Towards Automated Computational Auditing of mHealth Security and Privacy Regulations

Brian Tung*[†], Zhiyuan Yu*, Ning Zhang
Washington University in St. Louis, MO, USA
Harvard Law School, MA, USA

ABSTRACT

The growing complexity of our regulatory environment presents us with a hard problem: how can we determine if we are compliant with an ever-growing body of regulations? Computational legal auditing may help, as computational tools are exceptionally good at making sense of large amounts of data. In this research, we explore the possibility of creating a computational auditor that checks if mobile health (mHealth) apps satisfy federal security and privacy regulations. In doing so, we find that while it is challenging to convert open-ended, generally applicable, complicated laws into computational principles, the use of non-legal, authoritative, explanatory documents allows for computational operationalization while preserving the open-ended nature of the law. We test our auditor on 182 FDA/CE-approved mHealth apps. Our research suggests that the use of non-legal, authoritative, guidance documents may help with the creation of computational auditors, a promising tool to help us manage our ever-growing regulatory responsibilities.

CCS CONCEPTS

• **Security and privacy → Software security engineering**.

KEYWORDS

FDA regulations; computational law; auditing; mHealth; mobile app security

ACM Reference Format:
Brian Tung*[†], Zhiyuan Yu*, Ning Zhang. 2021. Towards Automated Computational Auditing of mHealth Security and Privacy Regulations. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21).* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3460120.3485342

1 INTRODUCTION

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 [2, 5, 21]. Our growing legal system is too complicated; it's impossible for people to know all the laws that apply to them. However, ignorance of the law is not an excuse; people are still subject to the law, even if they are unfamiliar with it. Therein lies the need for computational legal analysis. Can computational tools help us comply with an ever-growing body of regulation?

To explore this question, we ask if it is possible to automatically audit security and privacy compliance in a highly regulated industry – mobile health (mHealth). mHealth apps have been lauded for their potential to provide cheap, effective, and personalized healthcare to large groups of people [20]. In the past decade, mHealth apps have exploded in popularity, and they have been critical in the global response to COVID-19 [1, 23]. However, this explosive growth has made the industry challenging to regulate, and security concerns abound [19, 22]. Automated auditing is a scalable solution that would allow regulatory bodies like the Food and Drug Administration (FDA) to check if mHealth apps satisfy security and privacy regulations.

In this poster, we survey existing mHealth regulations and present key challenges that make it difficult to create computational auditors from legal requirements. Our key observation is that these challenges may be overcome by not using the legal text in the creation of an auditor, but by using a non-legal, authoritative, guidance document that explains regulatory intent. These documents may already exist; we use one to create an auditor that checks if 182 FDA/CE-approved apps violate FDA security and privacy recommendations. Our research suggests that the use of these non-legal, authoritative, guidance documents is a promising approach for computational auditing.

2 FEDERAL MHEALTH REGULATIONS

mHealth apps were grandfathered into the existing regulatory environment; therefore, they inherit the security and privacy regulations that apply to traditional healthcare devices. Accordingly, the FDA only plans to regulate mHealth apps that meet the legal definition of a medical device [15], defined in 21 U.S.C. § 321(h). The FDA often cites the United States Code (U.S.C.), the Code of Federal Regulations (C.F.R.), and FDA guidance documents when addressing mHealth regulations.

The U.S. Code is a codification of 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.

The C.F.R. is the official codification of the rules and regulations promulgated by departments and agencies of the executive branch

* Brian Tung and Zhiyuan Yu contributed equally. [†] Portions of this research appear in Brian Tung's Washington University in St. Louis master's thesis [24]. This work was supported in part by US National Science Foundation under the grant CNS-1916926.

CCS '21, November 15–19, 2021, Virtual Event, Republic of Korea.

ACM ISBN 978-1-4503-8454-4/21/11.
https://doi.org/10.1145/3460120.3485342

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. 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.

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 [18].

As an aside – our analysis is limited to federal mHealth security and privacy regulations. Due to the large number of regulations, state regulations (e.g., Illinois' Biometric Information Privacy Act) and case law are considered out-of-scope of our analysis.

3 CHALLENGES

There are a few key challenges when creating a computational auditor for legal affairs. Our research suggests that the use of non-legal, authoritative, guidance documents may address these challenges. This is illustrated in Figure 1.

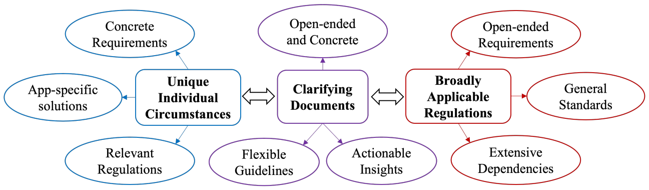

Figure 1: Non-legal, authoritative, guidance documents balance the challenges raised by broadly applicable regulations and unique individual circumstances

Open-ended vs. concrete requirements: One key challenge is that the law is often open-ended by design; by avoiding specificity, the law remains relevant, even as technology evolves. However, auditors require specific requirements to check. One solution to this problem is to write more-specific laws, though we are unlikely to realize greater specificity in all areas of the law. Another solution is to use authoritative, but non-legal, supporting documents that clarify regulatory intent, like FDA guidance documents. Because these documents do not have the full weight of the law behind them, they allow regulators like the FDA, standards organizations like NIST, and industry groups to provide more-concrete recommendations while leaving the law open-ended.

Application-specific solutions vs. general standards: There is a fundamental tension between giving OEMs more power to define their own defense profiles and requiring a general standard that applies to all applications. Applications face different risks due to their different operating requirements; it may be beneficial to allow individual product developers to provide their own application-specific security measures. However, a general standard may be easier to enforce from a regulatory standpoint. The publication of non-legal, more-specific, regulatory guidelines allows the best of both worlds: these recommendations are specific-enough to guide inexperienced developers, yet do not constrain the security solutions developed by more experienced teams.

Extensive dependencies vs. relevant regulations: Another challenge developers face is navigating the large collection of documents that define our complicated regulatory environment. While scholars are unsure about how to study legal complexity, one proposed method is to measure how many other documents each legal unit cites [3, 4]. We adopt 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.

Figure 2 maps the out-bound citations for eleven key FDA mHealth guidance documents [6–14, 16, 17]. 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. Furthermore, these other documents have their own dependency networks. The large number of regulations that apply to mHealth apps may weaken the explanatory purpose of these documents; it can be hard for developers to enumerate all of the relevant standards that apply to their application.

It may be helpful, then, to direct developers to a single, authoritative, non-legal, guidance document that distills many standards into a list of actionable insights. Developers satisfied with these insights may not need to explore other regulations or guidance documents.

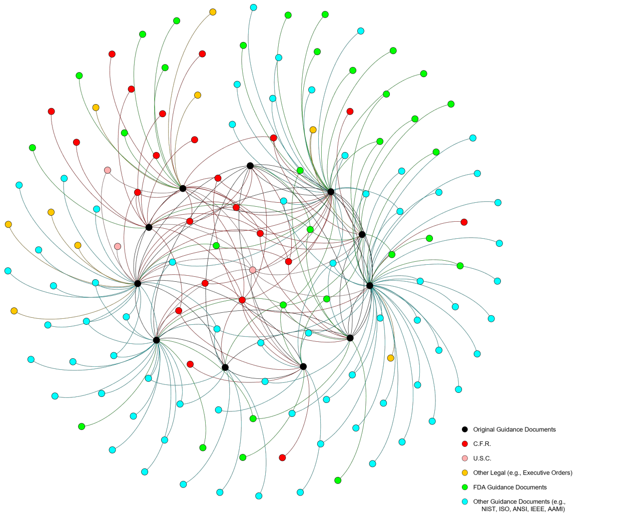
● Original Guidance Documents
● C.F.R.
● U.S.C.
● Other Legal (e.g., Executive Orders)
● FDA Guidance Documents
● Other Guidance Documents (e.g., NIST, ISO, ANSI, IEEE, AAMI)

Figure 2: Out-bound citations for eleven FDA guidance documents

Table 1: Interpreted FDA Recommendations

Category	Recommendation
A. Authenticate Users	A-1 User authentication
	A-2 Multi-factor authentication
	A-3 Strong passwords
B. Authenticate Commands	B-1 Safe crypto usage
	B-2 Authenticate connections
	B-3 Verify software
	B-4 Deny by default
C. Data Integrity	C-1 Verify incoming data
	C-2 Secure data transfer
	C-3 Protect essential local data
	C-4 Strong crypto algorithms
	C-5 Use unique key per device
D. Execution Integrity	D-1 Verify code integrity
E. Detect Security Events	E-1 Record security events
	E-2 Secure configuration

4 TRANSLATION AND COMPLIANCE VERIFICATION

As discussed in Section III, there are a few key challenges that make it difficult to translate regulatory requirements into a computation-friendly format. Existing work has achieved some success by applying ML and NLP [25]. However, these techniques may only be effective in some legal domains. Our key observation is that by translating not the legal text itself but "middleware" (i.e., non-legal, authoritative, guidance documents that clarify regulatory intent), we may be able to enjoy both the benefits of an open-ended, flexible legal framework and a descriptive regulatory environment that can be operationalized into a computational auditor. These clarifying documents may already be produced by regulatory bodies and can be independently developed by standards organizations and industry groups.

In the mHealth space, we operationalized the security recommendations in the 2018 FDA draft guidance document *Content of Premarket Submissions for Management of Cybersecurity in Medical Devices*. As shown in Table 1, the recommendations in this document can be grouped into five categories: user authentication, authentication of safety-critical functions, data integrity, execution integrity, and the detection of cybersecurity incidents. We collected 182 FDA/CE-approved mHealth apps from healthskouts[1] in February 2021 and used our analysis framework to audit their compliance with these FDA recommendations.

Warnings were then manually analyzed. One common type of warning was the presence of hard-coded keys. Twenty-five apps hard-coded OAuth secrets and seventeen apps hard-coded API keys; these keys can be extracted and used for malicious purposes. Seven apps hard-coded cryptographic keys for pairing or inter-device communication. Finally, four apps hard-coded Android Keystore passwords and three apps hard-coded SQLite database keys.

5 CONCLUSION

In this project, we propose a new way to build automatic auditors for legal compliance. Our key observation is that instead of operationalizing legal text itself, we may find more success by translating non-legal, authoritative, guidance documents that clarify regulatory intent. By doing so, we enjoy both the benefits of an open-ended, flexible legal framework and a descriptive regulatory environment that can be translated into a computational auditor. These documents may already be available from regulatory bodies. Alternatively, standards organizations and industry groups can create them independently. We applied our principles to create an auditor and tested it on existing FDA/CE-approved mHealth apps. Our research suggests that the use of these non-legal, authoritative, guidance documents is a promising approach for computational auditing.

RESPONSIBLE DISCLOSURE

We have not publicly disclosed the vulnerabilities and are working with individual app vendors to fix them.

REFERENCES

[1] Adans-Dester et al. 2020. Can mHealth technology help mitigate the effects of the COVID-19 pandemic? *IEEE Open Journal of Engineering in Medicine and Biology* 1 (2020), 243–248.

[2] National Archives. 2019. Code of Federal Regulations Total Pages 1936 - 1949, And Total Volumes and Pages 1950 - 2019.

[3] Michael James Bommarito and Daniel Martin Katz. 2009. Properties of the United States code citation network. *Available at SSRN 1502927* (2009).

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

[5] U.S. Congress. 1934. United States Code: Table of Contents (1934).

[6] FDA. 2002. General Principles of Software Validation.

[7] FDA. 2005. Cybersecurity for Networked Medical Devices Containing Off-the-Shelf (OTS) Software.

[8] FDA. 2005. Guidance for the Content of Premarket Submissions for Software Contained in Medical Devices.

[9] FDA. 2014. Content of Premarket Submissions for Management of Cybersecurity in Medical Devices.

[10] FDA. 2014. Distinguishing Medical Device Recalls from Medical Device Enhancements.

[11] FDA. 2016. Medical Device Reporting for Manufacturers.

[12] FDA. 2016. Postmarket Management of Cybersecurity in Medical Devices.

[13] FDA. 2017. Design Considerations and Pre-market Submission Recommendations for Interoperable Medical Devices.

[14] FDA. 2018. Content of Premarket Submissions for Management of Cybersecurity in Medical Devices.

[15] FDA. 2019. General Wellness: Policy for Low Risk Devices.

[16] FDA. 2019. Off-The-Shelf Software Use in Medical Devices.

[17] FDA. 2019. Policy for Device Software Functions and Mobile Medical Applications.

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

[19] Dongjing He, Muhammad Naveed, Carl A Gunter, and Klara Nahrstedt. 2014. Security concerns in Android mHealth apps. In *AMIA Annual Symposium Proceedings*, Vol. 2014. American Medical Informatics Association, 645.

[20] Kathryn Howley. 2018. Role of mHealth in PHC. https://dukepersonalizedhealth.org/2018/10/role-of-mhealth-in-phc/.

[21] U.S. House of Representatives. 2020. United States Code: Public Law 116-259 (12/23/2020).

[22] Miloslava Plachkinova, Steven Andrés, and Samir Chatterjee. 2015. A taxonomy of mhealth apps–Security and privacy concerns. In *2015 48th Hawaii International Conference on System Sciences*. IEEE, 3187–3196.

[23] Connor Stewart. 2020. 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/.

[24] Brian Tung. 2021. *The Challenges of Applying Computational Legal Analysis to mHealth Security and Privacy Regulations.* Master's thesis. Washington University in St. Louis.

[25] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel R Reidenberg, N Cameron Russell, and Norman Sadeh. 2019. Maps: Scaling Privacy Compliance Analysis to a Million Apps. *Proc. Priv. Enhancing Tech.* 2019 (2019), 66.

[1] https://apps.healthskouts.com/