

# Quotable Signatures for Authenticating Shared Quotes<sup>\*</sup>

Joan Boyar<sup>1</sup>[0000-0002-0725-8341], Simon Erfurth<sup>1</sup>[0000-0001-8862-2856], Kim S.  
Larsen<sup>1</sup>[0000-0003-0560-3794], and Ruben Niederhagen<sup>1,2</sup>

<sup>1</sup> University of Southern Denmark, Odense, Denmark

joan@imada.sdu.dk

simon@serfurth.dk

kslarsen@imada.sdu.dk

<sup>2</sup> Academia Sinica, Taipei, Taiwan

ruben@polycephaly.org

**Abstract.** Quotable signature schemes are digital signature schemes with the additional property that from the signature for a message, any party can extract signatures for (allowable) quotes from the message, without knowing the secret key or interacting with the signer of the original message. Crucially, the extracted signatures are still signed with the original secret key. We define a notion of security for quotable signature schemes and construct a concrete example of a quotable signature scheme, using Merkle trees and classical digital signature schemes. The scheme is shown to be secure, with respect to the aforementioned notion of security. Additionally, we prove bounds on the complexity of the constructed scheme. Finally, concrete use cases of quotable signatures are considered, using them to combat misinformation by bolstering authentic content on social media. We consider both how quotable signatures can be used, and why using them could help mitigate the effects of fake news.

**Keywords:** quotable signatures · digital signatures · Merkle trees · authenticity · fake news

## 1 Introduction

Digital signature schemes are a classical and widely used tool in modern cryptography (the canonical reference is [11], and [6] contains some current standards). A somewhat newer concept is *quotable signature schemes* [18], which are digital signature schemes with the additional property that signatures are *quotable* in the following sense. The *Signer* of a message  $m$  generates a quotable signature  $s$  for  $m$  using a private key  $sk$ . Given a message  $m$  and the quotable signature  $s$ , a

---

<sup>\*</sup> The first and third authors were supported in part by the Independent Research Fund Denmark, Natural Sciences, grant DFF-0135-00018B. All authors are currently associated with DDC – the Digital Democracy Center at the University of Southern Denmark.

*Quoter* (any third party) can extract a second quotable signature  $s'$  for a quote  $q$  from  $m$  without knowing  $sk$  or interacting with the original *Signer*. A quote can be any “allowable subsequence” of  $m$ . We write  $q \preceq m$  to indicate that  $q$  is a quote from  $m$ . This quotable signature  $s'$  is still signed with the private key  $sk$  of the *Signer* and hence authenticates the original *Signer* as the author of the quote. These signatures for quotes have the same required properties with respect to verification and security as a standard digital signature, in addition to allowing one to derive where content has been removed, relative to the quote. A signature for a quote is again a quotable signature with respect to sub-quotes of the quote, and neither authenticating a quote nor sub-quoting require access to the original message.

Quotable signatures can be used to mitigate the effects of fake news and disinformation. These are not new problems, and it is becoming increasingly apparent that they are posing a threat for democracy and for society. There is not one single reason for this, but one reason among many is a fundamental change in how news is consumed: a transition is happening, where explicit news products such as printed newspapers and evening news programs are still consumed, but are increasingly giving way for shorter formats and snippets of news on social media platforms [25]. However, people tend to be unable to recall from which news brand a story originated when they were exposed to it on social media [17]. This is problematic since the news media’s image is an important heuristic when people evaluate the quality of a news story [34]. In addition, according to the Reuters Institute Digital News Report 2022 [26], across markets, 54% of those surveyed say they worry about identifying the difference between what is real and fake on the Internet when it comes to news, but people who say they mainly use social media as a source of news are more worried (61%).

In recent years, a common approach to fighting back against fake news has been flagging (potentially) fake news, using either manual or automatic detection systems. While this might be a natural approach, research has shown repeatedly that flagging problematic content tends to have the opposite result, i.e., it increases the negative effects of fake news [13, 21, 29]. This indicates that flagging problematic content is not sufficient and alternative approaches need to be developed.

We present a method that complements flagging problematic content with the goal of mitigating the effect of fake news. Our idea builds on the observation that *which* news media published a news article is an important heuristic people use to evaluate the quality of the article [34]. However, since people get their news increasingly via social media, it is becoming more likely that they are not aware of who published the news they are consuming. To address this, we propose using quotable signatures to allow people on social media to find out and be certain of where the text they are reading originates from, and to verify that any modifications to the text were all allowed. Specifically for news, the proposed idea is that a news media publishing an article also publishes a quotable signature for the article signed with their private key. When someone shares a quote from the article, they then also include the signature for the quote that is derived from

the initial signature (without access to the private key), which we emphasize is signed with the same key. Finally, when one reads the quote, the signature can be checked, and it can be verified from where the quote originates.

The idea of mitigating the effects of fake news and misinformation, by using digital signatures to verify the source of media content, is one that has been addressed by others. One example is C2PA [9], which involves many companies, including Adobe, the BBC, Microsoft, and Twitter. C2PA focuses on providing a history of a published item, i.e., which device was used to capture it, how it has been edited and by whom, etc. Thus, quotable signatures could be of interest to their approach.

Another issue involving fake news is that news articles are perceived as more credible if they contain attributed quotes [32]. This is misused by fake news to appear more credible by providing attributions for their content [2, 30, 19, 27, 12], but can in turn be used to automatically detect fake news by considering the existence and quality of attributions [1, 33, 24] (among other things). Quotable signatures, in contrast, could be used to sign quotes to make a strong and verifiable connection between the original source and the quote. On the other hand, fake news would generally not be able to link their quotes to reputable sources, thereby providing another heuristic helping users to distinguish between authentic and fake content.

Without major changes to the system, it could be extended to further use cases such as signing Facebook and Twitter posts, official governmental rules and regulations, scientific publications, etc. For all of these instances, an important feature of our system that we have not used explicitly so far is that signing also binds the Signer, meaning that the signing party cannot later deny having signed the signed document.

We provide an overview over related work in Section 2. In Section 3, we give a more thorough introduction to and definition of quotable signatures, and we show how we can realize quotable signatures using Merkle trees [22, 23]. We define a notion of security for quotable signature schemes, and prove that the notion is satisfied by our construction. Additionally, we prove a number of bounds on the size and computational costs of quotable signatures obtained using Merkle trees. We revisit the application of quotable signatures to counter fake news in more detail in Section 4 and we conclude the paper with an outlook to future work in Section 5. In the full version of this paper we also describe concrete algorithms for our construction of quotable signatures from Merkle trees.

## 2 Related Work

Quotable signatures have been introduced in [18], which suggests constructing quotable signatures using Merkle trees and provides a rudimentary complexity analysis. The authors also suggest using quotable signatures to mitigate the effects of fake news. Compared to [18], we define a security model, and prove that our construction is secure in this security model. Additionally, we also provide proofs of our claims about the cost of using Merkle trees for quotable

signatures, and provide more in-depth considerations for why one could expect this to be a good approach.

A concept closely related to quotable signature schemes is *redactable signature schemes* (RSSs). Simultaneously introduced in [31] (as *Content Extraction Signatures*) and [16], RSSs essentially allow an untrusted redactor to remove (“redact”) parts of a signed message, without invalidating the signature. Often this requires modifying the signature, but crucially, it is still signed with the original key, despite the redactor not having access to the private key. Thus, quotable signatures share many similarities with RSSs; if one considers a quotation as a redaction of all parts of a text except for the quote, they are conceptually identical. Where quotable signatures and RSSs differ is in the security they must provide. Both signature schemes require a similar notion of unforgeability, but an RSS must also guarantee that the redacted parts remain private. A standard formulation is that an outsider not holding any private keys should “not be able to derive any information about redacted parts of a message”, and even stronger requirements, such as transparency or unlinkability, are not uncommon [7]. Quotable signatures have no such privacy requirements, allowing quotable signatures to be faster. In fact, it is worth noting that there are scenarios where RSSs’ notion of privacy would be directly harmful to a quotable signature. For instance, RSS would specifically make it impossible to tell if a quote is contiguous or not, something that we consider essential for a quotable signature scheme. To see the value of dropping the privacy requirement, we observe that some RSSs with  $O(n)$  performance may have  $O(n)$  expensive public key cryptography operations [8, 28], whereas quotable signatures can be obtained with  $O(n)$  (cheap) symmetric cryptographic operations (hashing), and only one expensive public key operation. There are approaches obtaining RSSs using only one expensive operation, but they either require many more cheap operations than quotable signatures do, or they result in considerably larger signatures, for example [14]. Early examples of RSSs had a weaker notion of privacy, but still stronger than what we require. They require only hiding of the redacted elements, not their location and number. Examples can be found in [16, 31]. Their approaches are similar to ours, also using Merkle trees, but we provide rigorous proofs of the claimed performance, and our lack of privacy requirements allows our scheme to be both more efficient and conceptually simpler. One consideration that is very relevant for quotable signatures, is how a quote (redaction) being contiguous will affect the complexity results. In a different setting [10] considers this question for Merkle trees, but provides no rigorous proof.

Considering the motivating example again, approaches to mitigate the impact of fake news, using either digital signatures or directly rating the source of the content, have been proposed and tried before. One approach, serving as inspiration for our approach, is [3]. They use digital signatures to verify the authenticity of images and other forms of multimedia. One drawback of their implementation is that it requires the media to be bit-for-bit identical to the version that was signed. Hence, the image can for instance not be compressed or resized, and thus their solution is not compatible with many platforms, e.g.,

Facebook compresses uploaded images, and many news websites resize images for different screen sizes. An example of directly rating the source of content, and flagging trustworthy sources, can be found in “NewsGuard Ratings” (NG), which provides a rating of trustworthiness for news sources. NG adds a flag that indicates if a news source is generally trustworthy (green) or not (red) to websites and outgoing links on websites. This approach has not been widely successful. For example, the study in [4] shows that NG’s labels have “limited average effects on news diet quality and fail to reduce misperceptions”. While this is somewhat related to our approach, there are two major differences. (1) NG only flags content that directly links to the source of the content with a URL. In contrast, our digital signature can be attached to any text quote. Hence, NG only adds additional information when it is already straightforward to figure out from where the content originates. Our approach also provides this information where there might otherwise be no clear context. (2) NG focuses on providing a rating for how trustworthy a news source is. This approach is similar to the typical approach of telling people when something might be problematic, which tends to have the opposite result. In contrast, we focus solely on providing and authenticating the source of a quote.


Summing up, the contributions of this paper is as follows. (1) We rigorously define the notion of security that quotable signature schemes must satisfy. (2) We rigorously prove the security of and analyze the complexity of, a quotable signature scheme constructed using Merkle trees. (3) This provides a scheme for quotable signatures that is more efficient than using an RSS for the same purpose.

### 3 Quotable Signatures

To construct a quotable signature scheme, we follow the approach suggested in [18] and use a combination of a classical digital signature scheme [11] and Merkle trees [22, 23].

Before getting into the construction, we summarize the setting of quotable signatures. In Section 3.1, we define the security notion that quotable signature schemes should satisfy. Then, in Section 3.2, we introduce Merkle trees, in Section 3.3 we construct a quotable signature scheme and show it is secure, and finally we analyze the complexity of the scheme in Section 3.4.

*General setting for quotable signatures.* A quotable signature scheme consists of four efficient algorithms,  $QS = (\text{KeyGen}, \text{Sign}, \text{Quo}, \text{Ver})$ . These four algorithms are essentially the standard three algorithms from a classical digital signature scheme for key generation, signing, and verification, with the added quoting algorithm **Quo**. To quote from a message, **Quo** allows extracting a valid signature for the quote from the signature of the message in such a way that it is still signed with the public key used to sign the original message. Additionally, it should be possible to derive from the signature of a quote where tokens from the original message have been removed relative to the quote.



**Fig. 1.** The general setting for a quotable signature.

We refer to the involved parties as the *Signer*, the *Quoter*, and the *Verifier*. We use  $\lambda$  to denote the security parameter. To summarize:

- $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\lambda)$  takes as input the security parameter  $1^\lambda$ . It outputs a public key pair. This is typically done by the Signer once, offline as part of the initial setup.
- $s \leftarrow \text{Sig}_{\text{sk}}(m)$  takes as input a secret key  $\text{sk}$  and a message  $m$ . It outputs a quotable signature for  $m$ . This is done by the Signer.
- $s' \leftarrow \text{Quo}(m, q, s)$  takes as input a message  $m$ , a quote  $q$  from  $m$ , and a quotable signature  $s$  for  $m$ . It outputs a quotable signature  $s'$  for  $q$ , that is still signed with the secret key used to generate  $s$ . Verifying  $s'$  does not require knowing  $m$ . Note that  $m$  and  $s$  could have been obtained via an earlier quote operation. This is done by the Quoter.
- $\top/\perp \leftarrow \text{Ver}_{\text{pk}}(q, s')$  takes as input a public key  $\text{pk}$ , a quote (message)  $q$ , and a signature  $s'$  for  $q$ . It outputs  $\top$  if  $s'$  is a valid signature for  $q$  with respect to  $\text{pk}$ , and  $\perp$  otherwise. This is done by the Verifier.

Figure 1 illustrates the typical interactions between the parties.

### 3.1 Security Model

Taking inspiration from the RSS notion of unforgeability, we define the security notion of quotable signatures schemes in Definition 1. At its core, this is the standard notion of unforgeability for digital signature schemes, with the additional requirement that the adversary’s chosen message cannot be a quote from any of the messages that the adversary sent to the signing oracle.

**Definition 1 (Unforgeability).** *Let  $\text{QS} = (\text{KeyGen}, \text{Sign}, \text{Quo}, \text{Ver})$  be a quotable signature scheme. We say that QS is existentially unforgeable, if for every probabilistic polynomial time adversary  $\mathcal{A}$ , the probability of the following experiment returning 1 is negligible:*

```

(pk, sk) ← KeyGen(1λ)
(m*, s*) ←  $\mathcal{A}^{\text{Sign}_{\text{sk}}(\cdot)}(\text{pk})$ 
// denote the queries that  $\mathcal{A}$  make to the signing oracle by  $m_1, m_2, \dots, m_Q$ .
if (Verpk(m*, s*) =  $\top$ ) ∧ (∀k ∈ {1, 2, ..., Q}: m*  $\not\leq$  mk)
  return 1
  
```

### 3.2 Merkle Trees

A Merkle tree (also known as a *hash tree*) allows one to efficiently and securely verify that one or more *tokens* are contained in a longer sequence of tokens, without having to store the entire sequence [22, 23]. Examples of this could be words forming a sentence, sentences forming an article, or data blocks making up a file.

Since our scheme will rely on hash functions, we assume that the tokens are binary strings. Equivalently, one could assume an implicitly used, well defined injective mapping from the token space to the space of binary strings. For data blocks, the identity mapping suffices and for words one such mapping could be the mapping of words to their UTF-8 representations.

The structure of a Merkle tree for a sequence of tokens is a binary tree, where each leaf corresponds to a token from the sequence, with the leftmost leaf corresponding to the first token, its sibling corresponding to the second token, and so on. Each leaf is labeled with the hash of its token and each internal node is labeled with the hash of the concatenation of the labels of its children. Hence, the  $i$ 'th internal node on the  $j$ 'th level will be labeled as

$$u_{j,i} = H(u_{j+1,2i} \parallel u_{j+1,2i+1}).$$


This way, one can show that any specific token is in the sequence by providing the “missing” hashes needed to calculate the hashes on the path from the leaf corresponding to the token to the root of the tree. Following established terminology, we call this the *verification path* for the token.<sup>3</sup>

Figure 2 shows the Merkle tree for a sequence of words forming the sentence “The quick brown fox jumps over the dog”. The verification path for the word “jumps” consisting of nodes  $u_{3,5}$ ,  $u_{2,3}$ , and  $u_{1,0}$  is highlighted in red. Similarly, one can obtain the verification path for a subsequence of more than just one token. In Figure 2, we also indicate the verification path for the contiguous subsequence “the quick” in blue. Note that the size of the verification path depends not only on how many tokens are chosen, but also on where in the sequence they are placed. In Section 3.4, we analyze how large the verification path can become, i.e., how many nodes need to be provided in the signature in the worst case.

In these examples, we have chosen a sequence of tokens where the length of the sequence, i.e., the number of tokens, is a power of two. If the sequence length is not a power of two, we require that the tree is *heap-shaped*, i.e., all levels are filled, except for possibly the lowest level, which is filled from the left up to some point, after which the lowest level is empty.

*Remark 1.* Observe that from the structure of the Merkle tree, one can see where in the sequence the quoted tokens are placed, and if they are sequential or discontinuous.

<sup>3</sup> This use of “path” is slightly counter intuitive, since it refers to the hashes needed to calculate the hashes on the path from the leaf to the root, and hence not the nodes on this path but their siblings.



**Fig. 2.** An example of a Merkle tree where the tokens are words and the sequence is a sentence. The verification path for the token “jumps” is highlighted in red ( $u_{1,0}, u_{2,3}, u_{3,5}$ ), and the verification path for the subsequence “The quick” is highlighted in blue ( $u_{1,1}, u_{2,1}$ ).

### 3.3 A Quotable Signature Scheme

Using a Merkle tree, we can now devise a scheme by which the Quoter can convince the Verifier that some quote is contained in a larger text, if the Verifier is already in possession of the root hash. The Quoter simply shares the verification path together with the quote, and the Verifier verifies that this indeed leads to the original root hash. In order to turn this into a quotable signature scheme, we include a classical digital signature for the root hash, signed by the Signer, with the verification path. Thus, letting  $DS = (\text{KeyGen}^{\text{DS}}, \text{Sign}^{\text{DS}}, \text{Ver}^{\text{DS}})$  be a classical digital signature scheme, our quotable signature scheme can be described as follows:

- **KeyGen:** Identical to  $\text{KeyGen}^{\text{DS}}$ .
- **Sign:** Find the root hash of the Merkle tree and sign it with  $\text{Sign}^{\text{DS}}$ .
- **Quo:** Find the verification path of the quote. Together with the signature of the root hash, this forms the signature for the quote.
- Find the root hash of the Merkle tree using the quote and its verification path. Use  $\text{Ver}^{\text{DS}}$  to verify the authenticity of the root hash.

**Proof of Security** We will show that the construction of the previous section is secure with respect to the notion of security introduced in Definition 1, when instantiated with a secure hash function and a secure classical signature scheme. Before doing so, we observe that currently, our scheme is trivially vulnerable to a forgery attack, as follows. An adversary obtains a quotable signature for a message from a signing oracle and then simply replaces the last two tokens on the lowest level with a single token, which is the concatenation of the tokens’ hashes. For example, using this attack on the message used in Figure 2, would give the message “The quick brown fox jumps over  $H(\text{the}) \parallel H(\text{dog})$ ”. However, there is an easy fix to this vulnerability. Noting that the problem is that an adversary can claim that an internal node is a leaf, we can prevent this by applying domain separation in the form of adding one value to the leaves before hashing, and another value to the internal nodes before hashing. Taking inspiration from RFC



6962 [20], the Merkle trees are modified by prepending 00 to the leaves before hashing and 01 to the internal nodes before hashing. From now on, we implicitly assume that this is done.

We can now argue that the construction is secure.

**Theorem 1.** *Under the assumption that*

- $H$  comes from a family of cryptographic hash function,
- $DS = (\text{KeyGen}^{\text{DS}}, \text{Sign}^{\text{DS}}, \text{Ver}^{\text{DS}})$  is an existentially unforgeable classical signature scheme,

$QS = (\text{KeyGen}, \text{Sign}, \text{Quo}, \text{Ver})$  constructed as described above, is an existentially unforgeable quotable signature scheme.

We have to show that no probabilistic polynomial time adversary can win the unforgeability experiment in Definition 1 with non-negligible probability.

*Proof.* Assume that  $\mathcal{A}$  is a probabilistic polynomial time adversary against the unforgeability of QS. We show that the probability of  $\mathcal{A}$  being successful is negligible. Let  $(m^*, s^*)$  be the output of  $\mathcal{A}$ , where  $s^* = (\text{Sig}_{\text{sk}}^{\text{DS}}(u_{0,0}^*), \{u_{i,j}^*\})$ , i.e., the classical digital signature of the root hash and a (possibly empty) verification path.

Consider first the case where the root hash  $u_{0,0}^*$  of  $m^*$  (found using  $\{u_{i,j}^*\}$ ) is different from the root hashes of the queries  $\mathcal{A}$  made to the signing oracle. In this case,  $(u_{0,0}^*, \text{Sig}_{\text{sk}}^{\text{DS}}(u_{0,0}^*))$  is a forgery against DS, and since DS is assumed to be existentially unforgeable, this can only happen with negligible probability. Denote this probability as  $\epsilon_{\text{DS}}$ .

If this is not the case, there must be an  $m_k$ , such that the root hash of  $m_k$  is  $u_{0,0} = u_{0,0}^*$ , but  $m^* \not\leq m_k$ . Denote by  $T^*$  the tree for  $m^*$  (constructed using the verification path, if one is included) and by  $T$  the tree for  $m_k$ .

Consider first the case where all leaves, corresponding to tokens, in  $T^*$  are at a location in the tree, where there is also a leaf, corresponding to a token, in  $T$ . Since  $m^* \not\leq m_k$  there must be tokens  $a^*, a$  such that  $a^* \in m^*$  and  $a \in m_k$  are at the same positions in their respective trees, and  $a^* \neq a$ . Observe that if  $H(00\|a^*) = H(00\|a)$ , we have found a collision to  $H$ . If  $H(00\|a^*) \neq H(00\|a)$ , let the nodes on the path between the leaf corresponding to  $a^*$  and the root of  $T^*$  be denoted by  $u_{i,j_i}^*, u_{i-1,j_{i-1}}^*, \dots, u_{1,j_1}^*, u_{o,o}^*$  and the nodes on the path between the leaf corresponding to  $a$  and the root of  $T$  by  $u_{i,j_i}, u_{i-1,j_{i-1}}, \dots, u_{1,j_1}, u_{o,o}$ . Since  $u_{i,j_i}^* \neq u_{i,j_i}$  and  $u_{o,o}^* = u_{o,o}$ , there exists a  $0 \leq \ell < j$  such that  $u_{\ell,j_\ell}^* = u_{\ell,j_\ell}$  and  $u_{\ell+1,j_{\ell+1}}^* \neq u_{\ell,j_{\ell+1}}$ . Thus,  $u_{\ell+1,j_{\ell+1}}^*$  and  $u_{\ell,j_{\ell+1}}$  (together with their siblings and 01) form a collision.

Consider now the case where there is a leaf, corresponding to a token, in  $T^*$  that is not at a location in the tree, where there is a leaf, corresponding to a token, in  $T$ . In this case there must be nodes  $u_{i,j}^* \in T^*$  and  $u_{i,j} \in T$  at the same position in their respective trees such that one of them is internal and the other corresponds to a token. If  $u_{i,j}^*$  and  $u_{i,j}$  do not have the same label, we can apply the method from the previous paragraph to find a collision. If they have the same

**Table 1.** Theoretical bounds on the performance of our version of a quotable signature. For the Quoter, we consider both if we allow quoting arbitrary tokens from the sequence, and when we require that the quoted tokens must be consecutive.

	Computation Time	Signature Size
<b>The Signer</b>	$2n - 1$ hashes and 1 classical signature	1 classical signature
<b>The Quoter</b>		
Arbitrary	$2n - 1$ hashes	1 classical signature, at most $t(\lceil \log n \rceil - \lceil \log t \rceil - 1)$ $+ 2^{\lceil \log t \rceil}$ hashes
Consecutive	$2n - 1$ hashes	1 classical signature, at most $2\lceil \log n \rceil - 2$ hashes
<b>The Verifier</b>	1 classical verification and up to $2n - 1$ hashes	—

label, we must have two nodes  $u_{i+1,2j}, u_{i+1,2j+1}$  in  $T$  or  $T^*$ , and a token  $a$  in  $m^*$  or  $m_i$  such that  $H(01 \parallel u_{i,j} \parallel u_{i,j+1}) = H(00 \parallel a)$ , and we have found a collision.

We observe that in all cases, we have found a collision for  $H$ . Since  $H$  is assumed to be secure, and hence collision resistant, this can happen only negligible probability. Denote this probability as  $\epsilon_H$ .

Hence,  $\mathcal{A}$ 's advantage of at most  $\epsilon_{DS} + \epsilon_H$  is negligible.

### 3.4 Performance

Table 1 shows the cost of our quotable signature scheme for each of the three parties. This is measured in terms of computation due to the number of required hash operations and classical signature operations as well as in terms of the size of the generated signature due to the required hash values and classical signatures, presumably the dominant operations. In all cases, we assume that the message  $m$  has length  $n$ , i.e.,  $m$  consists of  $n$  tokens. For the Quoter and the Verifier, we additionally assume that the quote has length  $t \leq n$ .

To put the results into context, running the command `openssl speed` on a modern laptop shows that it is capable of computing hundreds of thousands or even millions of hashes every second (depending on the size of the data being hashed and the hash algorithm being used). Additionally, a classical digital signature only takes a fraction of a second create or verify. Thus, it is nearly instantaneous to generate/quote/verify a quotable signature, even for sequences and quotes that are thousands of tokens long.

The cost for the Signer, the Quoter, and the Verifier is derived as follows.

**The Signer** To generate the Merkle tree, the Signer needs to compute  $2n - 1$  hashes. To create the quotable digital signature for  $m$ , she creates a classical digital signature for the root hash. This classical digital signature is the Signer's signature for her message  $m$ .

**The Quoter** The Quoter also has to generate the entire Merkle tree, from which he can extract the verification path for the quote he wishes to make. However, the size of the verification path (and hence the signature for the quote) depends on the size of the quote, and where in the text the quote is located. The most simple case is when just one token is quoted, in which case the size of the verification path is at most  $\lceil \log n \rceil$ , which, together with the classical signature for the root hash, forms the signature for the quote. Similarly, as shown in the following, the worst case can be obtained by quoting every second token, in which case the Quoter would need  $\lceil \frac{n}{2} \rceil$  hashes on the verification path.<sup>4</sup>

In Proposition 1 we quantify the worst-case size of the verification path (and hence the signature) for the quote in terms of message and quote lengths. In Proposition 2, we consider the special case where we require that the quote be contiguous.

**Proposition 1.** *For a message  $m$  of size  $n$  tokens and a quote of size  $t$  tokens, the worst-case size of the verification path of the quote is at most*

$$t(\lceil \log n \rceil - \lceil \log t \rceil - 1) + 2^{\lceil \log t \rceil}.$$

*Proof.* In Lemma 2, we consider the case where  $n$  is a power of two. In this case, we identify a worst-case set of  $t$  leaves of the Merkle tree on  $n$  tokens. In Lemma 3, we establish that it is sufficient to consider  $n$  a power of two.


To argue about the size of the signature, we consider what we call the *forest of independent trees* for a quote. To find the forest of independent trees for a quote, we do the following. For each token in the quote, consider the path between the node corresponding to that quote and the root (the root-token path). Define the *independent tree corresponding to that token* to be the subtree rooted in the highest node on the root-token path, which is not on the root-token path for any other token in the quote. The forest of independent trees for the quote is now the collection of the independent trees of all the tokens in the quote. In Figure 3, we consider a message of size  $n = 8$  and a quote of size  $t = 3$ , quoting the first, third, and fifth token. The red line indicates a separation between the independent trees and the nodes that are on multiple root-token paths. The forest of independent trees consists of the trees rooted in  $u_{2,0}$ ,  $u_{2,1}$ , and  $u_{1,1}$ .

**Lemma 1.** *If  $n$  is a power of two, the heights of the trees in the independent forest for a quote that maximizes the size of the signature can differ by at most 1.*


*Proof.* Assume towards a contradiction that  $Q$  is a quote that maximizes the size of the signature for  $Q$  such that the difference between the heights of the smallest and largest trees in the forest of independent trees for  $Q$  is at least 2. Let  $A$  be the root of a tree of minimal height in the forest of independent trees, and let  $B$  be its sibling. Note that  $B$  is also the root of a tree in the forest

---

<sup>4</sup> Of course, algorithms can be adapted to include the entire text instead in such (rare) cases where that might require less space.



**Fig. 3.** A Merkle tree for a sequence of size  $n = 8$  and a quote of size  $t = 3$ .



**Fig. 4.** Note that there might be trees rooted at  $A, B, D_1, D_2, D_3$ , and  $D_4$ , which we have omitted drawing, but by our assumption, the trees rooted at  $D_1, D_2, D_3$ , and  $D_4$  must be at least as high as the ones rooted at  $A$  and  $B$ .

of independent trees (otherwise the tree rooted at  $A$  would not be of minimal height). Additionally, let  $C$  be the root of a tree of maximal height in the forest of independent trees. We illustrate this in Figure 4.


Observe that we can now create a quote  $Q'$  requiring more hashes than  $Q$ , by changing  $Q$  in the following ways:

- Instead of quoting one token from the tree rooted at  $A$  and one token from the tree rooted at  $B$ ,  $Q'$  quotes only one token from the tree rooted at  $P$ .
- Instead of quoting just one token from the tree rooted at  $C$ ,  $Q'$  quotes one token from the tree rooted at  $L$  and one token from the tree rooted at  $R$ .

It is clear that  $Q$  and  $Q'$  quote equally many tokens and that the forest of independent trees for  $Q'$  is only changed from the forest for  $Q$  in the trees that involves  $A, B$ , and  $C$ . The new situation is illustrated in Figure 5.

If each of the trees rooted at  $A$  and  $B$  contributed with  $k$  hashes to  $Q$ , then the tree rooted at  $C$  contributed with  $k' + 2$  hashes, where  $k' \geq k$ . In total,  $A, B$ , and  $C$  contributed  $2k + k' + 2$  hashes. However, in  $Q'$  we see that the tree rooted at  $P$  contributes  $k + 1$  hashes, and each of the trees rooted at  $L$  and  $R$  contributes  $k' + 1$  hashes, for a total of  $k + 2k' + 3$  hashes. But since  $k' \geq k$ , we have that  $k + 2k' + 3 \geq 2k + k' + 3 > 2k + k' + 2$ , contradicting that  $Q$  maximizes the size of the signature.

**Lemma 2.** *When  $n$  is a power of two, we can assume that the quote generating the largest signature has the properties that*



**Fig. 5.** Note that there might be trees rooted at  $A, B, D_1, D_2, D_3,$  and  $D_4$ , which we have omitted drawing, but by our assumption, the trees rooted at  $D_1, D_2, D_3,$  and  $D_4$  must be at least as high as the ones rooted at  $A$  and  $B$ .

1. the heights of the trees in the independent forest for the quote differ by at most 1,
2. for each tree in the forest of independent trees, the left-most leaf corresponds to the token that is quoted, and
3. the trees in the forest of independent trees are arranged with the smallest trees first.

*Proof.* Claim 2.1 follows immediately from Lemma 1. Further, Claim 2.2 follows from observing that we can bring any tree to this form simply by swapping the children of some of the nodes on the path to the leaf corresponding to a quoted token (hereby changing which token is quoted, but not how many are quoted), and that these swaps do not affect the size of the signature. Finally, Claim 2.3 follows from observing that if two nodes are on the same level of the Merkle tree, and the labels of both are known, then we can “swap” the subtrees that they are roots of without affecting the size of the signature. By “swapping”, we mean that if the  $i$ 'th leaf in the first node's subtree corresponds to a quote before the swap, then the  $i$ 'th leaf in the second node's subtree corresponds to a quote after the swap, and so on. To see that this does not affect the size of the quote, note that outside of the two subtrees, nothing has changed; the hash of both nodes is still known. Additionally, from the first subtree we now get as many hashes as we got from the second subtree before the swap, and vice versa.

Lemma 2 implies that for any  $n$  a power of two and  $t \leq n$ , we need only consider one choice of which tokens are quoted. For example, Figure 3 shows the only quote of size  $t = 3$  in a tree of size  $n = 8$  that we need to consider.

**Lemma 3.** *For any message  $m$  of length  $n$  and quote  $Q$  of length  $t$ , there is a quote  $Q'$  of length  $t$  from a message  $m'$  of length  $2^{\lceil \log n \rceil}$  such that the signature for  $Q'$  is no smaller than the signature for  $Q$ .*

*Proof.* For fixed  $m$  and  $Q$ , we create  $m'$  by adding tokens to  $m$  until  $|m'| = 2^{\lceil \log n \rceil}$ . We now create  $Q'$  from  $Q$  by going over each quote  $q$  in  $Q$ .

1. If the leaf corresponding to  $q$  in the Merkle tree for  $m$  is on the deepest level, we quote the same token in  $m'$ .

2. If the leaf corresponding to  $q$  in the Merkle tree for  $m$  is not on the deepest level, there is an internal node in the Merkle tree for  $m'$  at the location of the leaf in  $m$ . We quote the token corresponding to its left child, which is a leaf.

Clearly, the tokens in  $Q'$  from case 1 contribute with the same number of hashes to the signature for  $Q'$  as the corresponding ones did to the signature for  $Q$ , and the tokens from case 2 contribute with exactly one more hash. Hence, the signature for  $Q'$  is at least as large as the signature for  $Q$ .

We are now ready to derive the claim in Proposition 1. For any message  $m$  and quote  $Q$  we can assume that  $|m| = n$  is a power of two, i.e.,  $n = 2^{\lceil \log n \rceil}$  (otherwise Lemma 3 allows us to instead consider an  $m'$  that is a power of two), and that  $Q$  has size  $|Q| = t$  and exactly the structure described in Lemma 2.

There are  $t$  trees in the forest of independent trees for the quote, and all the way up to (but not including) their roots, each of these trees provides one hash per level. The roots of the trees in the forest are on the deepest level with less than  $t$  nodes and the first level with more than  $t$  nodes (if  $t$  is a power of two, all roots are instead on the level with exactly  $t$  nodes). Hence, all levels that are at depth more than  $\lceil \log t \rceil$  contributes with 1 hash per tree, for a total of  $t(\lceil \log n \rceil - \lceil \log t \rceil)$  hashes. Additionally, we need to count how many hashes we get from the level at depth  $\lceil \log t \rceil$ . On this level, every node is either a root of an independent tree or a child of a root of an independent tree. In the first case, the hash of the node is calculable from information from lower levels. In the second case, for every pair of siblings, one of the nodes' hash is calculable from information from lower levels (the one on a root-token path for a token corresponding to a quoted token) and the other nodes' hash must be provided by the signature. Since there are  $2^{\lceil \log t \rceil}$  nodes on this level, and  $t$  independent trees, the signature must provide  $2^{\lceil \log t \rceil} - t$  hashes on this level.

In total, this shows that an upper bound on the number of hashes provided by the signature for a quote of  $t$  tokens from an  $n$  tokens sequence is

$$\begin{aligned} & t(\lceil \log n \rceil - \lceil \log t \rceil) + 2^{\lceil \log t \rceil} - t \\ & = t(\lceil \log n \rceil - \lceil \log t \rceil - 1) + 2^{\lceil \log t \rceil}, \end{aligned}$$

which finishes the proof of Proposition 1.

**Corollary 1.** *For a message of size  $n$  tokens and any quote, the worst-case size of the verification path of the quote is  $\lceil \frac{n}{2} \rceil$ .*

Another easy corollary to the proof of Proposition 1—and Lemma 3 in particular—we can bound the error when  $n$  is not a power of two (when  $n$  is a power of two, the bound is, of course, exact).

**Corollary 2.** *When  $n$  is not a power of two, the bound of Proposition 1 overcounts by at most  $t$  hashes.*

*Proof.* At each level of the Merkle tree, the signature needs to provide at most one hash for each quoted token. In the construction used in the proof of Proposition 1 when  $n$  is not a power of two, no levels are added to the Merkle tree, and hence the signature becomes no more than  $t$  hashes larger.

**Proposition 2.** *For a message of size  $n > 2$  tokens and a contiguous quote of  $t$  tokens, the worst-case size of the verification path of the quote is  $2\lceil \log n \rceil - 2$  hashes.*

*Proof.* We prove this proposition by induction on the height of the Merkle tree.

As the base case, we consider trees of height 2. Either picking just one token or picking one token among the first two tokens and one token among the last one or two tokens, gives a verification path of worst-case size  $2 \cdot 2 - 2 = 2$ .


Assume now that in a tree of height  $k$ , the largest possible size of the verification path for a contiguous quote is  $2k - 2$ . As our inductive step, we show that if the height of the Merkle tree of a message is  $k + 1$ , then the largest possible size of the verification path for a contiguous quote from the message is  $2(k + 1) - 2$ . For any contiguous quote  $Q$ , we consider two cases: (1)  $Q$  is either contained in the first  $2^k$  tokens or contains none of the first  $2^k$  tokens, and (2)  $Q$  contains both the  $2^k$ 'th and the  $(2^k + 1)$ 'st token.

**Case 1:** If  $Q$  corresponds to leaves that are completely contained in one of the subtrees of the root, it follows from the induction hypothesis that the verification path consists of at most  $2k - 2$  hashes from that subtree. The verification path contains only one additional hash, that of the root of the other subtree. Thus, the total number of hashes is at most  $2k - 2 + 1 < 2(k + 1) - 2$ .

**Case 2:** We make a few observations. Considering a level of the Merkle tree from left to right, the nodes with hashes that the Verifier calculates are consecutive. In Figure 6, we have illustrated this by highlighting in green all the nodes with labels that the Verifier calculates.

Additionally, observe that for any level of depth  $j \geq 2$ , the only nodes of depth  $j - 1$  with a label that the Verifier has to calculate and that, at the same time, (potentially) has a child outside the consecutive sequence of nodes that the Verifier calculated the labels for at depth  $j$ , are the parents of the leftmost and rightmost nodes in that consecutive sequence at depth  $j$ . All the nodes that might be characterized like this are on the two paths of black arrows in Figure 6. Hence, it follows that on each level, the verification path needs to provide at most 2 hashes. Clearly, the root's label will not need to be provided by the verification path, and the root's children will also not need to have their labels provided since the quote contains a token from each child's subtree. Finally, observing that there are a total of  $k + 2$  levels in a tree of height  $k + 1$ , allows us to conclude that the verification path needs to provide at most  $2 \cdot (k + 2 - 2) = 2 \cdot (k + 1) - 2$  hashes, completing the case and the proof.

**The Verifier** The Verifier has to verify one classical digital signature and to reconstruct the Merkle tree using the quote together with the verification path. Once again, the cost of this depends on where in the message the quote is located,



**Fig. 6.** Merkle tree with a contiguous quote divided between the left- and right subtree. The labels of all the nodes in the green area are calculated from the labels of their children and do not need to be part of the signature.

with the number of hashes generally going towards  $2n - 1$  as the quote gets closer to being the full message. For example, if all but one token has been quoted, the Verifier needs to compute  $2n - 2$  hashes, and if only one token has been quoted, the Verifier only needs to compute  $\lceil \log n \rceil + 1$  hashes.

## 4 Quotable Signatures and Fake News

In the introduction, we argued that the current approach to mitigating the effects of fake news, focusing on flagging problematic content, is not sufficient. As mentioned, one supplementary approach could be to bolster authentic content by authenticating the source of quotes, for example on social media, and the literature gives reason to believe this could have an impact. This approach could be implemented using a quotable signature scheme. Here, the message that is the original source of a quote would be an article and the creator or distributor of the article (a news agency, for instance) would act as the Signer, the one sharing the quote as the Quoter, and the one verifying the quote as the Verifier. For this approach to be effective, it would need to be widely adopted, both by news media and by users sharing and reading quotes from articles. We make the following observations on these problems.

Regarding the news media, there is wide interest in supporting initiatives to combat fake news, see for example [9]. Additionally, from our discussions with a news media company,<sup>5</sup> it is apparent that the current workflow employed by modern media companies is already highly automated, and it appears that it should be quite simple to integrate a process by which, when an article is published (or updated), it is automatically signed with the media company's public key. Regarding user adoption, there is the challenge of getting a sufficiently large proportion of users using the tool, but one would also have to teach users what a quote being authenticated means, i.e., that the source and integrity of

<sup>5</sup> Specifically, we talked with the editor in charge of the platforms and the editor in charge of the digital editorial office at a large media company that produces multiple newspapers for different regional areas, in both paper and digital versions.



the quote has been assessed, but not its truthfulness or the quality of its source, for example.

If news media and social media integrate this approach into their websites, our methods can be employed without any explicit user awareness. With such an integration, when a user copies a quote from a signed article, a signature for the quote is automatically generated, and an element including both quote and text is put into the clipboard, together with the plain text quote (in practise, this would be a `text/html` element and a `text/plain` element). When the user then pastes the quote, a website supporting signatures will use the clipboard element with a signature [35]. One challenge with this approach is that the verification is now performed by the websites, rather than a browser extension, for example. Thus, the user has to trust the website to perform the authentication correctly.

An essential choice is how to divide text into tokens, since any subsequence of the tokens is an allowable quote. Natural choices could be by word, sentence, or paragraph. As a more involved choice, one could also define the tokens at a per-token basis, and simply mark the tokens in the HTML code. A variation of this would be to have a default setting, but to allow the Signer to decide how to split the article into tokens when signing. As a variant, one could also consider using content extraction policies, as in [31], so the Signer can specify which subsequences of tokens are allowable quotes. A media company might want to disallow quotes of noncontiguous segments, for example, or disallow including only parts of a sentence containing a negative, such as “not”, “neither”, or “never”. Such restrictions could be handled efficiently using regular expressions.

We are implementing a prototype,<sup>6</sup> separated into two parts: a library that can be used by media companies to sign their articles and a browser extension that allows users to quote with signatures and to verify signatures for quotes. The library contains implementations of the proposed methods that each media companies can integrate into their publishing workflow. The browser extension modifies websites such that text (both full articles and quotes) with verified signatures is shown to be signed, and allows the user to make quotes from the signed text that include a signature for the quote. The browser extension also allows the user to get more information from the signature for a quote, e.g., who signed it, when it was signed, an indication of where text was removed, and a link to the original article.

One could further extend the system with different labels, depending on the quality of the source of a quote. For example, many countries have press councils enforcing press ethics, which includes providing correct information, e.g., by researching sufficiently and publishing errata when needed. Hence, it may make sense to mark quotes from articles written by news media certified as following press ethics and rulings of a national press council. One could even go so far as to authenticate only signatures signed by such sources.

To make a difference in the future, media companies and users on social media need to adopt these quotable signatures. To have the best effect, social

---

<sup>6</sup> To be made available at <https://serfurth.dk/research/archive/>

media platforms should directly support quotable signatures and the required extension should be natively integrated into browsers.

## 5 Future Work

With this paper, we have extended the theory on quotable signatures and presented an application of quotable signatures as a supplementary approach to mitigating the effect of fake news.

Further work on quotable signatures could include using methods similar to the ones employed in [15] and [5] to remove the requirement that the used hash function be collision-resistant, and thereby remedy a vulnerability against multi-target attacks against hash functions. Additionally, variants of quotable signatures optimized for different types of media should be developed and compared. Our current variant is in some sense optimized for cases where one will often wish to quote something contiguous in one dimension, such as text. If, instead, the goal is to crop an image, one would end up with a “quote” that is contiguous in two dimensions. We have not yet explored how to handle this case effectively. Finally, as discussed in Section 4, different policies for dividing text into tokens could be studied.

A natural next step towards using quotable signatures to combat misinformation would be to verify the effectiveness of the proposed method experimentally. In particular, the effects of using quotable signatures for verifying news shared on social media and elsewhere need to be investigated. A suggestion for a first study could be to investigate if the use of quotable signatures improves participants’ ability to recall from which news brand a story originated, which was an issue identified in [17]. Additional studies along the lines of [4], investigating the effects on the quality of the news diet of participants, would also be of interest.

## References

- [1] F. K. Abu Salem, R. Al Feel, S. Elbassuoni, H. Ghannam, M. Jaber, and M. Farah. “Meta-learning for fake news detection surrounding the Syrian war”. In: *Patterns* 2.11 (2021).
- [2] ADDO. *How hate speech trolls targeted Kenya’s 2022 elections*. [Online](#). Mar. 6, 2023.
- [3] E. L. Amoruso, S. P. Johnson, R. N. Avula, and C. C. Zou. “A Web Infrastructure for Certifying Multimedia News Content for Fake News Defense”. In: *Symposium on Computers and Communications (ISCC)*. IEEE Computer Society, 2022, pp. 1–7.
- [4] K. Aslett, A. M. Guess, R. Bonneau, J. Nagler, and J. A. Tucker. “News Credibility Labels Have Limited Average Effects on News Diet Quality and Fail to Reduce Misperceptions”. In: *Science Advances* 8.18 (2022).
- [5] J.-P. Aumasson et al. *SPHINCS+. Submission to the NIST Post-Quantum Project, v.3.1.1*. [Online](#). 2022.

- [6] E. Barker. *Digital Signature Standard (DSS)*. Federal Inf. Process. Stds. (NIST FIPS). Gaithersburg, MD, USA: National Institute of Standards and Technology, 2013.
- [7] A. Bilzhaue, H. C. Pöhls, and K. Samelin. “Position Paper: The Past, Present, and Future of Sanitizable and Redactable Signatures”. In: *Availability, Reliability and Security - ARES 2017*. ACM, 2017.
- [8] C. Brzuska et al. “Redactable Signatures for Tree-Structured Data: Definitions and Constructions”. In: *Applied Cryptography and Network Security - ACNS 2010*. Vol. 6123. LNCS. 2010, pp. 87–104.
- [9] C2PA. *Coalition for Content Provenance and Authenticity (C2PA)*. [Online](#). Mar. 3, 2023.
- [10] P. T. Devanbu, M. Gertz, C. U. Martel, and S. G. Stubblebine. “Authentic Third-party Data Publication”. In: *Data and Application Security - IFIP 2000*. Vol. 201. IFIP Conference Proceedings. Kluwer, 2000, pp. 101–112.
- [11] W. Diffie and M. Hellman. “New Directions in Cryptography”. In: *IEEE Trans. Inf.* 22.6 (1976), pp. 644–654.
- [12] R. Domingos. *É fake que g1 publicou reportagem afirmando que Lula disse que, se eleito, irá revogar o PIX*. [Online](#). Mar. 6, 2023.
- [13] C. Drummond, M. Siegrist, and J. Árvai. “Limited effects of exposure to fake news about climate change”. In: *Environmental Research Communications* 2.8 (2020). Article 081003.
- [14] S. Hirose and H. Kuwakado. “Redactable Signature Scheme for Tree-Structured Data based on Merkle Tree”. In: *Conference on Security and Cryptography - SECRYPT 2013*. IEEE, 2013, pp. 313–320.
- [15] A. Hülsing, J. Rijneveld, and F. Song. “Mitigating Multi-target Attacks in Hash-Based Signatures”. In: *Practice and Theory in Public-Key Cryptography (PKC)*. Vol. 9614. LNCS. Springer, 2016, pp. 387–416.
- [16] R. Johnson, D. Molnar, D. X. Song, and D. A. Wagner. “Homomorphic Signature Schemes”. In: *Topics in Cryptology - CT-RSA 2002*. Vol. 2271. LNCS. Springer, 2002, pp. 244–262.
- [17] A. Kalogeropoulos, R. Fletcher, and R. K. Nielsen. “News Brand Attribution in Distributed Environments: Do People Know Where They Get Their News?” In: *New Media & Society* 21.3 (2018), pp. 583–601.
- [18] M. Kreutzer, R. Niederhagen, K. Shrishak, and H. Simo Fhom. “Quotable Signatures using Merkle Trees”. In: *INFORMATIK 2019*. Vol. P-294. Lecture Notes in Informatik. 2019, pp. 473–477.
- [19] N. R. Kristensen. *Schmeichel om fabrikeret, mavesur udtalelse: ”Det er noget sludder”*. [Online](#). Mar. 6, 2023.
- [20] B. Laurie, A. Langley, and E. Käsper. “Certificate Transparency”. In: *RFC* 6962 (2013), pp. 1–27.
- [21] S. Lewandowsky, U. K. H. Ecker, C. M. Seifert, N. Schwarz, and J. Cook. “Misinformation and Its Correction: Continued Influence and Successful Debiasing”. In: *Psychological Science in the Public Interest* 13.3 (2012), pp. 106–131.

- [22] R. C. Merkle. “Protocols for Public Key Cryptosystems”. In: *IEEE Symposium on Security and Privacy 1980*. IEEE Computer Society, 1980, pp. 122–134.
- [23] R. C. Merkle. “A Certified Digital Signature”. In: *Advances in Cryptology - CRYPTO’ 89*. Vol. 435. LNCS. Springer, 1990, pp. 218–238.
- [24] M. D. Molina, S. S. Sundar, T. Le, and D. Lee. ““Fake News” Is Not Simply False Information: A Concept Explication and Taxonomy of Online Content”. In: *American Behavioral Scientist* 65.2 (2021), pp. 180–212.
- [25] N. Newman, R. Fletcher, A. Kalogeropoulos, and R. K. Nielsen. *Reuters Institute Digital News Report 2019*. Tech. rep. Reuters Institute for the Study of Journalism, 2019.
- [26] N. Newman, R. Fletcher, C. T. Robertson, K. Eddy, and R. K. Nielsen. *Reuters Institute Digital News Report 2022*. [Online](#). 2022.
- [27] Reuters Fact Check. *Fact Check-Screenshot of BBC News report on Russia is fake*. [Online](#). Mar. 6, 2023.
- [28] K. Samelin, H. C. Pöhls, A. Bilzhause, J. Posegga, and H. de Meer. “On Structural Signatures for Tree Data Structures”. In: *Applied Cryptography and Network Security - ACNS 2012*. Vol. 7341. LNCS. Springer, 2012, pp. 171–187.
- [29] L. Schaewitz and N. C. Krämer. “Combating Disinformation: Effects of Timing and Correction Format on Factual Knowledge and Personal Beliefs”. In: *Multidisciplinary International Symposium on Disinformation in Open Online Media (MISDOOM)*. Vol. 12259. LNCS. Springer, 2020, pp. 233–245.
- [30] A. H. Schiochet. *É falso que instituto alemão apontou fraude nas eleições do Brasil*. [Online](#). Mar. 6, 2023.
- [31] R. Steinfeld, L. Bull, and Y. Zheng. “Content Extraction Signatures”. In: *Information Security and Cryptology — ICISC 2001*. Vol. 2288. LNCS. Springer, 2001, pp. 285–304.
- [32] S. S. Sundar. “Effect of Source Attribution on Perception of Online News Stories”. In: *Journalism & Mass Communication Quarterly* 75.1 (1998), pp. 55–68.
- [33] T. Traylor, J. Straub, Gurmeet, and N. Snell. “Classifying Fake News Articles Using Natural Language Processing to Identify In-Article Attribution as a Supervised Learning Estimator”. In: *IEEE International Conference on Semantic Computing - ICSC 2019*. IEEE, 2019, pp. 445–449.
- [34] J. Urban and W. Schweiger. “News Quality from the Recipients’ Perspective”. In: *Journalism Studies* 15.6 (2014), pp. 821–840.
- [35] W3C. *Clipboard API and events, W3C Working Draft, 6 August 2021*. [Online](#). 2021.