

**DESIGN AND IMPLEMENTATION OF A THRESHOLD-BASED
LINKING SCHEME TO EXTEND BROWSER FINGERPRINT
LIFESPAN**

by
ELİF ECEM ŞAMLIOĞLU

Submitted to the Graduate School of Engineering and Natural Sciences
in partial fulfillment of
the requirements for the degree of Master of Science

Sabancı University
July 2025

**DESIGN AND IMPLEMENTATION OF A THRESHOLD-BASED
LINKING SCHEME TO EXTEND BROWSER FINGERPRINT
LIFESPAN**

Approved by:

Prof. ALBERT LEVİ
(Thesis Supervisor)

Prof. CEMAL YILMAZ

Assoc. Prof. KÜBRA KALKAN ÇAKMAKÇI

Date of Approval: July 9, 2025

ELİF ECEM ŞAMLIOĞLU 2025 ©

All Rights Reserved

ABSTRACT

DESIGN AND IMPLEMENTATION OF A THRESHOLD-BASED LINKING SCHEME TO EXTEND BROWSER FINGERPRINT LIFESPAN

ELİF ECEM ŞAMLIOĞLU

Computer Science and Engineering, M.Sc. Thesis, July 2025

Thesis Supervisor: Prof. Albert Levi

Keywords: browser fingerprinting, user identification, authentication, privacy

Browser fingerprinting is a powerful tool for user identification in financial and other security-critical applications that require strong authentication. However, due to the instability of browser attributes, fingerprints often change rapidly, reducing their lifespan and negatively impacting user convenience. We propose ThresholdFP, a novel linking algorithm designed to extend the durability of browser fingerprints without compromising precision.

Instead of replacing the fingerprint after even a small change, ThresholdFP computes a difference score between fingerprints. If the total change remains below a configurable threshold, the new fingerprint is linked to its predecessor, forming a lineage over time. This allows minor attribute fluctuations to be tolerated, improving persistence without triggering additional identity checks.

To ensure realistic performance evaluation, we collected and utilized two real-world datasets: one institutional dataset with over 235,000 fingerprints and a labeled volunteer dataset containing 5,774 fingerprints from known browser instances. The best-performing variant of ThresholdFP achieved an average tracking duration (i.e. fingerprint lifetime) of 55.7 days on the institutional dataset and 50.1 days on the volunteer dataset, while maintaining over 98% estimated precision and 99.5% actual precision, respectively. Compared to rival methods in the literature, ThresholdFP improves average fingerprint tracking duration by 24.33% to 106.30%, demonstrating its effectiveness and robustness under real-world conditions.

ÖZET

TARAYICI PARMAK İZİ ÖMRÜNÜ UZATMAK İÇİN EŞİK TABANLI BİR BAĞLANTILAMA YÖNTEMİNİN TASARIMI VE UYGULAMASI

ELİF ECEM ŞAMLIOĞLU

Bilgisayar Bilimi ve Mühendisliği, Yüksek Lisans Tezi, Temmuz 2025

Tez Danışmanı: Prof. Dr. Albert Levi

Anahtar Kelimeler: tarayıcı parmak izi, kullanıcı tanımlama, kimlik doğrulama, mahremiyet

Tarayıcı parmak izi alma, sağlam bir kimlik doğrulamasının gerekliliği olduğu finansal ve diğer güvenlik açısından kritik uygulamalarda kullanıcı tanımlama için güçlü bir araçtır. Ancak tarayıcı özniteliklerinin kararsızlığı nedeniyle parmak izleri sıklıkla hızlı bir şekilde değişmekte, bu da ömürlerini azaltmakta ve kullanıcı kolaylığını olumsuz etkilemektedir. Bu çalışmada, parmak izi doğruluğunu zedelemeden tarayıcı parmak izlerinin dayanıklılığını artırmak üzere tasarlanmış yeni bir bağlantılılama algoritması olan ThresholdFP önerilmektedir.

ThresholdFP, küçük bir değişiklikten sonra parmak izini tamamen değiştirmek yerine parmak izleri arasında bir fark skoru hesaplar. Eğer toplam değişiklik ayarlanabilir bir eşik değerinin altında kalıyorsa yeni parmak izi bir öncekine bağlanır ve zamanla bir soy zinciri oluşturur. Bu yaklaşım, küçük öznitelik dalgalanmalarının tolere edilmesini sağlar ve ek kimlik doğrulama kontrolleri tetiklenmeden izlenebilirliği artırır.

Gerçekçi bir performans değerlendirmesi sağlamak amacıyla biri 235.000'den fazla parmak izi içeren kurumsal bir veri kümesi, diğeri bilinen tarayıcı örneklerinden elde edilmiş 5.774 parmak izine sahip etiketli bir gönüllü veri kümesi olmak üzere iki gerçek dünya veri kümesi toplanmış ve kullanılmıştır. ThresholdFP'nin en iyi performans gösteren varyantı, kurumsal veri kümesinde ortalama 55.7 gün, gönüllü veri kümesinde ise 50.1 gün izleme süresi (yani parmak izi ömrü) elde etmiş; sırasıyla %98'in üzerinde tahmini doğruluk ve %99.5 gerçek doğruluk sağlamıştır. Liter-

atürdeki rakip yöntemlerle karşılaştırıldığında ThresholdFP ortalama parmak izi izlenebilirliğini %24.33 ile %106.30 arasında artırarak gerçek dünya koşulları altında etkinliğini ve dayanıklılığını ortaya koymaktadır.

ACKNOWLEDGEMENTS

Despite spending a long time searching for the right words, I am still not certain I have succeeded in expressing the depth of my gratitude to my supervisor, Prof. Albert Levi, as usual phrases feel too small to capture it. Without his unfaltering support, not only this thesis, but many important steps I have taken in my academic journey would not have been possible. During the many moments when I lost confidence in myself and struggled to stay motivated, it was his constant encouragement, endless patience, and invaluable guidance and wisdom that consistently grounded me and helped me find my way forward. I am and will remain deeply grateful to have had his guidance throughout this experience.

I would also like to express my sincere gratitude to Prof. Cemal Yilmaz, whose courses I had the opportunity to take during both my undergraduate and graduate studies, and with whom I had the privilege of working closely as a teaching assistant for three semesters, an experience that greatly contributed to my academic and professional growth.

Throughout every high and low, my family has been my unwavering foundation. Their constant support and encouragement, and their way of turning even the smallest achievement into a reason to smother me in love, lifted me and gave me strength in ways they may not even realize. I also feel incredibly lucky to have shared this path—both in life and in our field of study—with two brothers who are not only family, but also my lifelong best friends, and whose insight, steady presence, and companionship have been a constant source of inspiration and comfort.

Words feel too light to reflect Aysegül's place in my journey. Thank you for being there for me, truly, for everything. Your presence is quietly stitched into everything I have managed to accomplish, and I cannot even fathom doing this without you.

I would like to extend my thanks to Ada, Nil, Sinem, and Zülal for their wholehearted support and for believing in me with a conviction that I am still trying to live up to, through every step I have taken since high school.

Lastly, I would like to thank the members of FENS 2014 for providing both an intellectual and genuinely enjoyable environment that added warmth, support, and perspective throughout this journey, and for the many friendships it brought into my life.

*Dedicated to
my beloved family*

TABLE OF CONTENTS

LIST OF TABLES	xi
LIST OF FIGURES	xii
1. INTRODUCTION	1
2. BACKGROUND AND MOTIVATION	4
2.1. Fundamentals of Browser Fingerprinting	5
2.2. Related Work in Fingerprint Linking	9
2.3. Motivation	10
3. METHODOLOGY	11
3.1. General Workflow	11
3.2. Obtaining the Fingerprint Hash	15
3.3. Threshold-based Algorithm	15
3.4. Categorizing Fingerprints	18
3.4.1. Appearance of an Active Fingerprint	18
3.4.2. Reappearance of an Outdated Fingerprint	18
3.4.3. Newly Introduced Fingerprints	19
3.5. Remediation Process	21
4. DATASETS AND RESULTS	23
4.1. Evaluation Metrics	23
4.2. Datasets	24
4.2.1. Institutional Dataset	25
4.2.2. Volunteer Dataset	26
4.3. Results on the Institutional Dataset	29
4.3.1. Precision	29
4.3.2. Tracking Time of Fingerprints	29
4.4. Results on the Volunteer (Labeled) Dataset	31
4.4.1. Precision	31

4.4.2. Tracking Time of Fingerprints	33
4.5. Comparative Analysis.....	34
4.5.1. Tracking Time and Precision.....	36
4.5.2. Scalability	37
5. DISCUSSION AND ETHICAL ISSUES	39
5.1. Discussion.....	39
5.1.1. Addressing the Research Questions	40
5.1.1.1. Research Question 1	40
5.1.1.2. Research Question 2	41
5.1.1.3. Research Question 3	41
5.2. Ethical Considerations	42
5.3. Ethical Implications	43
5.4. Threats to Validity	44
6. CONCLUSION	45
BIBLIOGRAPHY.....	47

LIST OF TABLES

Table 2.1. Fingerprint attributes (part 1)	7
Table 2.2. Fingerprint attributes (part 2)	8
Table 4.1. Distribution of operating systems across two datasets.	25
Table 4.2. Distribution of browsers across two datasets.	25
Table 4.3. Normalized Shannon entropy and calculated scores of attributes in the institutional dataset.	27
Table 4.4. Normalized Shannon entropy and calculated scores of attributes in the volunteer dataset.	28
Table 4.5. Comparison of the average tracking time obtained in number of days by ThresholdFP and previous work	36

LIST OF FIGURES

Figure 3.1. Overview of the decision-making flow in ThresholdFP	12
Figure 3.2. Overview of the threshold algorithm	14
Figure 3.3. An example sequence of fingerprints from a single user, illustrating how the threshold algorithm updates the active set and lineage structure over time.	20
Figure 3.4. A remediation scenario	21
Figure 4.1. Estimated precision of three threshold algorithm variants on the institutional dataset	30
Figure 4.2. Average tracking time of three threshold algorithm variants on the institutional dataset	31
Figure 4.3. Precision and estimated precision of three threshold algorithm variants on the volunteer dataset	32
Figure 4.4. Average tracking time of three threshold algorithm variants on the volunteer dataset	33
Figure 4.5. Average decision time of ThresholdFP and FP-Stalker’s rule-based algorithm for one fingerprint	38

1. INTRODUCTION

Browser fingerprinting is a collection of methods used to gather various types of information from a client’s browser in order to acquire a fingerprint that can be associated with the user. Browser fingerprinting methods are considered harder to detect, as they are stateless, unlike cookies, and do not store data in the client browser. Many browsers have taken measures against third-party cookies to address the controversies surrounding tracking (Iqbal, Englehardt & Shafiq, 2021). As a result of these recent adjustments, stateless techniques such as browser fingerprinting have attracted significant popularity. Even though some browsers have also taken measures against browser fingerprinting to mitigate concerns about the potential threat it poses to user privacy, the use of browser fingerprinting by websites has been gaining momentum.

However, apart from its adversarial usage, browser fingerprinting has also been used for benign purposes such as bot detection (Jonker, Krumnow & Vlot, 2019; Vastel, Rudametkin, Rouvoy & Blanc, 2020; Wu, Sun, Zhao & Cao, 2023) and fraud detection (Iqbal et al., 2021). In particular, fraud detection is a prominent objective that organizations pursue using industry-standard tools (Durey, 2022). Furthermore, browser fingerprinting has begun to be integrated into authentication schemes as an additional security layer and typically as a triggering mechanism to induce complementary security procedures in case of any discrepancy concerning the fingerprints (Laperdrix, Avoine, Baudry & Nikiforakis, 2019; Lin, Ilia, Solanki & Polakis, 2022). These systems aim to find a balance between convenience and security, which underscores the need for browser fingerprints to be both unique and stable.

To accurately identify users, fingerprints must be sufficiently distinctive. At the same time, they must remain consistent over time to be useful for long-term tracking or authentication. On the other hand, Vastel, Laperdrix, Rudametkin, and Rouvoy (2018) concluded that nearly half of fingerprints do not even last a week,

and around 80% only last a few days longer than a week. The authors point to a series of reasons that can account for the evolution of fingerprints, including, but not limited to, software updates, relocations that cause a change in timezones, and some configuration changes made by users.

Although these problems have been widely acknowledged and amplifying the uniqueness of browser fingerprints has been a central focus of many works in the domain (Cao, Li & Wijmans, 2017; Gómez-Boix, Laperdrix & Baudry, 2018; Laperdrix, Rudametkin & Baudry, 2016), to the best of our knowledge, only a limited number of studies—most notably Eckersley’s Panopticlick (2010) and Vastel et al.’s FP-Stalker (2018)—have investigated the means of extending the tracking time, i.e., the useful time of fingerprints, by associating them with one another.

Furthermore, given that browser fingerprinting attributes may vary in reliability and also that some browsers have been interfering with or deprecating some of the attributes to combat browser fingerprinting, a study addressing these recent developments is needed.

To this end, in this thesis, we propose ThresholdFP, a threshold-based algorithm to link a user’s fingerprints to increase tracking time. ThresholdFP calculates a difference score between the fingerprints and tolerates changes up to a threshold. To put it differently, new fingerprints with acceptable differences from previous ones are linked and are not treated as unknown. The threshold is configurable, and a threshold value suitable for the needs of a website can be selected, taking into account the trade-off between precision and tracking time. Additionally, ThresholdFP carries out a remediation process if it detects that a previous decision was faulty.

This work was guided by the following research questions:

- RQ1.** Extending fingerprint lifespan via threshold-based linking: To what extent can a threshold-based algorithm extend the effective lifespan of browser fingerprints?
- RQ2.** Tolerance to attribute changes without precision loss: Can benign changes in fingerprint attributes be tolerated without compromising linking precision?
- RQ3.** Comparison with existing approaches: How does a threshold-based linking algorithm compare to existing methods in terms of tracking duration and precision?

To address these research questions, two separate datasets with different sizes, populations, frequencies, and collection intervals were employed. In addition to developing three variants of our algorithm, we adapted Vastel et al. (2018) and Eckersley’s work (2010) to conduct a comparative analysis. All threshold algorithm variants outperformed the other two algorithms, with optimal thresholds selected separately for each dataset. The best-performing variant achieved an average tracking time of 55.7 days on the first dataset and 50.1 days on the second dataset.

The organization of the thesis is as follows: Section 2 provides a brief literature review of the domain, elaborates on the state-of-the-art fingerprint linking studies, and outlines the motivation behind this study. Section 3 explains our approach in detail. Section 4 describes the datasets used in our study and presents the results obtained. Section 5 discusses the results presented in the previous section and addresses the threats to validity. Finally, Section 6 provides the concluding remarks.

2. BACKGROUND AND MOTIVATION

Despite ongoing improvements in web authentication, traditional systems remain constrained by usability and security trade-offs. The use of passwords is still unavoidable in most cases, even though they are vulnerable to reuse, leakage, phishing, and social engineering attacks (Abromaityte, Dubero, Kruger, Lucassen, Vos, van den Hout, Bouma, Treur & Roelofsma, 2025; Durey, Laperdrix, Rudametkin & Rouvoy, 2021). Typical two-factor authentication schemes are also generally considered inconvenient by end users (Marky, Ragozin, Chernyshov, Matviienko, Schmitz, Mühlhäuser, Eghtebas & Kunze, 2022). While biometrics may offer stronger assurance, they are still susceptible to replication attacks, raise user privacy concerns, and require customized hardware (Alrawili, AlQahtani & Khan, 2024), making them impractical for many web-based applications. SMS-based authentication introduces several attack surfaces as well. Messages are transmitted unencrypted over mobile networks, leaving them vulnerable to interception and man-in-the-middle attacks (Alzomai, Josang, McCullagh & Foo, 2008). SIM swapping is another documented threat (Mulliner, Borgaonkar, Stewin & Seifert, 2013).

Time-based One-Time Passwords (TOTP), by contrast, eliminate many of these network-related risks by generating codes locally on the user’s device. This makes TOTP more secure than SMS in many respects. However, it is less universally deployable, since it requires a smartphone and the setup of a dedicated app. In addition, the short validity window of each code can create usability problems and increase the likelihood of user error (Reese, Smith, Dutson, Armknecht, Cameron & Seamons, 2019).

Given these challenges, browser fingerprinting offers an alternative to traditional second factors for authentication. It requires no user interaction and can run silently in the background to complement traditional methods, helping to verify the user’s identity through their browser without disrupting the authentication experience. Additional security checks are triggered only when an unrecognized fingerprint is encountered (Alaca & van Oorschot, 2016; Senol, Ukani, Cutler & Bilotrevic, 2024).

The remainder of this section provides background on browser fingerprinting, presents related work in the literature, including both foundational studies and recent advancements, and outlines the motivation behind the approach proposed in this study.

2.1 Fundamentals of Browser Fingerprinting

Browser fingerprinting is fundamentally a technique for identifying users based on the configuration and behavior of their browsers. Instead of relying on stored identifiers, such as cookies, fingerprinting collects a combination of browser-exposed attributes to construct a unique identifier (i.e., a fingerprint) (Laperdrix, Bielova, Baudry & Avoine, 2020). These attributes may be gathered passively from network-level observations, such as HTTP headers, or actively through client-side JavaScript calls to browser APIs (Ajay & Guptha, 2022).

The origins of browser fingerprinting date back to Mayer's (2009) study, in which he demonstrated that by collecting the values of a few JavaScript objects, 96.23% of users could be uniquely identified. His small-sample work was followed by Eckersley (2010), who launched a large-scale experiment by announcing it on his website and inviting users to contribute their fingerprints. This effort resulted in the collection of 470,161 fingerprints and became a widely recognized study that helped bring browser fingerprinting into broader awareness.

Since then, numerous studies have been conducted in the field to point out the crucial role JavaScript APIs play in browser fingerprinting by allowing access to many browser-specific pieces of information such as installed fonts (Nikiforakis, Kapravelos, Joosen, Kruegel, Piessens & Vigna, 2013), plugins (Eckersley, 2010; Mayer, 2009), and screen attributes (Mayer, 2009; Nikiforakis et al., 2013). Moreover, JavaScript APIs also facilitate further identification through various APIs such as the Canvas API (Acar, Eubank, Englehardt, Juarez, Narayanan & Diaz, 2014; Mowery & Shacham, 2012), the WebGL API (Cao et al., 2017), and the Web Audio API (Englehardt & Narayanan, 2016) by making the browser render complex images and animations or process audio signals to produce specific outputs.

The fingerprint itself is typically formed by concatenating or hashing these attribute values into a single identifier string. Each attribute contributes differently: some offer high uniqueness but may change due to browser updates, for example, the user-agent, or system modifications (such as installed fonts), while others, such as the operating system platform, are more stable but less distinctive (Laperdrix et al., 2020). The effectiveness of a fingerprinting system depends on balancing the trade-off between entropy and stability (Andriamilanto, Allard, Le Guelvouit & Garel, 2021).

A widely adopted implementation of browser fingerprinting is FingerprintJS, an open-source library used in many production systems (FingerprintJS, 2023). It gathers a standardized set of fingerprinting attributes through both passive and active techniques. Tables 2.1 and 2.2 provide an overview of the attributes used by FingerprintJS, along with sample values.

Although the attributes used by libraries like FingerprintJS are effective for identification, the APIs that expose them are often employed for website functionality, making the intention behind their usage unclear and fingerprinting detection difficult in practice (Boussaha, Hock, Bermejo, Rumin, Rumin, Klein, Johns, Compagna, Antonioli & Barber, 2024). While users can disable JavaScript entirely, doing so often breaks site functionality and renders many websites unusable (Laperdrix et al., 2020).

However, it would be misleading to reduce browser fingerprinting solely to JavaScript APIs, as several studies have demonstrated that fingerprinting remains possible even when JavaScript is disabled, for instance, through CSS-based techniques (Laperdrix, Starov, Chen, Kapravelos & Nikiforakis, 2021; Lin, Araujo, Taylor, Jang & Polakis, 2023; Trampert, Weber, Gerlach, Rossow & Schwarz, 2025). Similarly, although JavaScript does not provide an API for retrieving browser extensions, several studies have reported various workarounds for fingerprinting browser extensions (Solomos, Ilia, Nikiforakis & Polakis, 2022; Starov & Nikiforakis, 2017; Takei, Saito, Takasu & Yamada, 2015).

Table 2.1 Fingerprint attributes (part 1)

Attribute	Source	Description	Example
User agent	HTTP header	User agent string identifying browser and OS	Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/135.0.0.0 Safari/537.36
Fonts	JavaScript	List of supported system fonts	Agency FB, Calibri...
Font preferences	JavaScript	Measured text widths using sample fonts	"default": 149.3125...
Audio	JavaScript	Sum of absolute sample values from a rendered audio buffer generated with Web Audio API	124.04
Screen resolution	JavaScript	Width and height of the screen	[720, 1280]
Screen frame	JavaScript	Dimensions of the visible browser window frame	[40, 0, 70, 0]
Color depth	JavaScript	Number of bits used per pixel for color	24
Timezone	JavaScript	Local time zone setting	Europe/Istanbul
Session storage	JavaScript	Availability of session-based storage	true
Local storage	JavaScript	Availability of persistent local storage	true
DOM blockers	JavaScript	Detected content blockers	adGuardBase
Platform	JavaScript	Platform identifier string	Win32
OS CPU	JavaScript	Reported operating system and version	Windows NT 10.0
Architecture	JavaScript	CPU architecture indicator	255
CPU class	JavaScript	CPU classification (legacy attribute)	x86
Hardware concurrency	JavaScript	Number of logical CPU cores	8
Device memory	JavaScript	Available device memory for browser	8
Indexed DB	JavaScript	Support for IndexedDB API	true
Open database	JavaScript	Support for Web SQL database	true

Table 2.2 Fingerprint attributes (part 2)

Attribute	Source	Description	Example
Plugins	JavaScript	Installed browser plugins with MIME info	{ "name": "PDF Viewer", "description": "Portable Document Format", "mimeTypes": ["type": "application/pdf", "suffixes": "pdf" , "type": "text/pdf", "suffixes": "pdf"] }
Languages	JavaScript	Preferred languages list	[en]
Canvas	JavaScript	Encoded image rendered on HTML canvas	"geometry": "data:image/png;base64,..."
Touch support	JavaScript	Touch capability indicators	{"maxTouchPoints": 0, "touchEvent": false, "touchStart": false}
Vendor	JavaScript	Reported browser vendor	"Google Inc."
Vendor flavors	JavaScript	Distinguishes browsers with same engine	chrome
Cookies enabled	JavaScript	Whether cookies are enabled	false
Color gamut	JavaScript	Supported color space	srgb
Inverted colors	JavaScript	Whether display colors are inverted	inverted
Forced colors	JavaScript	Whether a forced color palette is active	false
Monochrome	JavaScript	Whether device uses monochrome display	0
Contrast	JavaScript	Preferred contrast mode	more
Reduced motion	JavaScript	Whether reduced motion is enabled	reduce
HDR	JavaScript	High dynamic range display support	false
Math	JavaScript	Results of floating-point operations	"acos": 1.44..., "acosh": 709...
Video card	JavaScript	GPU vendor and renderer details	"vendor": "Google Inc.", ...
PDFViewer enabled	JavaScript	Support for in-browser PDF viewing	true

2.2 Related Work in Fingerprint Linking

The recent developments have escalated the privacy concerns of internet users and increased the demands for anti-tracking measures. Several major browsers took concrete steps in order to address these concerns. For example, Brave started to interfere with Canvas API output (Brave Privacy Team, 2020). As a result, fingerprints became more prone to changes. This shift has, in turn, led to greater significance for fingerprinting methods that introduce mitigation strategies against these kinds of anti-tracking measures.

One of the ways to improve the durability of fingerprints is by associating evolved fingerprints with matching candidates selected from the set of previous fingerprints. In his large-scale study, Eckersley (2010) formulated a basic fingerprint matching algorithm, which constituted the first attempt to link fingerprints. Each new fingerprint was compared with the set of previously encountered fingerprints. If a single candidate differed by only one attribute, the algorithm further checked whether the differing attribute belonged to either: (1) the set of attributes allowed to change regardless of value, or (2) the set of attributes allowed to change only if the string similarity ratio exceeded 85%. The study did not measure the tracking time; however, the accuracy and precision rates were reported as 65% and 99.1%, respectively.

Eight years later, Vastel et al. (2018) proposed FP-Stalker, which used 16 browser attributes and aimed to improve on Panopticlick through a rule-based and a hybrid fingerprint matching algorithm. They imposed seven constraints in the rule-based algorithm, deeming the first three rules non-negotiable. Although the remaining four were more flexible, the rules were collectively overly tailored to the available data. Moreover, covering all possibilities exhaustively with predetermined rules was found to be infeasible. Therefore, they employed a machine learning model in the second step of the hybrid algorithm, which was preceded by the enforcement of the first three rules of the rule-based algorithm as the first step. They also implemented Eckersley's (2010) simple algorithm to compare with their approaches. The hybrid algorithm attained the best performance, while the rule-based approach managed to outperform Eckersley's algorithm.

Later, one study evaluated the effect of optimizing feature sets and removing version information on fingerprint stability (Pugliese, Riess, Gassmann & Benenson, 2020). They were able to double the tracking time of FP-Stalker on average.

2.3 Motivation

However, the gains reported by Pugliese et al. (2020) came with a significant trade-off: tailoring the feature set to maximize tracking time rendered more than 10 browser types less trackable. Moreover, Li and Cao’s (2020) million-scale measurement study, in which they tested the performance of FP-Stalker, demonstrated that the hybrid algorithm suffered from significant scalability issues. Furthermore, although the rule-based algorithm was relatively more scalable, it was still infeasible for practical use.

The limitations and scalability issues of rule-based approaches and learning-based methods suggest that there is still room for a practical, flexible, and accurate linking mechanism to improve fingerprint durability. Rule-based algorithms often rely on fixed, handcrafted constraints that are tightly coupled to the attribute set used. As a result, their performance may degrade when applied to different datasets, attribute sets, or real-world conditions where certain attributes may be noisy, missing, or manipulated by privacy-preserving browsers. On the other hand, learning-based approaches such as classification models require large volumes of labeled training data and often incur significant computational overhead, making them difficult to deploy in time-sensitive or resource-constrained environments.

This motivates the development of a method that is both flexible and efficient, capable of tolerating natural variation without relying on rigid assumptions or resource-intensive models. To this end, we propose ThresholdFP, a threshold-based linking algorithm that allows for minor changes in browser attributes, provided they remain within a configurable threshold. This offers a more adaptable and context-aware linking strategy, enabling system designers to tune the trade-off between fingerprint persistence and linking precision. Unlike prior work, ThresholdFP is not tightly coupled to a specific attribute set and does not require fingerprint collection at fixed intervals, allowing it to operate effectively under realistic usage patterns. These design choices make ThresholdFP well-suited for use in large-scale, security-focused systems that require reliable recognition of browser instances while maintaining acceptable execution time.

3. METHODOLOGY

In this section, we describe the implementation of ThresholdFP and how it is employed to achieve a longer tracking time, that is, the period for which a fingerprint is traced.

3.1 General Workflow

In this section, the general workflow of the system will be explained. To provide a clear and comprehensive understanding of the workflow, a flowchart is shown in Figure 3.1, which summarizes the decision-making process in a compact manner. For a more detailed representation, Figure 3.2 provides a visual diagram that illustrates the interactions between the different stages of ThresholdFP.

The set of active fingerprints is the collection of a user’s up-to-date fingerprints and is initially empty. Once a fingerprint is introduced, our framework first checks whether it lies in the set of active fingerprints. If so, it is marked as safe. Else, we check whether it is a previously encountered fingerprint. If this is true, then it points out that the fingerprint was an active one before, but was later discarded from the set, meaning that until now, it was considered to be an outdated fingerprint that has been succeeded by a newer one. To follow up, a remediation process is applied to take corrective measures. The remediation process is explained in more detail in Section 3.5. If the fingerprint is neither an active nor an outdated one, then it must be recorded for the first time. It can be (a) an evolved version of a previous fingerprint or (b) an unknown fingerprint from a completely new browser instance. To decide if the fingerprint can be linked to a previous fingerprint, a threshold-based

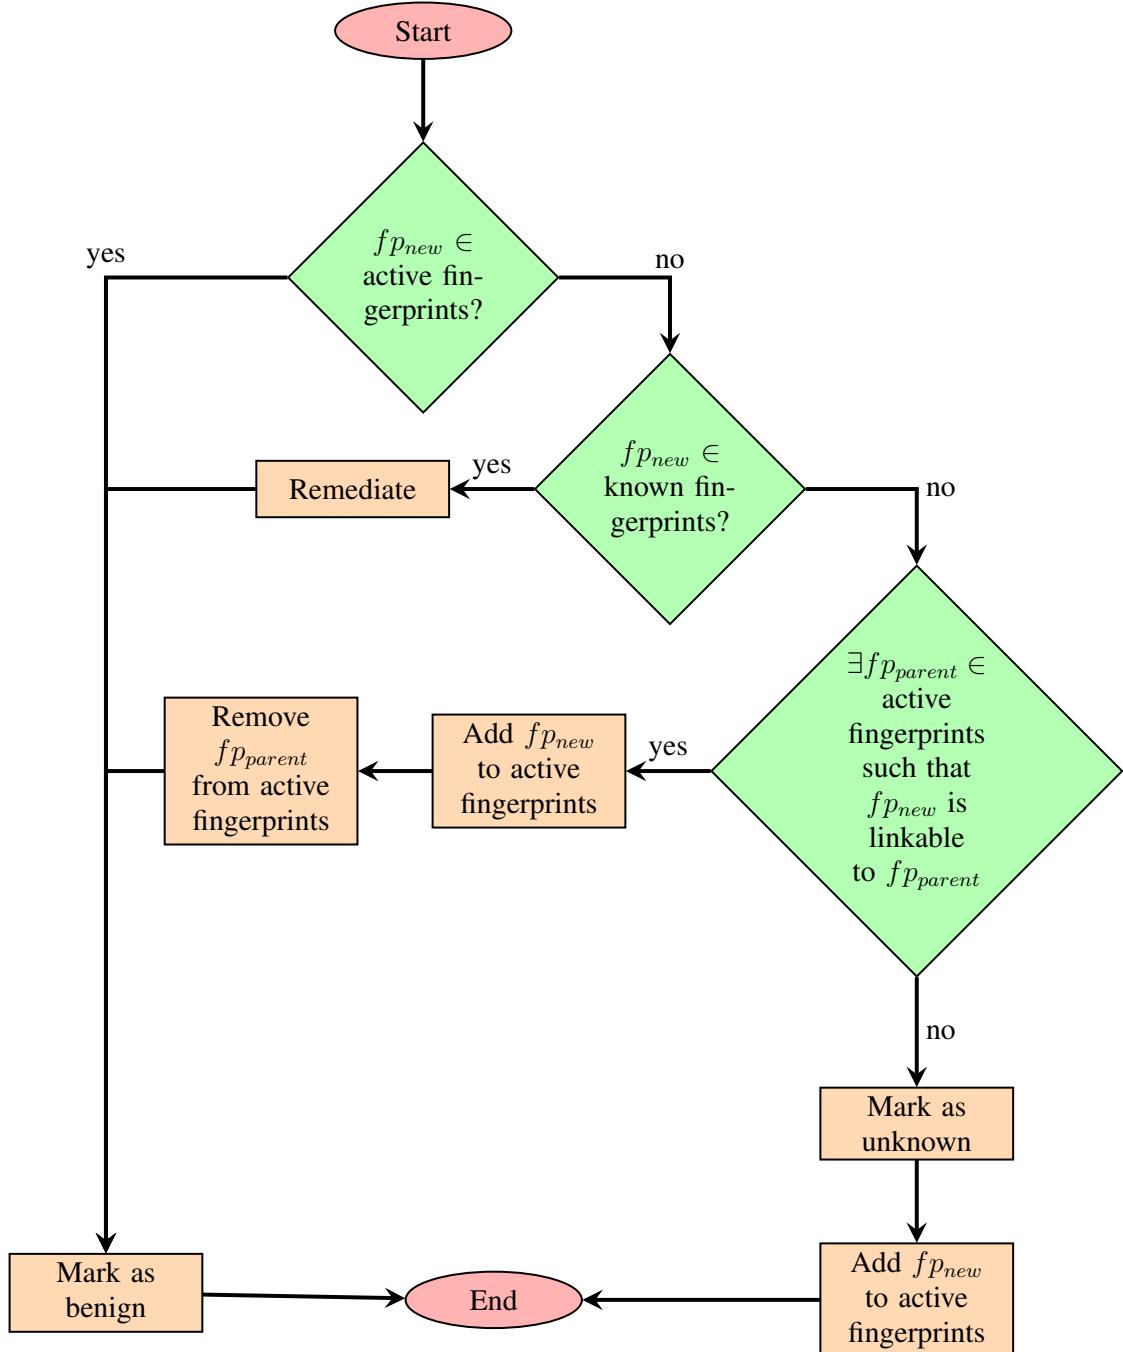


Figure 3.1 Overview of the decision-making flow in ThresholdFP

algorithm is employed. If the algorithm is able to find candidates, then the candidate that shares the greatest similarity with the fingerprint is regarded as the predecessor of the fingerprint. Thereafter, the predecessor fingerprint is removed from the set of active fingerprints and the successor fingerprint is added in its place. Otherwise, if no such candidate is found, the fingerprint is treated as a new one.

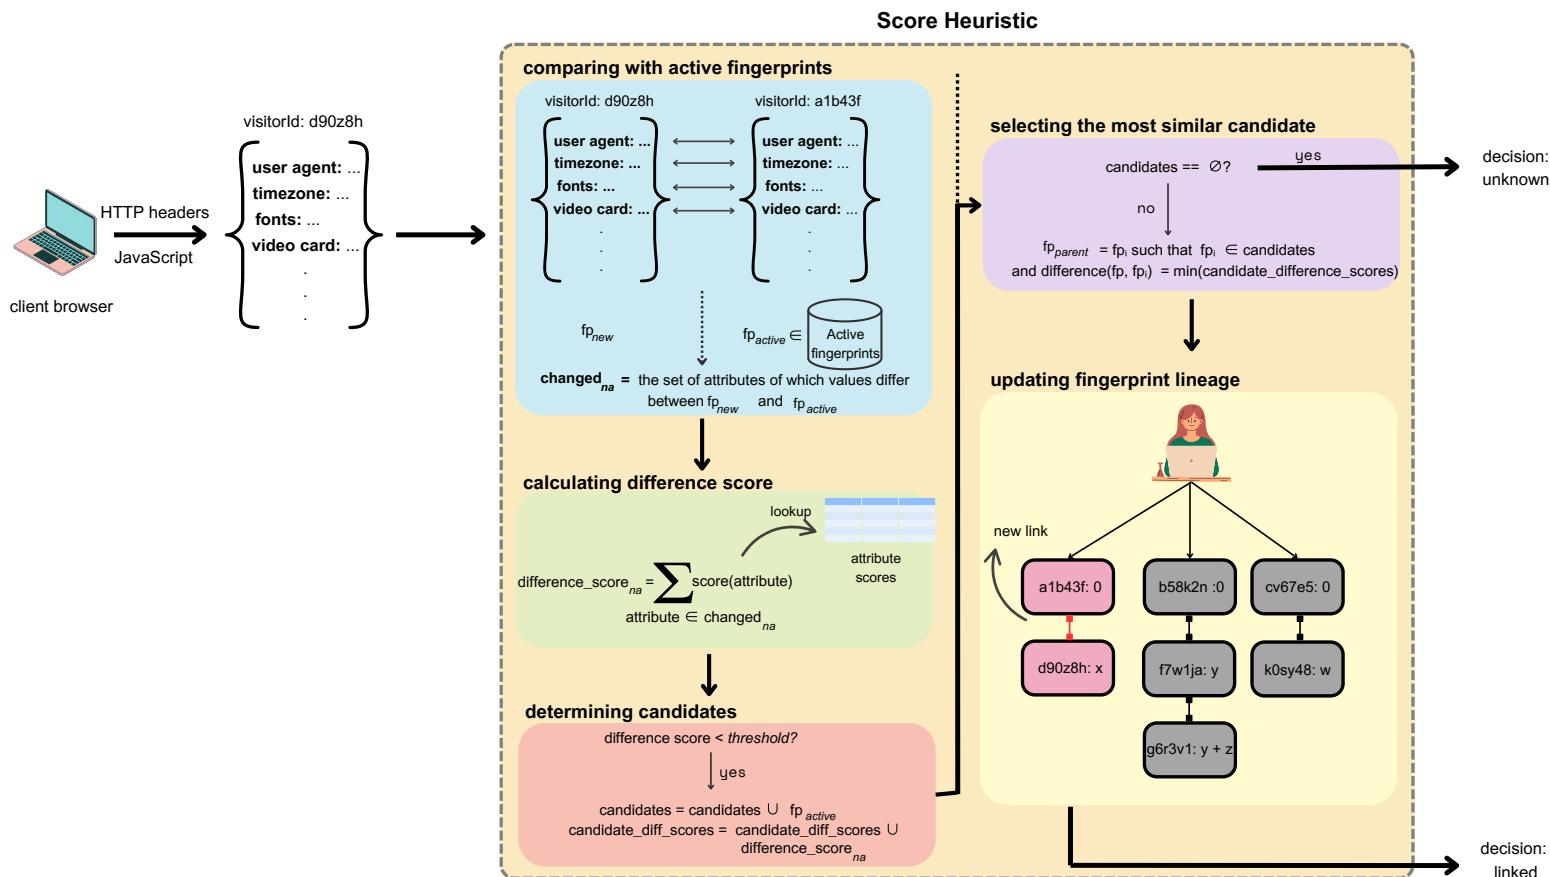


Figure 3.2 Overview of the threshold algorithm

3.2 Obtaining the Fingerprint Hash

First, FingerprintJS, which is an open-source JavaScript library for fingerprinting, is loaded into the client’s browser (FingerprintJS, 2023). The default settings do not include the user agent header in the hash; however, the user agent header holds crucial information regarding the browser brand, version, operating system, and operating system version. Furthermore, several large-scale studies such as Panopticlick by Eckersley (2010), AmIUnique by Laperdrix et al. (2016), and Hiding in the Crowd by Gómez-Boix et al. (2018) investigated the individual uniqueness of fingerprint attributes using Shannon’s entropy and suggested that high entropy values implied higher uniqueness. The user agent was ranked as the second, third, and second most unique attribute in these studies, respectively. In light of these factors, the user agent header was deemed valuable enough to be incorporated into the final hash in our work.

3.3 Threshold-based Algorithm

Fingerprint hashes offer an efficient way to compare client information. However, even a single change that stems from the unstable nature of a browser attribute can result in a completely new hash. Hence, hashes provide little information on the amount and significance of the changes that have occurred. Consequently, we formed a heuristic algorithm which is based on a scoring system that tolerates the changes until the threshold is exceeded. In other words, the algorithm connects new fingerprint instances to previously observed ones in order to extend the tracking time. Algorithm 1 is provided to display a high-level implementation of the approach.

To put this heuristic algorithm into practice, each attribute is associated with a dynamic score that is inversely proportional to the number of times it evolves throughout user logins. Each individual score is intended to quantify the instability of a given attribute. The exact calculation of the attribute scores is demonstrated in Algorithm 2.

Algorithm 1: Linking Fingerprints to the Most Relevant Predecessor

Function differenceScore(*fpAttributes*, *fpAttributes2*):

```
    sum  $\leftarrow$  0
    for attribute  $\in$  fpAttributes do
        | val1  $\leftarrow$  fpAttributes[attribute]
        | val2  $\leftarrow$  fpAttributes2[attribute]
        | if val1  $\neq$  val2 then
        |   | sum  $\leftarrow$  sum + attributeScores[attribute]
        | end
    end

    return sum
```

Data: *newFP* \neq null*minChangeScore* \leftarrow 100*parentFP* \leftarrow None**for** *activeFP* \in *ActiveFPs* **do**

```
    score  $\leftarrow$  currentScore[activeFP] + differenceScore(newFP, activeFP)
```

if score $<$ threshold **then**

```
    | if score  $<$  minChangeScore then
    |   | minChangeScore  $\leftarrow$  score
    |   | parentFP  $\leftarrow$  activeFP
    | end
```

end**end****if** *parentFP* is not None **then**

```
    | ActiveFPs  $\leftarrow$  ActiveFPs - {parentFP}
    | descendantDict[parentFP]  $\leftarrow$  newFP
    | currentScore[newFP]  $\leftarrow$  score
```

else

```
    | currentScore[newFP]  $\leftarrow$  0
```

end*ActiveFPs* \leftarrow *ActiveFPs* + {*newFP*}

Apart from the attribute scores, we also keep the scores of fingerprints. At the time of their introduction, the score of each fingerprint is initialized to zero. If the algorithm can link the new fingerprint to a preceding fingerprint, that is, if the sum of the current score of a preceding fingerprint and the difference score between that and the new fingerprint is smaller than the threshold, a parent-child relation is formed between the previous and the new fingerprint, and the score is passed on to the child fingerprint. Later on, the child fingerprint can have its own descendant fingerprint if a new fingerprint can be linked to it. Then, the new fingerprint inherits the accumulated score and joins the lineage. The chronologically linear sequence of

Algorithm 2: Calculating Individual Attribute Scores

Data: $attributes \leftarrow$ set of all fingerprinting attributes
 $attributeScores \leftarrow \{\}$
 $attributeChanges \leftarrow \{a : 0, \forall a \in attributes\}$
 $totalChanges \leftarrow 0$

for $fpAttributes1 \in userFPs$ **do**
 for $fpAttributes2 \in userFPs - \{fpAttributes1\}$ **do**
 $anyChanges \leftarrow false$
 for $attribute \in fpAttributes1$ **do**
 $val1 \leftarrow fpAttributes1[attribute]$
 $val2 \leftarrow fpAttributes2[attribute]$
 if $val1 \neq val2$ **then**
 $attributeChanges[attribute] \leftarrow attributeChanges[attribute] + 1$
 $anyChanges \leftarrow true$
 end
 end

 if $anyChanges$ **then**
 $totalChanges \leftarrow totalChanges + 1$
 end
 end
end

for $attribute \in attributes$ **do**
 $attributeChange \leftarrow attributeChanges[attribute]$
 $percentage \leftarrow \frac{attributeChange}{totalChanges} \times 100$
 $attributeScores[attribute] \leftarrow 100 - percentage$
end

fingerprints that were found related to each other will be addressed as a lineage throughout this thesis. If there are multiple candidates that satisfy the conditions, the one which produces the minimum difference score with the new fingerprint is selected. Formally, the transmission of score to descendants could be described as the following:

$$\text{score}_{\text{child}} = \text{score}_{\text{parent}} + f(\text{child}, \text{parent}) \quad (3.1)$$

where f is a function that corresponds to `differenceScore` in Algorithm 1. It finds the set of attributes whose values differed between the two fingerprints and calculates the overall change score by summing the individual scores of the attributes.

3.4 Categorizing Fingerprints

The fingerprints of a user are classified into one of the following three categories: (1) appearance of an active fingerprint, (2) reappearance of an outdated fingerprint, and (3) a new fingerprint.

3.4.1 Appearance of an Active Fingerprint

The set of active fingerprints is familiar to the system and is not considered a potential threat. They are marked as benign upon their arrival.

3.4.2 Reappearance of an Outdated Fingerprint

A fingerprint is a part of outdated fingerprints if it is among the set of fingerprints encountered so far but is not included in the set of active fingerprints. Such a case occurs if the threshold algorithm determines that fp_{new} is linked to some $fp_{previous}$ in the set of active fingerprints. The algorithm's finding relevance between a previous fingerprint and a more recent one has two implications:

- Both fingerprints have originated from the same browser instance.
- There has been a relatively small change in the attributes of the browser instance since the previous fingerprint last appeared. The more recent fingerprint is the newer version of the browser instance's fingerprint and there is a predecessor-successor relation between the two fingerprints. The terms "predecessor-successor" and "parent-child" are used interchangeably in this thesis.

In this work, we assume that each browser instance has at most one active fingerprint at any given time. However, if a fingerprint that was previously labeled as outdated reappears, it suggests that although a newer fingerprint had been marked as its successor based on similarity—leading to the original being marked as outdated—the two were not actually related. This indicates that the threshold algorithm reached a wrong conclusion about their connection. This category of fingerprints will be

addressed as `mislinked` fingerprints for the remainder of the thesis. Such cases may result from frequent or reversible changes, whether intentional or unintentional, that initially do not exceed the algorithm’s tolerance threshold, leading it to infer a successor relationship. However, the fingerprint may occasionally revert to its former state, resulting in a reappearance. Ultimately, this category requires a remediation process, but, similar to the first category, it does not raise concerns surrounding an unfamiliar new browser instance. The remediation process will be elaborated in Section 3.5.

3.4.3 Newly Introduced Fingerprints

This category includes fingerprints that have not previously been encountered. New fingerprints could result from two different reasons: a familiar browser instance producing a new fingerprint due to system evolution or an unfamiliar browser instance being introduced. Differentiating the former and the latter cases is not a straightforward task and demands a more complex approach than simple predetermined rules. Consequently, we employed the threshold algorithm to decide whether the new fingerprint can be associated with a previous fingerprint. In that case, the new fingerprint is marked as the successor of the previous one. These fingerprint pairs will be referred to as `linked` throughout this work. Otherwise, an unknown browser instance is believed to be used, and the website is informed in case any additional security measures need to be imposed. Finally, this new fingerprint is integrated into the set of active fingerprints of the user.

An illustrative example of how these fingerprint categories emerge in practice is provided in Figure 3.3. The figure demonstrates a sequence of fingerprints collected from a single user over five discrete time steps, and how each is classified into one of the three defined categories. It also shows how these classifications affect the set of active fingerprints and the lineage structure maintained by the threshold algorithm.

In both Figures 3.3 and 3.4, each box represents a fingerprint collected at time step t (indicated in the top-left corner), with the fingerprint hash shown inside. Below each box, *active* indicates the current active fingerprints after processing, and *lineages* shows the updated lineage structure.

- $t = 1$: Fingerprint “a” is observed for the first time. Since there are no previous fingerprints, it is classified as a **new fingerprint** (Category 3). It is added to the active set, and a new root node is created for it in the lineage structure.

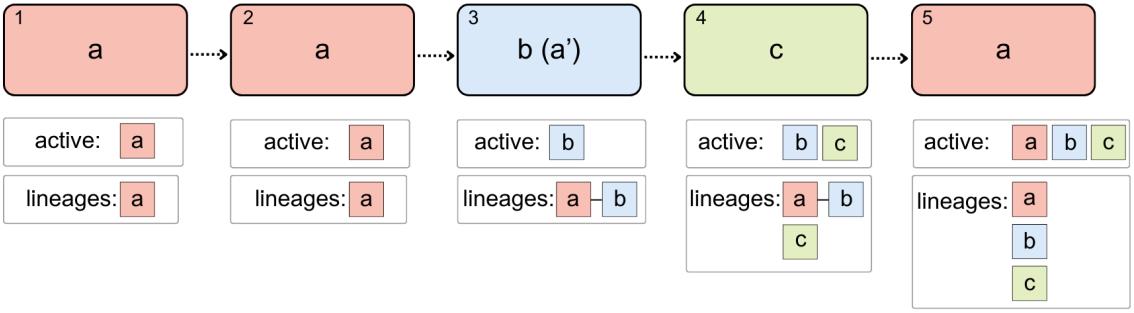


Figure 3.3 An example sequence of fingerprints from a single user, illustrating how the threshold algorithm updates the active set and lineage structure over time.

- $t = 2$: Fingerprint “a” appears again. This fingerprint exactly matches the current active fingerprint from the previous step. It is therefore classified as an **appearance of an active fingerprint** (Category 1). No updates are made to the active set or the lineage.
- $t = 3$: A **new fingerprint** (Category 3) “b” arrives. The algorithm determines that “b” is sufficiently similar to the active fingerprint “a”, and considers it a possible update. As a result, the active set is updated by replacing “a” with “b”, and the lineage structure is updated by linking “b” as a successor of “a”.
- $t = 4$: Another **new fingerprint** (Category 3), “c”, is received. However, it is not similar enough to any existing fingerprint in the active set. Since it cannot be linked to a known fingerprint, it is marked as *unknown*, and this status is returned to the caller. Then, it is added to the active set, and a new root node is created for it in the lineage structure.
- $t = 5$: Fingerprint “a” reappears. At this point, “a” was previously removed from the active set after being linked to “b” as its assumed successor. However, the reappearance of “a” indicates that this assumption was incorrect. If “b” had truly been a successor of “a”, then no further instances of “a” would be expected. Therefore, “a” is classified as a **reappearance of an outdated fingerprint** (Category 2). The algorithm adds “a” back to the active set and removes the previously established parent-child link between “a” and “b” in the lineage structure, treating them as independent root fingerprints.

3.5 Remediation Process

As mentioned in Section 3.4.2, the comeback of a fingerprint that was marked as outdated necessitates some corrections. Such a fingerprint can have:

- no ancestors but descendants
- both ancestors and descendants

No case without descendants is possible, because then, the reappeared fingerprint would not have been labeled as outdated in the first place.

The treatment of both cases is identical. Firstly, the parent-child relationship between the reappeared fingerprint and the descendant is broken. Then, the current score of the descendant, which is the sum of the scores of the changed attributes between the reappeared fingerprint and the descendant, is reverted back to 0. If the descendant has descendants itself, then throughout the whole lineage, the same amount is subtracted from the scores of remaining descendants to negate the effect of the parent fingerprint. The reason why the same resolution applies to the second case is that the relationship between the ancestor and the reappeared fingerprint is not a part of the erroneous decision. Consequently, ancestor fingerprints should preserve the link to the reappeared fingerprint.

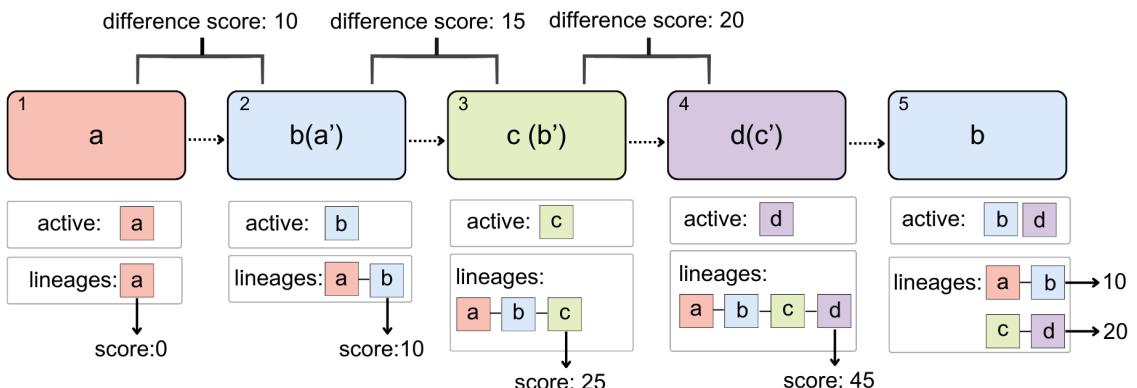


Figure 3.4 A remediation scenario

Figure 3.4 is provided to depict a scenario in which the reappeared fingerprint has both an ancestor and a chain of descendants. At $t = 2$, Fingerprint “b” is linked to Fingerprint “a” and is assigned the difference score between them, which is calculated as 10. After that, Fingerprint “c” arrives. Likewise, it is linked to Fingerprint “b”. It inherits the score of Fingerprint “b”, and a difference score of 15 is added on top of that, accumulating Fingerprint “c”’s score to 25 as demonstrated in the fig-

ure. This score is passed on to Fingerprint “d” after it gets linked to Fingerprint “c”, with an addition of 20 in regards to their difference score. At the end, Fingerprint “b” is observed to make a reappearance. This does not disrupt its logical relevance to Fingerprint “a”; therefore, its score and link to Fingerprint “a” are preserved. On the other hand, although the bond between Fingerprint “c” and Fingerprint “d” remains intact, they are revealed to be unrelated to Fingerprint “b” by this recent development. Thereafter, the link between Fingerprint “b” and Fingerprint “c” is broken, creating another lineage with Fingerprint “c” as the root and Fingerprint “d” as its child. Finally, the score of Fingerprint “c” is subtracted from the score of Fingerprint “d”, and then Fingerprint “c”’s score is reset to 0 in order to cancel the contribution of the erroneous link.

4. DATASETS AND RESULTS

In this section, a comprehensive view of the datasets is offered, and the results of the study are delivered. In this context, "linking" refers to the act of associating a newly received fingerprint with a previously seen one, forming a continuity chain or lineage of fingerprints that are presumed to belong to the same browser instance.

4.1 Evaluation Metrics

The evaluation metrics that are used to assess the performance of this work are as follows:

- **True Positive (TP):** Number of `linked` fingerprint pairs where the parent and child have the same browser label, in other words, number of accurate links established between fingerprints.
- **Estimated True Positive (ETP):** Number of `linked` fingerprints minus number of `mislinked` fingerprints.
- **False Positive (FP):** Number of `linked` fingerprint pairs where the parent and child have different browser labels, in other words, number of pairs between which the threshold algorithm formed an inaccurate link.
- **Precision:** The ratio of the number of true positives (TP) to that of all positive predictions, i.e., the sum of $TP + FP$.

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (4.1)$$

- **Estimated Precision:** The ratio of the number of estimated true positives (ETP) to that of all positive predictions, i.e., the sum of $TP + FP$.

$$\text{Estimated Precision} = \frac{\text{ETP}}{\text{TP} + \text{FP}} \quad (4.2)$$

- **Tracking Time:** Time passed between the first encounter of the root ancestor and the last encounter of the latest descendant in a fingerprint lineage. To put it differently, the amount of time that a fingerprint, specifically the root ancestor in a lineage, can be identified.
- **Matching Time:** Time to process and return a decision for a single fingerprint, that is, a decision on whether it can be linked to a previous one or not.

Our hypothesis is that as the threshold increases, the algorithm tolerates more changes, which in turn extends the tracking time. However, this increased tolerance might lead to the establishment of inaccurate links between fingerprints. As a result, both precision and estimated precision are expected to deteriorate. This trade-off highlights the importance of optimizing the threshold value to achieve a balance between the metrics. This, in fact, is what we analyze in the rest of this section.

4.2 Datasets

The same set of fingerprint attributes was collected for both of the datasets. Tables 2.1 and 2.2 provide an explanatory overview of the fingerprint attributes used. The descriptions and examples provided for the attributes in the table are prepared in accordance with the source code of FingerprintJS library and the "Web APIs" page of Mozilla Developer Network (2024), and may correspond to different information in different resources. In addition, Tables 4.1 and 4.2 summarize the operating system and browser distributions observed in the two datasets, offering further context about the user environments represented in the data.

Moreover, individual scores of attributes that are used in the threshold algorithm, as well as the normalized Shannon entropy for each of them, were computed separately for each dataset and are listed in Table 4.3 and Table 4.4.

The individual attribute scores are calculated based on how frequently each attribute changes across a user’s fingerprint visits. Attributes that change less frequently are assigned higher scores, reflecting their stability and their contribution to fingerprint durability.

The Shannon entropy values, in contrast, are included for descriptive purposes only and are not used in the scoring process. Entropy is computed from the distribution of attribute values across all users in the dataset and reflects attribute uniqueness. A loose inverse relationship between the attribute scores and the Shannon entropy values has been observed. This is an expected trend, though not guaranteed, as the two metrics capture different aspects: temporal stability for attribute scores and global uniqueness for entropy.

Table 4.1 Distribution of operating systems across two datasets.

Operating System	Institutional (%)	Volunteer (%)
Windows	79.0	64.6
MacOS	14.9	34.2
Linux	6.1	1.2

Table 4.2 Distribution of browsers across two datasets.

Browser	Institutional (%)	Volunteer (%)
Chrome	68.6	61.3
Firefox	14.2	3.4
Safari	5.1	21.1
Edge	7.6	5.1
Opera	3.8	9.1
Other	0.7	-

4.2.1 Institutional Dataset

This dataset was collected through the website of an ICT (Information and Communications Technology) company based in Türkiye. The dataset consists of 271,000 fingerprints of 32,052 users that were collected over a period of 5.5 months between May 2023 and October 2023. 35,578 out of 271,000 fingerprints that were retrieved from mobile devices were excluded from the study, leaving 28,467 users.

Every fingerprint instance in the dataset has a unique user identifier of the account owner; however, the dataset does not include browser identifiers such as cookies. Although many browser fingerprinting studies in the domain rely on cookies as ground truth (Eckersley, 2010; Gómez-Boix et al., 2018; Laperdrix et al., 2016), they were recently shown to be unreliable for use as identifiers (Li & Cao, 2020). Similarly, Pugliese et al.(2020) emphasized the inadequacy of cookies for constituting reliable identifiers and employed a user-level identifier instead. Nevertheless, due to the lack of browser identifiers, the performance of the threshold algorithm on the institutional dataset was evaluated via an estimation metric, which will be further elaborated in Section 4.3.1.

Table 4.3 features the scores and normalized Shannon entropy values calculated for attributes with respect to this dataset.

4.2.2 Volunteer Dataset

This labeled dataset of 5,774 fingerprints was collected from volunteers at an academic institution, most of whom were undergraduate students. The participants were instructed to visit a fingerprinting website that collects the same fingerprinting attributes as those present in the former dataset on a daily basis from two different devices between March 7 and May 19, 2024. Unlike the former dataset, each fingerprint was labeled with the name of the device by the participants themselves, in efforts to compile a dataset with which not the estimated but the exact performance of the threshold algorithm could be measured. The dataset was filtered to exclude 70 instances where different browsers were utilized on the same device, to ensure that the labels identify single browser instances and establish a reliable ground truth, with each device label also serving as a browser label/identifier.

The individual scores and normalized Shannon entropy values computed for the fingerprint attributes based on this dataset are presented in Table 4.4.

Table 4.3 Normalized Shannon entropy and calculated scores of attributes in the institutional dataset.

Attribute	Score	Shannon entropy (normalized)
architecture	92.48	0.49
audio	66.29	0.18
canvas	20.08	0.65
colorDepth	93.42	0.21
colorGamut	82.40	0.30
contrast	97.05	0.03
cookiesEnabled	0	0.00
cpuClass	0	0.00
deviceMemory	70.05	0.48
domBlockers	95.47	0.08
fontPreferences	68.96	0.22
fonts	40.95	0.64
forcedColors	98.69	0.08
hardwareConcurrency	57.33	0.48
hdr	89.35	0.31
indexedDB	99.99	0.00
invertedColors	96.39	0.29
languages	66.35	0.29
localStorage	0	0.00
math	78.63	0.20
monochrome	0	0.00
openDatabase	80.61	0.52
osCpu	81.41	0.10
pdfViewerEnabled	95.08	0.14
platform	76.62	0.24
plugins	82.77	0.10
reducedMotion	92.91	0.24
screenFrame	44.92	0.35
screenResolution	49.51	0.36
sessionStorage	0	0.00
timezone	94.63	0.10
touchSupport	93.82	0.10
userAgent	9.60	0.51
vendor	79.80	0.40
vendorFlavors	79.51	0.23
videoCard	39.84	0.65

Table 4.4 Normalized Shannon entropy and calculated scores of attributes in the volunteer dataset.

Attribute	Score	Shannon entropy (normalized)
architecture	70.42	0.87
audio	48.37	0.53
canvas	19.99	0.82
colorDepth	84.09	0.52
colorGamut	68.11	0.94
contrast	94.96	0.15
cookiesEnabled	0	0
cpuClass	0	0
deviceMemory	69.81	0.51
domBlockers	80.14	0.53
fontPreferences	58.77	0.55
fonts	44.91	0.68
forcedColors	90.57	0.31
hardwareConcurrency	51.75	0.71
hdr	67.63	0.58
indexedDB	0	0
invertedColors	80.14	0.79
languages	58.23	0.75
localStorage	0	0
math	71.77	0.55
monochrome	0	0
openDatabase	97.56	0.11
osCpu	95.93	0.28
pdfViewerEnabled	89.64	0.32
platform	64.388	0.69
plugins	80.11	0.21
reducedMotion	91.63	0.45
screenFrame	31.12	0.67
screenResolution	35.26	0.71
sessionStorage	0	0
timezone	89.48	0.15
touchSupport	92.85	0.20
userAgent	15.56	0.82
vendor	77.80	0.67
vendorFlavors	77.80	0.67
videoCard	36.45	0.82

4.3 Results on the Institutional Dataset

This section presents the performance of ThresholdFP on the institutional dataset measured by precision and tracking time.

4.3.1 Precision

As mentioned in Section 3.4.2, the threshold algorithm may occasionally establish incorrect bonds between fingerprints. However, since the fingerprints were received without any label or attribute that could act as ground truth, another way of measurement was needed to unravel the errors of the threshold algorithm. Hence, the number of reappeared outdated fingerprints, which account for errors that had observable consequences, was measured as an estimation of the number of errors.

In order to test how ancestor selection in cases of multiple candidates performs, three variants of the threshold algorithm were implemented. The `earliest` variant selects the fingerprint that comes first in chronological order, the `min` variant selects the fingerprint that has the minimum difference score with the successor, and lastly, the `max` variant selects the fingerprint that has the maximum difference score with the successor. Figure 4.1 displays the precisions of all three variants of the threshold algorithm.

The reason behind the similar results has been observed to be that there is a single parental candidate in most of the cases. Therefore, the method of ancestor selection does not matter significantly. Nevertheless, we proceeded with the `min` variant for the comparative analysis, as it was both the most intuitive and the best-performing variant.

4.3.2 Tracking Time of Fingerprints

The tracking time is calculated for fingerprint lineages consisting of at least two fingerprints. The root fingerprint's first recorded arrival time and the leaf fingerprint's last recorded arrival time constitute the boundaries for the calculation, and the duration between them is measured as the tracking time of the root fingerprint.

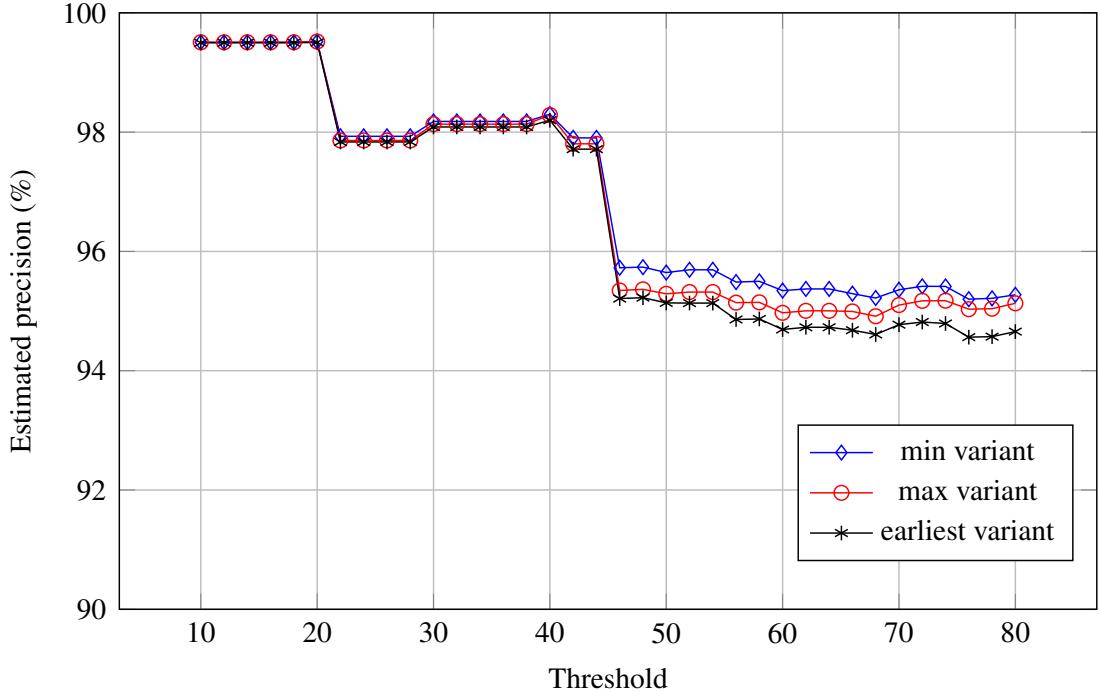


Figure 4.1 Estimated precision of three threshold algorithm variants on the institutional dataset

In order to compare the results of our work to a baseline, the average duration of the fingerprints without any algorithmic intervention was calculated. For every fingerprint, its initial appearance marked the kickoff time. After that, the time until a completely new fingerprint that was not encountered before arrived was measured. The time calculated was regarded as the original duration of the fingerprint. The process was repeated for every fingerprint, and the average time was obtained as 19.1 days. Here, we assumed that all of the previously seen fingerprints are remembered and not treated as new, allowing for the maximization of the measurements, which yields the best-case performance of the baseline.

Figure 4.2 displays the performance of all three variants of the threshold algorithm across different thresholds. The baseline duration is a constant value that is not dependent on the threshold.

In line with the precision results, the average duration of the fingerprints did not remarkably differ between the algorithm variants. The curve followed a generally increasing trend, and the best performance was obtained by the `min` variant. An average tracking time of 40.6, 55.7 and 57.6 days was recorded when threshold was set to 10, 40 and 80, respectively. All three versions managed to extend the tracking time to more than two to three times the original duration of 19.1 days, depending on the threshold.

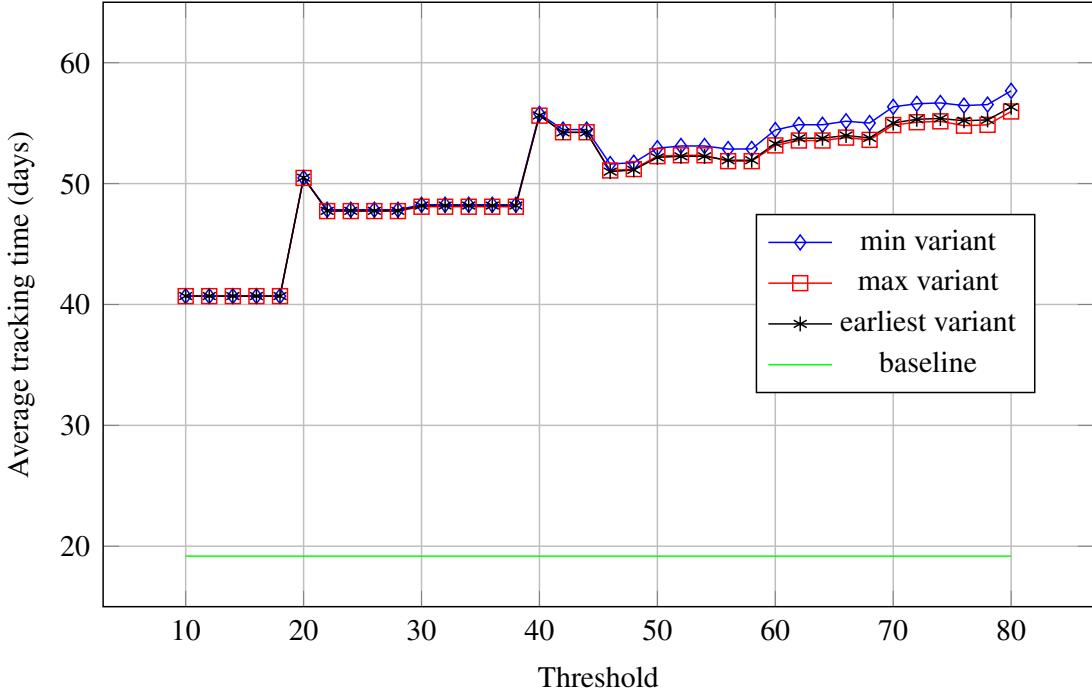


Figure 4.2 Average tracking time of three threshold algorithm variants on the institutional dataset

4.4 Results on the Volunteer (Labeled) Dataset

This section presents the performance of ThresholdFP on the volunteer dataset in terms of precision and tracking time. Unlike the figures for the previous dataset, the figures for the volunteer dataset start the threshold from 16 instead of 10. Since the smallest attribute score is 15.56, the difference scores calculated between fingerprints always exceed the limit when the threshold is set to values lower than this score, and consequently, no link gets established. Therefore, precision and tracking time were measured only for threshold values larger than 15.

4.4.1 Precision

Besides the precision estimation in Section 4.3.1, an exact measurement of the ThresholdFP's precision on this dataset was made possible by the availability of browser labels.

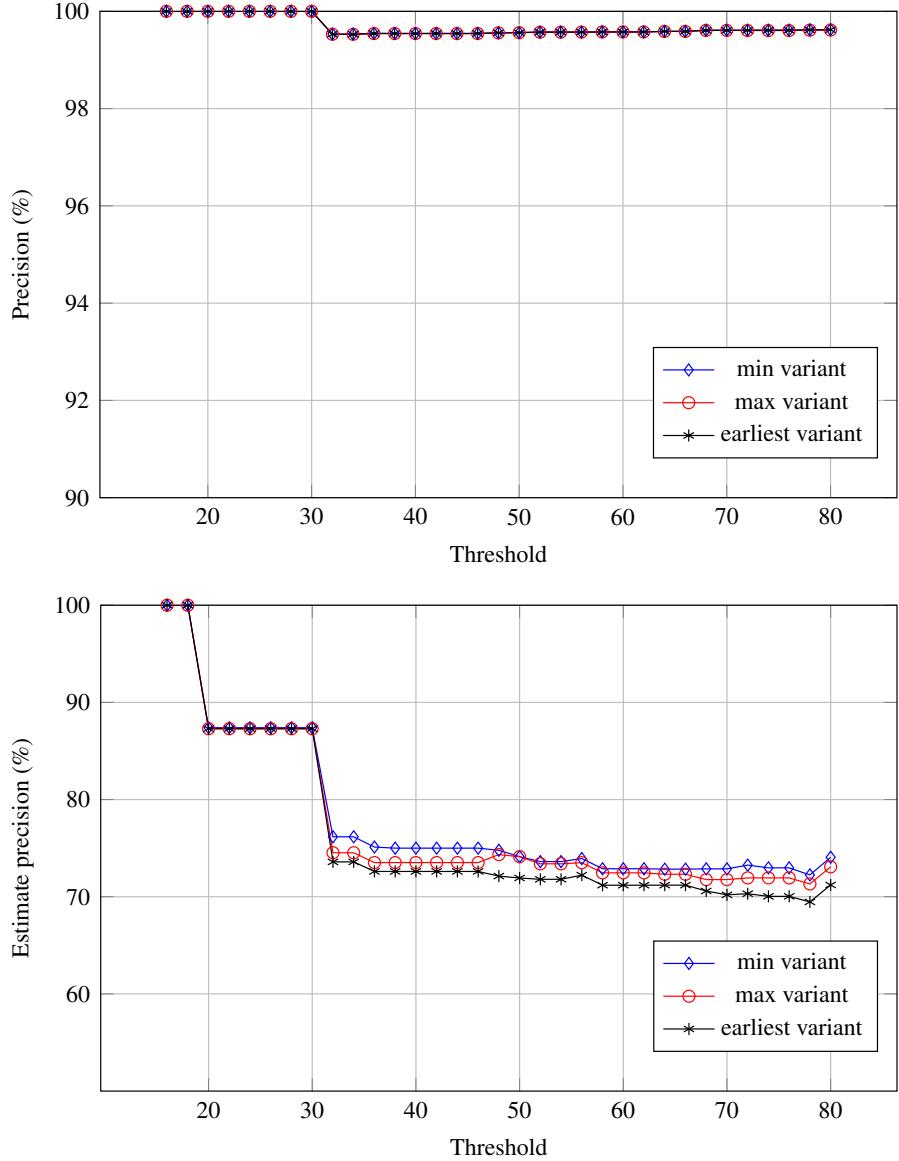


Figure 4.3 Precision and estimated precision of three threshold algorithm variants on the volunteer dataset

As evident in Figure 4.3, the actual precision and the estimated precision differed significantly. The actual precision never dropped below 99%, whereas the estimated precision experienced sharp declines before stabilizing around 70%. On the institutional dataset, however, the estimated precision remained consistently higher, exceeding 94% even at high thresholds, as depicted in Figure 4.1. A possible explanation for such dissimilar observations in the results of two datasets may lie in the population of the datasets, more specifically, the distribution of operating systems and browser brands, featured in Table 4.1 and Table 4.2, respectively. Consider-

ering that estimated precision decreases as the number of `mislinked` fingerprints increase, some operating systems or browser brands with built-in privacy features (e.g., Safari), which are more prevalent in this dataset, might be interfering with the fingerprinting results, leading to an escalated number of `mislinked` fingerprints. We leave this as future work.

4.4.2 Tracking Time of Fingerprints

The tracking time of ThresholdFP on the volunteer dataset was measured with the same approach discussed in Section 4.3.2.

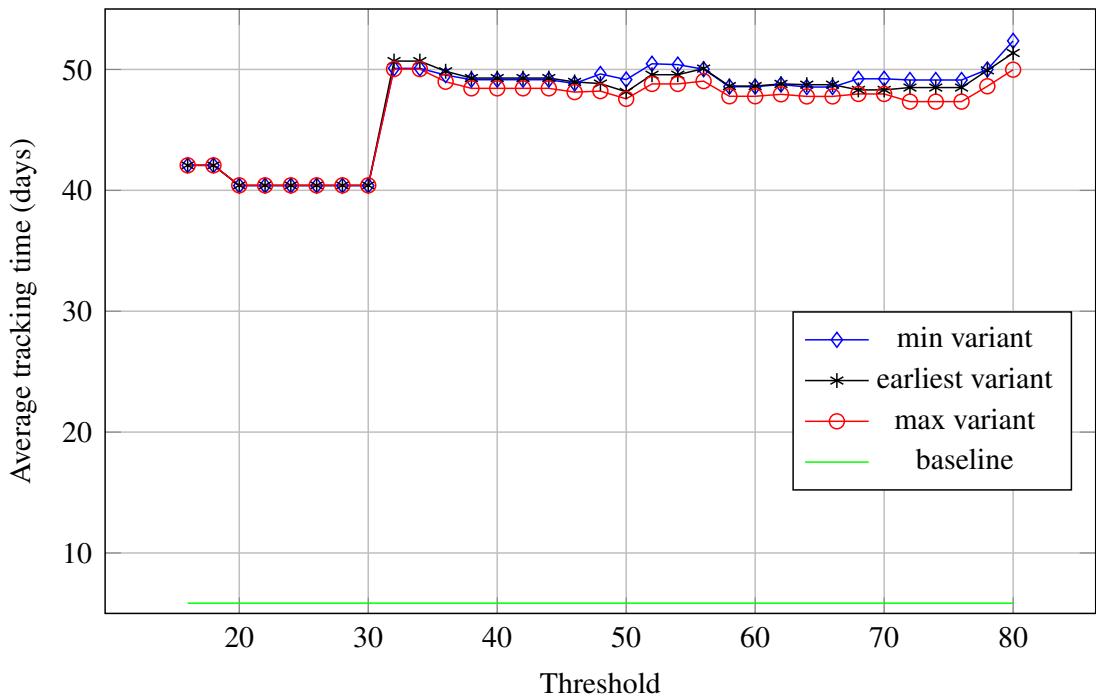


Figure 4.4 Average tracking time of three threshold algorithm variants on the volunteer dataset

Figure 4.4 displays that the `min` variant and the `earliest` variant alternately took the leading position, unlike the corresponding results for the first dataset where the `min` variant consistently maintained the top position. At the smallest and largest thresholds, the average tracking time attained with the `min` variant was 42.1 and 52.4 days. In between, it was observed to peak around threshold values 32 and 52, corresponding to 50.1 and 50.5 days of average tracking time, respectively.

4.5 Comparative Analysis

For our comparative analysis, we picked two major studies, namely FP-Stalker by Vastel et al. (2018) and Panopticlick by Eckersley (2010). This section outlines the process of applying their methods to our datasets and presents a comparison of the results. A common contributing characteristic of both studies is their exploration of methods for linking fingerprints to extend tracking duration. Although other studies on linkability exist (Li & Cao, 2020; Pugliese et al., 2020), they primarily focus on optimizations and empirical observations rather than introducing algorithmic novelties. While some commercial fingerprinting tools have been mentioned in prior work (Durey, 2022), they are not specifically designed for linkability analysis and are not replicable in practice, as their implementation details are proprietary.

Eckersley (2010) proposed a simple algorithm that determines whether an earlier fingerprint has evolved into a newly observed one. The algorithm used eight attributes, namely `user agent`, `accept`, `cookies enabled`, `screen resolution`, `timezone`, `plugins`, `fonts`, and `local storage`. It compared the previous fingerprints with the new fingerprint to check which attribute values had changed. If only one attribute had changed, it considered the previous fingerprint as a candidate. After all of the previous fingerprints were considered, if there was a single candidate, the new fingerprint was linked to it if one of the following held: (1) the changed attribute belonged to the set `cookies enabled`, `screen resolution`, `timezone`, `local storage`, or (2) the attribute values had a similarity score above 0.85.

Later, Vastel et al. (2018) proposed FP-Stalker, introducing one completely rule-based algorithm and one hybrid algorithm, which combines the rule-based one with machine learning methods to match new fingerprints with already encountered ones. The rule-based algorithm applies a predetermined set of seven rules, whereas the hybrid algorithm first applies the first three rules of the rule-based algorithm and then employs a random forest model. The hybrid algorithm was observed to perform better, but it lacked scalability (Li & Cao, 2020). Therefore, we proceeded with the rule-based algorithm in the comparative analysis.

There were a number of fundamental differences between the previous studies and our approach. Firstly, while a few attributes, namely `DoNotTrack`, `accept` and `encoding` are not available within our datasets, they include significantly more attributes than FP-Stalker and Panopticlick. Thus, not only do the studies not share the same set of attributes, but also the algorithms of Vastel et al. and Eckersley were heavily dependent on the collected attributes. For instance, a subset of attributes is

allowed to change, whereas another subset of attributes is permitted to change only if the change ratio is below a certain threshold, and for some attributes, the algorithm mandates no changes. Thus, assigning the extra attributes collected in our study to one of such subsets arbitrarily, in an effort to make a complete adaptation would have introduced subjectivity and risked inconsistency. As a result, in order to implement both algorithms accurately, the dataset was cleared of the extra attributes, leaving only the set of attributes present in each of the studies. Both of the algorithms were adapted from the GitHub repository referenced in the FP-Stalker study.¹

Secondly, Panopticlick leveraged HTTP cookies for browser identification that were persistent up to three months. In FP-Stalker, the users had installed browser extensions that uniquely identified the browser instance. As mentioned earlier, the institutional dataset includes a user identifier number, but unlike the volunteer dataset, it does not contain a device or browser identifier. As the estimated precision metric proposed earlier is specific to the threshold algorithm, the precision measurement of previous work was only applicable to the volunteer dataset. Additionally, under some conditions, the rule-based algorithm of FP-Stalker imposes an additional homogeneity check for the browser identifiers of the candidates, which was suitable for the volunteer dataset but not for the institutional dataset. Therefore, to be able to compare the performances of all algorithms on the institutional data, two variants of the rule-based algorithm were evaluated, one yielding an upper bound and the other yielding a lower bound on the tracking time. The first variant is configured to always act as if the candidates belong to the same browser. This results in a maximum tracking duration. Conversely, the second variant proceeds as if the homogeneity check always fails, which keeps the number of links and the tracking time at a minimum.

Finally, Eckersley and Vastel et al. compare the new fingerprint with the whole set of fingerprints received so far, whereas, due to our use case, we only compare the new fingerprint to the previous fingerprints of the corresponding user. Thus, the algorithms proposed by Eckersley and Vastel et al. were modified to incorporate this consideration.

¹Open source at: <https://github.com/Spirals-Team/FPStalker>

4.5.1 Tracking Time and Precision

Table 4.5 summarizes the average tracking time of the algorithms on both datasets. For the comparison, ThresholdFP’s threshold value was set to 40 and 32 for the institutional dataset and the volunteer dataset, respectively. These threshold values were determined independently for each dataset based on the empirical testing presented in the Datasets and Results section, balancing the trade-off between precision and tracking time to achieve optimal performance. The best performance for the average tracking time on the institutional dataset was achieved by ThresholdFP with 55.7 days and over 98% estimated precision. Since the estimated precision was specifically designed to account for the mistakes of the threshold algorithm, it was not applicable to the other linking algorithms. Panopticlick surpassed the lower bound of FP-Stalker by 7.5 days, but FP-Stalker’s upper bound managed to outperform Panopticlick by 10.3 days. Regarding the analysis with the volunteer dataset, the algorithms were observed to perform better than on the institutional dataset, relative to the span of the data collection. This might be rooted in the frequency and the regularity of the data collection: fingerprints in the volunteer dataset were collected almost every day. Similar to the results for the institutional dataset, ThresholdFP outperformed the other two studies with 50.1 days of average tracking time. In the volunteer dataset, as the ground truth was available, FP-Stalker could be adapted faithfully, allowing for the homogeneity check mentioned above and therefore, rendering the exact performance measurement possible. In contrast to the findings of Vastel et al. (2018), Panopticlick displayed better performance than FP-Stalker, exceeding it by 6.8 days. On the other hand, the findings demonstrated that FP-Stalker achieved a precision of 97.6%, whereas Panopticlick achieved a precision of 92.3%. Surpassing both, ThresholdFP obtained an almost perfect precision of 99.5%.

Table 4.5 Comparison of the average tracking time obtained in number of days by ThresholdFP and previous work

Dataset	ThresholdFP	Panopticlick (Eckersley, 2010)	FP-Stalker (Vastel et al., 2018)	
			lower bound	upper bound
Institutional	55.7	34.5	27	44.8
Volunteer	50.1	45.6		38.8

Across both datasets and all three prior methods (Panopticlick, FP-Stalker-upper, FP-Stalker-lower), ThresholdFP consistently achieved longer tracking durations. The relative increase ranged from 24.33% (compared to the FP-Stalker upper bound)

to 106.30% (compared to the FP-Stalker lower bound) on the institutional dataset, with Panopticlick comparisons and all corresponding results from the volunteer dataset also falling within this range. This thesis reports this as a representative performance range in the Abstract.

4.5.2 Scalability

Li and Cao (2020) demonstrated that FP-Stalker is not scalable in a large-scale setting. Their experiments show that with as few as 200,000 fingerprints, the rule-based algorithm struggles to meet the latency expectations of real-time bidding (RTB), where ad decisions are typically required in under 100 milliseconds (Wang, Zhang & Yuan, 2017).

In our setting, as the user identifiers are provided together with the fingerprints, a new fingerprint is only compared to the user's previous fingerprints. Therefore, we evaluated scalability based on the number of users instead of the number of fingerprints, since the latter did not consistently impact algorithm runtime in this setting. Matching time was measured in a realistic environment by adapting both algorithms to include database connections and associated overhead.

The experiments were run on an AWS EC2 m6i.2xlarge instance with 32 GB DDR4 memory and an Intel® Xeon® Platinum 8375C (Ice Lake) processor with up to 3.5GHz turbo boost frequency.

We used the lower-bound variant of FP-Stalker's rule-based algorithm, although this choice does not affect the matching time measurements reported in our scalability evaluation.

The average matching time remained consistently low for both algorithms, with values generally below 1 ms and only occasionally exceeding it, significantly under the 100 ms threshold required by RTB systems. To enable comparison with Li et al.'s (2020) findings, we note that at the 28k user mark on the x-axis, over 230,000 data points were processed in total for calculating the "average matching time" on the y-axis.

While our approach introduces only a marginal increase in latency compared to FP-Stalker, it delivers a significant improvement in performance, extending the average tracking time by at least 24.3%. This notable gain highlights the effectiveness of our algorithm, especially given that the average matching time remains consistently

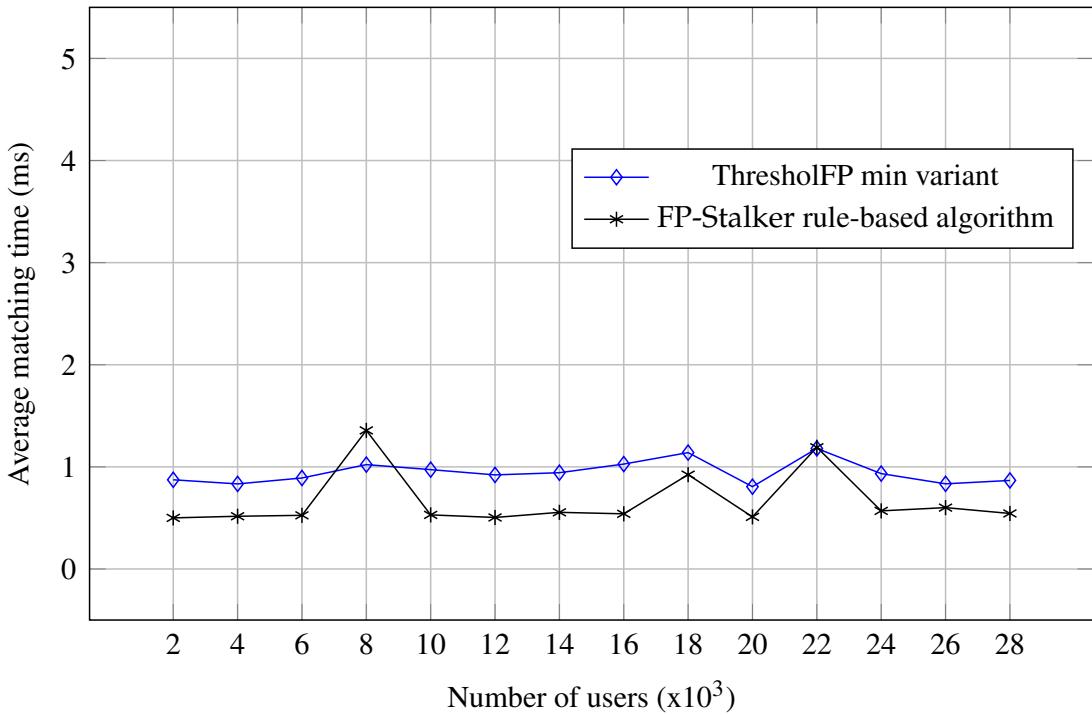


Figure 4.5 Average decision time of ThresholdFP and FP-Stalker's rule-based algorithm for one fingerprint

within just 0.3–0.4 ms of FP-Stalker. The slight latency overhead is thus a minimal cost for a substantial enhancement in tracking longevity, making our method a compelling alternative.

5. DISCUSSION AND ETHICAL ISSUES

This section discusses the implications of our findings, revisits the research questions, and addresses ethical considerations, implications and potential limitations of the study.

5.1 Discussion

Although it is intuitive to expect a monotonically increasing trend in average tracking time as the threshold increases, our models occasionally demonstrated fluctuating behavior. We identified several cases that contributed to such an outcome. For instance, larger thresholds bear an elevated chance of one fingerprint getting associated with another one that would not normally be considered similar. Consequently, as the difference scores that the threshold algorithm allows grow larger, the threshold merit of the lineage might get used up sooner, resulting in a shorter tracking time. Similarly, since the algorithm becomes more tolerant, the number of inaccurate links, which later get broken in remediation processes, increases and it gives rise to shorter fingerprint lineages. The combined effect of these cases may explain the formation of local minima in the curves.

An important observation was that, even though the volunteer dataset was gathered over a shorter period of time, the measurements of average tracking time on this dataset turned out to be close to, or even better than, those conducted on the institutional dataset for all of the algorithms. This outcome is in support of the finding of Vastel et al. (2018) that higher collection frequency results in a higher average tracking time for both Panopticlick and the rule-based variant of FP-Stalker.

Analyses across different datasets also show that individual attribute scores and optimal threshold values may vary. Thus, ThresholdFP is also dependent on data, just as the rival mechanisms it was compared to. However, this dependency is not continuous. When integrating ThresholdFP into a live system, default scores and a provisional threshold value can be used initially. As the system runs, the collected data can be used for fine-tuning as an offline process, and then the parameters can be updated. Alternatively, the system can first collect data to determine attribute scores and a threshold, and then initialize the algorithm using these values.

The expected trade-off between precision and tracking time is evident in the results. In other words, while tracking time increases, precision decreases with threshold. The non-monotonic behavior of both tracking time and precision metrics may even be helpful to choose optimal threshold values based on the system administrator's preference by picking threshold values of the right edges of the target metric value.

5.1.1 Addressing the Research Questions

This section discusses the findings in relation to the three research questions posed at the outset of the study. Each question is addressed in turn to show how the results contribute to the overall objectives of the research.

5.1.1.1 Research Question 1

To what extent can a threshold-based algorithm extend the effective lifespan of browser fingerprints?

The results in Section 4 show that our algorithm significantly increases the effective lifespan of browser fingerprints. In particular, ThresholdFP achieves an average tracking duration of 55.7 days on the institutional dataset and 50.1 days on the volunteer dataset, which represents a 24.33% to 106.30% improvement over prior schemes (Table 4.5). These findings indicate that a threshold-based linking approach can indeed extend fingerprint lifespan effectively under real-world conditions.

5.1.1.2 Research Question 2

Can benign changes in fingerprint attributes be tolerated without compromising linking precision?

ThresholdFP was able to tolerate benign changes across multiple fingerprinting attributes while maintaining high linking precision. As reported in Section 4, the algorithm achieved estimated precision above 98% on the institutional dataset. On the volunteer dataset, as browser instance-level ground truth was available, both estimated and actual precision could be measured. The results showed that the algorithm achieved an estimated precision of over 80% and an actual precision of 99.5%. The difference between actual and estimated precision suggests that the estimation method does not overstate performance and may, in fact, underestimate it in some cases. Moreover, the lower estimated precision observed on the volunteer dataset compared to the institutional dataset may stem from differences in the subject populations of the two datasets, as discussed in Section 4. These findings suggest that the approach can tolerate natural fluctuations in browser configurations without introducing significant mislinking.

5.1.1.3 Research Question 3

How does a threshold-based linking algorithm compare to existing methods in terms of tracking duration and precision?

When compared to adapted versions of existing algorithms by Vastel et al. (2018) and Eckersley (2010), ThresholdFP consistently outperformed them in tracking duration. On the institutional dataset, ThresholdFP achieved an average tracking time of 55.7 days, compared to an upper bound of 44.8 days and a lower bound of 27 days using the adapted FP-Stalker algorithm, and 34.5 days with the adapted Panopticlick algorithm. On the volunteer dataset, ThresholdFP reached 50.1 days, outperforming FP-Stalker (38.8 days) and Panopticlick (45.6 days).

In terms of precision, ThresholdFP demonstrated strong performance where browser instance-level ground truth was available. On the volunteer dataset, it achieved an actual precision of 99.5%. The baseline methods, evaluated only using actual precision due to metric constraints, achieved 97.6% (FP-Stalker) and 92.3% (Panopticlick). While the estimated precision metric was designed specifically for ThresholdFP and is not applicable to the baseline algorithms, it was shown to serve as

a conservative proxy for linking precision. On the volunteer dataset, the actual precision of ThresholdFP was consistently higher than its estimated precision (e.g., 99.5% vs. $\sim 80\%$), indicating that our method’s link decisions were more accurate than the estimation suggests. This further reinforces the reliability and robustness of ThresholdFP.

While prior measurement studies have raised concerns about the scalability of finger-print linking methods (Li & Cao, 2020), we note that these concerns were based on large-scale web-wide fingerprinting without access to user identifiers. In our experimental setup, user-level ground truth (user IDs) was available, which significantly reduced the candidate space and execution time for all methods. As a result, in our setting, all three algorithms demonstrated similar scalability in terms of runtime. We do not contradict earlier scalability findings, but show that in contexts where some user identity information is available (as is common in security-oriented applications), fingerprint linking, including ThresholdFP, can scale efficiently.

These results confirm that ThresholdFP not only improves tracking persistence significantly in a scalable way but also maintains high linking precision, making it a more robust and adaptive solution than existing rule-based or static fingerprint comparison schemes.

5.2 Ethical Considerations

The attributes were obtained via FingerprintJS, which is a well-known open-source browser fingerprinting library that is widely utilized by websites all around the world (Solomos et al., 2022). Apart from the attributes retrieved by FingerprintJS, no personal information was gathered from the users. Hence, none of the datasets contain any private information. The attributes were already being collected for a browser fingerprinting product employed by the ICT company due to security reasons. Data collection for our experiment was conducted under the existing consent agreements and setup. With regard to the volunteer dataset, the participants became a part of the experiment voluntarily after being informed about the data collection and providing their consent. Moreover, our system was designed in accordance with data minimization principles, and the fingerprint attributes collected are unrelated to the services provided on the websites, thereby not facilitating any behavioral or personal profiling.

5.3 Ethical Implications

This work aims to improve the persistence of browser fingerprints to support legitimate security use cases, such as fraud detection, device verification, and secondary authentication. By increasing the effective lifespan of a fingerprint under benign changes (e.g., software updates, configuration shifts), services can maintain continuity of user recognition without over-relying on fragile or user-resettable identifiers like cookies.

However, we acknowledge that improving fingerprint durability also raises privacy concerns. In particular, more stable fingerprints may be exploited to facilitate long-term user tracking without consent, especially in contexts outside security or authentication. In the wrong hands, these techniques could contribute to cross-site tracking, profiling, or surveillance that the user is not aware of.

To mitigate these concerns, we emphasize that this method should only be deployed in trusted, security-sensitive environments where users already expect some form of browser recognition. It should not be used for behavioral advertising or user profiling. Fingerprint histories should be stored securely, with limited retention and scope. Developers and organizations applying this method should ensure transparency and allow users to inspect or reset browser associations when applicable.

Additionally, users who wish to avoid fingerprint-based tracking may opt for browsers that actively interfere with or randomize fingerprint attributes. For example, Brave implements fingerprinting defenses by altering or reducing the stability of key attributes [Brave Privacy Team \(2020\)](#). Although we have not empirically tested the effectiveness of these browsers against our algorithm due to their absence in our dataset, their design philosophy aligns with efforts to reduce linkability over time and is therefore expected to reduce ThresholdFP’s durability.

We believe that publishing this technique contributes to the broader understanding of browser fingerprinting, both to improve defenses and to inform ethical deployment. Our intention is to support more robust and user-friendly authentication systems, not to enable covert tracking. Continued dialogue among browser developers, privacy advocates, and the research community is essential to ensure fingerprinting technologies are used responsibly.

5.4 Threats to Validity

In order to address possible limitations, we considered threats to internal, external and construct validity:

Threats to internal validity: For the reasons discussed in Section 4.5, the algorithms of Eckersley (2010) and Vastel et al. (2018), in their original shape, were not suitable for our analyses, rendering our intervention inevitable. Despite efforts to faithfully replicate the original algorithms, it is possible that unintentional errors were introduced during adaptation. Furthermore, the GitHub repository provided in Vastel et al.'s work was considered as the source material. If there was a bug in the source code, it might have propagated to our adaptations.

Threats to external validity: We analyzed the performance of all algorithms on two separate datasets that were collected independently. The institutional dataset was gathered by an ICT company, and the volunteer dataset had all of its participants from an academic environment. In addition to the population, the frequency of the data differed remarkably, reducing the chance of bias. The volunteer dataset consisted of fingerprints collected daily, with occasional exceptions, unlike the institutional dataset, which was composed of fingerprints collected at arbitrary times. Regardless of such differences, our approach performed the best in both of the comparative analyses employing either one of the datasets, and the performances of other algorithms were consistent with our results. Nevertheless, the participants in our study may not represent the broader user population, and our approach might not yield the same performance with different datasets.

Additionally, the time span of data collection was 2.5 months and 5.5 months, which might be deemed short in comparison to the prominent studies in the field. Finally, due to privacy reasons, we are unable to share the datasets, and that prevents the reproducibility of this work. Nonetheless, this is a common limitation in fingerprinting studies.

Threats to construct validity: Since the fingerprints in the institutional dataset did not contain browser identifiers, we formulated an "estimated precision" metric, as explained in Section 4.1 and Section 4.3.1. There is a possibility that this metric might not model the actual precision successfully or that it might cause inflated results. However, the analysis on the volunteer dataset revealed that actual precision was significantly better than the estimated precision, as illustrated in Figure 4.3, indicating that the estimated precision was not an overestimate of the actual precision.

6. CONCLUSION

In this thesis, we proposed ThresholdFP, a threshold-based fingerprint linking algorithm that associates new fingerprints with previously encountered ones. Fingerprints can change due to factors such as attribute instability, system evolution, and user intervention. ThresholdFP aims to mitigate these challenges in order to achieve an extended tracking time for the fingerprints.

To evaluate ThresholdFP’s performance, two datasets were generated and used. The first dataset was composed of 235,422 fingerprints collected over a span of 5.5 months. The second data collection phase, lasting a total of 2.5 months, was initiated to form a secondary dataset. Unlike the first dataset, this dataset included manually labeled browser identifiers as ground truth. In the end, the second dataset contained 5,774 fingerprints in total.

To decide which strategy to follow when multiple candidates are present to associate with a fingerprint, three variants of our threshold-based algorithm were created. Furthermore, all three variants were compared against two previous studies (Eckersley, 2010; Vastel et al., 2018), which also put forward rule-based algorithms to link fingerprints. A separate analysis was conducted for each of the datasets. In the analysis of the first dataset, instead of an exact measurement, upper and lower bounds were calculated for FP-Stalker. This is due to the fact that the original algorithm applied a check that was not applicable to our dataset. Therefore, we mimicked its behavior to maximize and minimize the average tracking time, obtaining the upper and lower bounds, respectively. The ThresholdFP variant that selects the candidate with the least difference to the new fingerprint yielded the best performance on the first dataset when the threshold was set to 40. At this threshold value, the algorithm achieved an average tracking time of more than 55 days with an estimated precision exceeding 98%, compared to 34.5 days by Panopticlick and 27 to 44.8 days by FP-Stalker.

The same variant of ThresholdFP also achieved the best performance on the second dataset, with 50.1 days and 99.5% precision, surpassing Panopticlick and FP-Stalker, which attained an average tracking time of 45.6 and 38.8 days, and a precision of 92.3% and 97.6%, respectively.

BIBLIOGRAPHY

Abromaityte, M., Dubero, E., Kruger, Q., Lucassen, B., Vos, T., van den Hout, N. J., Bouma, D., Treur, J., & Roelofsma, P. H. M. P. (2025). Computational analysis of user experience of password-based authentication systems. In Singh, D., van 't Klooster, J.-W., & Tiwary, U. S. (Eds.), *Intelligent Human Computer Interaction*, (pp. 341–352)., Cham. Springer Nature Switzerland.

Acar, G., Eubank, C., Englehardt, S., Juarez, M., Narayanan, A., & Diaz, C. (2014). The web never forgets: Persistent tracking mechanisms in the wild. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, (pp. 674–689).

Ajay, V. L. & Guptha, A. M. (2022). A defense against javascript object-based fingerprinting and passive fingerprinting. In *2022 International Conference on Computing, Communication, Security and Intelligent Systems (IC3SIS)*, (pp. 1–6).

Alaca, F. & van Oorschot, P. C. (2016). Device fingerprinting for augmenting web authentication: classification and analysis of methods. In *Proceedings of the 32nd Annual Conference on Computer Security Applications*, ACSAC '16, (pp. 289–301)., New York, NY, USA. Association for Computing Machinery.

Alrawili, R., AlQahtani, A. A. S., & Khan, M. K. (2024). Comprehensive survey: Biometric user authentication application, evaluation, and discussion. *Computers and Electrical Engineering*, 119, 109485.

Alzomai, M., Josang, A., McCullagh, A., & Foo, E. (2008). Strengthening sms-based authentication through usability. In *2008 IEEE International Symposium on Parallel and Distributed Processing with Applications*, (pp. 683–688).

Andriamilanto, N., Allard, T., Le Guelvouit, G., & Garel, A. (2021). A large-scale empirical analysis of browser fingerprints properties for web authentication. *ACM Trans. Web*, 16(1).

Boussaha, S., Hock, L., Bermejo, M., Rumin, R. C., Rumin, A. C., Klein, D., Johns, M., Compagna, L., Antonioli, D., & Barber, T. (2024). Fp-tracer: Fine-grained browser fingerprinting detection via taint-tracking and entropy-based thresholds. *Proceedings on Privacy Enhancing Technologies*, 2024(3), 540–560.

Brave Privacy Team (2020). Fingerprint randomization.

Cao, Y., Li, S., & Wijmans, E. (2017). (cross-) browser fingerprinting via os and hardware level features. In *Proceedings of the 2017 Network and Distributed System Security Symposium (NDSS)*, (pp. 1–15). Internet Society.

Durey, A. (2022). *Leveraging browser fingerprinting to strengthen web authentication*. PhD thesis, Université de Lille.

Durey, A., Laperdrix, P., Rudametkin, W., & Rouvoy, R. (2021). Fp-redemption:

Studying browser fingerprinting adoption for the sake of web security. In Bilge, L., Cavallaro, L., Pellegrino, G., & Neves, N. (Eds.), *Detection of Intrusions and Malware, and Vulnerability Assessment*, (pp. 237–257)., Cham. Springer International Publishing.

Eckersley, P. (2010). How unique is your web browser? In *Privacy Enhancing Technologies: 10th International Symposium, PETS 2010, Berlin, Germany, July 21-23, 2010. Proceedings 10*, (pp. 1–18). Springer.

Englehardt, S. & Narayanan, A. (2016). Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, (pp. 1388–1401).

FingerprintJS (2023). Fingerprintjs: The most advanced browser fingerprinting library. <https://github.com/fingerprintjs/fingerprintjs>.

Gómez-Boix, A., Laperdrix, P., & Baudry, B. (2018). Hiding in the crowd: an analysis of the effectiveness of browser fingerprinting at large scale. In *Proceedings of the 2018 world wide web conference*, (pp. 309–318).

Iqbal, U., Englehardt, S., & Shafiq, Z. (2021). Fingerprinting the fingerprinters: Learning to detect browser fingerprinting behaviors. In *2021 IEEE Symposium on Security and Privacy (SP)*, (pp. 1143–1161). IEEE.

Jonker, H., Krumnow, B., & Vlot, G. (2019). Fingerprint surface-based detection of web bot detectors. In *Computer Security—ESORICS 2019: 24th European Symposium on Research in Computer Security, Luxembourg, September 23–27, 2019, Proceedings, Part II 24*, (pp. 586–605). Springer.

Laperdrix, P., Avoine, G., Baudry, B., & Nikiforakis, N. (2019). Morellian analysis for browsers: Making web authentication stronger with canvas fingerprinting. In *Detection of Intrusions and Malware, and Vulnerability Assessment: 16th International Conference, DIMVA 2019, Gothenburg, Sweden, June 19–20, 2019, Proceedings 16*, (pp. 43–66). Springer.

Laperdrix, P., Bielova, N., Baudry, B., & Avoine, G. (2020). Browser fingerprinting: A survey. *ACM Transactions on the Web (TWEB)*, 14(2), 1–33.

Laperdrix, P., Rudametkin, W., & Baudry, B. (2016). Beauty and the beast: Diverting modern web browsers to build unique browser fingerprints. In *2016 IEEE Symposium on Security and Privacy (SP)*, (pp. 878–894). IEEE.

Laperdrix, P., Starov, O., Chen, Q., Kapravelos, A., & Nikiforakis, N. (2021). Fingerprinting in style: Detecting browser extensions via injected style sheets. In *30th USENIX Security Symposium (USENIX Security 21)*, (pp. 2507–2524).

Li, S. & Cao, Y. (2020). Who touched my browser fingerprint? a large-scale measurement study and classification of fingerprint dynamics. In *Proceedings of the ACM Internet Measurement Conference*, (pp. 370–385).

Lin, X., Araujo, F., Taylor, T., Jang, J., & Polakis, J. (2023). Fashion faux pas: Implicit stylistic fingerprints for bypassing browsers' anti-fingerprinting defenses. In *2023 IEEE Symposium on Security and Privacy (SP)*, (pp. 987–1004).

Lin, X., Ilia, P., Solanki, S., & Polakis, J. (2022). Phish in sheep's clothing: Ex-

ploring the authentication pitfalls of browser fingerprinting. In *31st USENIX Security Symposium (USENIX Security 22)*, (pp. 1651–1668).

Marky, K., Ragozin, K., Chernyshov, G., Matviienko, A., Schmitz, M., Mühlhäuser, M., Eghtebas, C., & Kunze, K. (2022). “nah, it’s just annoying!” a deep dive into user perceptions of two-factor authentication. *ACM Transactions on Computer-Human Interaction*, 29(5).

Mayer, J. R. (2009). Any person... a pamphleteer”: Internet anonymity in the age of web 2.0. *Undergraduate Senior Thesis, Princeton University*, 85.

Mowery, K. & Shacham, H. (2012). Pixel perfect: Fingerprinting canvas in html5. In *Proceedings of the Web 2.0 Security and Privacy Workshop (W2SP)*, (pp. 1–12).

Mozilla Developer Network (2024). Web apis. <https://developer.mozilla.org/en-US/docs/Web/API>. Accessed: 2024.

Mulliner, C., Borgaonkar, R., Stewin, P., & Seifert, J.-P. (2013). Sms-based one-time passwords: Attacks and defense. In Rieck, K., Stewin, P., & Seifert, J.-P. (Eds.), *Detection of Intrusions and Malware, and Vulnerability Assessment*, (pp. 150–159)., Berlin, Heidelberg. Springer Berlin Heidelberg.

Nikiforakis, N., Kapravelos, A., Joosen, W., Kruegel, C., Piessens, F., & Vigna, G. (2013). Cookieless monster: Exploring the ecosystem of web-based device fingerprinting. In *2013 IEEE Symposium on Security and Privacy*, (pp. 541–555). IEEE.

Pugliese, G., Riess, C., Gassmann, F., & Benenson, Z. (2020). Long-term observation on browser fingerprinting: Users’ trackability and perspective. *Proceedings on Privacy Enhancing Technologies*, 2020(2), 558–577.

Reese, K., Smith, T., Dutson, J., Armknecht, J., Cameron, J., & Seamons, K. (2019). A usability study of five Two-Factor authentication methods. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*, (pp. 357–370)., Santa Clara, CA. USENIX Association.

Senol, A., Ukani, A., Cutler, D., & Bilogrevic, I. (2024). The double edged sword: Identifying authentication pages and their fingerprinting behavior. In *Proceedings of the ACM Web Conference 2024, WWW ’24*, (pp. 1690–1701)., New York, NY, USA. Association for Computing Machinery.

Solomos, K., Ilia, P., Nikiforakis, N., & Polakis, J. (2022). Escaping the confines of time: Continuous browser extension fingerprinting through ephemeral modifications. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, (pp. 2675–2688).

Starov, O. & Nikiforakis, N. (2017). Xhoud: Quantifying the fingerprintability of browser extensions. In *2017 IEEE Symposium on Security and Privacy (SP)*, (pp. 941–956). IEEE.

Takei, N., Saito, T., Takasu, K., & Yamada, T. (2015). Web browser fingerprinting using only cascading style sheets. In *2015 10th International Conference on Broadband and Wireless Computing, Communication and Applications*

(*BWCCA*), (pp. 57–63). IEEE.

Trampert, L., Weber, D., Gerlach, L., Rossow, C., & Schwarz, M. (2025). Cascading spy sheets: Exploiting the complexity of modern css for email and browser fingerprinting. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, (pp. 1–20)., Reston, VA, USA. Internet Society.

Vastel, A., Laperdrix, P., Rudametkin, W., & Rouvoy, R. (2018). Fp-stalker: Tracking browser fingerprint evolutions. In *2018 IEEE Symposium on Security and Privacy (SP)*, (pp. 728–741). IEEE.

Vastel, A., Rudametkin, W., Rouvoy, R., & Blanc, X. (2020). Fp-crawlers: Studying the resilience of browser fingerprinting to block crawlers. In *Proceedings of the MADWeb’20—NDSS Workshop on Measurements, Attacks, and Defenses for the Web*, (pp. 1–13).

Wang, J., Zhang, W., & Yuan, S. (2017). Display advertising with real-time bidding (rtb) and behavioural targeting. *Foundations and Trends® in Information Retrieval*, 11(4-5), 297–435.

Wu, S., Sun, P., Zhao, Y., & Cao, Y. (2023). Him of many faces: Characterizing billion-scale adversarial and benign browser fingerprints on commercial websites. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, (pp. 1–16).