Cryptocurrency Scams: Analysis and Perspectives

MASSIMO BARTOLETTI[ID], **STEFANO LANDE**[ID], **ANDREA LODDO**[ID],
LIVIO POMPIANU, AND SERGIO SERUSI
Department of Mathematics and Computer Science, University of Cagliari, 09124 Cagliari, Italy

Corresponding authors: Massimo Bartoletti (bart@unica.it), Andrea Loddo (andrea.loddo@unica.it), and Livio Pompianu (livio.pompianu@unica.it)

ABSTRACT Since the inception of Bitcoin in 2009, the market of cryptocurrencies has grown beyond the initial expectations, as witnessed by the thousands of tokenised assets available on the market, whose daily trades exceed dozens of USD billions. The pseudonymity features of cryptocurrencies have attracted the attention of cybercriminals, who exploit them to carry out potentially untraceable scams. The wide range of cryptocurrency-based scams observed over the last ten years has fostered the study on their effects, and the development of techniques to counter them. The research in this field is hampered by various factors. First, there exist only a few public data sources about cryptocurrency scams, and they often contain incomplete or misclassified data. Further, there is no standard taxonomy of scams, which leads to ambiguous and incoherent interpretations of their nature. Indeed, the unavailability of reliable datasets makes it difficult to train effective automatic classifiers that can detect and analyse scams. In this paper, we perform an extensive review of the scientific literature on cryptocurrency scams, which we systematise according to a novel taxonomy. By collecting and homogenising data from different public sources, we build a uniform dataset of thousands of cryptocurrency scams. We build upon this dataset to implement a tool that automatically recognises scams and classifies them according to our taxonomy. We assess the effectiveness of our tool through standard performance metrics. We then analyse the results of the classification, providing key insights about the distribution of scam types, and the correlation between different types. Finally, we propose a set of guidelines that policymakers could follow to improve user protection against cryptocurrency scams.

INDEX TERMS Bitcoin, blockchain, cryptocurrency, frauds.

I. INTRODUCTION

According to the price-tracking website coinmarket-cap.com, the global market of cryptocurrencies features over 13,000 crypto assets, with a total capitalisation exceeding *USD* 2.5 trillion [1]. Over the years, cryptocurrencies have gained increasing attention from investors, entrepreneurs, regulators, and the general public. Besides the financial concerns about the instability of the cryptocurrency market and its suspected bubble dynamics [2], there is a major concern about cryptocurrency scams, where fraudsters try to deceive investors to gain an undue advantage.

Detecting cryptocurrency scams is not easy for the average user. Although a few websites allow users to report and search scams, they are not complete, and they often provide inconsistent, when not erroneous, information. A main issue of these websites is that scam reports must be inserted manually,

either by site administrators or by users, so they cannot keep up with the fast pace at which scams are created. To cope with this issue, tools that automatically recognise and track cryptocurrency scams would be in order. Classifying scams would also be crucial to understand their evolution and to study possible countermeasures. Closely related to this, a reference taxonomy of cryptocurrency scams would be instrumental to construct a reliable classification. Our aim is a technique for classifying scams, which can be used as the backbone of new detection tools.

However, automatically classifying scams is a difficult task, for various reasons:

- **Lack of reliable public data on scams**. Training a scam classifier requires reliable information about thousands of scams. Although a few public data sources for cryptocurrency scams exist, they contain a lot of spurious data, e.g. URLs pointing to empty web pages, or to web pages without enough usable data. The problem is even

more evident for scam reports posted by users, which often contain inaccurate or misleading scam descriptions. We discuss the issues with public data sources in more detail in Section IV.

- **Lack of a taxonomy of scams**. The online scam reporting systems, as well as most of the works in the literature, use a variety of different taxonomies of scams. These taxonomies usually associate a scam with a single category, thus failing to appropriately render the nature of a multitude of "hybrid" scams, which mix features of different categories. The absence of a standard, comprehensive taxonomy of cryptocurrency scams hinders the development of precise scam classifiers.
- **Wrong categorisation of scams**. Scam reports inserted by users are not always accurate, e.g. they often associate a scam with the wrong category, and sometimes classify legit sites as scams. For instance, blackmail scams are a common source of errors, since users often take as serious the fraudsters' claim about the presence of a malware on the victim's device, and thereby misclassify the scam as "malware". Overall, this inaccuracy makes the data extracted from scam reports inadequate for training classifiers.

Contributions: In this work, we address these issues in order to develop an effective tool to detect and classify cryptocurrency scams. More specifically, our main contributions are the following:

1) We survey the existing body of computer science literature on crypto scams, organising them in a novel taxonomy. Our taxonomy consists of seven main features, which make it possible to represent "pure" scams, that exhibit only one feature, and "hybrid" scams, which mix two or more features.
2) We construct a dataset of thousands of scams, by collecting and homogenising data from public sources.
3) We use this dataset to train a classifier of scams, which we evaluate according to standard performance metrics.
4) We then apply our tool to take a snapshot of the current status of scams, by classifying them according to our taxonomy. This allows us to analyse various aspects of scams, e.g. their distribution by type, and the most common combinations of features for hybrid scams.
5) Based on the knowledge developed during our work and on the results of our analysis, we propose a set of guidelines that policymakers could follow to enhance the protection of users against crypto scams.

To foster the reproducibility of our results, we make publicly available the toolchain we have developed for constructing our dataset and for performing our analyses [3], as well as the obtained dataset and the analyses results [4]–[6].

The rest of the paper is organised as follows. In Section II we survey the scientific literature on crypto scams, along with we introduce our taxonomy. In Section III we describe the public data sources we use in this work. In Section IV we detail our methodology to collect and classify scams. In Section V we detail the design of our scam classifier, and the

methodology we use to validate it. In Section VI we apply our tool to analyse the scams in our dataset. In Section VII we conclude, by synthesising our findings into a set or proposals to improve the protection for users, and by proposing some directions for future works.

II. CRYPTOCURRENCY SCAMS IN THE LITERATURE

The scientific literature on crypto scams includes works from different research areas, ranging from computer science to economics, finance and law [7], [8]. In our survey, we focus on literature from the area of computer science, where we cover, to the best of our knowledge, the most relevant works from 2013 (when the first works on crypto scams were published) until the end of 2020.

Before starting, it is convenient to circumscribe the meaning of the term *scam* as it is used in this work. We define scam as any unlawful behaviour of one or more persons who intentionally deceive people to obtain something illegal or unfair. More specifically, in a *crypto* scam, fraudsters exploit one the peculiar features of blockchain technologies, i.e. the availability of crypto-assets that can be anonymously (or pseudonymously) exchanged for fiat currencies. Note that this definition rules out several illegal activities that can be carried out through blockchains. For instance, it rules out the publication of illegal material (like e.g. pornographic or blasphemous content) on the blockchain, which is possible by embedding such material in transactions [9], [10]. Further, we do not consider "pump and dump" schemes, where fraudsters intentionally create hype on some crypto-assets in order to pump their prices and sell their stocks right away. Although pump and dump schemes are illegal in many countries, they are hardly distinguishable from legit investment marketing campaigns, so to stay on the safe side we prefer to omit them from our analysis.

In literature, scams are categorised in several ways, according to different criteria [11], [12]. Some categorizations are also induced by scam reporting services. For example, the Anti-Phishing Working Group [13] features an online reporting system which categorizes scam data into: (i) phishing URLs; (ii) phishing emails; (iii) malicious IPs; (iv) malicious domains; (v) cryptocurrency (suspicious exchanges, wallet providers, trading platforms, and investment fund).

A drawback of the existing categorizations is that they fail to capture the multi-faceted nature of crypto scams, most of which feature various traits. Indeed, strict classifications of scams into isolated categories is inadequate, since the boundary between the types of a scam is often blurred. For instance, fake mining scams can be quite similar to Ponzi schemes, in that both advertise investment schemes which promise to return a larger sum to the investor. For this reason, we propose a taxonomy of crypto scams which consists of seven main features, which capture different types of illegal activity, and can be mixed to characterize *hybrid* scams.

Since most of the research works focus on *pure* scams, in this section we partition the literature survey according to the principal feature addressed by the surveyed works.

TABLE 1. Number of publications by scam type.

Scam type	2013	2014	2015	2016	2017	2018	2019	2020	Total
Ponzi schemes	0	0	1	0	0	6	5	4	16
Fake crypto services	1	0	1	0	0	2	2	2	8
Malware	0	1	1	1	1	2	1	0	7
Blackmail	2	0	0	0	0	0	1	1	4
Advance-fee scams	0	0	0	0	0	0	0	1	1
Money laundering	1	0	1	0	0	1	2	1	6
Fake ICO	-	-	-	-	0	1	1	1	3

A. PONZI SCHEMES

Ponzi schemes are scams that advertise themselves as high-yield investment programmes (HYIPs) [14], [15]. They typically lure users with the promise of high profits in return for their investments by paying high levels of interest (see e.g. Figure 1). In practice, Ponzi schemes only pay users with the funds invested by new users, and therefore implode as soon as new investors stop joining. As a result, most investors in Ponzi schemes just lose their money. The Securities and Exchange Commission's Office of Investor Education and Advocacy [16] has provided some common red flags for understanding whether an investment is a scam:

- it advertises high investment returns with little or no risk (every investment carried some degree of risk);
- it claims to generate a steady return regardless of general market conditions;
- the investment is not registered, and the sellers are not licensed;
- it presents secret or complex investment strategies;
- it is hard to find complete information on it;
- there are problems with the documentation or excuses as to why it is not possible to review the information on the investment;
- there are non-payments or difficulties in collecting the investment;
- a few respected leaders can be enlisted to spread "investment" around the world.

Several works investigate Ponzi schemes operating on Bitcoin [17], [18], [20], [22], [25], [30] and on Ethereum [22], [31]. Among them, [18], [20], and [30] aim at automatically detecting Ponzi schemes by analysing their transactions.

Toyoda *et al.* [18] collected from *bitcointalk.com* a dataset of 1,500 Bitcoin addresses between HYIP and non-HYIP, later increased to 2,026 HYIPs and 26,976 non-HYIPs addresses [30]. Bartoletti *et al.* [20] collected 1211 Ponzi scam addresses from *Reddit* and *bitcointalk.org*, and 6,400 not-scam addresses from other sources. Both Bartoletti *et al.* and Toyoda *et al.* used addresses clustering techniques [33], [34] (in particular, the multi-input heuristic) to extend their datasets of scam addresses, and both applied supervised machine learning techniques to detect Ponzi schemes. In particular, Toyoda *et al.* used XGboost and Random Forest classifiers, correctly recognising 83% of HYIP, with a false

positive rate of 4,4%. Bartoletti *et al.* used a cost-sensitive Random Forest classifier, which correctly classified 31 Ponzi schemes out of 32. Vasek and Moore [17], [25] analysed the behaviour of Ponzi schemes on Bitcoin. In [17], they compiled a list of 349 scams from various sources, including *bitcointalk.org*, *badbitcoin* and the HYIP-tracking website *cryptohyip.com*, and analysed them according to various metrics, e.g. their lifetime and the gain for fraudsters. In [25] they found empirical evidence that scams advertised with newly created accounts die quicker than those with older accounts.

Bartoletti *et al.* [31], Chen *et al.* [22], [27] and Jung *et al.* [28] analysed Ponzi schemes developed as smart contracts on Ethereum. The work [31] proposed a set of criteria for determining when a smart contract implements a Ponzi scheme: 1) the contract distributes money among investors, 2) the contract receives money only from investors, 3) each investor makes a profit if enough investors invest enough money in the contract afterwards, 4) the later an investor joins the contract, the greater the risk of losing his investment. By examining the source code of several smart contracts, they categorized them according to the mechanism used to redistribute money. The same work also proposed a tool to detect hidden Ponzi schemes on Ethereum, exploiting the fact that many Ponzi schemes share code fragments, and so the Levenshtein distance between the bytecode of a hidden Ponzi scheme is usually close to that of a known one. Chen *et al.* [22] proposed a machine learning technique to detect Ponzi schemes on Ethereum, based on a dataset of 131 Ponzi schemes and 1,251 non-Ponzi contracts collected from *etherscan.io*. The features used to train the classifier are based mostly on the contract bytecode, e.g. the type of opcodes used. The authors found that SLOAD and CALLER are the most used opcodes in Ponzi schemes (used to get the caller's address) and LT. Using a regression tree model (XGBoost) classifier, they obtained a precision of 97% and a recall of 81%. Based on the results of the classification, the authors estimated that almost 434 Ponzi schemes have operated on Ethereum before May, 2017. A subsequent work of the same authors [27] improved the dataset and the classification technique, reaching a precision of 95% with a Random Forest classifier. Jung *et al.* [28] analysed both behaviour and code frequency of contracts to detect Ponzi schemes, obtaining a precision of 99% and recall of 97%. In 2017, Bartoletti *et al.* manually analysed the source code of smart contracts on

FIGURE 1. Cryptory Ponzi scheme (from https://bitcointalk.org/index.php?topic=633822.0).

Ethereum [35] to spot Ponzi schemes, finding 120 scam instances.

B. MALWARE

The alleged untraceability of cryptocurrencies has been extensively exploited by malware developers. There are two types of malware closely related to cryptocurrencies:

Ransomware After infecting the victim's device, this kind of malware encrypts the data on the device, and locks it until the user pays a ransom (usually in Bitcoin). Figure 2 shows a screenshot of a device infected by Wannacry, along with the instructions to pay the ransom.

FIGURE 2. Wannacry screen requests payment.

Crypto loggers This kind of malware tries to steal information about the victim's accounts on crypto services (like, e.g., wallets). In particular, crypto loggers try to obtain the private key needed to transfer crypto-assets from the victim's account to the fraudster's. They often work as

a transparent interface while the user is surfing the web, or is searching for password files.

Liao et al. [36] realised a scam detection framework based on the multi-input and change-addresses heuristics. They collected data from the blockchain and automatically identified payments to Bitcoin addresses related to the *CryptoLocker* ransomware. Moreover, they estimated the financial damages of the scam from September, 2013, to January, 2014. The authors in [37] developed a framework to identify, collect, and analyse addresses that belong to the same user. They also proposed an approach for classifying payments as ransoms. Finally, they explored the economic impact of twenty ransomware that have been operated before December, 2017. Huang et al. [38] performed a two-year analysis of the ransomware ecosystem using the multi-input heuristic. In particular, they collected 25 seed ransom addresses from actual victims across eight ransomware families, like e.g. *Locky*, *Cerber*, and *Spora*. The authors discovered that these scams generate unique addresses for every victim. The work [39] implemented a tool for identifying, extracting and analysing Bitcoin transactions related to ransomware attacks. Moreover, it traced the monetary flows of these attacks by computing the graph of transactions between the clustered addresses. They also realised a dataset (gathering data from *walletexplorer.com* and *blockchain.info*) for improving the interpretation of the monetary flows and associating them to actors (e.g., exchanges and gambling sites). This dataset consists of 7,118 addresses related to 35 ransomware families.

Liao et al. [36] collected a dataset of 968 Bitcoin addresses related to the *CryptoLocker* cluster. They estimated the payments valued at *USD* 310,473 in the period 2013-2014, which is not too far from the estimated by Conti et al. [37] (*USD* 449,274, period 2013-2017) and by Paquet-Clouston et al. [39] (*USD* 519,991, period 2013-2017). Huang et al. [38] traced *USD* 16 millions in 19,750 ransom payments belonging to five ransomware families over 22 months. Most ransom addresses have only one

TABLE 2. Summary of literature on Ponzi schemes.

Ref.	Year	CCY	Highlights	In USD	Dataset
[17]	2015	BTC	- dataset of 9 bridge Ponzi schemes and 23 Bitcoin-only schemes - measures payout of scams and daily volume of payments	7.2M	N/A
[18]	2017	BTC	- dataset of 43 Ponzi addresses and 1523 non-Ponzi addresses - automatic classification of scam addresses with RF and XGBoost	—	[19]
[20]	2018	BTC	- dataset of 32 Ponzi schemes, expanded via address clustering - cost-sensitive *vs.* sampling-based approaches to deal with class imbalance - among tested classifiers, the best one is Random Forest with CM20	9.5M	[21]
[22]	2018	ETH	- dataset of 131 Ponzi schemes and 1251 not Ponzi schemes - automatic classification of Ponzi schemes, XGBoost achieved the best results	—	[23]
[24]	2018	BTC	- time series analysis to detect anomalies in a transaction history - experiments with Pirate@40's HYIP, which raised 700K BTC in 2011-13	—	N/A
[25]	2018	BTC	- analyses the supply and demand for Bitcoin-based Ponzi schemes - 11902 victims from 89439 comments on 2629 threads on 1779 scams	—	N/A
[26]	2019	BTC	- analysis of MMM, one of the oldest and most popular Ponzi schemes - based on 423K transactions involving 16K addresses	—	N/A
[27]	2019	ETH	- automatic classification of smart Ponzi schemes - 3,780 open source smart contracts of which 200 are smart Ponzi schemes - among tested classifiers, the best one is Random Forest	—	N/A
[28]	2019	ETH	- automatic detection of Ponzi smart contract - unbalance dataset with 3,203 non-Ponzi instances and 172 Ponzi instances	—	N/A
[29]	2019	ETH	- investigate the prevalence of honey-pot smart contracts - study of 690 honeypot smart contracts as well as 240 victims in the wild	90K	N/A
[30]	2019	BTC	- automatic detection of Ponzi schemes on Bitcoin - dataset of 26976 non-HYIPs and 2026 HYIPs	—	N/A
[31]	2020	ETH	- analyses scam impact, source code, lifetime, inequality - detects hidden Ponzi schemes by bytecode similarity	630K	[32]

inflow, while others have two or more, probably because some victims did not consider the transactions fees. The authors also measured that the bitcoins remained in *Wannacry* for 79.8 days on average, while in *Cerber* and *Locky* for 5.3 and 1.6 days, respectively. Conti *et al.* [37] discovered that 83% of Bitcoin addresses used by *Cryptolocker* received at most two payments, out of 804 ransom payments which contributed to extorting 1,403 BTC. *CryptoDefense* infected computers through spam and received 108 ransom payments corresponding to BTC 126, equivalent to *USD* 63,859; *WannaCry* received over 238 payments and extorted 47 BTC. Paquet-Clouston *et al.* [39] estimated as *USD* 12 millions the overall market of ransomware payments from 2013 to mid 2017. Furthermore, they pointed out a few players were responsible for most payments, and that ransomware authors tended to keep their money in one or several addresses. One of the most successful ransomware, *Locky*, received payments worth *USD* 8 million, more than 50% of the total ransomware payments. Moreover, the first three families (*Locky*, *CryptXXX* and *DMALockerv3*) accounted collectively for 86% of the market. Although a single ransomware can infect many computers, the number of payments received by victims is much lower. Kshetri and Voas [40] measured that *WannaCry* received a payout rate of about 0.14%. Several factors cause this, but one of the main factors could be that there is no guarantee that data are unlocked after the random is paid.

Moreover, most insurance policies deny paying ransomware without pre-approval. Finally, since the KYC verification processes can take several days, the three-day deadline for *WannaCry* was too short for many victims. Liao *et al.* [36] developed Gephi, a software to visualise the transaction graph and discover evidence that connects popular crypto services, like exchanges, to other cybercrimes, like e.g. the Sheep Marketplace scam.

C. FAKE CRYPTO SERVICES

In the cryptocurrency ecosystem, there are multiple services to simplify their use and management. They include exchange services, wallets and mixers. However, numerous criminals develop these types of services in the form of fraud, as described below.

Fake exchange Fake exchange frauds deceive users by offering incredibly competitive market prices for purchasing cryptocurrencies. Indeed, they trick users with quick and easy access to some cheap currency. Figure 3 shows a fake exchange service at paybillsbitcoin.com, mirrored from BtcToPal.com.

Fake wallet Wallet services allow users to manage, send and receive cryptocurrencies. In this scenario, users can run into wallet scams characterised by various types of fraudulent behaviour. For example, some wallets steal the entire amount indiscriminately, while others take a

TABLE 3. Summary of literature on malware.

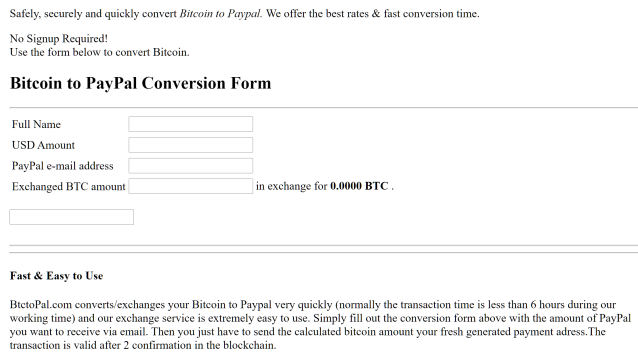| Ref.	Year	CCY	Highlights	In USD	Dataset
[36]	2016	BTC	- analysis of 968 payments to the **ransomware CryptoLocker** - they discover connections between Bitcoin services and other cybercrimes	310,473	N/A
[40]	2017	BTC	- analysis on **WannaCry ransomware** public reports - authors discover WannaCry payout rate of 0.14% as of May, 23, 2017	$\overline{86,000}$	N/A
[37]	2018	BTC	- framework to collect and analyse addresses belonging to the same user - system to classify a payment as ransom	449,274	N/A
[38]	2018	BTC	- analysis of 25 seed ransom addresses from 8 **ransomware families** - they obtained 19,750 ransom payments for 5 ransomware families	16,322,006	N/A
[39]	2019	BTC	- Bitcoin transactions analysis, related to **ransomware** attacks	519,991	N/A

FIGURE 3. An example of fake exchange service.

FIGURE 4. Fake mixing at https://bitcoinmixer.eu/.

FIGURE 5. Slush-pool is an example of fake mining which impersonates the real service at Slushpool [41].

small percentage of the daily deposit. Finally, others withdraw when the deposit exceeds a certain threshold.

Fake mixing Transactions in blockchain systems (e.g. Bitcoin) are linked together, so it is possible to inspect cryptocurrency movements between addresses. However, mixing services make it possible to erase the links between initial and final addresses, randomise the number of transactions, add delays to transactions, and use other extraneous addresses. On the other hand, fake mixers receive the money and steal it without sending it to the client. Figure 1 was a website of a fake mixer called "BitcoinMixer", online up to December 30, 2020. Several users sent money to this fake mixer and never received it back, as reported on BitcoinTalk. Moreover, a BitcoinTalk topic shows a list of fake mixers.

Fake mining pool Cryptocurrencies based on the Proof Of Work (POW) mechanism require a computational effort to create blocks. Therefore, users who create blocks, called miners, receive a reward. This type of scam asks users to participate by investing money to buy mining hardware. Despite what the scam promised, the money invested is not used to buy new hardware but rather to pay interest to previously registered users.

Fake donation Usually, donors make donations to projects or people for good. Fraudsters exploit people's virtue, creating fake donation campaigns, and instead of giving the money as a donation, they steal the money and disappear.

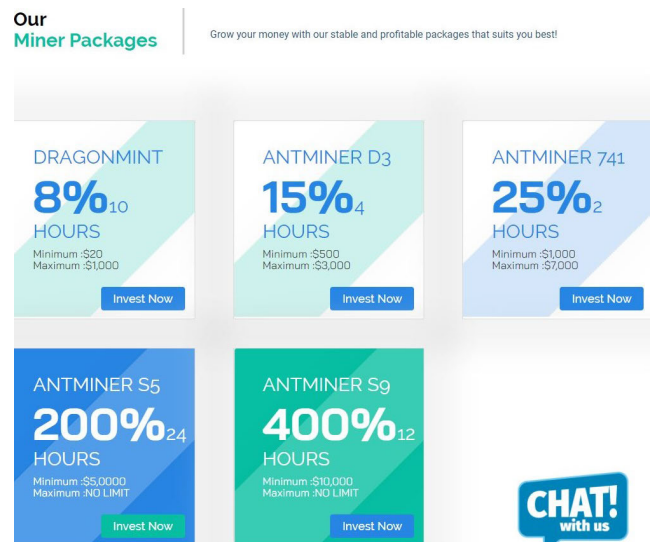Vasek and Moore [17] identified 192 scams out of 349 candidates from *bitcointalk.com*, a list of suspected fraudulent services and a website that tracks Bitcoin-based HYIPs called *cryptohyips.com*. They categorised them into four groups:

1) Ponzi schemes, 2) mining scams, 3) scam wallets, 4) and scam exchanges. They also discovered that two *Bitcointalk* users put over ฿10 into their account but remained with 0.099 ฿. According to their calculations, scam wallet revenue through September 11, 2014, was 4,100 ฿, near one million *USD*. Moreover, they analysed four Bitcoin exchanges scams: 1) BTC Promo, 2) btc-Quick, 3) CoinOpend, 4) and Ubitex. These scams offered features other exchanges do not provide, for example, PayPal payments, Credit Card processing, and better exchange rates. The longest-lived scam exchanges survived for three months while the quicker less than one month, and the total revenue of this type of scam was *USD* 0.36 million.

Moore and Christin [42] studied the risk investors face from Bitcoin exchanges. In particular, they tracked 40 Bitcoin exchanges in the last three years and found that 18 have closed. Eleven exchanges reimbursed the customers, while the other seven did not. Next, for each exchange, they calculated the average trade volume and the lifetime (they considered an exchange as closed if it has not made any trade in the last two weeks). They calculated that the median lifetime of exchanges is 381 days.

Andryukhin [43] reviewed and classified the prominent frauds that occurred in blockchains, how the fraudsters implemented them, and how to mitigate scams. While analysing social engineering attacks, the authors found "Clone", a clone website to create IOTA wallet keys, and estimated that *USD* 4 million worth of MIOTA tokens had been stolen.

Holub and O'Connor [7] tracked a Bitcoin phishing campaign called "Coinhoarder" for over six months. The campaign theft *USD* 10 million by launching typosquatting domains containing clones of crypto exchange websites. For example, they find "block-chain.info" and "blockchian.info/wallet".

Chen *et al.* [44] proposed a classifier to detect phishing account based on blockchain transaction. The authors suggested integrating it into users' cryptocurrency wallets to alert users of potential risks when interacting with dangerous accounts. They developed and tested the method on more than 7 million Ethereum transactions, more than 534 thousand addresses, 323 of which are phishing addresses. Their best classifier is DElightGBM, which obtained 82% of precision, 80% of recall, and 81% of AUC.

The authors in [45] and [46] investigated and proposed fake mining detection tools. The first one [45] inspects a system defined as the Bitcoin "GeneratorScam" (BGS), where the scammers promise to "generate" new bitcoins using the ones users sent to them. In particular, the attackers claim that they will provide free bitcoin in exchange for a small mining fee, using dubious claims such as the possibility of "hacking the blockchain ledger". The authors built a dataset of BGS, composed of 500 unique scam domains. These addresses received a total of *USD* 5,098,178, with an average of *USD* 47.3 per transaction. Instead, the second one [46] identified 3,000 websites with active mining activities and five extensive campaigns with more than 50 websites infected

through each. Moreover, they proposed *MiningHunter*, a web crawling framework that can identify mining activities on websites through WebSocket traffic and other indicators.

Bijmans *et al.* [47] also analysed the relationships between websites that perform crypto mining and the actors behind them. They studied 1,136 top-level domains (TLDs), resulting in 48.9M websites. They compared the installation base with actual mining traffic on the Internet using NetFlow data, finding that the most prominently installed miner is not the one that generates the most mining activity in practice. Moreover, they estimated that cryptojacking is present on 0.011% of all domains.

As regards fake mixing services, Wu *et al.* [48] extracted the transaction data of Bitcoin by using *WalletExplorer*, a smart Bitcoin block explorer that provides label information of addresses. The authors analysed the time interval between November 2014 and January 2016, discovering that the labelled addresses belonged to three mixing services: Bitcoin Fog, BitLaunder, and Helix. Moreover, they proposed an address detection tool that can distinguish the addresses belonging to mixing services by analysing the critical transactions involved.

D. ADVANCE-FEE SCAMS

According to the FBI [49], "An advance fee scheme occurs when the victim pays money to someone in anticipation of receiving something of greater value—such as a loan, contract, investment, or gift—and then receives little or nothing in return."

In an advance-fee scam, scammers typically contact the victim via e-mail or social media using a fake e-mail address or social media account. They promise the victim a significant amount of money in exchange for a small upfront payment that scammers claim will use to obtain a substantial sum as a reward. If the victim makes the payment, the scammer either disappears or adds several additional charges that the victim has to pay. Recently, several of these scams campaigns have taken place on Twitter, using well-known personalities such as Elon Musk, as shown in Figure 6 example.

Phillips and Wilder [50] analysed public and blockchain-based data to understand advance-fee scams. The authors use CryptoScamDB, that provides a list of cryptocurrency scam websites, even though most of the listed scam websites are no longer online. With the tool *URLScan*, they produced detailed snapshots containing the raw HTML, IP address and page content. They also captured additional information, for example, who registered the domain. Next, they extracted blockchain addresses from the snapshots of the scammed web pages through regular expression pattern matching. They found that 18.1% of the websites include a Bitcoin address and 38.7% contain an Ethereum address. Then, they used clustering techniques to analyse both the scam websites content and the details on registration and ownership. They found that the same entities run different types of prepayment scams. In addition, clustering of website content identified 171 clusters with 24 websites per cluster. The largest

TABLE 4. Summary of literature on fake crypto services.

| Ref.	Year	CCY	Highlights	In USD	Dataset
[42]	2013	BTC	- dataset of 40 **fake exchanges** until January 16, 2013 - average daily trade volumes measurements - 33 fiat currencies analysed	—	N/A
[17]	2015	BTC	- dataset of 4 **fake exchanges** and 5 **fake mining** services - weekly measures and overall payouts computation	3.8M	N/A
[7]	2018	BTC	- tracking of a **Bitcoin phishing campaign** for six months - identification of the built network and phishing websites related - reconstruction of the campaign structure, via Google Adwords	10M	N/A
[46]	2018	BTC	- identified 3,000 websites with active **mining activities** - tracking of 5 big campaigns with more than 50 websites infected	—	N/A
[47]	2019	BTC	- analysis 48.9M websites with relationship with **cryptomining** - analysis of NetFlow traffic for the popularity of cryptojacking services	—	N/A
[43]	2019	—	- analysis of four **fake wallets** on the Google Play Store	—	N/A
[44]	2020	ETH	- dataset of 323 **phishing** addresses out of 534,820 - comparison of GBM, SVM, DT and DE models for detection	—	N/A
[45]	2020	BTC	- tracking of a **Bitcoin fake mining campaign** for four months - dataset of 6,395 addresses associated to 500 scam domains	5M	N/A
[48]	2021	BTC	- extraction of **transaction data of Bitcoin** between Nov 2014 and Jan 2016 - Bitcoin Fog, BitLaunder Helix were the spreadest mixing services	—	N/A
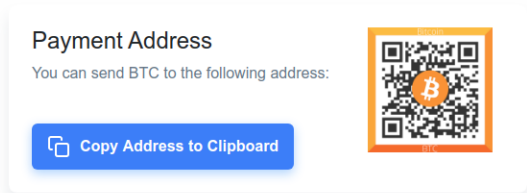
FIGURE 6. Fake Bitcoin airdrop at elonmuskdrop.com.

contains 1,359 prepayment websites. Finally, they used registration details and property identifiers in clustering by campaign, resulting in 25 identified campaigns, with an average of 17 websites per campaign cluster. Statistically, 29% of the identified campaigns run more than one type of prepayment scam, and 54% have US IP addresses. The authors also found that 55.4% of cryptocurrency transfers occurred in the first week after domain registration, 71.8% after two weeks and 91% after 30 days. Finally, by analysing funds sources, the authors discovered that more than 94% of revenue comes from exchanges, and exchanges receive more than 67% of the target funds.

E. BLACKMAIL

In blackmail scams, fraudsters usually claim to have hacked the victim's device, and installed a key logger, or recorded the victim with the webcam. The mail typically asks for a ransom in Bitcoin to delete the material, threatening that they will otherwise sell it, or publish on social networks. Figure 7 shows an actual example of blackmail that the scammer sent to a potential victim. More advanced fraudsters tailor personalised emails to victims, by exploiting databases of emails and hacked passwords.

Paquet-Clouston *et al.* [51] studied blackmail scams in the Bitcoin ecosystem. To this purpose, they used a dataset of 4,340,736 blackmails, from which they extracted the Bitcoin addresses to which the ransoms was supposed to be paid. The authors observed that frausters intentionally use grammar mistakes, synonyms, and language translations in emails to avoid spam filters. Moreover, they proposed an heuristic to cluster spam emails by textual similarity, obtaining 96 buckets. The buckets contain 12,533 Bitcoin addresses (before address clustering), even though only 245 received payments. In an 11-month analysis, the authors found that blackmail scams resulted in a revenue of *USD* 1,3 million. Oggier *et al.* [52] studied "sextortion" scams operated on Bitcoin from February to April, 2019. The authors proposed a technique, based on social networks analysis, to find relation among the Bitcoin addresses used by fraudsters, e.g. to detect which addressed are controlled by the same entity, and which addressed are used before exchanging or laundering Bitcoins.

F. FAKE ICO

An Initial Coin Offering (ICO) is a way for blockchain-related currencies to raise funds before their official launch,

From: <john_doe36@lilacmantle.com>

You may not know me, and you are probably wondering why you are getting this email, right? I'm a Hacker who cracked your devices. I setup a malware on the adult video (porn) website and guess what, you visited this site to have fun (you know what I mean). While you were watching videos, your internet browser started out functioning as a "HRDP" – Hidden Remote Desktop Protocol having a keylogger which gave me accessibility to your screen and webcam. After that, my software program obtained all your contacts and files. You entered a password on the websites you visited, and I intercepted it. Of course, you can change it, or already changed it. But it doesn't matter, my malware updated it every time.

What did I do? I generated a backup of your every system (private document files, video, photos, all files). I created a double-screen video. 1st part shows the video you were watching (you've got a good taste ha ha . . .), and 2nd part shows the recording of your webcam. Do not try to find and destroy my virus! (All your data is already uploaded to a remote server)

Do not try to contact me. Various security services will not help you; formatting a disk or destroying a device will not help either, since your data is already on a remote server. I guarantee you that I will not disturb you again after payment, as you are not my single victim. This is a hacker code of honor. Don't be mad at me, everyone has their own work.

Exactly what should you do? Well, in my opinion, $500 (USD) Dollars is a fair price for our little secret. You'll make the payment by Bitcoin (search "Buy Bitcon" in Google) Make a deposit to your wallet. After that, transfer it to my wallet. My Bitcoin (BTC) wallet address: 1N3rRvBVgCTB4D9bYdeCGCAzoyZuheTPxA.

Important: You have 2 days in order to make the payment. (I've a Facebook pixel in this mail, and at this moment I know that you have read through this email message). To track the reading of a message and the actions in it, I use the Facebook pixel. Thanks to them. (Everything that is used for the authorities can help us.) If I do not get the Bitcoin, I will certainly send out your video recording to all your contacts including relatives, coworkers and all contacts. Having said that, if I receive the payment, I'll destroy the video immediately. If you need evidence, reply with "Yes" and I will certainly send out your video recording to your 6 contacts. It is a non-negotiable offer, don't waste my personal time and yours by responding to this message.

FIGURE 7. A typical blackmail email.

similar in many ways to Initial Public Offerings for shares. Fake ICO scams apply the same strategy by luring users into buying fake coins. Typically, a cryptocurrency company releases a predetermined number of coins on the open market in the same way that shares are issued when a company goes public. Many ICOs are legitimate cryptocurrencies that have the potential to make an investor as much money as any other stock [53]. Fake cryptocurrencies advertise themselves with peculiar features that others do not have via brand new websites. To try to mitigate this problem, the SEC launched a parody website in 2018 that mocks ICO [54], along with a fake eight-page white paper, fake celebrity endorsements, and a fake team working on the ICO. According to a study realised by Satis Group in 2018, approximately 80% of ICOs conducted in 2017 were scams, with no actual product to offer [55]. Moreover, in 2017, ICO gained *USD* 1.6 billion, of which *USD* 150 million belongs to fake ones [43].

The general categorisation of fake ICOs strongly varies according to the definition of different authors. Indeed, Andryukhin [43] placed them in the more general context of social engineering phishing schemes, while Conlon and McGee framed them as Ponzi schemes [55].

To cite some examples of Fake ICO scams [53], [55], *Pincoin* was launched in 2018 and raised *USD* 660 million. *PlexCoin* in 2017 raised *USD* 8.5 million, while *Bitconnect* in 2016, reached a market cap of over *USD* 2.6 billion. *OneCoin* is perhaps the most famous. It was launched in 2014 as a mined cryptocurrency even though it was a Ponzi scheme. The FBI discovered it raised to *USD* 4 billion in income. *Savedroid* was funded in 2015 and raised *USD* 50 million until 2018. At the time of writing, it is still listed on exchanges. Finally, *AriseCoin* was an ICO attempt by a fake bank named *AriseBank*. The SEC stopped it in January 2018.

G. MONEY LAUNDERING

Money laundering consists of making large amounts of money obtained from illegal activities appear to come from legitimate sources. It consists of three stages: placement, layering and integration. In the first stage, dirty money is introduced into the legitimate financial system. Then, the money is moved several times to create confusion, moving through numerous accounts. Finally, it is integrated into the financial system through further transactions until the process is completed. The main problem associated with this criminal activity is to make the proceeds legal without arousing the suspicion of law enforcement.

Moser *et al.* [56] focus on money laundering by analysing three mixers based on the transaction graph extracted from the blockchain and trying to establish relationships between inputs and outputs. The authors found that *BitcoinFog* and *Blockchain.info* make difficult to relate input and output transactions. Indeed, they could not find any direct connections in their transaction graphs.

Brenig *et al.* [57] focused on money laundering implemented in cryptocurrencies from an economic perspective. They present the structure of the money laundering process and the primary anti-money laundering control and analyse the contextual and transactional factors that facilitate money laundering. In conclusion, they argue that cryptocurrencies may actually encourage the exploitation of money launderers.

Fanusie and Robinson [58] used Elliptic's forensic analysis tool, which uses blockchain data with a proprietary dataset of bitcoin addresses associated with 102 known illicit entities. The authors discovered that darknet marketplaces, such as Silk Road or AlphaBay, were the source of almost all illegal bitcoin laundered through the conversion services (exchanges, mixers, ATMs, online gaming sites) identified

Do you want to earn more than $500 per day? You've heard of Crypto currency, Ethereum, and Bitcoin. I want to show you a new way to achieve this profit!

1) First of all you need Bitcoin and Ethereum wallets. I prefer to use https://blockchain.com wallet but you can use any wallets that you like. https://login.blockchain.com/signup/ - link for registration.

2) You need to have at least $50 in Bitcoin (0.012 BTC) at your wallet

3) Then you need to swap 0.012 BTC to Ethereum using exchange https://cryptone.eu on site 0.012 BTC = 0.46355 ETH

How to exchange here:

1. Fill the form and click Exchange

2. You will see the BTC address from the exchange where you should send specified amount of BTC from your wallet (don't forget about commission for transaction in BTC network about $0.08)

3. You will receive Ethereum after the 1st confirmation of BTC network, it takes about 10 - 20 minutes

4) Now you exchange Ethereum to Bitcoin using exchangers https://en.changelly.com/ or https://www.alfacashier.com/ on course 0.46355 ETH = 0.017279 BTC

5) So you have 0.005279 BTC (that's about $22) profit just from one exchange!!!

6) Repeat this procedure 20+ times and your profit will be more than $500!!! If you start with a larger amount of BTC profit will be higher.

FIGURE 8. Fake exchanges linked by a fraudulent website suggesting how to earn money with cryptocurrencies.

in their study. Looking at geographical patterns, conversion services based in Europe received the most significant share of illicit bitcoins among the identifiable regions.

Hu *et al.* [59] analysed money laundering activities through the Bitcoin network. They used the data collected from July 2014 and May 2017 to differentiate money laundering transactions from regular transactions, implementing different classifiers based on deepwalk embeddings.

The work of Weber *et al.* [60] aimed to classify money laundering transactions. The authors also proposed a novel dataset of Bitcoin transactions, provided with handcrafted features. They realised it with the data provided by Elliptic, a cryptocurrency intelligence company focused on safe-guarding cryptocurrency ecosystems against criminal activity.

Lorenz *et al.* [61] aimed to detect money laundering patterns. In particular, they propose a detection system using minimal access to labelled transactions by utilising a dataset [62] composed of 200K transactions.

III. MATERIALS AND METHODS
In this section we describe the data sources from which we have collected scams, and the open-source toolchain that we have developed to support our work.

A. DATA SOURCES
To the best of our knowledge, no public dataset encompassing all the scams identified in Section II seems to exist. Therefore, we construct our dataset of scams by gathering and homogenizing data from various sources:

- BadBitcoin, a website which collects crypto scams since 2014. Scams are reported by users through a contact form (see Figure 9). Scam reports are then moderated by the site administrators.
- EtherAddressLookup, a browser add-on that alerts users when trying to access known scam domains in the

FIGURE 9. Form for reporting scams on BadBitcoin.

Ethereum realm. To this purpose, the add-on exploits a blacklist of scams curated by the authors and publicly available on GitHub.

- BitcoinAbuse, a public database of Bitcoin addresses used by scammers. The database is built from users reports, apparently without moderation. The report form (see Figure 10) allows users to specify the address of the scam and its type, among the following: (i) ransomware; (ii) blackmail scam; (iii) sextortion; (iv) darknet market; (v) Bitcoin tumbler; (vi) other. Further, users can specify the abuser (e.g., the email address from which the victim of the scam has been contacted) and a description (e.g., the body of a blackmail email). The database is accessible through public APIs.

Besides those listed above, there are other websites that collect crypto scams, but they have drawbacks that make them unusable for our purposes. CryptoScamDB is a public

TABLE 5. Summary of literature on money laundering.

Ref.	Year	CCY	Highlights	In USD	Dataset
[56]	2013	BTC	- analysis of three mixers based on the transaction graph - for BitLaundry found direct connections in transaction graph	—	N/A
[57]	2015	BTC	- economic point of view of money laundering implemented in cryptocurrencies - blockchain properties act as incentives for money laundering	—	N/A
[58]	2018	BTC	- analysis of conversion services - proprietary dataset of 102 illicit entities	—	N/A
[59]	2019	BTC	- automatic classification of money laundering transactions - analysis using data collected over three years	—	N/A
[60]	2019	BTC	- automatic classification of money laundering transactions - tested Logistic Regression, Random Forest, Multilayer Perceptrons, and Graph Convolutional Networks	—	N/A
[61]	2020	BTC	- active learning to detect money laundering patterns - used dataset of 200K labelled transactions	—	[62]

FIGURE 10. Form for reporting scams on BitcoinAbuse.

database of malicious URLs, but we exclude it because it is a strict subset of EtherAddressLookup. ScamAlert and BitcoinWhosWho collect users reports about crypto scams and make them accessible through a search form. Since there is no way to download reports from these websites, we cannot include them in our collection.

We partition scams in two main categories, according to the way they are indexed by the data sources:

URL-reported scams are those which are indexed by their URLs, like in BadBitcoin and EtherAddressLookup;

address-reported scams are those which are indexed by the Bitcoin address reported by users, like in BitcoinAbuse.

B. TOOLCHAIN
To support our work we make available a toolchain, which comprises the following tools:

- scripts that download lists of scams from our data sources, using the BadBitcoin, EtherAddressLookup, and BitcoinAbuse APIs (see Sections IV-A and IV-B);
- a script that downloads a list of snapshots for each URL reported as a scam from the Wayback Machine. The script follows potential redirect snapshots until it obtains an error response code or an HTML source, and classifies each scam accordingly (Section IV-B);

- a script that crawls *BitInfoCharts* and *Vivigle* to build a list of Bitcoin addresses that belong to reputable entities (e.g. wallets) that we use for filtering out our list of address-reported scams (Section IV-A);
- a script that classifies a text according to our taxonomy. We classify address-reported scams using the scam descriptions in BitcoinAbuse, and URL-reported frauds by inspecting snapshots of the websites (Section V).

IV. COLLECTION OF SCAM DATA
In this section we illustrate the methodology followed to collect crypto scams from the web, and draw some initial statistics on the resulting dataset.

A. COLLECTING ADDRESS-REPORTED SCAMS
The reports downloaded from BitcoinAbuse often contain spurious data that we need to fix in order to construct a reliable dataset. Indeed, the scam type reported by users does not always reflect the actual one: for instance, users often report blackmail scams as malware, just because the blackmail they receive claims that their computer has been hacked. Further, some users report as scam Bitcoin addresses that belong to exchanges and miners. For instance, a report for the address 1CK6KHY6MHgYvmRQ4PAafKYDrg1ejbH1cE, describes the scam as "*Shipping scam Fake logistics company used as front to blackmail people out of money and personal information [...]*". This report seems completely misleading, since the address is controlled by Slush Pool, the oldest Bitcoin mining pool. We discuss below how we filter out such imprecise or misleading reports.

1) We start by downloading all the reports through the BitcoinAbuse APIs. This procedure results in 163,840 reports (up to date 2020-05-15), which refer to 47,099 unique addresses. We discard the scam type from the reports.

2) To filter out erroneous reports, we construct a safe-list of addresses owned by reputable entities, like, e.g. exchanges and mining pools. For this purpose, we query BitInfoCharts and vivigle, two websites that tag Bitcoin addresses with their owners. We add to our safelist the addresses whose tag refers to a known mining pool or exchange (e.g., Coinbase).

TABLE 6. Statistics on the address-reported scams dataset.

	Addresses	Reports
Data extracted	47,099	163,840
Safe elements	24	63
Total Scams	**47,075**	**163,777**

3) We remove from the dataset the 167 safe addresses collected in the previous step. The filtered dataset contains a total of 47,075 scam addresses.

Table 6 summarizes the results of these steps, while Figure 11 shows the distribution of reports in the BitcoinAbuse dataset. The figure represents, in the y-axis, the number of Bitcoin addresses reported N times, where N belongs to the interval in the x-axis. Although the figure does not display the addresses with less than 20 reports, we count 30,191 addresses with only one report, and 15,928 addresses with $2 < N < 20$ reports. We also count 2 addresses with more than 500 reports. Although the figure displays only 978 addresses, the number of associated reports is relevant, and indeed it amounts to $\sim 40\%$ of the total. Overall, $\sim 70\%$ of reports concern a Bitcoin address that has been already reported.
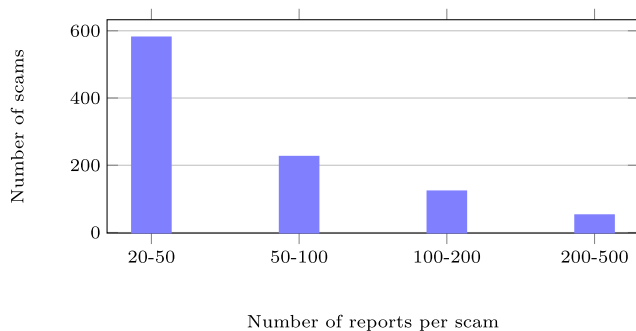
FIGURE 11. Distribution of reports among scams.

B. COLLECTING URL-REPORTED SCAMS

We collect URL-reported scams from BadBitcoin and EtherAddressLookup. Since we identify several issues in the datasets items, our final list of frauds does not comprehend all the URLs published by BadBitcoin and EtherAddressLookup. More precisely, we first download all datasets URLs, and then, for each issue we identify, we perform an ad hoc analysis that filters out the affected URLs. Some of our analyses require inspecting the HTML pages of scams websites, but both datasets focus on URLs and do not provide any HTML page. Moreover, most scam websites are not online anymore and, perhaps, have been online for a limited lifetime. For these reasons, we crawl Web Archive, a digital archive of the World Wide Web, also containing snapshots (i.e., archived copies) of now offline sites. We can summarise the pipeline for collecting domain scams in the below steps.

1) *Datasets crawling*: download scams from datasets;
2) *URL filtering*: remove scams with an invalid URL;
3) *Web Archive crawling*: crawl scams snapshots from Web Archive;
4) *HTML filtering*: remove scams which do not have at least one valid snapshot;

In the following, we discuss our methodology for collecting URL-reported scams. The results are shown in Table 7. Each row focuses on a different analysis, showing for both BadBitcoin and EtherAddressLookup the total amount of snapshots (*Snapshots* columns) and URLs (*URLs* columns) affected by the current analysis.

1) DATASETS CRAWLING

We start by crawling BadBitcoin and EtherAddressLookup to extract the list of reported scams (note that BadBitcoin no longer makes available the scam list). This results in a list of 6,721 URLs reported by BadBitcoin, and 12,338 URLs reported by EtherAddressLookup.

2) URL FILTERING

The first challenge we encounter is that some reported URLs are not valid. For instance, they do not contain the protocol or parts of the hostname, include spaces within the hostname, or are just nonsense (e.g., *War* and *Religion* are reported as scam URLs in BadBitcoin). Whenever possible, we try to fix these malformed URLs, e.g. by adding the missing *http://* or *https://* prefix. We remove from the dataset the malformed URLs which we did not manage to fix. Both BadBitcoin and EtherAddressLookup reported some URLs: we mark them as *duplicates* in Table 7, counting them only in the EtherAddressLookup column.

3) WEB ARCHIVE CRAWLING

EtherAddressLookup preemptively reports some URLs as scams. For instance, this is the case of *myetherieumwallet.com*, *myetheresswallet.com*, and many other URLs resembling the URL of MyEtherWallet, a widespread Ethereum wallet. It seems that the administrators of EtherAddressLookup uploaded many variants of the original URL in a single commit to contrast phishing attacks beforehand. The commit is available at this URL.

To detect the URLs that actual scams have exploited, we inspect the websites that scams have operated. Since most of these websites have been closed over time, we retrieve their snapshots on Web Archive. This procedure results in 516,699 snapshots of the websites reported by BadBitcoin and 113,457 snapshots for EtherAddressLookup. We note that downloading the snapshots from Web Archive sometimes produces internal errors, either for the whole URL or some of its snapshots. In Table 7 we show the total number of URLs/snapshots we exclude from the dataset because of internal errors. The "total items crawled" row in Table 7 displays the number of URLs/snapshots that remain in the dataset after this filtering step.

TABLE 7. Analysis of URL-reported scams datasets.

Methodology steps		BadBitcoin		EtherAddressLookup	
		URLs	Snapshots	URLs	Snapshots
DATASETS	Scam datasets elements	6,721	-	12,338	-
URL	Malformed	22	-	7	-
	Duplicates	-	-	20	-
WEB ARCHIVE	Total items available	6,699	516,699	12,311	113,457
	Internal errors	33	315,548	28	88,153
	Total items crawled	**6,666**	**201,151**	**12,283**	**25,304**
HTML	Empty/error/domain on sale	395	30,664	1,377	8,387
	No snapshots	861	-	8,250	-
	Total URLs/snapshots	**5,410**	**170,487**	**2,656**	**16,917**
Total			Scams		Snapshots
			8,066		**187,404**

4) HTML FILTERING

As the last step, we analyse the HTML content of snapshots to remove from the dataset those without relevant content. More specifically, we remove the snapshots of the following forms:

- the empty HTML pages, containing minimal HTML tags, but without a real content;
- the error pages, containing messages like, e.g. *"The requested URL was not found on this server"* or *"Settings are broken. Contact developers please"*;
- the domain advertisement pages, containing messages like e.g. *"Domain expired. Contact hosting provider"*.

We then exclude the URLs for which all snapshots are of any of the three forms above (including the URLs without any snapshot). Note from Table 7 that URLs without snapshots are predominant in EtherAddressLookup, coherently with our conjecture that this service used to report URLs preemptively.

The last row of Table 7 aggregates the data of the two datasets of URL-reported scams.

V. DESIGN OF THE SCAM CLASSIFIER

In this section we introduce our scam classification technique. This technique is based on the detection of characterizing keywords in reports (for address-reported scams) or websites (for URL-reported scams). Since the details vary for the two kinds of scams, we discuss them separately. We conclude the section by evaluating the effectiveness of our technique with respect to relevant metrics.

A. CLASSIFYING ADDRESS-REPORTED SCAMS

To classify address-reported scams we develop a text-based analysis of scam reports. The analysis is based on a preprocessing of reports. We start by manually inspecting hundreds of scams' reports to understand their typical structure and identify the keywords characterizing each type. Based on this, we associate each pair (scam type, keyword) with a weight, reflecting the keyword's relevance to identify the scam type.

We then develop a text-based classifier based on this dictionary of weighted keywords. For each scam report, the

TABLE 8. Classification of address-reported scams.

Category	#Addresses
Scam	33,570
Other	299
Not Enough Data	13,206
Total	**47,075**

classifier computes the *score* for each scam type. This score is the summation of the weights of all the keywords occurring in the report (considering only the keywords associated with the scam type against which we are computing the score).

We denote an address as a *scam* if at least one of its reports has a score above a given threshold for some scam type. In such a case, the scam types associated with the address are those for which the score is above the threshold.

For the addresses which are not identified as scams according to this condition, we further distinguish between those for which all the reports are too short to enable a classification (*"not enough data"*), and the *"other"* addresses. Note that the address of an actual scam may not be recognized as such when all its reports contain shallow or uninformative descriptions.

Table 8 counts the addresses in each category. A comparison with URL-reported scams is postponed to Section VI.

B. CLASSIFYING URL-REPORTED SCAMS

The classification technique for URL-reported scams extends that for address-reported scams shown in Section V-A, with respect to which it must cope with additional issues. A first issue is that some of the datasets' websites are not related to cryptocurrencies: this is the case e.g. of some Ponzi schemes that accept payments only in fiat currencies. Our goal is to classify these items as scams, specifying that they are *fiat* scams. Another issue is that the dataset also contains some non-fraudulent websites: this is the case of, e.g. of discussions websites on cryptocurrencies, and "spamdexing" websites that attempt to tamper with search engines algorithms by containing long pieces of uninformative crypto-related text. Since spamdexing does not fit our definition of scams, we do not want to classify these websites as scams.

To classify URL-reported scams we proceed as follows:
1) We first inspect the timestamp of snapshots to rule out the websites for which all the snapshots are older than the Bitcoin release date (3 January 2009).
2) We construct a dictionary of keywords related to cryptocurrencies. We use this dictionary to tag each snapshot as "fiat" or "crypto": more precisely, we count the occurrences of the keywords in snapshots, and we label as crypto the snapshots that reach a given threshold.
3) We then apply the classifier of Section V-A on each snapshot to detect whether it is a scam or not and the associated scam types.
4) We classify websites by aggregating the results of their snapshots. More precisely, we classify a website as:

TABLE 9. Classification of URL-reported scams.

Category	Crypto	Fiat	Total
Scam	5,329	729	6,058
Other	381	165	546
Not Enough Data	787	675	1,462
Total	**6,497**	**1,569**	**8,066**

- *"crypto scam"*, if at least one of its snapshots is tagged as "crypto" and it is associated with at least one scam type. Common examples of these scams are Ponzi schemes accepting bitcoins, and advance-fee scams pretending that some celebrity rewards users with cryptocurrencies.
- *"fiat scam"*, if none of its snapshots is related to cryptocurrencies, but it still conducts fraudulent activities. Common examples of these scams are Ponzi schemes accepting fiat currencies and not supporting cryptocurrencies.

Similarly to address-reported scams, the classifier tags as *"not enough data"* the websites for which no snapshot has enough content, and as *"other"* those for which no snapshot reaches the threshold of the scam-related score for any scam type. Unlike address-reported scams, the *"other"* category now contains several no-scam items. We classify websites mentioning cryptocurrencies without scam behaviour as *"crypto no-scam"*. This is the case, e.g., of websites reporting news about blockchains. Similarly, we classify websites which neither mention cryptocurrencies nor conduct scam activities as *"fiat no-scam"*. This is often the case of legit websites wrongly reported by users.

Table 9 summarises the results obtained by the classifier. The scam types measured (associated with the *crypto-scam* and *fiat-scam* categories) are discussed in Section VI, compared with address-reported scams.

C. METRICS

We use five standard metrics to evaluate the performance of our classifier: precision, sensitivity, specificity, accuracy and F-measure. All these metrics are defined in terms of true/false positives/negatives. A true positive (TP) is a result in which the classification model correctly predicts the positive class, while a true negative (TN) is an outcome in which it correctly predicts the negative class. On the contrary, a false positive (FP) is an outcome in which the model incorrectly predicts the positive class, and, finally, when the model incorrectly predicts the negative class, we have a false negative (FN). *Accuracy* indicates how many of the predictions are correct:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

Although this is a commonly used metric and a good general evaluator, it is not able to give a fair insight when the problem has an unbalanced distribution of classes, like our scenario. *Precision* is the proportion of true positives (TP)

in the set of positive results, and measures how good the classifier is when the prediction is positive:

$$Precision = \frac{TP}{TP + FP}$$

Specificity is the proportion of negative results that are correctly identified, and *Sensitivity* (called *Recall* in multiclass classification), is the proportion of positive results that are correctly identified:

$$Specificity = \frac{TN}{TN + FP}$$
$$Sensitivity = \frac{TP}{TP + FN}$$

F-measure is the weighted average of precision and recall:

$$F - measure = \frac{2 \times Precision \times Recall}{Precision + Recall}$$

Since the multiclass classification realised in this work is multi-label (ML), we use the following metrics to take into account that we have a TP when, for a scam, the real labels are a subset T of the labels produced by the classification model R, as proposed in [63]:

$$MLPrecision = \frac{T \cap R}{R}$$
$$MLRecall = \frac{T \cap R}{T}$$
$$MLF - measure = \frac{2 \times MLPrecision \times MLRecall}{MLPrecision + MLRecall}$$

We compute each of the presented metrics for every class included in both binary and multiclass classification evaluation to represent the per-class performance. In addition, we provide the weighted average for both evaluations as follows. We compute metrics for each class and define their average weighted by the number of actual instances for each class. In this way, we take into account the class imbalance.

Finally, we also represent the Receiver Operating Characteristic (ROC) curve and computed the Area Under the ROC (AUC) directly from the ROC curve itself. The ROC curve depicts the trade-off between False Positives (FP) and True Positives (TP); therefore, it represents the True Positive Rate (TPR), that is the *Specificity*, on the y-axis and False Positive Rate (FPR), computed as $1 - Sensitivity$ on the x-axis.

D. CLASSIFIER EVALUATION

We evaluate our scam classification tool as follows:
1) We sample 200 random elements from our collection of 8,066 *potential* scams constructed in Section IV.
2) Each author (unaware of the output of the tool) independently classifies each potential scam in the sample. To this purpose, the authors can inspect the same data used by the classifier, i.e. the HTML snapshots for URL-reported scams, and the user reports for address-reported frauds. The possible outputs are *No scam*

(which aggregates the *Other* and *Not Enough Data* categories produced by our tool) or *Scam*. In the latter case, the authors also indicate a set of scam types among those defined in Section II.

3) We share the results of the independent manual classifications, and we discuss each element until we find an agreement on its classification.

4) Finally, we compare the manual classification with the output of the tool, against:

   a) a *binary classification*, which distinguishes between scams and no scams. More precisely, by further specifying the classic definition of TP given in Section V-C, we consider an item to be a TP when both the manual and the automatic classification agree on classifying the item as a scam, regardless of the scam types provided.

   b) a *multi-label classification*, that evaluates the performance in the multiclass scenario, where each scam could potentially belong to more than one type. Here, we define an item as TP when the scam types determined by the manual classification are a subset of those calculated by the tool, and the manual and automatic classification agree on classifying the item as a scam. For instance, if an item is manually classified as Ponzi and the tool determines that both Ponzi and Fake Services meet the minimum threshold, we say that the item is classified correctly.

Tables 10 and 11 show the confusion matrix and the classification metrics for the binary classification. Table 12 shows the multi-label classification metrics.

1) ON THE BINARY CLASSIFICATION

With regard to binary classification, Table 10 works as a general indicator of the effectiveness of our classification tool. In particular, 170 items out of 200 were correctly classified (133 as scam, and 37 as no scam), while only 30 were misclassified. Concerning the latter point, we observe that only 9 scams were misclassified as no scams. This aspect is of particular relevance because, considering the goal of the tool, false positives are not a big issue. This trend is confirmed by the metrics reported in Table 11 and, in particular, by the AUC value in Figure 12a.

Table 11 shows that precision and sensitivity (both oriented towards the positive class) have high values on the Scam class, while the specificity is definitely lower on the Scam class since 21 No scams out of 30 were misclassified as Scams. Further, as depicted in Figure 12a, the binary classification reaches a true positive rate of 89%, which is quite high and permits correctly addressing a huge set of real scams.

In general, our tool achieves satisfactory results considering the importance of Scam class recognition for our classification.

2) ON THE MULTI-LABEL CLASSIFICATION

In Table 12 we report a specific evaluation of every single scam type addressed in the multi-label classification.

TABLE 10. Binary classification confusion matrix.

		Predicted	
		Scam	No Scam
Actual	Scam	**133**	9
	No Scam	21	37

In particular, for the sake of completeness, we show the metric values computed on the single classes, without taking into account the possibility that a scam may belong to more than one type. This aspect is more clearly represented by the ROC curve in Figure 12b.

From Table 12 it is clear that the most misclassified type of scam is the Fake Service, with a precision value of 44%. The reason is clearly due to the fact that there is a correlation between Fake Services and Ponzi schemes and, therefore, they are frequently classified as Ponzi schemes. In general, however, actual Fake Services are predicted correctly with a recall of 81%. For the same reason, Ponzi and Advance-Fee scams have a low precision of 73% and 63%, respectively. This is because Ponzi scams are often misclassified as Fake Services, while Advance-Fee scams are misclassified as Ponzi and Malware, and as Fake Services for the most part. The detail of the correlations between the scam types is deeply illustrated in the following Section VI.

The table does not contain any malware because none was detected from our sampling of 200 different scams.

Even in this case, we report the weighted average of the classes, in order to consider the class imbalance, and the samples average, and to calculate the metrics for each item, and then their average. Considering the average metric results, the classification tool proposed have obtained satisfactory values, even though they give room for improvements towards the accurate recognition of the most significant scam class.

Finally, in Figure 1b we show the ROC curve for each type of scam addressed in the multi-label classification evaluation, and the micro and macro average as terms of comparison. In contrast to the results shown in Table 12, which are targeted at individual classes, here we can give a more general evaluation of multi-label classification, taking into account the possibility that a scam may belong to more than one type. Specifically, we can see that the AUC of each of the scam types addressed is between 78% of the Fake Services category and 89% of the Advance-Fee scam category.

The average AUC values reported also represent a satisfactory classification result, as the macro average, which gives equal weight to the classification of each label, and the micro average are 85%, 86% respectively.

VI. APPLYING THE CLASSIFIER IN THE WILD

We now classify the scams collected in Section IV according to our taxonomy. Our tool associates each fraud to zero, one or more types of fraud. We exclude all frauds that have no associated type from the following evaluations,
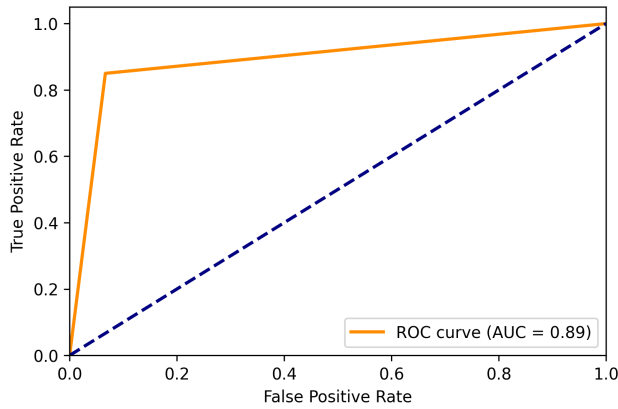
(a) Binary scam classification ROC Curve. The number in brackets indicates the AUC value for the Scam class, considered as positive.
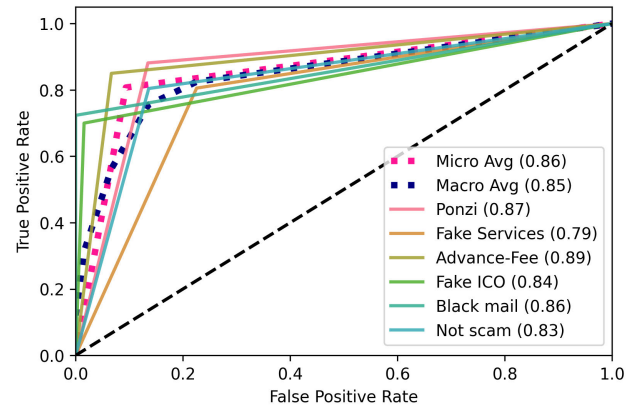
(b) Multi-label classification ROC Curve for each scam type. The numbers in brackets indicates the AUC value. Micro Avg AUC and Macro Avg are reported as comparison terms.

FIGURE 12. ROC curves of binary and multi-label classification.

TABLE 11. Evaluation of binary classification. "Scam" and "No Scam" columns represent the metric values for each of the two classes. The last columns indicates the weighted average metric values, computed independently for each class and then averaged taking into account the number of samples per class.

	Scam (%)	No Scam (%)	Weighted Avg (%)
Precision	86.36	80.43	84.64
Sensitivity	93.66	63.79	85.00
Specificity	63.79	93.66	72.46
Accuracy	85.00	85.00	85.00
F-measure	89.87	71.15	84.44

TABLE 12. Evaluation of multi-label classification. The metrics shown are calculated without taking into account that a scam may belong to more than one type, in order to have a more specific evaluation per class.

	ML Precision (%)	ML Recall (%)	ML F-measure (%)	# Samples
Ponzi	73.0	88.0	80.0	59
Fake Services	44.0	81.0	57.0	36
Advance-Fee Scam	59.0	85.0	69.0	20
Fake ICO	70.0	70.0	70.0	10
Black mail	100.0	72.0	84.0	47
No Scam	64.0	80.0	71.0	46
Weighted Avg	71.0	81.0	74.0	218

which corresponds to all the items we previously classified as Other and Not enough data. In summary, our analysis involves 39,628 elements, which are the subset identified as scams in Section V-A (33,570 address-reported), and Section V-B (6,058 URL-reported). The analysis graphs (i.e. Figures 13 to 15) show three different distributions of the scams identified above in relation to the total number of scams. For the sake of clarity, we aggregate all subtypes within their primary type, i.e. the *Fake Services* type also includes exchange, wallet, mixing, mining, donation, while the *Malware* type includes ransomware and crypto loggers. Consequently, the final list of types we focus on includes:

Ponzi, *Malware*, *Advance-Fee*, *Fake Services*, *Black Mail*, and *Fake ICO*.

A. PURE VERSUS HYBRID SCAMS

We define frauds associated with only one type as *pure*, and *hybrid* those that have more than one type. The former is deeply studied in the literature, while the latter is more recent and often more difficult to detect. Nevertheless, hybrid scam constitute an important percentage of frauds, as evidenced by Figure 13 which represents the distribution of both types. In particular, in Figure 13a we observe that 28% of address-based scams and almost half of URL-reported scams are hybrid.

To better understand hybrid scams, first we quantify how many types of scams they generally involve (Figure 13b). In particular, the first column represents the percentage of scams related to only one type, i.e. pure scams, while the other columns detail the structure of hybrid frauds. As it can be seen, few scams have three or more types. Most of them involve exactly two types. Specifically, in URL-reported scams, two-types scams (40%) are close to pure scams (49%). In the following analysis we provide more information on the nature of hybrid scams by understanding which specific types they involve.

B. DISTRIBUTION OF SCAM TYPES

Hybrid scams do not involve all types of fraud in the same way. In Figure 14, for each type of scam we examine the distribution of pure (light bars) and hybrid (dark bars) scams, also showing the distinction between URL-reported (blue bars) and address-reported (red bars) scams. Since we show the results as percentage values, for each type the sum of all four bars is exactly 100%. We note that the hybrid scams are counted in multiple bars (once for each type involved).

To be specific, the Black Mail type is the only one that consists mainly of pure scams (75%). In contrast, the other types are largely hybrid scams: Ponzi, Advance Fee, and
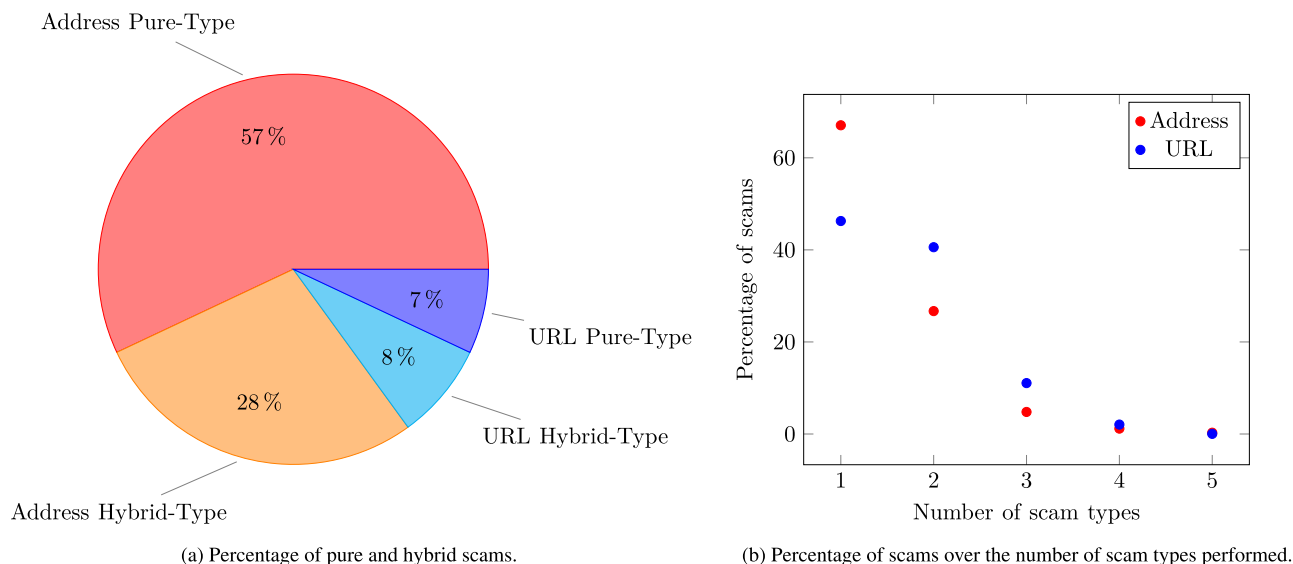
(a) Percentage of pure and hybrid scams.

(b) Percentage of scams over the number of scam types performed.

FIGURE 13. Distribution of pure-type and hybrid-type scams.
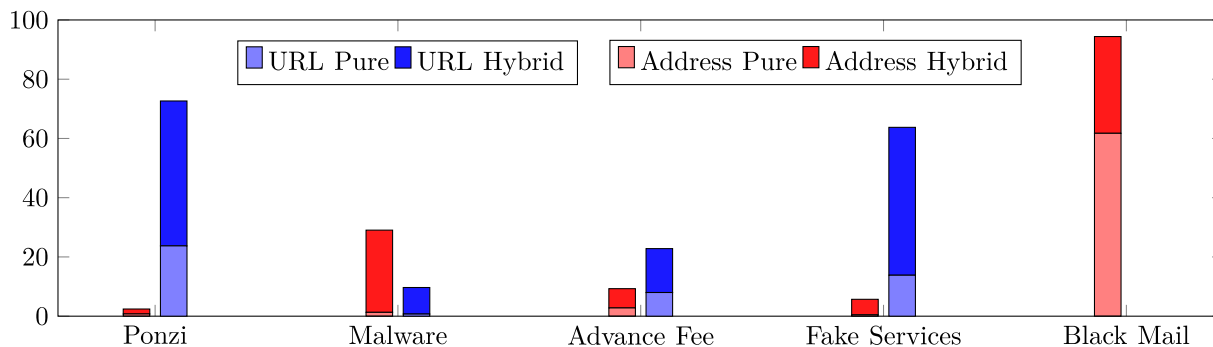

FIGURE 14. For each type, percentage of pure and hybrid scams.

Fake ICO with ∼ 70%; Fake Service more than 80%; and Malware more than 90%. Although all the types analysed are well defined in the literature, the above values suggest that, in practice, many of them are mostly used in combination with others rather than alone.

The distinction between blue and red bars also allows us to look at fraud types from a different perspective, finding that fraud types are not uniformly distributed between address-based and URL-reported scams. This peculiarity is related both to the dataset in which a scam is reported, and to the medium that scams use to reach end-users: indeed, scams reported by URL are websites, whereas scams reported by address are mainly distributed by mail. The Black Mail, Malware and Advance Fee types are usually address-based scams, while most Ponzi schemes and Fake ICOs are URL-reported scams. In particular, the Fake Services type evenly affects both categories.

Figure 13b shows that hybrid frauds usually involve two types and in Figure 14 we observe that some types have similarities. For instance, Malware and Black mail are mainly address-based scams, while Ponzi schemes and fake ICOs

often use websites and also have similar percentages of hybrid scams. In the following, we perform further analysis to check whether there is indeed any correlation between the various pairs.

C. CORRELATION BETWEEN SCAM TYPES
In Figure 15 we report a heatmap matrix to further investigate the correlations of scam types. Diagonal values indicate the correlation of a type with itself, which corresponds to a pure scam.

The main observation that can be made is that most Ponzi schemes are pure scams: often their websites claim to be HYIP, while sometimes they explicitly claim to be Ponzi schemes. We note that usually pure Ponzis are very simple websites with a minimal user interface, which only allows the user to calculate the expected profit and send money: a subset of these websites originated before Bitcoin and still use fiat currencies. Since pure Ponzi fraud can be easily recognised, fraudsters have created more elaborate websites, pretending to be sites offering blockchain services that reward users: for this reason, Ponzi has a clear correlation with Fake Services.
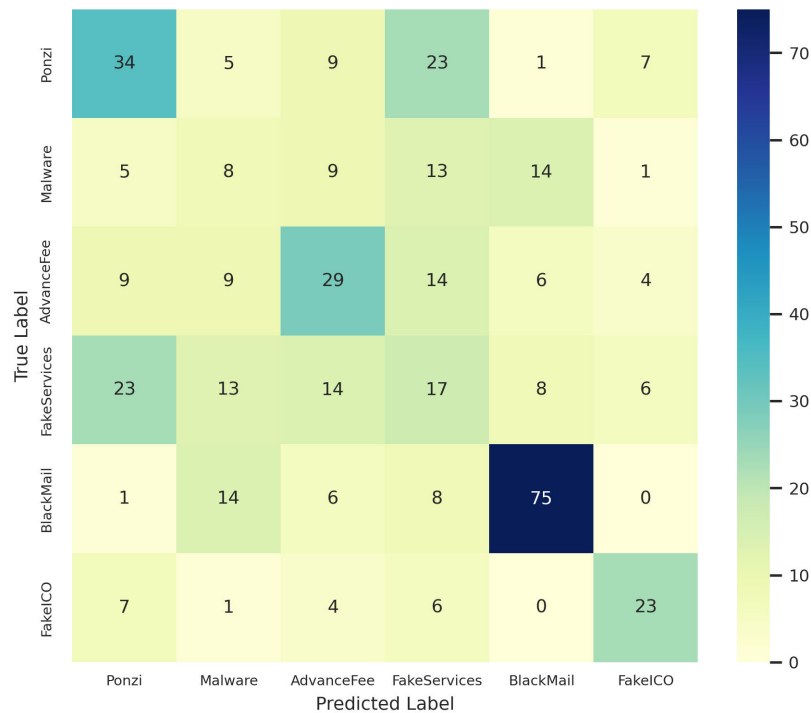
FIGURE 15. Heatmap of the confusion matrix obtained from the classifier proposal, showing the correlation between the types of scams.

They use various strategies, for example pretending to be: i) portfolios where you earn over time; ii) exchanges with really good rates; iii) mining pools that reward participants over time.

Consequently, we observe that some types such as Fake Services and Malware have a weak correlation with themselves: they are mostly used as support for other types of scams (such as Ponzi and Black Mails). On the other hand, Black mail type has a high correlation with itself, because they do not hide their true purpose but immediately make their request for money clear. Sometimes, in order to convince users that the sender really has compromising material belonging to the victim, black mail pretends to have used malware. This is why we notice a slight correlation between black mail and malware types. Since we collected less malwares than black mails, this correlation causes a spike in the hybrid malware count.

Fake ICO scams can be easily recognised from other types of scams because they are websites that describe a new cryptocurrency project, provide a roadmap and present profiles of the founders. Although their website sometimes mentions exchanges, mining systems, and other services, they are not fake services even though the primary fraud mechanism is the same: all addresses provided to end-users belong directly to the scammers, and they will never return the funds.

Advance-Fee scams also use this technique. However, we recognise them because they pretend to be some famous entrepreneur or company.

The widespread correlations existing between the analysed scam types, and provided in Figure 15, more clearly motivate the poor accuracy metric values obtained in the multi-label evaluation of some scam types, given in Section V-D2 and previously denoted in Table 12.

VII. CONCLUSION

To the best of our knowledge, this work is the first comprehensive research on cryptocurrency scams, ten years after their emergence. We have started by surveying the scientific literature on scams. Then, we have built a collection of scams by fetching data from various sources. Coarsely, we have distinguished between two types of scams, according on how they are indexed by websites: *address-reported scams* are Bitcoin addresses reported by potential victims; *URL-reported scams* are fraudulent websites, which we have retrieved from Web Archive. We have publicly released a dataset of scams [4]–[6] containing 47,075 address-reported scams (with 163,777 reports), and 8,066 URL-reported scams (with 187,404 snapshots). We have developed an open-source tool [3] for classifying scams according to our taxonomy. We have evaluated our tool using standard techniques and metrics. Finally, we have used our tool to perform a multi-label classification of the collected scams, based on which we have analysed the distribution and correlation between scam types, and the distinction between pure and hybrid scams.

We observe that, although most scams in our dataset are related to Bitcoin, almost all scams do not rely on the features

of that specific blockchain, but just use the native cryptocurrency of the blockchain as a means of payment. Therefore, it is possible to see in the future the emergence of scams operating on other blockchains besides Bitcoin, as their native cryptocurrencies become more popular. Indeed, our dataset already contains some scams related to Ethereum.

Based on our experience, we provide the following recommendations for counteracting crypto scams.

A. RECOMMENDATION #1. IMPROVE SCAM REPORTING SYSTEMS

As we have noted in Section III, the existing public data sources on crypto scams are heterogeneous, and not completely reliable. This hampers the development of effective scam detection and classification tools, for which it would be crucial to have a uniform and reliable dataset of scam. To overcome these issues, our recommendation is to construct a scam reporting system which is comprehensive (i.e., it allows users to report scams of any type, targeted to any blockchain), with a uniform taxonomy of scams (like, e.g., the one we have proposed in Section II) and moderated, so to reduce the amount of spurious or incorrect data. Further, this system should guide users towards the correct self-classification of scams, e.g. by providing an interactive questionnaire, and by showing scam templates of the various types.

B. RECOMMENDATION #2. DEVELOP A BROWSER EXTENSION TO WARN AGAINST SCAMS

A success factor of cryptocurrency scams, besides users' greediness, is that non-technical users often find it difficult to distinguish fraudulent websites from legit ones. Accordingly, we recommend the implementation of a browser extension which inspects websites, alerts users when it detects potential scams, and advises users what to do if they have already been scammed. Similarly, the browser extension could alert users when they try to send money to blockchain addresses related to scams, or when they read blackmails. Our toolchain can be the basis for developing such a browser extension.

C. FUTURE WORKS

Our work can be the basis for several future developments. First, we believe that the scam detection metrics observed on our tool (Section V-D) could be improved by more sophisticated machine learning techniques. For instance, our relatively high FP rate is the consequence of our design choice of preferring to report legit items as scams, rather than not reporting potentially fraudulent ones. A possible strategy for reducing the FP rate is to extend our dataset with "safelists" of legit items. To this purpose, one could exploit public datasets of blockchain-related websites, like e.g. those related to social good projects [64], digital coupons [65], and so on. Another possible development is to extend the set of features used by the detection engine, including e.g. the HTML structure of websites, the transactions related to the Bitcoin addresses possibly linked to a scam, etc.

The analysis of scams described in Section VI is focussed on classifying scams according to our taxonomy. Based on this, one could perform further analyses, e.g. to measure (and aggregate by type) the lifetime of scams, their evolution over time, and their economic impact, when Bitcoin addresses are available (as for address-reported scams). To this purpose, one would need to combine the tools developed in this work with blockchain query engines, like e.g. BlockAPI [66].

REFERENCES

[1] W. Chen, T. Zhang, Z. Chen, Z. Zheng, and Y. Lu, "Traveling the token world: A graph analysis of Ethereum ERC20 token ecosystem," in *Proc. Web Conf.*, Apr. 2020, pp. 1411–1421.

[2] N. Kyriazis, S. Papadamou, and S. Corbet, "A systematic review of the bubble dynamics of cryptocurrency prices," *Res. Int. Bus. Finance*, vol. 54, Dec. 2020, Art. no. 101254.

[3] *Cryptocurrency Scam Classifier Repository*. Accessed: Oct. 26, 2021. [Online]. Available: https://github.com/blockchain-unica/cryptoscam-classifier

[4] A. Loddo, "Badbitcoin dataset," Univ. Cagliari, Cagliari, Italy, Tech. Rep., 2020, doi: 10.7910/DVN/MOARX1.

[5] S. Serusi, "BitcoinAbuse dataset," Univ. Cagliari, Cagliari, Italy, Tech. Rep., 2020, doi: 10.7910/DVN/SMPQBQ.

[6] A. Loddo, "EtherAddressLookup dataset," Univ. Cagliari, Cagliari, Italy, Tech. Rep., 2020, doi: 10.7910/DVN/XCP6KF.

[7] A. Holub and J. O'Connor, "COINHOARDER: Tracking a Ukrainian bitcoin phishing ring DNS style," in *Proc. APWG Symp. Electron. Crime Res. (eCrime)*, San Diego, CA, USA, May 2018, pp. 1–5.

[8] S. Kethineni and Y. Cao, "The rise in popularity of cryptocurrency and associated criminal activity," *Int. Criminal Justice Rev.*, vol. 30, no. 3, pp. 325–344, Sep. 2020, doi: 10.1177/1057567119827051.

[9] M. Bartoletti and L. Pompianu, "An analysis of Bitcoin OP_RETURN metadata," in *Proc. Financial Cryptogr. Data Secur. Workshops*, in Lecture Notes in Computer Science, vol. 10323. New York, NY, USA: Springer, 2017, pp. 218–230.

[10] M. Bartoletti, B. Bellomy, and L. Pompianu, "A journey into bitcoin metadata," *J. Grid Comput.*, vol. 17, no. 1, pp. 3–22, Mar. 2019.

[11] E. Badawi and G.-V. Jourdan, "Cryptocurrencies emerging threats and defensive mechanisms: A systematic literature review," *IEEE Access*, vol. 8, pp. 200021–200037, 2020.

[12] A. Higbee, "The role of crypto-currency in cybercrime," *Comput. Fraud Secur.*, vol. 2018, no. 7, pp. 13–15, Jul. 2018.

[13] APWG. *The APWG Ecrime Exchange (ECX)*. Accessed: Oct. 26, 2021. [Online]. Available: https://apwg.org/

[14] T. Moore, J. Han, and R. Clayton, "The postmodern Ponzi scheme: Empirical analysis of high-yield investment programs," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.*, vol. 7397. Berlin, Germany: Springer, 2012, pp. 41–56.

[15] J. Neisius and R. Clayton, "Orchestrated crime: The high yield investment fraud ecosystem," in *Proc. APWG Symp. Electron. Crime Res. (eCrime)*, Sep. 2014, pp. 48–58.

[16] *Ponzi Schemes Using Virtual Currencies*, SEC, Washington, DC, USA, 2013.

[17] M. Vasek and T. Moore, "There's no free lunch, even using Bitcoin: Tracking the popularity and profits of virtual currency scams," in *Financial Cryptography and Data Security* (Lecture Notes in Computer Science), vol. 8975. Berlin, Germany: Springer, 2015, pp. 44–61.

[18] K. Toyoda, T. Ohtsuki, and P. T. Mathiopoulos, "Identification of high yielding investment programs in bitcoin via transactions pattern analysis," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2017, pp. 1–6.

[19] K. Toyoda, T. Ohtsuki, and P. T. Mathiopoulos. (2017). *Dataset for Identification of High Yielding Investment Programs in Bitcoin Via Transactions Pattern Analysis*. [Online]. Available: https://bitbucket.org/kentaroh_toyoda/research-for-hyip-classification-on-bitcoin/src/master/

[20] M. Bartoletti, B. Pes, and S. Serusi, "Data mining for detecting bitcoin Ponzi schemes," in *Proc. Crypto Valley Conf. Blockchain Technol. (CVCBT)*, Jun. 2018, pp. 75–84.

[21] M. Bartoletti, B. Pes, and S. Serusi. (2018). *Dataset for Data Mining for Detecting Bitcoin Ponzi Schemes*. [Online]. Available: https://github.com/bitcoinponzi/BitcoinPonziTool

[22] W. Chen, Z. Zheng, J. Cui, E. Ngai, P. Zheng, and Y. Zhou, "Detecting Ponzi schemes on ethereum: Towards healthier blockchain technology," in *Proc. World Wide Web Conf. World Wide Web (WWW)*, 2018, pp. 1409–1418.

[23] W. Chen, Z. Zheng, J. Cui, E. Ngai, P. Zheng, and Y. Zhou. (2018) *Dataset for Detecting Ponzi Schemes on Ethereum: Towards Healthier Blockchain Technology*. [Online]. Available: http://ibase.site/scamedb/

[24] K. Toyoda, T. Ohtsuki, and P. T. Mathiopoulos, "Time series analysis for bitcoin transactions: The case of Pirate@40's HYIP scheme," in *Proc. IEEE Int. Conf. Data Mining Workshops (ICDMW)*, H. Tong, Z. J. Li, F. Zhu, and J. Yu, Eds., Nov. 2018, pp. 151–155.

[25] M. Vasek and T. Moore, "Analyzing the Bitcoin Ponzi scheme ecosystem," in *Proc. Bitcoin Workshop*, in Lecture Notes in Computer Science, vol. 10958. New York, NY, USA: Springer, 2018, pp. 101–112.

[26] Y. Boshmaf, C. Elvitigala, H. Al Jawaheri, P. Wijesekera, and M. A. Sabah, "Investigating MMM Ponzi scheme on bitcoin," in *Proc. 15th ACM Asia Conf. Comput. Commun. Secur.*, Oct. 2020, pp. 519–530.

[27] W. Chen, Z. Zheng, E. C.-H. Ngai, P. Zheng, and Y. Zhou, "Exploiting blockchain data to detect smart Ponzi schemes on ethereum," *IEEE Access*, vol. 7, pp. 37575–37586, 2019.

[28] E. Jung, M. L. Tilly, A. Gehani, and Y. Ge, "Data mining-based ethereum fraud detection," in *Proc. IEEE Int. Conf. Blockchain (Blockchain)*, Jul. 2019, pp. 266–273.

[29] C. F. Torres, M. Steichen, and R. State, "The art of the scam: Demystifying honeypots in Ethereum smart contracts," in *Proc. 28th USENIX Secur. Symp.*, N. Heninger and P. Traynor, Eds. Berkeley, CA, USA: USENIX Association, 2019, pp. 1591–1607.

[30] K. Toyoda, P. Takis Mathiopoulos, and T. Ohtsuki, "A novel methodology for HYIP operators' bitcoin addresses identification," *IEEE Access*, vol. 7, pp. 74835–74848, 2019.

[31] M. Bartoletti, S. Carta, T. Cimoli, and R. Saia, "Dissecting Ponzi schemes on Ethereum: Identification, analysis, and impact," *Future Gener. Comput. Syst.*, vol. 102, pp. 259–277, Jan. 2020.

[32] M. Bartoletti, S. Carta, T. Cimoli, and R. Saia. (2020). *Dataset for Dissecting Ponzi Schemes on Ethereum*. [Online]. Available: https://github.com/blockchain-unica/ethereum-ponzi

[33] F. Reid and M. Harrigan, "An analysis of anonymity in the Bitcoin system," in *Security and Privacy in Social Networks*. New York, NY, USA: Springer, 2013, pp. 197–223.

[34] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage, "A fistful of bitcoins: Characterizing payments among men with no names," in *Proc. Conf. Internet Meas. Conf.*, Oct. 2013, pp. 127–140.

[35] M. Bartoletti and L. Pompianu, "An empirical analysis of smart contracts: Platforms, applications, and design patterns," in *Financial Cryptography and Data Security Workshops* (Lecture Notes in Computer Science), vol. 10323. New York, NY, USA: Springer, 2017, pp. 494–509.

[36] K. Liao, Z. Zhao, A. Doupé, and G.-J. Ahn, "Behind closed doors: Measurement and analysis of CryptoLocker ransoms in bitcoin," in *Proc. APWG Symp. Electron. Crime Res. (eCrime)*, Jun. 2016, pp. 1–13.

[37] M. Conti, A. Gangwal, and S. Ruj, "On the economic significance of ransomware campaigns: A Bitcoin transactions perspective," *Comput. Secur.*, vol. 79, pp. 162–189, Nov. 2018.

[38] D. Y. Huang, M. M. Aliapoulios, V. G. Li, L. Invernizzi, E. Bursztein, K. McRoberts, J. Levin, K. Levchenko, A. C. Snoeren, and D. McCoy, "Tracking ransomware end-to-end," in *Proc. IEEE Symp. Secur. Privacy (SP)*. New York, NY, USA: IEEE Computer Society, May 2018, pp. 618–631.

[39] M. Paquet-Clouston, B. Haslhofer, and B. Dupont, "Ransomware payments in the bitcoin ecosystem," *J. Cybersecurity*, vol. 5, no. 1, p. tyz003, Jan. 2019.

[40] N. Kshetri and J. Voas, "Do crypto-currencies fuel ransomware?" *IT Prof.*, vol. 19, no. 5, pp. 11–15, 2017.

[41] *Bitcoin Mining Scams to Avoid*. Accessed: Oct. 26, 2021. [Online]. Available: https://braiins.com/blog/bitcoin-mining-scams#:~:text=Slush%20Pool%20is%20the%20oldest,require%20payments%20from%20our%20users

[42] T. Moore and N. Christin, "Beware the middleman: Empirical analysis of Bitcoin-exchange risk," in *Financial Cryptography and Data Security* (Lecture Notes in Computer Science), vol. 7859. Berlin, Germany: Springer, 2013, pp. 25–33.

[43] A. A. Andryukhin, "Phishing attacks and preventions in blockchain based projects," in *Proc. Int. Conf. Eng. Technol. Comput. Sci. (EnT)*, Mar. 2019, pp. 15–19.

[44] W. Chen, X. Guo, Z. Chen, Z. Zheng, and Y. Lu, "Phishing scam detection on Ethereum: Towards financial security for blockchain ecosystem," in *Proc. 29th Int. Joint Conf. Artif. Intell.*, Jul. 2020, pp. 4506–4512.

[45] E. Badawi, G.-V. Jourdan, G. Bochmann, and I.-V. Onut, "An automatic detection and analysis of the bitcoin generator scam," in *Proc. IEEE Eur. Symp. Secur. Privacy Workshops (EuroS PW)*, Sep. 2020, pp. 407–416.

[46] J. Rauchberger, S. Schrittwieser, T. Dam, R. Luh, D. Buhov, G. Pötzelsberger, and H. Kim, "The other side of the coin: A framework for detecting and analyzing web-based cryptocurrency mining campaigns," in *Proc. 13th Int. Conf. Availability, Rel. Secur.*, Aug. 2018, pp. 18:1–18:10.

[47] H. L. J. Bijmans, T. M. Booij, and C. Doerr, "Inadvertently making cyber criminals rich: A comprehensive study of cryptojacking campaigns at Internet scale," in *Proc. USENIX Secur. Symp.* Berkeley, CA, USA: USENIX Association, 2019, pp. 1627–1644.

[48] J. Wu, J. Liu, W. Chen, H. Huang, Z. Zheng, and Y. Zhang, "Detecting mixing services via mining bitcoin transaction network with hybrid motifs," *IEEE Trans. Syst., Man, Cybern., Syst.*, early access, Jan. 21, 2021, doi: 10.1109/TSMC.2021.3049278.

[49] *FBI's Advance-Fee Scam Definition*. Accessed: Oct. 26, 2021. [Online]. Available: https://www.fbi.gov/scams-and-safety/common-scams-and-crimes/advance-fee-schemes

[50] R. Phillips and H. Wilder, "Tracing cryptocurrency scams: Clustering replicated advance-fee and phishing websites," in *Proc. IEEE Int. Conf. Blockchain Cryptocurrency*, May 2020, pp. 1–8.

[51] M. Paquet-Clouston, M. Romiti, B. Haslhofer, and T. Charvat, "Spams meet cryptocurrencies: Sextortion in the bitcoin ecosystem," in *Proc. 1st ACM Conf. Adv. Financial Technol.*, Oct. 2019, pp. 76–88.

[52] F. Oggier, A. Datta, and S. Phetsouvanh, "An ego network analysis of sextortionists," *Social Netw. Anal. Mining*, vol. 10, no. 1, p. 44, Dec. 2020.

[53] S. C. Baum, "Cryptocurrency fraud: A look into the frontier of fraud," Honors theses, School Accountancy, Georgia Southern Univ., Statesboro, GA, USA, 2018.

[54] *Sec's Fake Ico Parody*. Accessed: Oct. 26, 2021. [Online]. Available: https://www.howeycoins.com/index.html

[55] T. Conlon and R. McGee, "ICO fraud and regulation," in *Batten-Corbet-Lucey Handbooks in Alternative Investments*. Berlin, Germany: De Gruyter, 2021.

[56] M. Möser, R. Böhme, and D. Breuker, "An inquiry into money laundering tools in the bitcoin ecosystem," in *Proc. APWG eCrime Researchers Summit*, Sep. 2013, pp. 1–14.

[57] C. Brenig, R. Accorsi, and G. Müller, "Economic analysis of cryptocurrency backed money laundering," in *Proc. Eur. Conf. Inf. Syst. (ECIS)*, 2015, pp. 1–19.

[58] Y. J. Fanusie and T. Robinson, "Bitcoin laundering: An analysis of illicit flows into digital currency services," Center Sanctions Illicit Finance, Elliptic, London, U.K., Tech. Rep., 2018. Accessed: Oct. 26, 2021.

[59] Y. Hu, S. Seneviratne, K. Thilakarathna, K. Fukuda, and A. Seneviratne, "Characterizing and detecting money laundering activities on the Bitcoin network," *CoRR*, vol. abs/1912.12060, pp. 1–17, Dec. 2019.

[60] M. Weber, G. Domeniconi, J. Chen, D. K. I. Weidele, C. Bellei, T. Robinson, and C. E. Leiserson, "Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics," *CoRR*, vol. abs/1908.02591, pp. 1–7, Jul. 2019.

[61] J. Lorenz, M. I. Silva, D. Aparício, J. T. Ascens ao, and P. Bizarro, "Machine learning methods to detect money laundering in the Bitcoin blockchain in the presence of label scarcity," in *Proc. Int. Conf. AI Finance*, 2020, pp. 23:1–23:8.

[62] (2019). *Elliptic*. [Online]. Available: https://www.kaggle.com/ellipticco/elliptic-data-set

[63] S. Godbole and S. Sarawagi, "Discriminative methods for multi-labeled classification," in *Advances in Knowledge Discovery and Data Mining*, H. Dai, R. Srikant, and C. Zhang, Eds. Berlin, Germany: Springer, 2004, pp. 22–30.

[64] M. Bartoletti, T. Cimoli, L. Pompianu, and S. Serusi, "Blockchain for social good: A quantitative analysis," in *Proc. 4th EAI Int. Conf. Smart Objects Technol. Social Good (Goodtechs)*, 2018, pp. 37–42.

[65] A. S. Podda and L. Pompianu, "An overview of blockchain-based systems and smart contracts for digital coupons," in *Proc. IEEE/ACM 42nd Int. Conf. Softw. Eng. Workshops*, Jun. 2020, pp. 770–778.

[66] M. Bartoletti, S. Lande, L. Pompianu, and A. Bracciali, "A general framework for blockchain analytics," in *Proc. 1st Workshop Scalable Resilient Infrastructures Distrib. Ledgers*, Dec. 2017, pp. 7:1–7:6.

MASSIMO BARTOLETTI is currently an Associate Professor with the Department of Mathematics and Computer Science, University of Cagliari. He is also the Co-Founder and the Co-Director of the Trustworthy Computational Societies Research Group, University of Cagliari (http://tcs.unica.it), and the Founder of the Blockchain@Unica Laboratory (http://blockchain.unica.it), one of the first academic research groups on these topics in Italy. He is also the Director of the node of the Cyber Security National Laboratory, University of Cagliari, and a Core Member of the Italian Working Group in Distributed Ledger Technologies (http://dltgroup.dmi.unipg.it). His research activity concerns, in general, the development of tools and techniques for the specification, analysis and verification of properties of software and systems, both from a foundational viewpoint and from an applicative one. The specific problem addressed by his research is how to guarantee secure interactions among mutually distrusting participants, by the means of behavioural contracts. This problem has been tackled by developing various models for contracts (such as blockchain-based smart contracts, session types, logics, event structures, and Petri nets), and by realising tools to support programmers (such as domain-specific languages, static analyses, and middleware for contract-oriented interactions). He has published over 90 research papers on refereed journals and international refereed conferences and workshops.

Prof. Bartoletti served as the chair, a program committee member, or an organizer for several international conferences and workshops, and an external reviewer for several journals and international conferences.

ANDREA LODDO received the B.Sc., M.Sc., and Ph.D. degrees from the University of Cagliari, in 2012, 2014, and 2019, respectively. His PhD thesis faced blood cells image analysis and classification issues to create new tools for automatic diagnosis as a support to medical analysis. He is currently an Assistant Professor with the Department of Mathematics and Computer Science, University of Cagliari. He is the author of 20 scientific manuscripts in peer-reviewed journals and international conference proceedings related to this task. His research interests include computer vision, biomedical image analysis, pattern recognition, and machine learning. Currently, he is pursuing a research activity for crypto scams, human activity recognition, and biomedical image analysis for diagnosis support systems.

LIVIO POMPIANU received the Ph.D. degree in mathematics and informatics from the University of Cagliari, in 2018. He is currently a Postdoctoral Researcher with the Department of Mathematics and Computer Science, University of Cagliari. He is the coauthor of several scientific publications. His research is currently focused on information security, blockchain analytics, smart contracts, and cryptocurrency scams.

STEFANO LANDE received the B.Sc. and M.Sc. degrees (Hons.) in computer science, and the Ph.D. degree in mathematics and computer science from the University of Cagliari.

SERGIO SERUSI received the B.Sc. and M.Sc. degrees (Hons.) in computer science, and the Ph.D. degree in mathematics and computer science from the University of Cagliari.

● ● ●