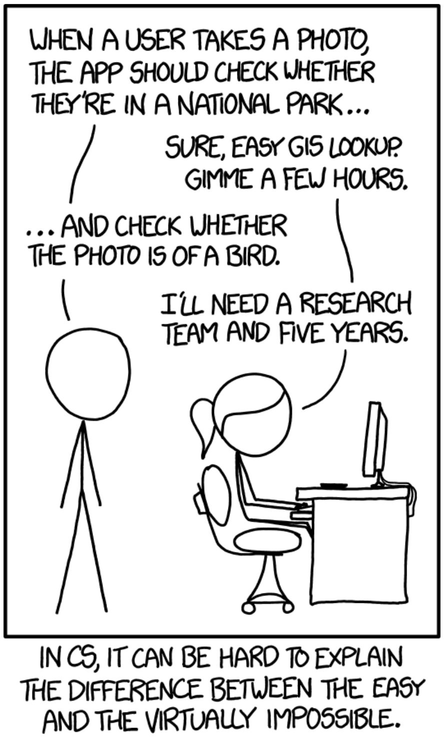
WHEN A USER TAKES A PHOTO, THE APP SHOULD CHECK WHETHER THEY'RE IN A NATIONAL PARK...
SURE, EASY GIS LOOKUP. GIMME A FEW HOURS.
...AND CHECK WHETHER THE PHOTO IS OF A BIRD.
I'LL NEED A RESEARCH TEAM AND FIVE YEARS.
IN CS, IT CAN BE HARD TO EXPLAIN THE DIFFERENCE BETWEEN THE EASY AND THE VIRTUALLY IMPOSSIBLE.

Figure 2.3: Moravec's Paradox.
Image from xkcd [169]

Chapter 3

An Applied Example: Automated Auditing of mHealth Security and Privacy Regulations can Increase Trust and Reduce Inequalities in Healthcare

While this thesis is broadly concerned with the application of computation to legal analysis, the task of teaching computers to understand and reason about all areas of the law is too broad in scope for my thesis. Therefore, I focus on a narrow application of computational legal analysis; I study if it is possible to teach computers to understand the privacy and security regulations for mobile health (mHealth) applications. I chose this narrow subsection of the law because I believe computational understanding of mHealth security and privacy regulations can ultimately improve the quality and equity of our healthcare system. This section gives an example of how computational legal analysis can have a significant, practical impact on the population at large.

3.1 Why mHealth?

mHealth applications have been lauded for their potential to provide cheap, effective, and personalized healthcare to large groups of people. Medical professionals are optimistic that mHealth adoption can increase access to high-quality healthcare and improve public health [94, 165, 31]. Patients are similarly optimistic, hoping that mHealth will improve affordability, convenience, and the quality of healthcare they will receive [151]. In a survey taken near the beginning of the COVID-19 pandemic, almost half of Americans reported they were users of an mHealth application [153]. The COVID-19 pandemic further accelerated mHealth adoption; during the first year of the pandemic mHealth app downloads surged by over 50% worldwide [152].

mHealth is also important because it presents us with an opportunity to increase equitable outcomes in healthcare. Current healthcare disparities in the United States and around the world are well documented [14, 140, 122, 95, 15, 97]. Experts hope that the low-cost and accessibility of mHealth applications can improve health outcomes in underserved communities [158, 35, 149, 115].

Despite this optimism, or perhaps because of it, others have been quick to point out that mHealth applications may exacerbate existing healthcare inequalities. Previous research on technology inequality shows that society's most vulnerable have the lowest technology adoption rates and are disproportionally harmed by new technologies [109, 110]. Effective mHealth applications, if only adopted by the healthiest and most privileged members of society, may increase health disparities.

Recent research supports the assertion that mHealth applications may be disproportionately benefiting society's most privileged, while leaving the most vulnerable members of society behind [24, 150, 41, 8]. Organizations have taken notice, and are actively studying how to tackle health disparities within mHealth [43, 164].

My investigation of the mHealth regulatory landscape is motivated by my belief that by automating a firm's or the FDA's ability to check an application's compliance with FDA regulations, we can ultimately improve the equitable adoption of mHealth applications. To make this argument, I contribute a novel model of technology adoption that explicitly links technology adoption to trust in the technology itself. This trust is predicated on many factors, among which include the security and privacy of the application. The ability to efficiently check if an application satisfies the FDA's security and privacy regulations will ultimately improve public trust in and the adoption of mHealth.

This mHealth case study is just one narrow example of how bringing computation to legal analysis can directly benefit ordinary people.

3.2 A Trust-Based Model for Technology Adoption

mHealth technologies are new and not evenly adopted by the population. If we are trying to equitably increase mHealth adoption, we need a theory for technology adoption. Here, I put forward my own trust-based model for technology adoption. This model is adapted from a theory Lawfare's Benjamin Wittes gave in a 2021 lecture at Duke University's Sanford School of Public Policy, in conversation with Professor David Hoffman, Intel's associate general counsel and global privacy officer [93].

Notice that the first layer of my model lists key characteristics that contribute to technology adoption. Trust is vitally important, and I elaborate on why trust matters for technology adoption in a later section. Because my analysis focuses on why trust is important and how computational auditing of FDA regulations can improve trust, I skim over the other characteristics of my model. These other characteristics are:

1. *Momentum* – people are reluctant to change usage patterns. For example, Facebook, Instagram, and WhatsApp usage in the EU remains high, though trust is low, and alternatives exist (e.g., Signal, Telegram). The continued use of old, unpatched software is another example of technological momentum.

2. *Market Conditions* – economic conditions influence adoption. For example, Huawei products are used in countries even when public trust is low because of their competitive price.

3. *Access* – people cannot adopt new technology if they cannot access it. For example, global demand for mRNA COVID-19 vaccines far outpace supply.

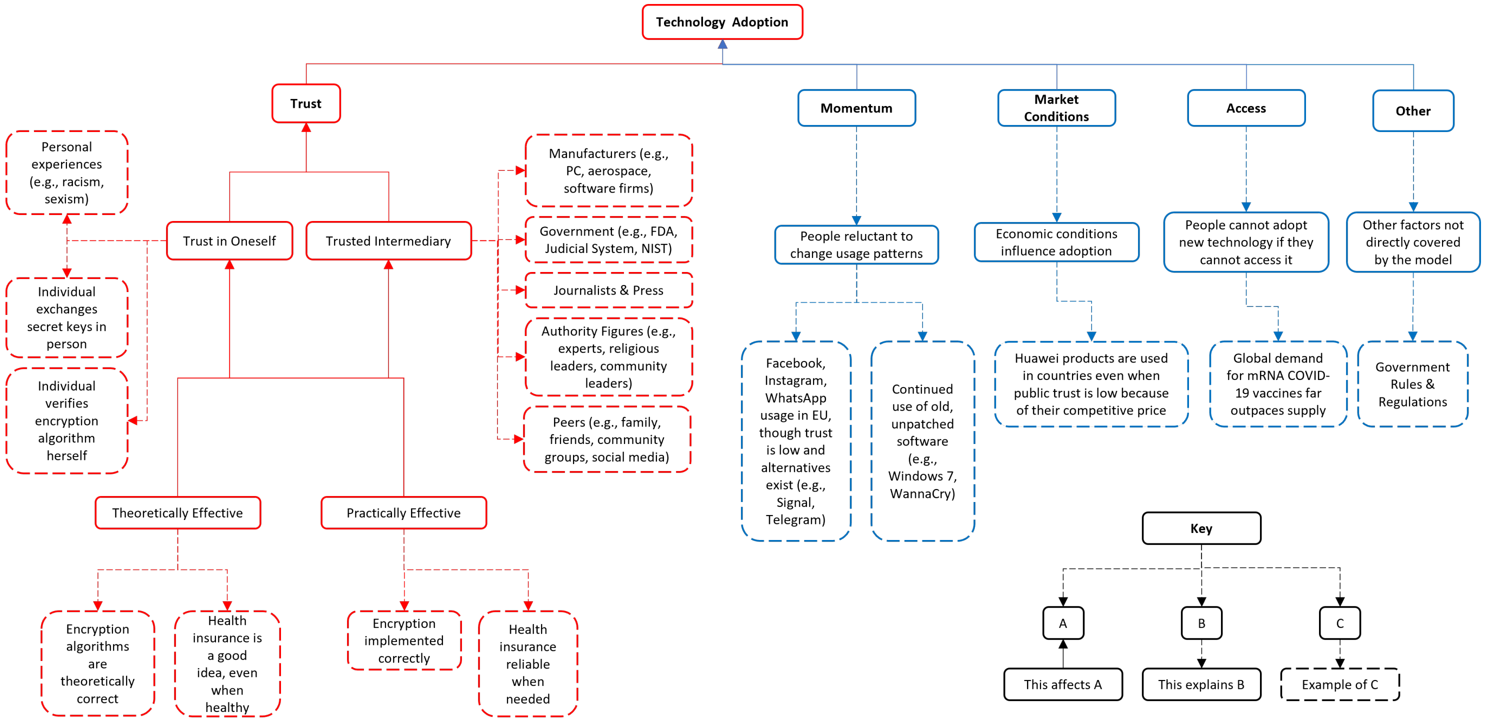
Technology Adoption
Trust
Momentum
Market Conditions
Access
Other
Personal experiences (e.g., racism, sexism)
Trust in Oneself
Trusted Intermediary
Manufacturers (e.g., PC, aerospace, software firms)
Government (e.g., FDA, Judicial System, NIST)
Journalists & Press
Authority Figures (e.g., experts, religious leaders, community leaders)
Peers (e.g., family, friends, community groups, social media)
Individual exchanges secret keys in person
Individual verifies encryption algorithm herself
Theoretically Effective
Practically Effective
Encryption algorithms are theoretically correct
Health insurance is a good idea, even when healthy
Encryption implemented correctly
Health insurance reliable when needed
People reluctant to change usage patterns
Economic conditions influence adoption
People cannot adopt new technology if they cannot access it
Other factors not directly covered by the model
Facebook, Instagram, WhatsApp usage in EU, though trust is low and alternatives exist (e.g., Signal, Telegram)
Continued use of old, unpatched software (e.g., Windows 7, WannaCry)
Huawei products are used in countries even when public trust is low because of their competitive price
Global demand for mRNA COVID-19 vaccines far outpaces supply
Government Rules & Regulations
Key
A
B
C
This affects A
This explains B
Example of C

Figure 3.1: A Trust-Based Model for Technology Adoption

4. *Other* – other factors not directly covered by the model. For example, specific government rules, regulations, and requirements.

3.2.1 How Security Affects Trust

My model includes two ways to build trust, each of which require two necessary conditions. All four of these areas are ultimately affected by security and privacy.

Trust in Oneself

The first way to build trust is to have trust in one's own first-hand experiences. For example, it is well established that many minorities experience racism in our healthcare system. These experiences result in greater medical mistrust and lower patient engagement [103, 148, 167]. Those who have experienced medical racism and distrust the medical industry may also mistrust mHealth applications. On the other hand, those who have only had positive experiences with the medical industry may be especially trusting of mHealth applications.

Another example is that victims of ransomware may need to see stricter assurances that their data will be protected from hackers before feeling comfortable with mHealth. Recall Vastaamo, the Finnish psychotherapy company whose hack became public knowledge in 2020. In a shocking and cruel move, criminals posted thousands of patient's private psychotherapy notes online [139, 13, 89]. Those who have been the victim of security and privacy violations in the past may require extra reassurance that their security and privacy will be protected in an mHealth application.

Trusted Intermediary

If we don't have any first-hand experience with a new technology, we can still rely on the experiences and wisdom of trusted intermediaries. These people or organizations may have first-hand experience with this technology or be in a better position to judge the technology's merits.

Relying on someone else's judgement is not a new concept. A traditional example of this kind of trust relationship exists in our public key infrastructure, in which users must ultimately trust their certificate authorities. However, these trust networks are also common in other areas of everyday life. For example, we trust aerospace and medical manufactures to make safe products. We also trust government entities (like the Federal Aviation Administration and the Food and Drug Administration) to certify the safety of these products. Trusted authority figures (e.g., experts, religious and community leaders), journalists, friends, and social groups may also weigh-in on new technology.

mHealth applications (or other technologies) with strong security and privacy practices can eventually build a strong reputation throughout one's network of trusted intermediaries. In this case, even if one has not used an mHealth application, other people one trusts may have used the application and can vouch for its trustworthiness.

Theoretically Effective

For an application to establish trust through first-hand experiences or through trusted intermediaries, it is necessary for this application to be theoretically effective. This means that, in theory, an application or technology can meet the demands of its users.

Health insurance is a good example of the necessity of this theoretical requirement. Health insurance is predicated on the belief that purchasing health insurance is a good idea, even when one is healthy. This belief, regardless of whether it's right or wrong, is necessary for the health insurance industry. If only those seeking expensive healthcare procedures sought out health insurance, insurers would not be able to spread out healthcare costs among many people.

An example closer to mHealth applications is that individuals must believe that mHealth apps can protect their sensitive health data. If people had absolutely no confidence in an app's privacy and security, this application would be unable to inspire confidence through first-hand experiences or trust networks.

Practically Effective

Finally, applications and new technologies must also be effective in practice. This means that applications must satisfy user demands in practice, not only in theory.

Again, health insurance is a good example of the difference between theoretical effectiveness and practical effectiveness. Not only must people believe that health insurance is a good idea when they're healthy (theoretical effectiveness), health insurance must actually be reliable when a person is in need (practical effectiveness). If someone's health insurance consistently refuses to insure their medical expenses, they might lose faith in the theoretical effectiveness of health insurance. This loss of faith and bad experience directly affects either one's first-hand experiences (trust in oneself) or the experiences of a trusted intermediary.

In the mHealth security and privacy sphere, practical effectiveness is required to convince patients that their private medical data can be safely entrusted to mHealth applications. Even if patients believe that it is theoretically possible to create an mHealth app that safely guards their medical information, if mHealth apps are regularly leaking personal data, patients may view mHealth apps with suspicion. Failures to protect medical data may directly affect a patient (trust in oneself) or affect someone in their trust network (via a trusted intermediary).

3.2.2 Computational Auditing of mHealth Security and Privacy Regulations Can Build Trust

I believe that computational auditing of mHealth security and privacy regulations can improve both necessary conditions required to build trust. If these regulations truly do improve the security and privacy of mHealth applications, the existence of these regulations may convince people that mHealth applications can theoretically meet their security and privacy demands.

The ability to enforce these regulations at scale through computational auditing would also improve mHealth's privacy and security in practice. Applications that fail federal privacy and security requirements would be prevented from entering the market, ensuring that all mHealth applications meet minimum privacy and security guarantees, allowing apps to be "effective in practice".

Furthermore, a computational audit would require the operationalization of mHealth privacy and security requirements. Therefore, firms with failed mHealth applications could be told what specific operationalized requirements they failed. By identifying specific points of improvement, this auditing process would allow firms to quickly fix their mistakes and build more private and secure applications.

Because computational auditing can help mHealth applications improve their theoretical and practical effectiveness, patients may be more willing to try out these applications. Other people's positive experiences with mHealth apps can improve the reputation of these applications in one's network of trusted intermediaries. Finally, if one is convinced that an mHealth app is theoretically effective, effective in practice, and trusted by one's network, one may be willing to test out the mHealth app oneself, resulting in first-hand experiences. These experiences, if positive, further increase trust in the application.

This improved trust, then, encourages the adoption of mHealth technologies, and is especially important for groups with lower health outcomes and greater medical mistrust.

Chapter 4

Can We Automatically Check if mHealth Applications Satisfy Security and Privacy Regulations?

To study if we can apply the tools of computation to legal analysis, I narrowed my focus to a small subsection of our legal landscape – the regulation of mHealth applications. This narrower focus dovetails nicely with my host lab, as one research area in my lab is the security and privacy of mHealth devices. In my thesis, instead of looking for new classes of vulnerabilities in existing mHealth applications, or evaluating defenses for existing vulnerabilities, I ask a different question – can we create an automated test to detect if mHealth apps satisfy current regulations for security and privacy?

To answer this question, we must first be able to identify what are the current mHealth regulations for security and privacy.

4.1 Background: Federal Rules, Regulations, and Recommendations for mHealth Security and Privacy

If we wish to design cybersecurity tests that are compatible with federal rules, regulations, and recommendations, it is imperative to be familiar with these federal guidelines. Therefore, I give a high-level overview of the federal regulatory landscape. My overview is not comprehensive; I limit my analysis to the federal U.S. regulatory environment. Other important non-federal regulations that mHealth apps may be subject to, like state-based biometric identification laws [138] or California's Consumer Privacy Act, are considered out-of-scope of this analysis. Federal case law is similarly considered out-of-scope.

4.1.1 mHealth Apps Inherit The Existing Device Regulatory Framework

The FDA regulatory framework generally treats mHealth apps like traditional medical devices. Importantly, if an mHealth app does not meet the definition of a medical device, its general effectiveness is not regulated by the FDA (though it is still subject to other laws, like those related to fraud) [64]. An mHealth app is a medical device when it meets the definitions in the FD&C Act in 21 U.S.C. § 321(h): a medical device is "intended for use in the diagnosis of disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease, in man or other animals, or intended to affect the structure or any function of the body of man or other animals" and does not include the software functions excluded in 21 U.S.C. § 360j(o). These software functions include a handful of exceptions; one notable exception is that software with an intended use of "maintaining or encouraging a healthy lifestyle and is unrelated to the diagnosis, cure, mitigation, prevention, or treatment of a disease or condition" is not considered a medical device. The FDA provides numerous resources to help developers determine if their mHealth app is a medical device. If an application is not a medical device, the FDA does not regulate its effectiveness [64, 44, 45, 66].

Even if an mHealth app is recognized as a medical device by the FDA, the FDA may choose to not regulate low-risk apps by exercising "enforcement discretion." This means that the FDA may regulate the app in the future, but does not intend to right now, as the "FDA intends to apply its regulatory oversight to only those software functions that are medical devices and whose functionality could pose a risk to a patient's safety if the device were to not function as intended" [66]. To determine if an mHealth app that is a medical device will be regulated by the FDA, one can reference Appendix B of the guidance document titled "Policy for Device Software Functions and Mobile Medical Applications" [66], view the FDA's online resources [44], or contact the FDA directly.

Medical devices are further sorted by their risk and functionality. Class I devices are low risk and may not require a premarket 510(k) notification. Class II devices are moderate risk, may have special controls, and often require a premarket 510(k) notification. Class III devices are high risk, are strictly regulated, and require a premarket approval (PMA) [69, 53, 67]. Additionally, the FDA assigns all devices a device type, based on the device's functionality (e.g., blood pressure alarm, electrocardiograph software for over-the-counter use). These device types may be found in 21 C.F.R. § 800-1050 and are searchable online [75]. In the Code of Federal Regulations each device type is described, classified as a Class I-III device, and given special controls (if applicable). Medical devices must satisfy all the FDA regulations for their device type, regardless if they are an mHealth app or a traditional medical device. Manufacturers can ask the FDA to formally classify their device by filing a 513(g) request [63]. Manufacturers with a medical device that does not fit into an existing device type can ask the FDA to create a new device type through a De Novo request, an alternative to premarket 510(k) notification and premarket approval (PMA) [61].

Finally, all devices are subject to FDA General Controls. These controls include regulations pertaining to Establishment Registration (21 C.F.R. Part 807), Medical Device Listing (21 C.F.R. Part 807), Premarket Notification 510(k) (21 C.F.R. Part 807 Subpart E) or Premarket Approval (PMA) (21 C.F.R. Part 814), Investigational Device Exemption (IDE) (21 C.F.R. Part 812), Quality System Regulation (QS Regulation) (21 C.F.R. Part 820), Labeling (21 C.F.R. Part 801), and Medical Device Reporting (21 C.F.R. Part 803).

4.1.2 Cybersecurity Specifics for mHealth

Many regulations that affect medical devices also affect their cybersecurity, as cybersecurity involves all parts of the technology stack. Nevertheless, some key principles affecting mHealth application developers include:

1. The FDA splits cybersecurity risks into two categories: acceptable ("controlled") risks, and unacceptable ("uncontrolled") risks. To determine if a risk is acceptable or unacceptable, manufacturers must consider the exploitability of the vulnerability and the severity of patient harm if the vulnerability is exploited. Manufacturers are required to mitigate unacceptable risks either by controlling them entirely or by reducing them to become acceptable risks. Manufacturers then have to document any residual acceptable risks and explain why they only pose a "sufficiently low" risk of patient harm. There is some flexibility when defining acceptable and unacceptable risks; the FDA guidance document "Postmarket Management of Cybersecurity in Medical Devices" provides further guidance [56].

2. mHealth application manufacturers are required to manage the cybersecurity risk of their entire application throughout the lifecycle of the app. This requirement is more difficult than it appears, as it requires application manufacturers to manage cybersecurity risks of "off-the-shelf" software (e.g., open-source software, proprietary third-party software). Manufacturers may also be required to provide the FDA with "continuity plans" that address how the manufacturer will handle no-longer supported "off-the-shelf" software. [65, 4, 56, 49].

3. There may be regulatory hurdles when trying to patch cybersecurity vulnerabilities. Generally, vulnerabilities that present controlled risks can be fixed without making substantial changes to the device. However, if an update consists of "substantial changes" or could "significantly affect the safety or effectiveness" of the device, the manufacturer may need to re-submit a premarket submission (e.g., PMA Supplement, 510(k)) to the FDA. [56, 48, 52, 57, 50].

4. Device manufacturers may be required to notify the FDA upon discovery of an uncontrolled cybersecurity risk. Reporting requirements depend on the device's class and

type. Generally, vulnerabilities that present only a “controlled risk of patient harm” do not need to be reported to the FDA. [56, 55, 52].

Most FDA cybersecurity guidance is general advice. However, a 2018 draft guidance titled “Content of Premarket Submissions for Management of Cybersecurity in Medical Devices” contains more specific technical recommendations [60]. The FDA also recommends the more-technical NIST “Framework for Improving Critical Infrastructure Cybersecurity” [118].

4.1.3 Demystifying the U.S.C., C.F.R., and FDA Guidance Documents

When learning about federal guidelines for mHealth cybersecurity, the FDA often cites the United States Code, the Code of Federal Regulations, and FDA Guidance Documents. It is important to understand what these publications are and how they relate to one another.

The United States Code (U.S. Code or U.S.C.) is a codification of the current laws passed by the legislative branch of the United States. The U.S. Code contains 53 Titles and is organized by subject; for example, Title 21 of the U.S.C. generally covers the laws applicable to food and drugs. The Food, Drug, and Cosmetic Act (FD&C Act), which gives the FDA many of its powers, is housed within Chapter 9 of 21 U.S.C. To keep the U.S.C. up-to-date, new main editions of the U.S.C. are published every six years by the legislative branch of the U.S. government, with annual cumulative supplements published in-between main editions.

The Code of Federal Regulations (C.F.R.) is the official codification of the rules and regulations promulgated by departments and agencies of the executive branch of the United States. While laws in the U.S.C. give executive bodies (like the FDA) their powers, the C.F.R. describes how these bodies choose to exercise their powers. Accordingly, it is published by the executive branch of the U.S. government. One well-known example of the C.F.R. is the Health Insurance Portability and Accountability Act of 1996 (HIPAA). While the laws that comprise HIPAA are scattered throughout the U.S.C., the rules that describe what HIPAA means to the public are located in Title 45 of the C.F.R. Part 160, Part 162, and Part 164. The official C.F.R. is updated annually, though there is an unofficial online version (e-CFR) that is updated daily.

FDA Guidance Documents often explain, elaborate, or complement parts of the C.F.R. These documents describe the FDA's current thinking about topics under its authority, including how the FDA plans to regulate (or not regulate) mHealth applications. It's important to note that these guidance documents are only recommendations, with the exception of sections that cite specific regulatory or statutory requirements. However, these documents offer invaluable advice; for example, those who are required to submit their mHealth application for premarket evaluation should familiarize themselves with what the FDA says to include in premarket submissions. These guidance documents cover many topics, and all 2,500+ guidance documents are searchable online [76].

4.1.4 Labeling Requirements for mHealth Applications

mHealth labeling requirements are not significantly different from labeling requirements for traditional medical devices. This means that mHealth labeling, like traditional medical device labeling, is tightly regulated. Unfortunately, these regulations are not contained in one place. Common labeling regulations can be found in 21 C.F.R. Parts 801, 809, 812, 820, 830, and 1010. FDA Guidance Documents also explain labeling requirements [46, 58, 54]. Because labeling requirements are often uniquely tailored for each medical device, this section only clarifies a few general requirements applicable to all medical devices.

mHealth application manufacturers must first understand that a device's label and its labeling are different legal constructs. A label is more narrowly applicable, and it is defined in the FD&C Act as "a display of written, printed, or graphic matter upon the immediate container of any article" (21 U.S.C. § 321(k)). Labeling is defined in the FD&C Act as "all labels and other written, printed, or graphic matter (1) upon any article or any of its containers or wrappers, or (2) accompanying such article" (21 U.S.C. § 321(m)). Historically, labeling has been interpreted broadly to encompass material like advertising, posters, pamphlets, instructions for use, user manuals, quick-start guides, and packaging [2, 5, 46, 54].

Manufacturers should note that labeling must address both the intended use of the product and its current use. For example, if the manufacturer knows that their device is to be used for purposes other than its original purpose, the manufacturer must update its labeling (including its directions for use) to accommodate these other use cases. [66, 3].

As pursuant to the FD&C Act, labeling must not be misbranded (21 U.S.C. § 352) or contain any false or misleading statements [46]. Labeling, including advertisements, are considered misbranded if it is false or misleading to the user. This obviously includes both willful and accidental deception from the manufacturer, but also includes truths that may "lead to a false impression in the mind", as "labeling can be deemed by the FDA to be in violation of the law if it proves deceptive to the customer" [46]. Labeling characteristics common in tech and other industries like "expressions of opinion or subjective statements" and "failure to reveal material facts, consequences that may result from use, or the existence of difference of opinion", are categorized as misleading in the medical context [46].

4.1.5 Apps Must Be HIPAA Compliant If Used By Covered Entities

In short, the Health Insurance Portability and Accountability Act (HIPAA) is a piece of legislation that governs how "covered entities" (e.g., health care providers like hospitals and other organizations that bill patients, health plans like insurance providers, and health care clearinghouses like billing services) and their "business associates" (an organization, person, or subcontractor that creates, receives, maintains, or transmits "protected health information") handle "protected health information" (individually identifiable health information). Legal definitions of "covered entities", "business associates" and "protected health information" can be found at 45 C.F.R. § 160.103. The HITECH Act was passed in 2009 to encourage the adoption of electronic medical health records and to strengthen parts of HIPAA. These two Acts were eventually merged into the 2013 HIPAA Omnibus Rule, and today HIPAA usually refers to the merged regulation.

Note that HIPAA only applies to "covered entities" and their "business associates". Therefore, an mHealth app that is not directly used by a healthcare provider, health plan, or healthcare clearinghouse is exempt from HIPAA. For example, an mHealth app that monitors a user's sleep that is not intended for use by "protected entities" may collect protected health information not covered by HIPAA. However, if a hospital instructs a patient to download the same app, then the protected health information collected by the app is protected by HIPAA. A good rule of thumb is that a manufacturer should make their mHealth app HIPAA compliant if the manufacturer wants "protected entities" to be able to use their app.

mHealth developers can think of HIPAA as consisting of three main Rules, the Privacy Rule (45 C.F.R. Part 160 and Subparts A and E of Part 164), the Security Rule (45 C.F.R. Part 160 and Subparts A and C of Part 164), and the Breach Notification Rule (45 C.F.R. Part 164 Sections 400-414). The Privacy Rule is generally concerned with when protected health information (PHI) can be transferred to another entity. The Security Rule is generally concerned with the security requirements needed to protect PHI. And the Breach Notification Rule covers notification requirements when there is a suspected breach of PHI. Developers might also be interested in the Enforcement Rule (45 C.F.R. Part 160 in Subparts C, D, and E), which details how the U.S. Department of Health and Human Services' Office of Civil Rights (OCR) enforces HIPAA, including legal procedure and monetary penalties for non-compliance. The Department of Health and Human Services has extensive online resources for HIPAA on its website [121].

4.1.6 New Frameworks on The Horizon

The FDA, working with other national health regulators, has defined Software as a Medical Device (SaMD) as “software intended to be used for one or more medical purposes that perform these purposes without being part of a hardware medical device”. SaMD includes mHealth applications, also called Mobile Medical Applications (MMAs) [87].

The FDA is working on developing regulations for AI/ML-based SaMDs. Although currently lacking clear Guidance Documents, the FDA has released an action plan that details what kind of regulations they will likely pursue in the future. These regulations include a “tailored regulatory framework for AI/ML-based SaMD”, a “good machine learning practice (GMLP)”, a “patient-centered approach incorporating transparency to users”, “regulatory science methods related to algorithm bias & robustness”, and “real-world performance (RWP) monitoring” [74]. The FDA regularly publishes discussion papers and requests for feedback when developing guidelines, and this practice is expected to continue for AI/ML-based SaMDs [68].

The COVID-19 crisis also affects the regulation of mHealth applications. The FDA temporarily relaxed some mobile health regulations for the length of the COVID-19 emergency

[72, 71]. Although these changes are temporary, they may become permanent when the public health emergency has been lifted. Mobile medical applications is an especially innovative field, and those working in this area are encouraged to keep up with the latest publications from the FDA.

Finally, developers may be interested in the FDA's Breakthrough Devices Program or Digital Health Software Precertification (Pre-Cert) Program, new programs that may help expedite the approval of novel treatments and technologies [59, 70].

4.1.7 Where to Learn More

The FDA has many useful online materials for those who wish to learn about how devices are regulated. Most questions have been answered somewhere on their website. The FDA has many useful, detailed webpages, like their "Overview of Device Regulation" and "Classify Your Medical Device" pages [73, 69]. The FDA's Center for Devices and Radiological Health (CDRH) is responsible for regulating medical devices. CDRH Learn is their online video platform that explains their current policies.

FDA Guidance Documents explain official FDA policies, and all 2,500+ guidance documents are searchable online [76]. Withdrawn medical device guidance documents are also searchable online [77]. Unfortunately, it can be hard to determine when a guidance document has not been withdrawn but instead has been replaced or merged into a newer document.

Those with further questions may contact the FDA's Division of Industry and Consumer Education (DICE) [62]. Each guidance document also contains contact information for questions pertinent to its subject area. Additionally, other organizations may provide consultants with specialized expertise. Finally, developers should contact a lawyer, as each development situation is unique.

4.2 It Would Be Very Challenging to Automatically Check if mHealth Applications Satisfy Security and Privacy Regulations

After grappling with the current state of mHealth regulations, I conclude that it would be very challenging to automatically detect if mHealth apps satisfy their FDA security and privacy regulations. This is due to a few main reasons: (1) the open-ended design of the law makes it difficult to operationalize, (2) the FDA's reliance on "self-certification" of security and privacy compliance, and (3) a complicated regulatory environment that makes it difficult to know what regulations apply.

4.2.1 Current Regulations are too Open-Ended to be Meaningfully Operationalized

If an auditor wants to check if an mHealth application technically satisfies its security and privacy regulations, then the auditor needs specific, technical requirements to check. However, existing regulations are often too open-ended to be used for auditing security and privacy compliance. From a philosophical perspective, the law may be open-ended by design; by avoiding specificity the law remains relevant, even as technology evolves.

Furthermore, because mHealth applications were grandfathered into the existing device regulatory environment, these existing regulations are often focused on securing the safety of manufactured hardware devices. The creation of mHealth applications, sometimes called the "manufacturing of software devices", doesn't quite fit into the existing regulatory framework. Because of this misfit, there may be many ways to interpret how existing regulations apply to mHealth applications.

For example, the Quality System Regulations, located at 21 C.F.R. Section 820, is commonly cited as an important federal regulation that ensures the safety of medical devices. These regulations are also commonly cited when discussing the cybersecurity requirements of healthcare devices, as cybersecurity risks can directly affect the safety of a medical device.

However, the current Quality System Regulations are too open-ended to ensure strong cyber safety. While these regulations may be satisfactory for experienced developers, inexperienced developers may interpret these regulations in an unsafe way. Some sections of the Quality System Regulations most relevant to cybersecurity read:

> *Design validation.* Each manufacturer shall establish and maintain procedures for validating the device design. Design validation shall be performed under defined operating conditions on initial production units, lots, or batches, or their equivalents. Design validation shall ensure that devices conform to defined user needs and intended uses and shall include testing of production units under actual or simulated use conditions. Design validation shall include software validation and risk analysis, where appropriate. The results of the design validation, including identification of the design, method(s), the date, and the individual(s) performing the validation, shall be documented in the DHF. – 21 C.F.R. 820.30(g)
>
> *Quality audit.* Each manufacturer shall establish procedures for quality audits and conduct such audits to assure that the quality system is in compliance with the established quality system requirements and to determine the effectiveness of the quality system. Quality audits shall be conducted by individuals who do not have direct responsibility for the matters being audited. Corrective action(s), including a reaudit of deficient matters, shall be taken when necessary. A report of the results of each quality audit, and reaudit(s) where taken, shall be made and such reports shall be reviewed by management having responsibility for the matters audited. The dates and results of quality audits and reaudits shall be documented. – 21 C.F.R. 820.22

As the reader can see, these sections are so open-ended that firms may choose to satisfy these regulations in a myriad of unsafe ways. “Design validation shall include software validation and risk analysis, where appropriate” does not specify how rigorous “software validation” or “risk analysis” must be, nor does the regulation specify what cyber risks firms need to look for to be in compliance.

Similarly, while the Quality Audit provision requires that firms perform quality audits from time to time, this section raises more questions than it answers. How often must audits be performed? What should the audit consist of? Should firms hire a red team to look for cyber vulnerabilities, or simply check that users are using new passwords each year? What are these “established quality system requirements” the regulation references? What does it mean to “be in compliance” with “established quality system requirements”? And how is the “effectiveness” of the quality system to be determined? Without specificity, it’s possible to have firms satisfy these requirements without actually improving their cyber defenses.

Another reason why mHealth security and privacy regulations may be open-ended is that open-ended requirements gives OEMs (original equipment manufactures) greater freedom to design a cybersecurity strategy that is tailored to the individual needs of their application. Because applications have different operating requirements, their security needs will also be different. Detailed regulation that applies to all applications may ignore these differences, which may actually weaken application cybersecurity. By leaving cybersecurity requirements open-ended, the FDA allows experienced teams to develop cybersecurity strategies that are tailored to application requirements.

Not all FDA mHealth security and privacy documentation is open-ended, though. In my investigation, I found one document with very specific, operationalizable cybersecurity recommendations, a 2018 FDA draft guidance document titled “Content of Premarket Submissions for Management of Cybersecurity in Medical Devices” [60]. The recommendations in this document, an FDA guidance document, are not legally binding. Also note this document is a “draft guidance”, which means that it’s not an official guidance document. Draft guidance documents are often issued by the FDA to solicit feedback; they show us what the FDA is currently thinking about. The content in this 20-page draft guidance may someday replace the current 2014, 7-page “Content of Premarket Submissions for Management of Cybersecurity in Medical Devices” guidance document [51].

This 2018 draft guidance splits devices into two categories. Tier 1 “Higher Cybersecurity Risk” Devices, networked medical devices in which a cybersecurity incident could directly lead to patient harm in multiple patients; and Tier 2 “Standard Cybersecurity Risk” Devices, which includes all devices not in Tier 1. From an implementation standpoint, this guidance document has 22 specific recommendations to “identify and protect devices assets and functionality”, and 15 specific design expectations to “detect, respond, [and] recover” from a cybersecurity incident.

Security researchers will recognize that these recommendations are helpful. Some recommendations include:

- “Employ a layered authorization model by differentiating privileges based on the user role (e.g., caregiver, patient, health care provider, system administrator) or device functions”.

- “Authenticate firmware and software. Verify authentication tags (e.g., signatures, message authentication codes (MACs)) of software/firmware content, version numbers, and other metadata. The version numbers intended to be installed should themselves be signed/have MACs. Devices should be electronically identifiable (e.g., model number, serial number) to authorized users”.

- “Ensure capability of secure data transfer to and from the device, and when appropriate, use methods for encryption and authentication of the end points with which data is being transferred”.

- “Use current NIST recommended standards for cryptography (e.g., FIPS 140-2, NIST[26] Suite B[27]), or equivalent-strength cryptographic protection for communications channels”.

- “Implement design features that allow for security compromises to be detected, recognized, logged, timed, and acted upon during normal use”.

- “Ensure the design enables forensic evidence capture. The design should include mechanisms to create and store log files for security events. . . . Examples of security events include but are not limited to configuration changes, network anomalies, login attempts, and anomalous traffic (e.g., sending requests to unknown entities)”.

- “Implement device features that protect critical functionality and data, even when the device’s cybersecurity has been compromised”.

Notice that these recommendations are more specific than the Quality System Regulations, yet remain intentionally open-ended so teams can choose to satisfy these recommendations based on their application’s operational requirements. For example, “layered authorization” is generally a good idea. So are “device features that protect critical functionality and data”. How these recommendations should be met, however, is not specified – the FDA leaves implementation details up to the developers.

Thus, these more-specific recommendations, like the more open-ended Quality System Regulations, leave applications vulnerable to poor implementation errors. However, it's critical to note the difference between the recommendations in the 2018 draft guidance and the Quality System Regulations: the Quality System Regulations is so open-ended that it is hard to operationalize. On the other hard, the recommendations in the 2018 draft guidance document are much more well-defined.

This specificity makes a world of difference. Previous research suggests that as long as security requirements are well-defined, we may be able to audit poor security implementations (e.g., see CryptoGuard [1]). Although these auditors will never be 100% accurate, they may be able to catch a significant number of security-related implementation errors while an application is still in development.

The 2018 draft guidance document "Content of Premarket Submissions for Management of Cybersecurity in Medical Devices" is enormously helpful. However, most cybersecurity regulations and recommendations are more similar in substance to the Quality System Regulations. Finally, I add that while the draft guidance was released on October 18, 2018, its 150-day comment period closed on March 17, 2019. This draft guidance has not yet been replaced by a new, official guidance document.

4.2.2 Current Regulations Rely on "Self-Certification"

Currently, the FDA primarily relies on firms to "self-certify" that their software devices are "secure enough". This collaborative system of "self-certification" makes it more difficult for the FDA to determine when an application fails to meet its security and privacy requirements. This is because a system of "self-certification" relies on trust; the FDA must trust that firms' self-evaluations are correct. However, given the large number of mHealth applications, it is inevitable that the FDA will end up trusting some inexperienced developers or bad actors that produce an unsatisfactory products. There's no way for the FDA to be completely sure of the security and privacy of every mHealth application without testing them all.

The FDA splits cybersecurity concerns into two general categories: (1) controlled risks, and (2) uncontrolled risks. Each firm must try to turn all "uncontrolled risks" into "controlled risks", explain how they did so, and argue that any remaining risks, controlled or not, are acceptable.

The determination of what counts as a "controlled risk" or an "uncontrolled risk" depends on the risk's exploitability and the severity of patient harm if exploited. To evaluate exploitability, the FDA suggests some tools like the "Common Vulnerability Scoring System Version 3.0", the "AAMI TIR57: Principles for Medical Device Security – Risk Management Framework", and the "IEC 80001: Application of Risk Management for IT Networks Incorporating Medical Devices Framework", though firms may use whatever system they please (including their own). The FDA's definitions for the severity of patient harm are more focused. Severity ranges include negligible (inconvenience or temporary discomfort), minor (results in temporary injury or impairment not requiring professional medical intervention), serious (results in injury or impairment requiring professional medical intervention), critical (results in permanent impairment or life-threatening injury), and catastrophic (results in patient death). A diagram helps explain these relationships [56].

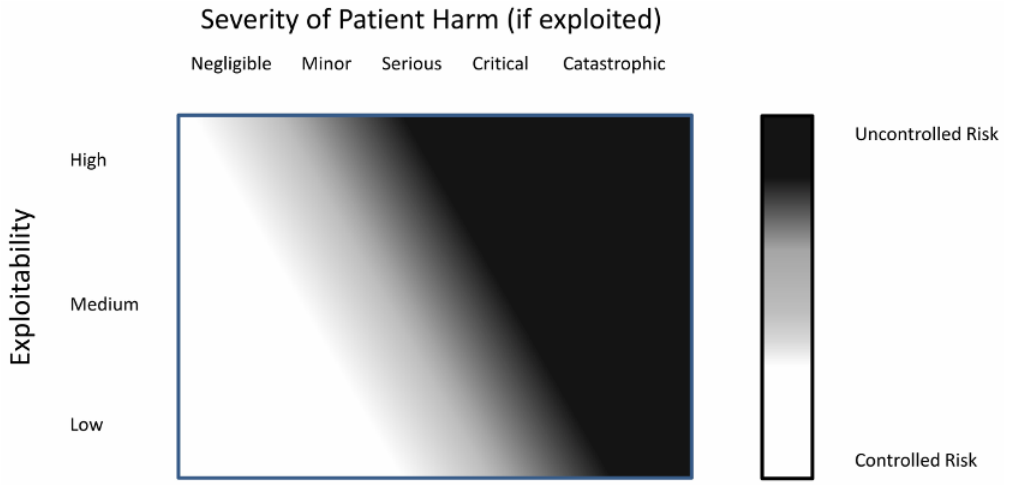

Figure 4.1: Controlled vs. Uncontrolled Cybersecurity Risks in mHealth.
Image from Postmarket Management of Cybersecurity in Medical Devices (2016) [56]

If a particular medical application requires premarket or postmarket submissions, firms must identify all cybersecurity risks and tell the FDA why these risks are all "controlled risks". This system, however, relies on firms' "self-evaluation". The FDA may claim that a firm's actions to remediate cybersecurity risks is inadequate, but for the most part, it appears that the FDA must trust that a firm's self-evaluation and implementation was done correctly.

I suspect this system of “self-evaluation” is in-place because it can be difficult for a third-party, even a third-party with governmental authority, to effectively assess a firm’s cybersecurity risk. This is especially acute because many of the recommendations the FDA gives are difficult for a third-party to assess.

For example, the FDA writes “manufactures may wish to deploy an additional control(s) as a part of a ‘defense-in-depth’ strategy” to manage “controlled risks” [56]. Defense-in-depth strategies, and their correct implementation, are left up to each firm. There are benefits to this approach, as it gives firms the flexibility to design a “defense-in-depth” strategy that is tailored to the needs of their application. This flexibility, though, makes it difficult for the FDA to audit the strength of all “defense-in-depth” strategies. There are many firms creating mHealth applications and each firm will have their own “defense-in-depth” strategy. It is unreasonable to expect the FDA to have the expertise or resources to evaluate each of these firm’s strategies.

Let’s look at another example. For newly discovered “uncontrolled risks”, the FDA recommends that “[if] fixing the vulnerability may not be feasible or immediately practicable, manufactures should identify and implement risk mitigations and compensating controls to adequately mitigate the risk” and recommends that “customers and the user community should be provided with relevant information on recommended controls and residual cybersecurity risks so that they can take appropriate steps to mitigate the risk and make informed decisions regarding device use” [56]. Some vulnerabilities may need to be reported to the FDA, though other vulnerabilities (even critical vulnerabilities) may not need to be reported [56]. Again, because the FDA gives firms great flexibility in how they choose to satisfy their cybersecurity requirements, the FDA may not be well placed to determine if fixing a new vulnerability is “feasible or immediately practicable” – firms may be better positioned to make this judgement. While this framework allows the people with the best understanding of an application (i.e., the firms) to determine the severity of a cybersecurity risk, it also requires that regulatory bodies trust these firms’ self-evaluation of their products. This trust may be well-placed for experienced teams; however, inexperienced teams may unintentionally make mistakes in their cybersecurity implementations.

The practice of relying on a system of “self-certification” is not new for government agencies. In fact, this system is pretty common, even in highly regulated industries like the defense industry. One can argue that it may not be appropriate or efficient for the FDA to check implementation details, like if all health firms' cloud configurations are "secure-enough". One can also argue that more prescriptive guidelines would stifle innovation and lead to less secure, suboptimal outcomes. Another argument can be made that open-ended requirements give expert teams the flexibility they need to design cybersecurity controls that suit an application's operational requirements. I'm not disagreeing with any of these arguments. Instead, I claim that open-ended regulations and a system of "self-certification" comes with inherent trade-offs. While allowing experienced teams to design their own cybersecurity solutions may be helpful, this same system gives inexperienced teams or bad actors the leeway to design inadequate solutions.

Finally, I note that there are also alternatives to this trust-based regulatory framework. For example, the California Air Resources Board adopts an antagonistic approach to vehicle emission regulatory compliance. Instead of a trust-based system of "self-certification", all vehicles sold in California must be tested at a facility to verify that their emissions satisfy California regulations. One can argue that this antagonistic approach may be more effective than a trust-based approach in ensuring that all vehicles meet their emission requirements. However, this antagonistic approach may require more resources (e.g., for the development of complicated testing procedures [22], then to fund these extensive tests). And finally, this antagonistic approach requires a set of criteria that is testable. Currently, the open-ended nature of mHealth security and privacy regulations makes these regulations difficult to test at scale.

4.2.3 It's Unclear What Regulations Apply

Finally, it's important to note that the current regulatory environment is very complicated. So complicated, in fact, that it can be hard to determine what regulations and what FDA guidance documents even apply to a product. Without knowing what regulations or recommendations to check for, we cannot automate a process to check for legal compliance.

mHealth application developers are subject to federal regulations in the U.S.C and the C.F.R., in addition to federal case law. Title 21 of the C.F.R. alone (the part of the C.F.R. devoted to rules for the FDA) was a whopping 4,907 pages in 2020 [127, 128, 129, 129, 130, 131, 132, 133, 134, 135]. Furthermore, mHealth app developers are also subject to state laws and state case law.

In addition to the vast quantity of regulations mHealth application developers must satisfy, the FDA's guidance documents present an additional challenge. While these guidance documents usually do not represent legal requirements, they are often important because they explain how the FDA will exercise its legal authority. When the law is ambiguous or leaves enforcement details up to the FDA, these guidance documents explain the FDA's position.

However, understanding these FDA guidance documents is itself a challenge. There are over 2,500 guidance documents currently active; each of these documents are searchable online [76]. Thankfully, all of the 200+ withdrawn guidance documents are also searchable via an FDA database [77]. However, it can be difficult to determine when a guidance document has been updated. One must manually investigate (e.g., through Google) to determine if a guidance document has been replaced with a newer version of the same name, and it's especially difficult to determine if a guidance document has been updated or merged into a newer document with a different name.

It is also difficult to determine what the FDA's policies are when multiple guidance documents apply to same application. Which guidance should the developer follow? As an example of the complexity of this environment, I've compiled a table that features a few of the most important regulations and FDA guidance documents, organized by category. Note how these documents often overlap.

Table 4.1: A Partial Listing of Important FDA Guidance Documents, by Subject

Legal	Software	Cybersecurity	Networks	Labeling	Organizational	Premarket Submissions	Reporting
Food, Drug, and Cosmetic Act (FD&C Act) [21 U.S.C. Chapter 9]	Policy for Device Software Functions and Mobile Medical Applications	Content of Premarket Submissions for Management of Cybersecurity in Medical Devices (2018)	Cybersecurity for Networked Medical Devices Containing Off-the-Shelf (OTS) Software	Labeling - Regulatory Requirements for Medical Devices (FDA 89-4203)	QS Regulations (21 C.F.R. Part 820)	Premarket Approval of Medical Device (21 C.F.R. 814)	Postmarket Management of Cybersecurity in Medical Devices
Food and Drugs (Title 21 of C.F.R.)	Guidance for the Content of Premarket Submissions for Software Contained in Medical Devices	Content of Premarket Submissions for Management of Cybersecurity in Medical Devices (2014)	Design Considerations and Premarket Submission Recommendations for Interoperable Medical Devices	Labeling (21 CFR 801)	HIPAA (Health Insurance Portability and Accountability Act)	Premarket Notification 510(k) (21 CFR Part 807 Subpart E)	Medical Device Reporting (21 CFR Part 803)
QS Regulations (21 C.F.R. 820)	Off-The-Shelf Software Use in Medical Devices	Cybersecurity for Networked Medical Devices Containing Off-the-Shelf (OTS) Software	Medical Device Data Systems, Medical Image Storage Devices, and Medical Image Communications Devices	In Vitro Diagnostic Products for Human Use - Labeling (21 CFR 809 Subpart B)	Breakthrough Devices Program	Guidance for the Content of Premarket Submissions for Software Contained in Medical Devices	Medical Devices; Reports of Corrections and Removals (21 CFR 806)
HIPAA (Health Insurance Portability and Accountability Act)	Medical Device Data Systems, Medical Image Storage Devices, and Medical Image Communications Devices	Postmarket Management of Cybersecurity in Medical Devices Radio Frequency Wireless Technology in Medical Devices	Radio Frequency Wireless Technology in Medical Devices	Unique Device Identification (21 CFR 830)	Developing the Software Precertification Program: Summary of Learnings and Ongoing Activities	Content of Premarket Submissions for Management of Cybersecurity in Medical Devices (2018)	QS Regulations (21 C.F.R. 820)
Labeling (21 C.F.R. 801)	Clinical Decision Support Software	Multiple Function Device Products: Policy and Considerations		Labeling of Investigational Devices (21 CFR § 812.5)	De Novo Classification Process (Evaluation of Automatic Class III Designation)	Content of Premarket Submissions for Management of Cybersecurity in Medical Devices (2014)	Medical Device Reporting for Manufacturers
Establishment Registation (21 CFR Part 807)	Changes to Existing Medical Software Policies Resulting from Section 3060 of the 21st Century Cures Act			QS Regulations - Labeling and Package Control (21 CFR 820 Subpart K)	Humanitarian Device Exemption (HDE) Program	Design Considerations and Premarket Submission Recommendations for Interoperable Medical Devices	Postmarket Surveillance Under Section 522 of the Federal Food, Drug, and Cosmetic Act
Medical Device Listing (21 CFR Part 807)	General Wellness: Policy for Low Risk Devices			General Electronic Products - Identification (21 CFR § 1010.3)	Modifications to Devices Subject to Premarket Approval (PMA) - The PMA Supplement Decision-Making Process	Appropriate Use of Voluntary Consensus Standards in Premarket Submissions for Medical Devices	Distinguishing Medical Device Recalls from Medical Device Enhancements
Premarket Approval of Medical Device (21 C.F.R. 814)	Multiple Function Device Products: Policy and Considerations			Multiple Function Device Products: Policy and Considerations	Deciding When to Submit a 510(k) for a Change to an Existing Device	Developing the Software Precertification Program: Summary of Learnings and Ongoing Activities	Deciding When to Submit a 510(k) for a Change to an Existing Device
Premarket Notification 510(k) (21 CFR Part 807 Subpart E)	General Principles of Software Validation				FDA and Industry Procedures for Section 513(g) Requests for Information under the Federal Food, Drug, and Cosmetic Act	De Novo Classification Process (Evaluation of Automatic Class III Designation)	Modifications to Devices Subject to Premarket Approval (PMA) - The PMA Supplement Decision-Making Process
Investigational Device Exemption (IDE) (21 CFR Part 812)	Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan				Requests for Feedback and Meetings for Medical Device Submissions: The Q-Submission Program	FDA and Industry Procedures for Section 513(g) Requests for Information under the Federal Food, Drug, and Cosmetic Act	Developing the Software Precertification Program: Summary of Learnings and Ongoing Activities
Medical Device Reporting (21 CFR Part 803)	Developing the Software Precertification Program: Summary of Learnings and Ongoing Activities				Applying Human Factors and Usability Engineering to Medical Devices		
Medical Devices; Reports of Corrections and Removals (21 CFR 806)					General Principles of Software Validation		

Another challenge is that the health regulatory environment is very convoluted. While scholars are unsure about how to study legal complexity, one proposed method is to measure how many other documents each legal unit cites [26, 25]. I utilize this method to study the out-bound citation network for key mHealth FDA guidance documents.

These citation networks are critical. FDA guidance documents, laws, and other policies do not exist in a vacuum. Oftentimes, these documents rely upon principles and regulations from other documents. Therefore, to understand a document, one may need to be familiar with its citations. However, my investigation suggests that the complexity of FDA guidance documents and the many other documents they cite may weaken the explanatory purpose of these guidance documents.

As a part of my investigation, I created a graphic that maps the out-bound citations for eleven key FDA mHealth guidance documents. These eleven original documents ultimately reference 133 other documents a total of 239 times. On average, each of the original FDA guidance documents cites 21.72 other documents.

The eleven original FDA guidance documents are: “Policy for Device Software Functions and Mobile Medical Applications”, “Content of Premarket Submissions for Management of Cybersecurity in Medical Devices (2018)”, “Content of Premarket Submissions for Management of Cybersecurity in Medical Devices (2014)”, “Guidance for the Content of Premarket Submissions for Software Contained in Medical Devices”, “Off-The-Shelf Software Use in Medical Devices”, “Cybersecurity for Networked Medical Devices Containing Off-the-Shelf (OTS) Software”, “Postmarket Management of Cybersecurity in Medical Devices”, “General Principles of Software Validation”, “Medical Device Reporting for Manufacturers”, “Distinguishing Medical Device Recalls from Medical Device Enhancements”, and “Design Considerations and Premarket Submission Recommendations for Interoperable Medical Devices” [66, 60, 51, 49, 65, 48, 56, 47, 55, 52, 58].

This graphic shows the distributed and convoluted nature of the regulatory landscape. To understand these eleven FDA Guidance Documents, one may need to be familiar with over 200 other documents. All of these other documents (especially the legal documents and other FDA Guidance Documents) have their own dependencies. The complexity of this citation network grows exponentially.

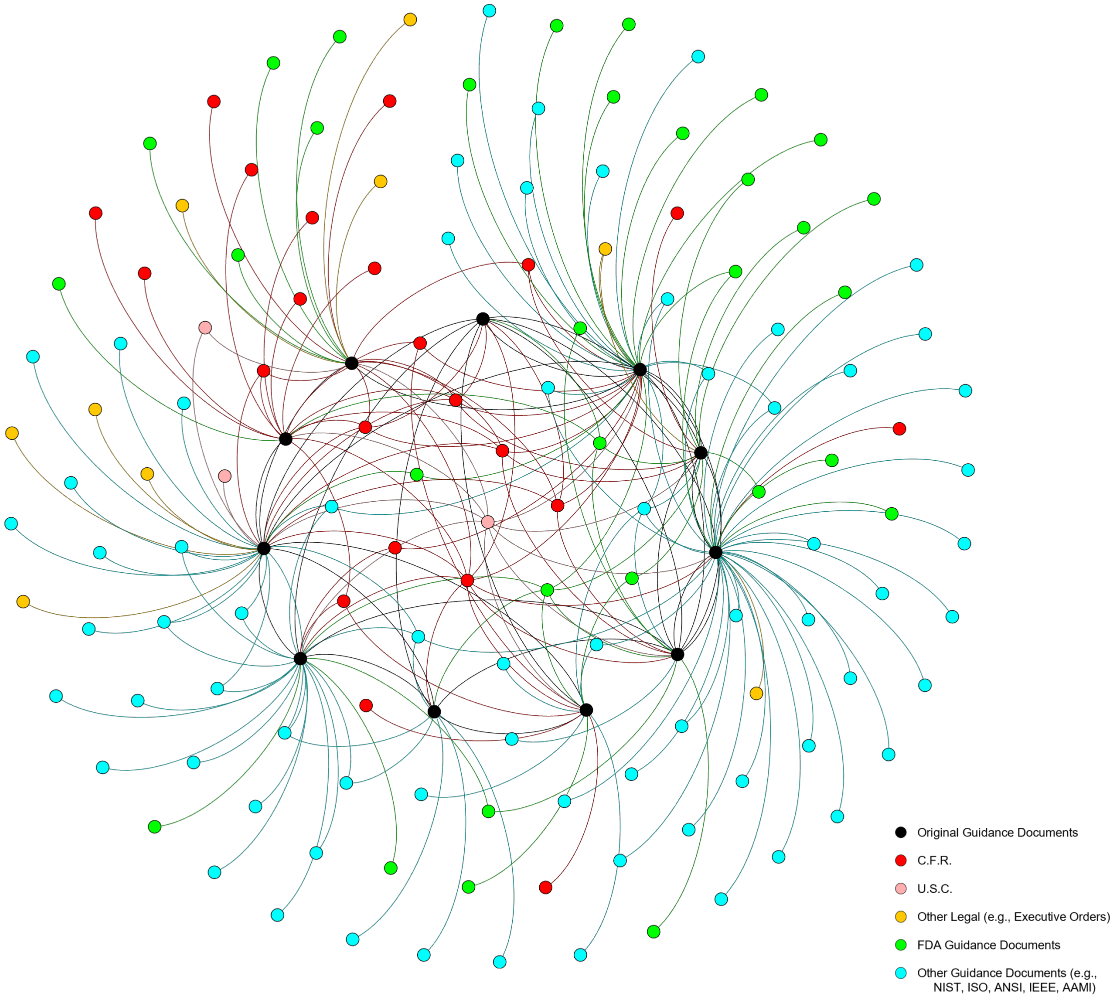
Original Guidance Documents
C.F.R.
U.S.C.
Other Legal (e.g., Executive Orders)
FDA Guidance Documents
Other Guidance Documents (e.g., NIST, ISO, ANSI, IEEE, AAMI)

Figure 4.2: Out-Bound Citations for Eleven FDA Guidance Documents

4.3 The Use of Non-legal, Authoritative Guidance Documents That Explain Regulatory Intent May Aid in the Creation of Automated Auditors

The previous sections describe some key challenges that make it difficult to computationally audit mHealth security and privacy regulatory compliance: the open-ended nature of the law, the FDA's reliance upon "self-certification", and the complicated regulatory environment. In previous sections, I speculate why our regulatory environment was designed this way and explain some of the pros and cons of our current approach.

While our current regulatory approach has many benefits, I do note that its current design makes it hard to audit application compliance. The open-ended nature of the law is an especially large challenge, as auditors (especially computational auditors that can scale to millions of apps) require specific, technical specifications to check against. If a regulation can be satisfied in thousands of different ways, there is no good way for a computational auditor to check every possible compliance strategy. There are simply too many. An auditor that tries to cover many regulations, each of which have thousands of compliance strategies, may be too difficult to create. Furthermore, even if we could create such an auditor, it may suffer from too many false-positives or false-negatives.

Because of these challenges, I wonder if it would be better to create the auditor in another way – can we use non-legal, authoritative guidance documents that explain regulatory intent to create an auditor? This solution may not be intuitive; the conventional wisdom is to operationalize the law itself if we are trying to create an auditor that checks for compliance with the law. It seems counter-intuitive to willingly set aside the law and instead rely upon non-legal documents.

But, the use of non-legal documents brings important benefits. Because these documents do not carry the full weight of the law, authoritative bodies (e.g., the FDA, NIST) can give more-specific recommendations while preserving the open-ended nature of the law. Furthermore, because these recommendations are not mandatory, they can be specific-enough to aid inexperienced developers, yet not prescriptive, giving experienced developers the freedom to design their own security solutions. And, a comprehensive guidance document can also help clarify the current regulatory framework. If this document distills many convoluted rules into list of actionable insights, developers may not feel the need to parse through other regulations and their dependency networks.

It's also important to note that these documents must be authoritative. Not all interpretations of regulatory intent may satisfy the government. The good news, however, is that many authoritative bodies may have the legal sophistication to create these guidance documents. These documents can come from the government itself (e.g., the FDA), but can also come from other authoritative bodies (like industry groups, research organizations, or non-profits). Technologists do not need to wait for the government to make an "official" document that clarifies regulatory intent. This operationalization can come from many places, as long as their legal interpretation is trustworthy.

Related work in the Security and Privacy Laboratory at Washington University in St. Louis suggests that the use of these non-legal, authoritative guidance documents that clarify regulatory intent is a promising way to create automated, computational auditors. This work uses the 2018 FDA draft guidance document "Content of Premarket Submissions for Management of Cybersecurity in Medical Devices" as the non-legal, authoritative guidance document that is operationalized into an auditor [60]. The recommendations in this draft guidance document are specific-enough such that we were able to create an auditor that uses static analysis techniques to detect potential violations of FDA security and privacy recommendations.

While this auditor did not include all FDA security and privacy recommendations for mHealth applications, the auditor's creation does show the potential of using non-legal, authoritative guidance documents as the basis of computational auditors. Guidance documents, because they are non-legal, can be written to arbitrary degrees of specificity. As long as these guidance documents are specific-enough to computationally operationalize, we can use them to create an auditor that checks if applications satisfy their legal requirements, all while preserving the open-ended nature of the law.

However, it's important to note that by using non-legal, authoritative guidance documents that clarify regulatory intent to create auditors, we are not teaching computers to understand and reason about the law. The hard problem of legal interpretation (going from law to operationalizable items) must still be done by experts. This proposal suggests that human interpretations can be operationalized for computers, which can potentially enable the creation of automated auditors. Getting computers to reason about the law remains an open question.

Chapter 5

Related Work

We can organize other work that combines computer science and legal analysis into a few categories.

5.1 Artificial Intelligence, Machine Learning, and Natural Language Processing

Many projects that intersect legal academia and the computer science security and privacy community have focused on the analysis of privacy policies. Many projects, like the Usable Privacy Project, seek to use natural language processing and crowdsourcing to semi-autonomously extract key provisions from privacy policies and organize them in a user-friendly way [142, 141].

Bhatia et al. (2016) use Tregex, a tree regular expression language, to identify hyponyms (i.e., specific examples) present in privacy policies. They found that the vocabulary used in privacy policies is very different than the vocabulary used in popular lexical databases [21]. Kumar et al. (2020) use supervised machine learning to automatically extract opt-out hyperlinks from privacy policies [18]. Zimmeck et al. (2019) use supervised machine learning and natural language processing to automatically extract privacy practice descriptions from Android apps. Static analysis of apps' Android APIs, strings, class structures, and permissions checks if data (e.g., location data accessed through an API) is used without being disclosed in a privacy policy [171].

Other research outside of the computer science security and privacy community focuses on legal information retrieval and prediction [80, 170, 20]. One example of legal information retrieval involves search engines – given a legal query, can a search engine find relevant legal documents (e.g., judicial decisions, legal statutes). Participants in the Competition on Legal Information Extraction/Entailment (COLIEE) have used a variety of techniques to improve legal information retrieval. Kim and Goebel (2017) use TF-IDF and language model-based information retrieval [101]. Tran, Nguyen, and Satoh (2019) apply encoded summarization and lexical matching [156]. Shao et al. (2020) use word-entity duet representations [147]. And Gain et al. (2021) use BM25 and Bidirectional Encoder Representations from Transformers (BERT) [81]. Information retrieval in the legal domain is different than general information retrieval; the many differences between legal and general language may complicate the application of traditional NLP techniques to the legal domain [168, 170, 108].

Another area of research is using computation to predict the outcome of judicial decisions. Legal judgement prediction has long been of interest of the legal community [102, 116]. Recently, researchers have applied AI and NLP to improve predictive results. Katz, Bommarito, and Blackman (2017) use random forests to predict US Supreme Court voting behavior with 70% accuracy [99]. Aletras et al. (2016) use a support vector machine to predict judicial decisions from the European Court of Human Rights with 80% accuracy [9]. Wang and Jin (2020) combine convolutional neural networks and recurrent neural networks to predict judicial decisions from the Supreme People's Court of China [163].

5.2 Logical Frameworks, Privacy Languages, and Legal Ontologies

One branch of research tries to take legal documents and precisely represent them with knowledge representation languages. Some researchers attempt to represent privacy policies in logical frameworks. Breaux and Rao (2013) use Description Logic to model actors, data, and data use purpose hierarchies in Facebook, Zynga, and AOL-Advertising privacy policies. Their investigation is motivated by the fact that digital products (e.g., Zynga's Farmville game) often depend on many technological actors (e.g., Farmville was playable on Facebook, where ads were then served by AOL). However, the privacy policies of Zynga, Facebook, and AOL-Advertising may contain contradictory requirements. [30].

Instead of presenting privacy policies as knowledge representation languages, other researchers have developed new languages to enforce privacy and security requirements (e.g., access control frameworks). EPAL includes support for hierarchies, obligations, and conditions a firm may specify in its privacy policy [16]. XACML is another attribute-based access control system, though it lacks adequate support for the temporal conditions often found in privacy regulations [19, 40]. Other privacy languages and proposals (like P3P and Do-Not-Track) can be too domain-specific and therefore cannot effectively be applied to the legal system [39, 124, 40]. Researchers affiliated with CodeX, the Stanford Center for Legal Informatics, have found more success applying structured programming languages and formal logic to contract law, as the domain-specific qualities of contract law are more amenable to computation [155, 83].

This domain-specific approach is not unique; researchers who study how to apply logical programming to specific subsections of the law (instead of privacy policies) often focus on very narrow applications. In the 1980s, Sergot et al. modeled the British Nationality Act of 1981 in the Prolog logic language; one of their conclusions was that the vagueness of the law precluded its nuanced analysis [146]. More recently in 2005, Hilty, Basin, and Pretschner use distributed temporal logic to model future access controls, an idea often found in security and privacy regulation [92]. In 2010, DeYoung et al. use least fixed point logic to create a logical framework called PrivacyLFP to formalize the transmission-related portions of the Health Insurance Portability and Accountability Act (HIPAA) and the Gramm-Leach-Bliley Act (GLBA). PrivacyLFP includes support for disclosure purposes, real-time constructs, and self-reference via fixed points. The authors argue that their logical specification reveals that the law cannot be completely mechanized; they find that a large portion (e.g., almost all HIPAA clauses) are "inherently subjective", and well-over-half of the clauses require human involvement (often expert human involvement) at the time of data access or audit [40].

Perhaps because of this subjectivity, since the 1990s much research in applying logic to the law has been focused on ontologies [20]. While ontologies have a historical meaning in philosophy and other related fields, ontologies in logical legal research follows the AI definition – the explicit specification of a conceptualization (i.e., model) of the legal domain.

This includes explicit definitions of the logic, structure, and relationships in our legal system and its many sub-domains [159]. Ontologies are critical in AI, and provide a shared vocabulary and allow for the creation and development of domain-specific knowledge [159, 88]. Without these explicit definitions and shared vocabulary, it is challenging to "translate" legal concepts and statues into a logical framework. I encountered this problem in my own research; the open-ended nature of the law complicates its translation into a well-defined, computer-friendly format.

The creation of legal ontologies remains an active, unsettled, and pressing research area. However, research in creating legal ontologies is very difficult. From a measurement perspective, no one fully knows what the law looks like or how it is used. In fact, the U.S. Justice Department once tried to count the number of federal criminal laws in the U.S. Code; they were unsuccessful. Current estimates peg the number from 10,000 to 300,000 [79]. Others have tried to understand the structure of the legal system by analyzing the hierarchical structure, citation network, associated text function, and between-word Shannon Entropy of legal collections like the U.S. Code [26]. Others have measured when key words (e.g., "privacy") entered legal documents, page length, paragraph count, and word count [161, 112, 32, 107]. However, while these studies can help us better understand our legal structure and the effects laws have in our society, these discoveries are also reminders about how little we currently know about the legal system. For example, Liebman et al. (2017) reviewed tens of millions of historical judicial judgements from China's Henan province and in a random sample found that roughly half of the administrative lawsuits in their corpus reflect an underlying civil dispute. This directly defies conventional wisdom, as we expect civil disputes to be resolved through civil lawsuits [108]. It's hard to create a model and ontology of the legal system when we're currently unclear about how exactly the legal system is used and structured.

Perhaps an even greater challenge to creating legal ontologies is that there are currently many philosophically unknown questions about the fundamental nature of the law. For example, those who have been able to accurately predict judicial decisions have used this predictive power to argue that the law in practice is more similar to legal realism (in which social interests and public policy influence the outcome of a case) than legal formalism (in which the rules of the law uniquely determine a judicial outcome) [9]. This question of how laws are interpreted and applied remains an unanswered question in legal philosophy; it may never be answered. I previously mentioned that there are other key, unanswered legal questions that are fundamental to our legal system (e.g., federalism). The philosophical and political realities of the law prevent its simple classification. I, too, note this difficulty in my research in mHealth: it's hard to know what and how regulations apply when the relationship among legal documents is convoluted and imprecise. Some modeling techniques, like mapping the outbound citation network, may give structure to some relevant documents; however, the relationship among these documents remain unclear.

5.3 Modeling Legal and Philosophical Ideas

Other work, instead of directly trying to understand the law, focuses on modeling legal and philosophical ideas. The law is ultimately grounded in philosophy (e.g., our conceptions of what is right and wrong, our definition of privacy, our beliefs about the role of the state in private affairs) and the trade-offs inherent in the political decision-making process. Modeling these larger philosophical ideas is like creating a philosophical superset; a subset of these ideas comprises our current laws. Furthermore, once this subset of ideas has been identified, we can try to formalize these ideas into a computation-friendly representation. This kind of meta-analysis, therefore, is related to the creation of an ontology used to study computation and the law.

One example of this kind of philosophical inquiry is Nissenbaum's theory of Contextual Integrity [117]. Contextual Integrity builds upon the rich, historical, social-scientific study of the "spheres" – or the distinct roles, places, and fields that comprise our society [85, 27, 28, 162, 111]. Nissenbaum identifies two norms that categorize our expectations of privacy (i.e., the "norms of appropriateness" and the "norms of flow or distribution") and argues that we feel that our privacy has been violated when one of these norms have been transgressed. These norms are specific to each "field" we inhabit, and therefore depend on the context of the transmission, information included in the transmission, the agents who are sending and receiving the information, and other terms placed on the transmission. This theory is descriptive and focuses on when people feel their privacy has been violated through information transmission (n.b. its well-defined scope).

The development of such a theory allows for its formalization in a logical context. Barth et al. (2006) formalize Contextual Integrity to focus on who the information is about, how the information is transmitted, and what past and future actions are allowed and disallowed by both the subject and the receivers of the information. They then apply linear temporal logic (LTL) to this more-concrete definition of Contextual Integrity to model some information transmission restrictions present in the Health Insurance Portability and Accountability Act (HIPAA), the Children's Online Privacy Protection Act (COPPA), and the Gramm-Leach-Bliley Act (GLBA). Unlike privacy languages that represent transmission rules using functions, their formalization of Contextual Integrity represents transmission rules as logical formulas. Their representation of these transmission restrictions in a logical framework instead of as functions means that their logical implementation is fundamentally different than the approach taken by privacy languages; their research provides a summary of the different capabilities and limitations of different privacy languages (RBAC, XACML, EPAL, P3P) and their linear temporal logic formalization of Contextual Integrity [19]. However, this formalization, while impressive, remains incomplete. For example, it is limited to information transmission, does not cover group actors, and does not touch upon case law. Some of these limitations are due to the inherent restrictions of the Contextual Integrity framework.

It's important to mention that Contextual Integrity has been formalized in other ways as well. Apthorpe, Varghese, and Feamster (2019) interpret Contextual Integrity as comprising of five parameters: "(1) the subject of the information being transferred, (2) the sender of this information, (3) the attribute or type of information, (4) the recipient of the information, and (5) the transmission principle or condition imposed on the transfer of information from the sender to the recipient". They then surveyed 200 parents with children from ages 3-13 and examine if the privacy transmission restrictions in COPPA for IoT smart-toys aligned with parental expectations. Importantly, to formalize COPPA's transmission principles, they used the Federal Trade Commission's (FTC's) Six Step Compliance Plan for COPPA [10]. The FTC's formalization of COPPA is important, as it gives an authoritative summary of COPPA, which by itself may be very complicated and poorly understood [37]. This FTC explanation is similar to the FDA's guidance documents – they are non-legal, authoritative documents that complement formal law.

Others have sought to model legal and philosophical ideas in other ways. Garg, Goldwasser, and Vasudevan (2020), in response to "right to be forgotten" laws, ask the question "what does it truly mean to delete data". They model data collectors as Interactive Turing Machines, define what qualifies as an insert, lookup, and deletion, and theorize as to when a data collector is statistically deletion-compliant [82].

Cohen and Nissim (2020) attempt to formalize the EU's General Data Protection Regulation (GDPR)'s concept of "singling out". While protection from "singling out" is required by the GDPR, what this legal requirement means is much less clear. Cohen and Nissim attempt to capture the intent of this regulation mathematically by defining a privacy attack called "predicate singling out". A design that is secure against "predicate singling out" (PSO secure) is a necessary (but perhaps not sufficient) condition to be compliant with the GDPR's protections from "singling out". Finally, they show that k-anonymity is not PSO secure, but differential privacy implies PSO security [36]. Because the GDPR's definition of "singling out" was insufficiently detailed, they had to model what they believed was the intention behind this rule to represent this rule in a mathematical framework.

On the case law side, Sheffler and Varia (2021) explain the foregone conclusion doctrine, a legal doctrine in the United States center to the unresolved question about the limits of government-compelled decryption and disclosure. They create foregone conclusion games to explicitly define their definition, then walk through a few cryptographic systems (e.g., one-way functions, symmetric encryption) and show in what situations the government can and cannot compel the defendant to decrypt her data [143].

In all of these examples, researchers had to create new models of legal concepts. Sometimes these models were supersets of legal rules (e.g., Contextual Integrity for information transmission), while other models were a smaller refinement of a larger concept (e.g., "predicate singling out" is a necessary but not sufficient condition for "singling out"). In all of these examples, the original legal idea was insufficient in some way for precise analysis (e.g., too vague, unsettled philosophical question). In these situations, researchers need to extrapolate what they believe to be the intent behind the law and model it in a computation-friendly way.

Chapter 6

Conclusion

In this thesis, I asked the two following questions: (1) can we teach computers to understand the law, and (2) can we automate checking for legal compliance? To investigate these questions, I pursued a case-study in mHealth security and privacy regulations.

To show the benefits of applying computation to legal analysis, I argue that by automating the auditing of mHealth security and privacy regulations, we can ultimately encourage equitable mHealth adoption and reduce inequalities in healthcare.

To make this argument, I introduced a theoretical model of technology adoption. In this model, technology adoption is directly affected by one's trust in the technology, technological momentum, market conditions, technological access, and other factors. One's trust in technology is ultimately built through two pathways, each of which require two necessary conditions. These two pathways are one's trust in oneself (e.g., through first-hand experiences) and through one's network of trusted intermediaries (e.g., government entities, authority figures, journalists, friends, social groups). The two necessary conditions for trust are the technology's theoretical effectiveness (i.e., the ability for an application or technology to theoretically meet the demands of its users) and practical effectiveness (i.e., applications must satisfy user demands in practice, not only in theory).

I then examined if it is possible to automate compliance checking given existing mHealth security and privacy regulations. Because of the breadth of these regulations, I limit my investigation in scope to federal rules and regulations and exclude case law. As a part of this investigation, I give a comprehensive summary of the current federal mHealth regulatory environment.

I conclude that it is very challenging to automatically check if mHealth applications satisfy current security and privacy regulations. This conclusion relies upon three observations. The first observation is that current regulations are too open-ended to be meaningfully operationalized, which prevents the translation of these regulations into a precise, machine-readable format. My second observation is that regulations currently rely too much on "self-certification", which makes objective measurements of an mHealth app's security and privacy difficult. The FDA does not ask for the detailed information required to verify satisfactory implementations of security and privacy controls. My third observation is that it's unclear what federal regulations or recommendations even apply to one's mHealth application. These regulations and recommendations often overlap and rely on other documents, which quickly snowballs into an unmanageable network of legal documents one must know and satisfy.

However, I note that the use of non-legal, authoritative guidance documents that explain regulatory intent may help with the creation of computational auditors. Because these documents do not carry the full weight of the law, they can be arbitrarily specific, all the while allowing the law to remain open-ended. Furthermore, because these documents are not mandatory, they can be specific-enough to help inexperienced developers, while not being prescriptive, giving experienced developers the freedom to design their own security solutions. And finally, if these documents distill many convoluted rules into actionable insights, developers may not feel the need to parse through other regulations and their dependency networks.

Related work in the realm of computational legal analysis generally focuses on: (1) utilizing artificial intelligence, machine learning, and natural language processing to identify key statements, improve legal search, and predict the outcome of judicial decisions; (2) the creation of logical frameworks, privacy languages, and legal ontologies to represent and communicate legal ideas; and (3) modeling legal and philosophical ideas by extrapolating the intent or philosophical basis behind laws, and translating this extrapolation into a computationally-friendly representation.

The challenges present in conducting computational legal research suggests that future work should focus on conducting basic research in legal computation. Because significant basic questions about the structure and usage of the law remains a mystery, it is hard to model the law in a computation-friendly format; it is hard to model what we do not know. The lack of computational legal ontologies makes it difficult to communicate in this research area and slows the development of incremental knowledge. Open-ended legal statements, legal contradictions, and unsettled areas of the law (e.g., federalism, evolving case law) present additional challenges. Computational legal analysis is still a field in its infancy, and basic research is needed to establish the buildings blocks that must undergird future developments.

References

[1] cryptoguard: High precision detection of cryptographic vulnerabilities in massive-sized java projects.

[2] Kordel v. United States, 335 U.S. 345 (1948).

[3] Labeling, 21 C.F.R. 801.

[4] Quality Systems Regulation, 21 C.F.R. 820.

[5] United States v. Research Laboratories, 126 F.2d 42 (9th Cir. 1942).

[6] United States v. Wong Kim Ark, 169 U.S. 649, 18 S. Ct. 456, 42 L. Ed. 890 (1898). Available at: `https://www.law.cornell.edu/supremecourt/text/169/649`.

[7] Central Intelligence Agency. Legal System — CIA - The World Factbook. `https://www.cia.gov/the-world-factbook/field/legal-system/`. Accessed on: July 18, 2021.

[8] Tanvir Ahmed, Syed Jafar Raza Rizvi, Sabrina Rasheed, Mohammad Iqbal, Abbas Bhuiya, Hilary Standing, Gerald Bloom, and Linda Waldman. Digital Health and Inequalities in Access to Health Services in Bangladesh: Mixed Methods Study. *JMIR mHealth and uHealth*, 8(7):e16473, 2020. Available at: `https://mhealth.jmir.org/2020/7/e16473/`.

[9] Nikolaos Aletras, Dimitrios Tsarapatsanis, Daniel Preoţiuc-Pietro, and Vasileios Lampos. Predicting judicial decisions of the European Court of Human Rights: A natural language processing perspective. *PeerJ Computer Science*, 2:e93, 2016. Available at: `https://peerj.com/articles/cs-93/`.

[10] Noah Apthorpe, Sarah Varghese, and Nick Feamster. Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus {COPPA}. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, pages 123–140, 2019. Available at: `https://www.usenix.org/system/files/sec19-apthorpe.pdf`.

[11] Samuel Arbesman. Measuring the Complexity of the Law — WIRED. `https://www.wired.com/2014/06/scienceblogs0602law/`, 06 2014.

[12] National Archives. Code of Federal Regulations Total Pages 1936 - 1949, And Total Volumes and Pages 1950 - 2019. https://uploads.federalregister.gov/uploads/2020/04/01123111/cfrTotalPages2019.pdf, 2019.

[13] Pieter Arntz. Vastaamo psychotherapy data breach sees the most vulnerable victims extorted — Malwarebytes Labs. https://blog.malwarebytes.com/cybercrime/2020/10/vastaamo-psychotherapy-data-breach-sees-the-most-vulnerable-victims-extorted/, October 2020.

[14] American Medical Association. Reducing Disparities in Health Care. https://www.ama-assn.org/delivering-care/patient-support-advocacy/reducing-disparities-health-care, 2018.

[15] American Medical Student Association. Global Health Equity. https://www.amsa.org/about/mission-aspirations/global-health-equity/.

[16] Michael Backes, Birgit Pfitzmann, and Matthias Schunter. A Toolkit for Managing Enterprise Privacy Policies. In *European Symposium on Research in Computer Security*, pages 162–180. Springer, 2003. Available at: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.4.3700&rep=rep1&type=pdf.

[17] The World Bank. Key Features of Common Law or Civil Law Systems. https://ppp.worldbank.org/public-private-partnership/legislation-regulation/framework-assessment/legal-systems/common-vs-civil-law, May 2021.

[18] Vinayshekhar Bannihatti Kumar, Roger Iyengar, Namita Nisal, Yuanyuan Feng, Hana Habib, Peter Story, Sushain Cherivirala, Margaret Hagan, Lorrie Cranor, Shomir Wilson, et al. Finding a Choice in a Haystack: Automatic Extraction of Opt-Out Statements from Privacy Policy Text. In *Proceedings of The Web Conference 2020*, pages 1943–1954, 2020. Available at: https://www.usableprivacy.org/static/files/kumar_iyengar_www_2020.pdf.

[19] Adam Barth, Anupam Datta, John C Mitchell, and Helen Nissenbaum. Privacy and Contextual Integrity: Framework and Applications. In *2006 IEEE symposium on security and privacy (S&P'06)*, pages 15–pp. IEEE, 2006. Available at: https://theory.stanford.edu/people/jcm/papers/barth-datta-mitchell-nissenbaum-2006.pdf.

[20] Trevor Bench-Capon, Michał Araszkiewicz, Kevin Ashley, Katie Atkinson, Floris Bex, Filipe Borges, Daniele Bourcier, Paul Bourgine, Jack G Conrad, Enrico Francesconi, et al. A History of AI and Law in 50 Papers: 25 Years of the International Conference on AI and Law. *Artificial Intelligence and Law*, 20(3):215–319, 2012. Available at: https://link.springer.com/article/10.1007/s10506-012-9131-x#citeas.

[21] Jaspreet Bhatia, Morgan C Evans, Sudarshan Wadkar, and Travis D Breaux. Automated Extraction of Regulated Information Types Using Hyponymy Relations. In *2016 IEEE 24th International Requirements Engineering Conference Workshops (REW)*, pages 19–25. IEEE, 2016. Available at: `https://www.cs.cmu.edu/~jbhatia/papers/jbhatia_aire2016.pdf`.

[22] California Air Resources Board. Low-Emission Vehicle Regulations & Test Procedures. `https://ww2.arb.ca.gov/our-work/programs/advanced-clean-cars-program/lev-program/low-emission-vehicle-regulations-test`.

[23] Giovanni Bognetti, Matthew F. Shugart, David Fellman, and C. Neal Tate. Constitutional Law - Unitary and Federal Systems — Encyclopædia Britannica. `https://www.britannica.com/topic/constitutional-law/Unitary-and-federal-systems`.

[24] Nadine Bol, Natali Helberger, and Julia CM Weert. Differences in mobile health app use: a source of new digital inequalities? *The Information Society*, 34(3):183–193, 2018. Available at: `https://pure.uva.nl/ws/files/25256057/Differences_in_mobile_health_app_use.pdf`.

[25] Michael James Bommarito and Daniel Martin Katz. Properties of the United States code citation network. *Available at SSRN 1502927*, 2009. Available at: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1502927`.

[26] Michael J Bommarito II and Daniel M Katz. A Mathematical Approach to the Study of the United States Code. *Physica A: Statistical Mechanics and its Applications*, 389(19):4195–4200, 2010. Available at: `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1578094`.

[27] P. Bourdieu and R. Nice. *Distinction: A Social Critique of the Judgement of Taste*. Polity Short Introductions. Harvard University Press, 1984.

[28] Pierre Bourdieu. The Social Space and the Genesis of Groups. *Theory and society*, 14(6):723–744, 1985. Available at: `https://link.springer.com/content/pdf/10.1007/BF00174048.pdf`.

[29] Ted Brackemyre. America's First Failure at Government — US History Scene. `https://ushistoryscene.com/article/articles-of-confederation/`, November 2020.

[30] Travis D Breaux and Ashwini Rao. Formal Analysis of Privacy Requirements Specifications for Multi-Tier Applications. In *2013 21st IEEE International Requirements Engineering Conference (RE)*, pages 14–23. IEEE, 2013. Available at: `https://www.cs.cmu.edu/~agrao/paper/Analysis_of_Privacy_Requirements_Facebook_Google_Zynga.pdf`.

[31] Jennifer Bresnick. 50% of Execs Think mHealth will be Key to Patient Engagement. https://ehrintelligence.com/news/50-of-execs-think-mhealth-will-be-key-to-patient-engagement, February 2015.

[32] Katelyn Brown. Tax Code Is So Long That Nobody's Really Sure of Its Length — POLITIFACT. https://www.politifact.com/factchecks/2017/oct/17/roy-blunt/tax-code-so-long-nobodys-really-sure-its-length/, October 2017.

[33] Jeanine Cali. Frequent Reference Question: How Many Federal Laws Are There? — In Custodia Legis: Law Librarians of Congress. https://blogs.loc.gov/law/2013/03/frequent-reference-question-how-many-federal-laws-are-there/, March 2012.

[34] National Constitution Center. Brown v. Board: When the Supreme Court Ruled Against Segregation. https://constitutioncenter.org/blog/on-this-day-the-supreme-court-rules-against-segregation, May 2021.

[35] Arul Chib. The Promise and Peril of mHealth in Developing Countries. *Mobile Media & Communication*, 1(1):69–75, 2013. Available at: https://journals.sagepub.com/doi/full/10.1177/2050157912459502.

[36] Aloni Cohen and Kobbi Nissim. Towards Formalizing the GDPR's Notion of Singling Out. *Proceedings of the National Academy of Sciences*, 117(15):8344–8352, 2020. Available at: https://www.pnas.org/content/pnas/117/15/8344.full.pdf.

[37] Federal Trade Commission. Children's Online Privacy Protection Rule: A Six-Step Compliance Plan for Your Business. https://www.ftc.gov/tips-advice/business-center/guidance/childrens-online-privacy-protection-rule-six-step-compliance, June 2017.

[38] U.S. Congress. United States Code: Table of Contents (1934). https://www.loc.gov/item/uscode1934-001000002/, 1934.

[39] Lorrie Faith Cranor. P3P: Making Privacy Policies More Useful. *IEEE Security & Privacy*, 1(6):50–55, 2003. Availabe at: https://users.ece.cmu.edu/~adrian/630-f05/readings/cranor-p2p.pdf.

[40] Henry DeYoung, Deepak Garg, Limin Jia, Dilsun Kaynar, and Anupam Datta. Experiences in the Logical Specification of the HIPAA and GLBA Privacy Laws. In *Proceedings of the 9th Annual ACM Workshop on Privacy in the Electronic Society*, pages 73–82, 2010. Available at: https://dl.acm.org/doi/pdf/10.1145/1866919.1866930.

[41] Lauren A Eberly, Michael J Kallan, Howard M Julien, Norrisa Haynes, Sameed Ahmed M Khatana, Ashwin S Nathan, Christopher Snider, Neel P Chokshi, Nwamaka D Eneanya, Samuel U Takvorian, et al. Patient Characteristics Associated With Telemedicine Access for Primary and Specialty Ambulatory Care During the COVID-19 Pandemic. *JAMA network open*, 3(12):e2031640–e2031640, 2020. Available at: https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2774488?resultClick=3.

[42] Eric Engle. Europe Deciphered: Ideas, Institutions, and Laws. *Fletcher F. World Aff.*, 33:63, 2009. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1336490.

[43] EuroHealthNet. Mhealth Provides Opportunities, But Risks Widening Inequalities. https://eurohealthnet.eu/media/news-releases/mhealth-provides-opportunities-risks-widening-inequalities.

[44] FDA. Device Software Functions Including Mobile Medical Applications.

[45] FDA. Examples of Mobile Apps That Are NOT Medical Devices.

[46] FDA. Labeling - Regulatory Requirements for Medical Devices (FDA 89-4203). https://www.fda.gov/regulatory-information/search-fda-guidance-documents/labeling-regulatory-requirements-medical-devices-fda-89-4203, 1989.

[47] FDA. General Principles of Software Validation. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/general-principles-software-validation, 2002.

[48] FDA. Cybersecurity for Networked Medical Devices Containing Off-the-Shelf (OTS) Software. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/cybersecurity-networked-medical-devices-containing-shelf-ots-software, 2005.

[49] FDA. Guidance for the Content of Premarket Submissions for Software Contained in Medical Devices. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/guidance-content-premarket-submissions-software-contained-medical-devices, 2005.

[50] FDA. Modifications to Devices Subject to Premarket Approval (PMA) - The PMA Supplement Decision-Making Process. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/modifications-devices-subject-premarket-approval-pma-pma-supplement-decision-making-process, 2008.

[51] FDA. Content of Premarket Submissions for Management of Cybersecurity in Medical Devices. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/content-premarket-submissions-management-cybersecurity-medical-devices, 2014.

[52] FDA. Distinguishing Medical Device Recalls from Medical Device Enhancements. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/distinguishing-medical-device-recalls-medical-device-enhancements, 2014.

[53] FDA. The 510(k) Program: Evaluating Substantial Equivalence in Premarket Notifications [510(k)]. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/510k-program-evaluating-substantial-equivalence-premarket-notifications-510k, 2014.

[54] FDA. Applying Human Factors and Usability Engineering to Medical Devices. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/applying-human-factors-and-usability-engineering-medical-devices, 2016.

[55] FDA. Medical Device Reporting for Manufacturers. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/medical-device-reporting-manufacturers, 2016.

[56] FDA. Postmarket Management of Cybersecurity in Medical Devices. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/postmarket-management-cybersecurity-medical-devices, 2016.

[57] FDA. Deciding When to Submit a 510(k) for a Change to an Existing Device. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/deciding-when-submit-510k-change-existing-device, 2017.

[58] FDA. Design Considerations and Pre-market Submission Recommendations for Interoperable Medical Devices. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/design-considerations-and-pre-market-submission-recommendations-interoperable-medical-devices, 2017.

[59] FDA. Breakthrough Devices Program. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/breakthrough-devices-program, 2018.

[60] FDA. Content of Premarket Submissions for Management of Cybersecurity in Medical Devices. https://www.fda.gov/regulatory-information/search-fda-guidance-

documents/content-premarket-submissions-management-cybersecurity-medical-devices, 2018.

[61] FDA. De Novo Classification Request. https://www.fda.gov/medical-devices/premarket-submissions/de-novo-classification-request, November 2019.

[62] FDA. Division of Industry and Consumer Education. https://www.fda.gov/about-fda/cdrh-offices/division-industry-and-consumer-education, August 2019.

[63] FDA. FDA and Industry Procedures for Section 513(g) Requests for Information under the Federal Food, Drug, and Cosmetic Act. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/fda-and-industry-procedures-section-513g-requests-information-under-federal-food-drug-and-cosmetic, 2019.

[64] FDA. General Wellness: Policy for Low Risk Devices. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/general-wellness-policy-low-risk-devices, 2019.

[65] FDA. Off-The-Shelf Software Use in Medical Devices. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/shelf-software-use-medical-devices, 2019.

[66] FDA. Policy for Device Software Functions and Mobile Medical Applications. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/policy-device-software-functions-and-mobile-medical-applications, 2019.

[67] FDA. Premarket Approval (PMA). https://www.fda.gov/medical-devices/premarket-submissions/premarket-approval-pma, May 2019.

[68] FDA. Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD). https://www.fda.gov/media/122535/download, 2019.

[69] FDA. Classify Your Medical Device. https://www.fda.gov/medical-devices/overview-device-regulation/classify-your-medical-device, February 2020.

[70] FDA. Developing the Software Precertification Program: Summary of Learnings and Ongoing Activities. https://www.fda.gov/media/142107/download, 2020.

[71] FDA. Enforcement Policy for Digital Health Devices For Treating Psychiatric Disorders During the Coronavirus Disease 2019 (COVID-19) Public Health Emergency. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/enforcement-policy-digital-health-devices-treating-psychiatric-disorders-during-coronavirus-disease, 2020.

[72] FDA. Enforcement Policy for Non-Invasive Remote Monitoring Devices Used to Support Patient Monitoring During the Coronavirus Disease 2019 (COVID-19) Public Health Emergency (Revised). https://www.fda.gov/regulatory-information/search-fda-guidance-documents/enforcement-policy-non-invasive-remote-monitoring-devices-used-support-patient-monitoring-during, 2020.

[73] FDA. Overview of Device Regulation. https://www.fda.gov/medical-devices/device-advice-comprehensive-regulatory-assistance/overview-device-regulation, September 2020.

[74] FDA. Artificial Intelligence and Machine Learning (AI/ML) Software as a Medical Device Action Plan. "https://www.fda.gov/media/145022/download, 2021.

[75] FDA. Product Classification. https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfPCD/PCDSimpleSearch.cfm, July 2021.

[76] FDA. Search for FDA Guidance Documents. https://www.fda.gov/regulatory-information/search-fda-guidance-documents, July 2021.

[77] FDA. Withdrawn Guidance. https://www.fda.gov/medical-devices/guidance-documents-medical-devices-and-radiation-emitting-products/withdrawn-guidance, June 2021.

[78] Kevin Featherly. Bernstein v. the U.S. Department of State — Encyclopædia Britannica. https://www.britannica.com/event/Bernstein-vs-the-US-Department-of-State, May 2016.

[79] Gary Fields and John R. Emshwiller. The Many Failed Efforts to Count Nation's Federal Criminal Laws — WSJ. https://www.wsj.com/articles/SB10001424052702304319804576389601079728920, July 2011.

[80] Jens Frankenreiter and Michael A Livermore. Computational Methods in Legal Analysis. *Annual Review of Law and Social Science*, 16:39–57, 2020. Available at: https://www.annualreviews.org/doi/abs/10.1146/annurev-lawsocsci-052720-121843.

[81] Baban Gain, Dibyanayan Bandyopadhyay, Tanik Saikh, and Asif Ekbal. IITP in COLIEE@ ICAIL 2019: Legal Information Retrieval using BM25 and BERT. *arXiv preprint arXiv:2104.08653*, 2021. Available at: https://arxiv.org/pdf/2104.08653.pdf.

[82] Sanjam Garg, Shafi Goldwasser, and Prashant Nalini Vasudevan. Formalizing Data Deletion in the Context of the Right to Be Forgotten. *Advances in Cryptology–EUROCRYPT 2020*, 12106:373, 2020. Available at: https://eprint.iacr.org/2020/254.pdf.

[83] Michael Genesereth. Computational Law The Cop in the Backseat, 2015. Available at: http://logic.stanford.edu/publications/genesereth/complaw.pdf.

[84] Andrew R Glencross. Federalism, Confederalism and Sovereignty Claims: Understanding the Democracy Game in the EU. *SOVEREIGNTY GAMES: INSTRUMENTALIZING STATE SOVEREIGNTY IN EUROPE AND BEYOND, R. Adler-Nissen and T. Gammeltoft-Hansen, eds., New York: Palgrave*, 2008. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1004326.

[85] Erving Goffman. *The Presentation of Self in Everyday Life*, volume 21. Harmondsworth London, 1978.

[86] Ann Griffiths, Rupak Chattopadhyay, John Light, and Carl Stieren. *The Forum of Federations Handbook of Federal Countries 2020*. Springer, 2020.

[87] IMDRF SaMD Working Group. Software as a Medical Device (SaMD): Key Definitions. http://www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-131209-samd-key-definitions-140901.pdf, 2013.

[88] Thomas R Gruber. A Translation Approach to Portable Ontology Specifications. *Knowledge acquisition*, 5(2):199–220, 1993. Available at: https://tomgruber.org/writing/ontolingua-kaj-1993.pdf.

[89] The Guardian. 'Shocking' Hack of Psychotherapy Records in Finland Affects Thousands — The Guardian. https://www.theguardian.com/world/2020/oct/26/tens-of-thousands-psychotherapy-records-hacked-in-finland, October 2020.

[90] Tivas Gupta. The Future of Federalism — Harvard Political Review. https://harvardpolitics.com/the-future-of-federalism/, September 2019.

[91] Mark A Hall, Elizabeth Dugan, Beiyao Zheng, and Aneil K Mishra. Trust in Physicians and Medical Institutions: What Is It, Can It Be Measured, And Does It Matter? *The Milbank Quarterly*, 79(4):613–639, 2001. Available at: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0009.00223?casa_token=c0ZT76iUglEAAAAA:0lbY5P7sNBZ3ftByZ_asTmUSmBb4ThCDDJ3WfnBLiPBkZ6XGlLrbtZ1jbRYAtiPRBGMXU-BOjzI3yQ.

[92] Manuel Hilty, David Basin, and Alexander Pretschner. On Obligations. In *European Symposium on Research in Computer Security*, pages 98–117. Springer, 2005. Available at: https://link.springer.com/content/pdf/10.1007/11555827_7.pdf.

[93] Jen Patja Howell. The Lawfare Podcast: Trust, Software and Hardware. https://www.lawfareblog.com/lawfare-podcast-trust-software-and-hardware, February 2021.

[94] Kathryn Howley. Role of mHealth in PHC — Duke Personalized Health Care. https://dukepersonalizedhealth.org/2018/10/role-of-mhealth-in-phc/, October 2018.

[95] Jordyn Imhoff. Health Inequality Actually Is a "Black and White Issue", Research Says. https://healthblog.uofmhealth.org/lifestyle/health-inequality-actually-a-black-and-white-issue-research-says, June 2020.

[96] Washington University in St. Louis School of Law. What is the Difference Between Common Law and Civil Law? https://onlinelaw.wustl.edu/blog/common-law-vs-civil-law/, January 2014.

[97] Healthcare Information and Management Systems Society. Global Health Disparities Infographic. https://www.himss.org/resources/global-health-disparities-infographic, July 2020.

[98] Justia. US Case Law. https://law.justia.com/cases/.

[99] Daniel Martin Katz, Michael J Bommarito, and Josh Blackman. A General Approach For Predicting the Behavior of the Supreme Court of the United States. *PloS one*, 12(4):e0174698, 2017. Available at: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0174698.

[100] Daniel Martine Katz and Michael J. Bommarito II. Measuring the Complexity of the Law: The United States Code. https://computationallegalstudies.com/2010/08/02/measuring-the-complexity-of-the-law-the-united-states-code/, August 2010.

[101] Mi-Young Kim and Randy Goebel. Two-Step Cascaded Textual Entailment for Legal Bar Exam Question Answering. In *Proceedings of the 16th edition of the International Conference on Articial Intelligence and Law*, pages 283–290, 2017. Available at: https://sites.ualberta.ca/~miyoung2/Papers/2017_cascaded.pdf.

[102] Fred Kort. Predicting Supreme Court Decisions Mathematically: A Quantitative Analysis of the "Right to Counsel" Cases. *American Political Science Review*, 51(1):1–12, 1957. Available at: https://www.jstor.org/stable/1951767.

[103] Thomas A LaVeist, Lydia A Isaac, and Karen Patricia Williams. Mistrust of Health Care Organizations Is Associated with Underutilization of Health Services. *Health services research*, 44(6):2093–2105, 2009. Available at: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2796316/.

[104] John Law. How Can We Define Federalism? *Perspectives on Federalism*, 5(3):88–120, 2013. Available at: http://www.on-federalism.eu/attachments/169_download.pdf.

[105] LexisNexis. Perfect 10, Inc. v. Amazon.com, Inc. https://www.lexisnexis.com/community/casebrief/p/casebrief-perfect-10-inc-v-amazon-com-inc.

[106] LexisNexis. Zeran v. Am. Online, Inc. https://www.lexisnexis.com/community/casebrief/p/casebrief-zeran-v-am-online-inc.

[107] William Li, Pablo Azar, David Larochelle, Phil Hill, and Andrew W. Lo. Law is Code: a Software Engineering Approach to Analyzing the United States Code. *J. Bus. & Tech. L.*, 10:297, 2015. Available at: http://digitalcommons.law.umaryland.edu/jbtl/vol10/iss2/6.

[108] Benjamin L Liebman, Margaret E Roberts, Rachel E Stern, and Alice Z Wang. Mass Digitization of Chinese Court Decisions: How to Use Text as Data in the Field of Chinese Law. *Journal of Law and Courts*, 8(2):177–201, 2020. Available at: https://scholarship.law.columbia.edu/cgi/viewcontent.cgi?article=3040&context=faculty_scholarship.

[109] Mary Madden. Privacy, Security, and Digital Inequality — Data & Society. https://datasociety.net/library/privacy-security-and-digital-inequality/, September 2017. Available at: https://datasociety.net/library/privacy-security-and-digital-inequality/.

[110] Mary Madden, Michele Gilman, Karen Levy, and Alice Marwick. Privacy, poverty, and big data: A matrix of vulnerabilities for poor Americans. *Wash. UL Rev.*, 95:53, 2017. Available at: https://openscholarship.wustl.edu/law_lawreview/vol95/iss1/6/.

[111] John Levi Martin. What is Field Theory? *American journal of sociology*, 109(1):1–49, 2003.

[112] Dylan Matthew. The Myth of the 70,000-Page Federal Tax Code — Vox. https://www.vox.com/policy-and-politics/2017/3/29/15109214/tax-code-page-count-complexity-simplification-reform-ways-means, March 2017.

[113] Alex McBride. Dred Scott v. Sandford (1857) — Educational Broadcasting Corporation and PBS. https://www.thirteen.org/wnet/supremecourt/antebellum/landmark_dred.html, December 2006.

[114] Sabrina McCubbin. Summary: The Supreme Court Rules in Carpenter v. United States. https://www.lawfareblog.com/summary-supreme-court-rules-carpenter-v-united-states, June 2018.

[115] Elizabeth Mynatt, Gregory D Hager, Santosh Kumar, Ming Lin, Shwetak Patel, Jack Stankovic, and Helen Wright. Research Opportunities and Visions for Smart and

Pervasive Health. *arXiv preprint arXiv:1706.09372*, June 2017. Available at:https://arxiv.org/abs/1706.09372.

[116] Stuart Nagel. Predicting Court Cases Quantitatively. *Michigan Law Review*, 63(8):1411–1422, 1965. Available at: https://repository.law.umich.edu/mlr/vol63/iss8/6/.

[117] Helen Nissenbaum. Privacy as Contextual Integrity. *Wash. L. Rev.*, 79:119, 2004. Available at: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.79.1603&rep=rep1&type=pdf.

[118] NIST. Framework for Improving Critical Infrastructure Cybersecurity. https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf, April 2018.

[119] Library of Congress. Road to the Constitution - Creating the United States — Library of Congress. https://www.loc.gov/exhibits/creating-the-united-states/road-to-the-constitution.html, April 2008.

[120] The Editors of Encyclopædia Britannica. Unitary State — Encyclopædia Britannica. https://www.britannica.com/topic/unitary-state, August 2019.

[121] U.S. Department of Health and Human Services. Resources for Mobile Health Apps Developers. https://www.hhs.gov/hipaa/for-professionals/special-topics/health-apps/index.html, September 2020.

[122] U.S. National Library of Medicine. Health Disparities — U.S. National Library of Medicine: MedlinePlus. https://medlineplus.gov/healthdisparities.html, May 2021.

[123] University of Ottawa Juriglobe. JuriGlobe. http://www.juriglobe.ca/eng/index.php. Archived at: https://web.archive.org/web/20210415050722/http://www.juriglobe.ca/eng/index.php.

[124] Future of Privacy Forum. All About DNT. https://allaboutdnt.com/.

[125] U.S. House of Representatives. United States Code: Public Law 116-259 (12/23/2020). https://uscode.house.gov/download/releasepoints/us/pl/116/259/usc-rp@116-259.htm, 12 2020.

[126] University of South Carolina Law Library. Guide to International and Foreign Law Research - A Quick Primer on the World's Legal Systems. https://guides.law.sc.edu/c.php?g=315476&p=2108388, June 2018.

[127] Office of the Federal Register and National Archives and Records Administration. Title 21 - Food and Drugs - Volume 1. https://www.govinfo.gov/app/details/CFR-2020-title21-vol1, March 2020.

[128] Office of the Federal Register and National Archives and Records Administration. Title 21 - Food and Drugs - Volume 2. https://www.govinfo.gov/app/details/CFR-2020-title21-vol2, March 2020.

[129] Office of the Federal Register and National Archives and Records Administration. Title 21 - Food and Drugs - Volume 3. https://www.govinfo.gov/app/details/CFR-2020-title21-vol3, March 2020.

[130] Office of the Federal Register and National Archives and Records Administration. Title 21 - Food and Drugs - Volume 4. https://www.govinfo.gov/app/details/CFR-2020-title21-vol4, March 2020.

[131] Office of the Federal Register and National Archives and Records Administration. Title 21 - Food and Drugs - Volume 5. https://www.govinfo.gov/app/details/CFR-2020-title21-vol5, March 2020.

[132] Office of the Federal Register and National Archives and Records Administration. Title 21 - Food and Drugs - Volume 6. https://www.govinfo.gov/app/details/CFR-2020-title21-vol6, March 2020.

[133] Office of the Federal Register and National Archives and Records Administration. Title 21 - Food and Drugs - Volume 7. https://www.govinfo.gov/app/details/CFR-2020-title21-vol7, March 2020.

[134] Office of the Federal Register and National Archives and Records Administration. Title 21 - Food and Drugs - Volume 8. https://www.govinfo.gov/app/details/CFR-2020-title21-vol8, March 2020.

[135] Office of the Federal Register and National Archives and Records Administration. Title 21 - Food and Drugs - Volume 9. https://www.govinfo.gov/app/details/CFR-2020-title21-vol9, March 2020.

[136] Oyez. Reno v. ACLU. https://www.oyez.org/cases/1996/96-511.

[137] Oyez. Roe v. Wade. https://www.oyez.org/cases/1971/70-18.

[138] Natalie A. Prescott. The Anatomy of Biometric Laws: What U.S. Companies Need To Know in 2020. https://www.natlawreview.com/article/anatomy-biometric-laws-what-us-companies-need-to-know-2020, January 2020.

[139] William Ralston. They Told Their Therapists Everything. Hackers Leaked It All — WIRED. https://www.wired.com/story/vastaamo-psychotherapy-patients-hack-data-breach/, May 2021.

[140] Wayne J Riley. Health Disparities: Gaps in Access, Quality and Affordability of Medical Care. *Transactions of the American Clinical and Climatological Association*, 123:167, 2012. Available at: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3540621/.

[141] Norman Sadeh, Alessandro Acquisti, Travis D Breaux, Lorrie Faith Cranor, Aleecia M McDonald, Joel Reidenberg, Noah A Smith, Fei Liu, N Cameron Russell, Florian Schaub, et al. Towards Usable Privacy Policies: Semi-Automatically Extracting Data Practices from Websites' Privacy Policies. *Poster Proceedings, SOUPS*, pages 9–11, 2014. Available at: https://cups.cs.cmu.edu/soups/2014/posters/soups2014_posters-paper20.pdf.

[142] Norman Sadeh, Alessandro Acquisti, Travis D Breaux, Lorrie Faith Cranor, Aleecia M McDonald, Joel R Reidenberg, Noah A Smith, Fei Liu, N Cameron Russell, Florian Schaub, et al. The Usable Privacy Policy Project: Combining Crowdsourcing, Machine Learning and Natural Language Processing to Semi-Automatically Answer Those Privacy Questions Users Care About. 2013. Available at: http://reports-archive.adm.cs.cmu.edu/anon/isr2013/CMU-ISR-13-119.pdf.

[143] Sarah Scheffler and Mayank Varia. Protecting Cryptography Against Compelled Self-Incrimination. In *30th {USENIX} Security Symposium ({USENIX} Security 21)*, 2021. Available at: https://www.usenix.org/system/files/sec21summer_scheffler.pdf.

[144] Cornell Law School. Case Law — Cornell Legal Information Institute. https://www.law.cornell.edu/wex/case_law, May 2020.

[145] SCOTUSblog. Obergefell v. Hodges. https://www.scotusblog.com/case-files/cases/obergefell-v-hodges/.

[146] Marek J. Sergot, Fariba Sadri, Robert A. Kowalski, Frank Kriwaczek, Peter Hammond, and H Terese Cory. The British Nationality Act As a Logic Program. *Communications of the ACM*, 29(5):370–386, 1986. Available at: https://www.doc.ic.ac.uk/~rak/papers/British%20Nationality%20Act.pdf.

[147] Yunqiu Shao, Bulou Liu, Jiaxin Mao, Yiqun Liu, Min Zhang, and Shaoping Ma. THUIR@ COLIEE-2020: Leveraging Semantic Understanding and Exact Matching for Legal Case Retrieval and Entailment. *arXiv preprint arXiv:2012.13102*, 2020. Available at: https://arxiv.org/pdf/2012.13102v1.pdf.

[148] Vanessa B Sheppard, Darren Mays, Kenneth P Tercyak, and Thomas LaVeist. Medical Mistrust Influences Black Women's Level of Engagement in brca1/2 Genetic Counseling and Testing. *Journal of the National Medical Association*, 105(1):17–22, 2013. Available at: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3721431/.

[149] Katie Siek, Tiffany Veinot, and Beth Mynatt. Research opportunities in Sociotechnical Interventions for Health Disparity Reduction. *arXiv preprint arXiv:1908.01035*, August 2019. Available at: https://arxiv.org/abs/1908.01035.

[150] Karandeep Singh, Kaitlin Drouin, Lisa P Newmark, JaeHo Lee, Arild Faxvaag, Ronen Rozenblum, Erika A Pabo, Adam Landman, Elissa Klinger, and David W Bates. Many Mobile Health Apps Target High-Need, High-Cost Populations, But Gaps Remain. *Health Affairs*, 35(12):2310–2318, 2016. Available at: https://pubmed.ncbi.nlm.nih.gov/27920321/.

[151] Statista. Patient Optimism of mHealth Improvements to Overall Healthcare as of 2012, By Market. https://www.statista.com/statistics/328695/expectation-of-patients-on-mhealth-improvement-to-healthcare/, July 2014.

[152] Statista. Growth in the Number of Medical Apps Downloaded during the COVID-19 Pandemic by Country in 2020. https://www.statista.com/statistics/1181413/medical-app-downloads-growth-during-covid-pandemic-by-country/, October 2020.

[153] Statista. Percentage of Health App Users in Selected Countries as of 2020. https://www.statista.com/forecasts/1181436/share-of-health-app-users-by-country, October 2020.

[154] Matt Stroud. These Six Lawsuits Shaped the Internet — The Verge. https://www.theverge.com/2014/8/19/6044679/the-six-lawsuits-that-shaped-the-internet, August 2014.

[155] Harry Surden. Computable Contracts. *UCDL Rev.*, 46:629, 2012. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2216866.

[156] Vu Tran, Minh Le Nguyen, and Ken Satoh. Building Legal Case Retrieval Systems with Lexical Matching and Summarization Using a Pre-trained Phrase Scoring Model. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law*, pages 275–282, 2019. Available at: https://dl.acm.org/doi/pdf/10.1145/3322640.3326740.

[157] American Civil Liberties Union. Windsor v. United States. https://www.aclu.org/cases/lesbian-and-gay-rights/windsor-v-united-states, April 2014.

[158] Tara Van Veen, Sophia Binz, Meri Muminovic, Kaleem Chaudhry, Katie Rose, Sean Calo, Jo-Ann Rammal, John France, and Joseph B Miller. Potential of Mobile Health Technology to Reduce Health Disparities in Underserved Communities. *Western Journal of Emergency Medicine*, 20(5):799, 2019. Available at: https://pubmed.ncbi.nlm.nih.gov/31539337/.

[159] Pepijn RS Visser and Trevor JM Bench-Capon. A Comparison of Four Ontologies for the Design of Legal Knowledge Systems. *Artificial Intelligence and Law*, 6(1):27–57, 1998. Available at: https://intranet.csc.liv.ac.uk/~tbc/publications/4ontologies.pdf.

[160] Robert J Walsh. An History of US Tax Code Complexity within Computer-Based Return Preparation. *Accounting & Taxation*, 11(1):47–57, 2019. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3480705.

[161] Bernhard Waltl and Florian Matthes. Towards Measures of Complexity: Applying Structural and Linguistic Metrics to German Laws. In *Legal knowledge and information systems*, pages 153–162. IOS Press, 2014. Available at: https://www.semanticscholar.org/paper/Towards-Measures-of-Complexity%3A-Applying-Structural-Waltl-Matthes/e72028689ad140c42de8eda77c8991c77d5782e3?p2df.

[162] Michael Walzer. *Spheres of Justice: A Defense of Pluralism and Equality*. Basic books, 2008.

[163] Chenlu Wang and Xiaoning Jin. Study on Prediction of Legal Judgments Based on the CNN-BiGRU Model. In *Proceedings of the 2020 6th International Conference on Computing and Artificial Intelligence*, pages 63–68, 2020. Available at: https://dl.acm.org/doi/pdf/10.1145/3404555.3404573.

[164] Eric Wicklund. New Coalition Plans to Apply mHealth to Tackling Health Disparities. https://mhealthintelligence.com/news/new-coalition-plans-to-apply-mhealth-to-tackling-health-disparities, September 2020.

[165] Eric Wicklund. Personalized mHealth Tools Help Providers Improve Care Management. https://mhealthintelligence.com/news/personalized-mhealth-tools-help-providers-improve-care-management, February 2020.

[166] Pete Williams. Does the Constitution Guarantee Citizenship to Anyone Born in U.S.? https://www.nbcnews.com/news/us-news/does-constitution-guarantee-citizenship-all-born-here-n411451, August 2015.

[167] Lillie D Williamson and Cabral A Bigman. A Systematic Review of Medical Mistrust Measures. *Patient education and counseling*, 101(10):1786–1794, 2018. Available at: https://www.sciencedirect.com/science/article/pii/S0738399118302040?via%3Dihub.

[168] Chaojun Xiao, Haoxi Zhong, Zhipeng Guo, Cunchao Tu, Zhiyuan Liu, Maosong Sun, Tianyang Zhang, Xianpei Han, Zhen Hu, Heng Wang, et al. Cail2019-Scm: A Dataset of Similar Case Matching in Legal Domain. *arXiv preprint arXiv:1911.08962*, 2019. Available at: https://arxiv.org/pdf/1911.08962.pdf.

[169] xkcd. Tasks. https://xkcd.com/1425/.

[170] Haoxi Zhong, Chaojun Xiao, Cunchao Tu, Tianyang Zhang, Zhiyuan Liu, and Maosong Sun. How Does NLP Benefit Legal System: A Summary of Legal Artificial Intelligence. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5218–5230, Online, July 2020. Association for Computational Linguistics. Available at:https://arxiv.org/pdf/2004.12158.pdf.

[171] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel R Reidenberg, N Cameron Russell, and Norman Sadeh. Maps: Scaling Privacy Compliance Analysis to a Million Apps. *Proc. Priv. Enhancing Tech.*, 2019:66, 2019. Available at: https://www.usableprivacy.org/static/files/popets-2019-maps.pdf.

Comput. Legal for mHealth S&P, Tung, M.S. 2021