DOI:10.1145/3299768

Tracing the tangled web of unsolicited and undesired email and possible strategies for its demise.

BY EMILIO FERRARA

The History of Digital Spam

SPAM! THAT'S WHAT Lorrie Faith Cranor and Brian LaMacchia exclaimed in the title of a popular call-to-action article that appeared 20 years ago in *Communications*.¹⁰ And yet, despite the tremendous efforts of the research community over the last two decades to mitigate this problem, the sense of urgency remains unchanged, as emerging technologies have brought new dangerous forms of digital spam under the spotlight. Furthermore, when spam is carried out with the intent to deceive or influence at scale, it can

alter the very fabric of society and our behavior. In this article, I will briefly review the history of digital spam: starting from its quintessential incarnation, spam emails, to modern-days forms of spam affecting the Web and social media, the survey will close by depicting future risks associated with spam and abuse of new technologies, including artificial intelligence (AI), for example, digital humans. After providing a taxonomy of spam, and its most popular applications emerged throughout the last two decades, I will review technological and regulatory approaches proposed in the literature, and suggest some possible solutions to tackle this ubiquitous digital epidemic moving forward.

An omni-comprehensive, universally acknowledged definition of digital spam is hard to formalize. Laws and regulation attempted to define particular forms of spam, for example, email (see 2003's Controlling the Assault of Non-Solicited Pornography and Marketing Act.) However, nowadays, spam occurs in a variety of forms, and across different techno-social systems. Each domain may warrant a slight different definition that suits what spam is in that precise context: some features of spam in a domain, for example, volume in mass spam campaigns, may not apply to

» key insights

- Throughout the Internet's history, digital spam has pervaded all techno-social platforms and it is constantly evolving. This article provides a taxonomy of digital spam, from its inception to current spam techniques.
- Since the email spam epidemic of the early 1990s, new forms of spam have emerged, including search engine spam, fake reviews, spam bots, and false news. In its latest incarnation, spam threatens to pollute AI systems making them biased and ultimately dangerous for our society.
- By illustrating some of the risks posed by digital spam in all its forms, including AI spam, we provide policy recommendations and technical insights to tackle old and new forms of spam.

Figure 1. Timeline of the major milestones in the history of spam, from its inception to modern days.

The History of Digital Spam

1898–2018

others, for example, carefully targeted phishing operations.

In an attempt to propose a general taxonomy, I here define digital spam as the attempt to abuse of, or manipulate, a techno-social system by producing and injecting unsolicited, and/or undesired content aimed at

steering the behavior of humans or the system itself, at the direct or indirect, immediate or long-term advantage of the spammer(s).

This broad definition will allow me to track, in an inclusive manner, the evolution of digital spam across its most popular applications, starting

from spam emails to modern-days spam. For each highlighted application domain, I will dive deep to understand the nuances of different digital spam strategies, including their intents and catalysts and, from a technical standpoint, how they are carried out and how they can be detected.

Wikipedia provides an extensive list of domains of application:

“While the most widely recognized form of spam is email spam, the term is applied to similar abuses in other media: instant messaging spam, Usenet newsgroup spam, Web search engine spam, spam in blogs, wiki spam, online classified ads spam, mobile phone messaging spam, Internet forum spam, junk fax transmissions, social spam, spam mobile apps, television advertising and file sharing spam.” (<https://en.wikipedia.org/wiki/Spamming>)

The accompanying table summarizes a few examples of types of spam and relative context, including whereas there exist machine learning solutions (ML) to each problem. Email is known to be historically the first example of digital spam (see Figure 1) and remains uncontested in scale and pervasiveness with billions of spam emails generated every day.¹⁰ In the late 1990s, spam landed on instant messaging (IM) platforms (SPIM) starting from AIM (AOL Instant Messenger) and evolving through modern-days IM systems such as WhatsApp, Facebook Messenger, and WeChat. A widespread form of spam that emerged in the same period was Web search engine manipulation: content spam and link farms allowed spammers to boost the position of a target Website in the search result rankings of popular search engines, by gaming algorithms like PageRank and the like. With the success of the social Web,²² in the early 2000s we witnessed the rise of many new forms of spam, including Wiki spam (injecting spam links into Wikipedia pages¹), opinion and review spam (promoting or smearing products by generating fake online reviews²⁷), and mobile messaging spam (SMS and text messages sent directly to mobile devices³). Ultimately, in the last decade, with the increasing pervasiveness of online social networks and the sig-

Detecting Spam Email

Email spam detection is an arms race between attackers (spammers) and defenders (service providers). Two decades of research in the data mining and machine learning communities produced troves of techniques to tackle this problem. Some milestones include:

SMTP solutions. SMTP is the protocol at the foundation of the email exchange infrastructure. Blacklists were introduced to keep track of spam propagators.⁷ Mail servers can consult blacklisting services to determine whether to route emails to their destination. A softer version of blacklisting is greylisting. Greylists keep track of triplets of IP addresses (sender, receiver, SMTP host) involved into an email exchange. The first time a triplet involving a dubious SMTP host appears, the exchange is denied, but the triplet is stored to authorize future exchanges. This is based on the rationale that spammers rarely retry sending spam through the same relay, and was proven effective in reducing early spam circulation.⁷

Another approach is keyword-based filtering: whenever the subject or the body of an email contains flagged terms (belonging to a keyword list), the SMTP service provider would not route it to its intended recipient, and flag the sending offender—multiple offenses would lead to permanent bans. Other strategies like DomainKeys Identified Mail (DKIM) and digital signatures are authentication methods designed to detect email spoofing and assess email provenance.

Supervised learning. In their seminal work, Drucker et al.¹³ proposed one of the first machine learning systems for spam detection, based on support vector machines (then the state of the art in terms of supervised learning). The success of supervised learning over traditional keyword-based filters demonstrated by Drucker et al.¹³ motivated the first wave of machine learning research in email spam detection. Shortly after, Androutsopoulos et al.⁴ showed the power of naive Bayesian anti-spam filtering: Bayesian systems yielded state-of-the-art spam detection performance for many years. The advent of more sophisticated learning models, like boosting trees, set the accuracy bar higher but paradigm shifts lagged for nearly a decade.

Hybrid neural systems. More recently, Wu³⁷ proposed behavior-based spam detection using combinations of simple association rules and neural networks. Given their ability to naturally handle visual information, neural network methods to detect spam were extended to multimedia content. For example, Wu et al.³⁸ and Fumera et al.¹⁷ proposed methods exploiting visual cues to detect spam content injected in images embedded into emails.

Dedicated hardware. Networking companies are developing anti-spam appliances. Dedicated hardware can detect various types of spam, including phishing, malware, and ransomware, guaranteeing high efficiency and accuracy. For example, Cisco advertises that their Email Security Appliance (ESA) detects over 99.9% of incoming spam email with lower than one in a million false positive rate.

Examples of types of spam and relative statistics.

Spam Type	Start	Today's Volume	ML	Ref
Email	1978	Billions x day	✓	10
Instant Messaging	1997	Millions x day	✓	20
Search Engine	1998	Unknown	✓	31
Wiki	2001	Thousands x day	—	1
Opinion and Reviews	2005	Millions across platforms	✓	11
Mobile Messaging	2007	Millions x day	✓	3
Social Bots	2010	Millions across platforms	✓	16
False News	2016	Thousands across websites	—	36
Multimedia	2018	Unknown	—	25

nificant advancements in AI, new forms of spam involve social bots (accounts operated by software to interact at scale with social Web users¹⁶), false news websites (to deliberately spread disinformation³⁶), and multi-media spam based on AI.²⁵

In the following, I will focus on three of these domains: email spam, Web spam (specifically, opinion spam and fake reviews), and social spam (with a focus on social bots). Furthermore, I will highlight the existence of a new form of spam that I will call AI spam. I will provide examples of spam in this new domain, and lay out the risks associated with it and possible mitigation strategies.

Flooded By Junk Email

The 1998 article by Cranor and LaMacchia¹⁰ in *Communications*, characterized the problem of junk email messages, or email spam, as one of the earliest forms of digital spam.

Email spam has mainly two purposes, namely advertising (for example, promoting products, services, or contents), and fraud (for example, attempting to perpetrate scams, or phishing). Neither ideas were particularly new or unique to the digital realm: advertisement based on unsolicited content delivered by traditional post mail (and, later, phone calls, including more recently the so-called “robo-calls”) has been around for nearly a century. As for scams, the first reports of the popular advance-fee scam (in modern days known as 419 scam, a.k.a. the Nigerian Prince scam), called the Spanish Prisoner scam were circulating in the late 1800s.^a

The first reported case of digital spam occurred in 1978 and was attributed to Digital Equipment Corporation, who announced their new computer system to over 400 subscribers of ARPANET, the precursor network of modern Internet (see Figure 1). The first mass email campaign occurred in 1994, known as the USENET green card lottery spam: the law firm of Canter & Siegel advertised their immigration-related legal services simultaneously to over 6,000 USENET newsgroups. This event con-

Email spam has mainly two purposes: advertising and fraud.

tributed to popularizing the term spam. Both the ARPANET and USENET cases brought serious consequences to their perpetrators as they were seen as egregious violations of common code of conduct in the early days of the Internet (for example, Canter & Siegel ran out of business and Canter was disbarred by the Arizona Bar Association.) However, things were bound to change as the Internet became an increasingly more pervasive technology in our society.

Email spam: Risks and challenges.

The use of the Internet for distributing unsolicited messages provides unparalleled scalability, and unprecedented reach, at a cost that is infinitesimal compared to what it would take to accomplish the same results via traditional means.¹⁰ These three conditions created the ideal conjecture of economical incentives that made email spam so pervasive.

In contrast to old-school post mail spam, digital email spam introduced a number of unique challenges:¹⁰ If left unfiltered, spam emails can easily outnumber legitimate ones, overwhelming the recipients and thus rendering the email experience from unpleasant to unusable; email spam often contains explicit content that can hurt the sensibility of the recipients—depending upon the sender/recipient country’s laws, perpetrating this form of spam could constitute a criminal offense;^b by embedding HTML or JavaScript code into spam emails, the spammers can emulate the look and feel of legitimate emails, tricking the recipients and eliciting unsuspecting behaviors, thus enacting scams or enabling phishing attacks;²³ finally, mass spam operations pose a burden on Internet service providers (ISPs), which have to process and route unnecessary, and often large, amounts of digital junk information to millions of recipients—for the larger spam campaigns, even more.

The Internet was originally designed by and for tech-savvy users: spammers quickly developed ways to take advantage of the unsophisticated ones. Phishing is the practice of using

a See *New York Times*, Mar. 20, 1898; <https://nyti.ms/2DD6oIn>

b For example, see the U.S. Federal Law on Obscenity; <https://bit.ly/2wfpDgt>

deception and social engineering strategies by which attackers manage to trick victims by disguising themselves as a trusted entity.^{9,23} The end goal of phishing attacks is duping the victims into revealing sensitive information for identity theft, or extorting funds via ransomware or credit card frauds. Email has been by far and large the most common vector of phishing attacks. In 2006, Indiana University carried out a study to quantify the effectiveness of phishing email messages.²³ The researchers demonstrated that a malicious attacker impersonating the university would have a 16% success rate in obtaining the users' credentials when the phishing email came from an unknown sender; however, success rate arose to 72% when the email came from an attacker impersonating a friend of the victim.

Fighting email spam. Over the course of the last two decades, solutions to the problem of email spam revolved around implementing new regulatory policies, increasingly sophisticated technical hurdles, and combinations of the two.¹⁰ Regarding the former, in the context of the U.S. or the European Union (EU), policies that regulate access to personal information (including email addresses), such as the EU's General Data Protection Regulation (GDPR) enacted in 2018, hinder the ability of bulk mailers based in EU countries to effectively carry out mass email spam operations without risks and possibly serious consequences. However, it has become increasingly more obvious that solutions based exclusively on regulatory affairs are ineffective: spam operations can move to countries with less restrictive Internet regulations. However, regulatory approaches in conjunction with technical solutions have brought significant progress in the fight against email spam.

From a technical standpoint, two decades of research advancements led to sophisticated techniques that strongly mitigate the amount of spam email ending up in the intended recipients' inboxes. A number of review papers have been published that surveyed data mining and machine learning approaches to detect and filter out email spam,⁷ some with a specific focus on scams and phishing spam.²¹

From a technical standpoint, two decades of research advancements led to sophisticated techniques that strongly mitigate the amount of spam email ending up in the intended recipients' inboxes.

In the sidebar “Detecting Spam Email,” I summarize some of the technical milestones accomplished in the quest to identify spam emails. Unfortunately, I suspect that much of the state-of-the-art research on spam detection lies behind close curtains, mainly for three reasons: First, large email-related service providers, such as Google (Gmail), Microsoft (Outlook, Hotmail), Cisco (IronPort, Email Security Appliance—ESA) devote(d) massive R&D investments to develop machine learning methods to automatically filter out spam in the platforms they operate (Google, Microsoft, among others) or protect (Cisco); the companies are thus often incentivized to use patented and close-sourced solutions to maintain their competitive advantage. Secondly, related to the former point, fighting email spam is a continuous arms-race: revealing one's spam filtering technology gives out information that can be exploited by the spammers to create more sophisticated campaigns that can effectively and systematically escape detection, thus calling for more secrecy. Finally, the accuracy of email spam detection systems deployed by these large service providers has been approaching nearly perfect detection: a diminishing return mechanism comes into play where additional efforts to further refine detection algorithms may not warrant the costs of developing increasingly more sophisticated techniques fueling complex spam detection systems; this makes established approaches even more valuable and trusted, thus motivating the secrecy of their functioning.

Web 2.0 or Spam 2.0?

The new millennium brought us the Social Web, or Web 2.0, a paradigm shift with an emphasis on user-generated content and on the participatory, interactive nature of the Web experience.²² From knowledge production (Wikipedia) to personalized news (social media) and social groups (online social networks), from blogs to image and video sharing sites, from collaborative tagging to social e-commerce, this wealth of new opportunities brought us as many

new forms of spam, commonly referred to as social spam.

Differently from spam emails, where spam can only be conveyed in one form (such as email), social spam can appear in multiple forms and modi operandi. Social spam can be in the form of textual content (for example, a secretly sponsored post on social media), or multimedia (for example, a manufactured photo on 4chan); social spam can aim at pointing users to unreliable resources, for example, URLs to unverified information or false news websites;³⁶ social spam can aim at altering the popularity of digital entities, for example, by manipulating user votes (upvotes on Reddit posts, retweets on Twitter), and even that of physical products, for example, by posting fake online reviews (say, for example, about a product on an e-commerce website).

Spammy opinions. In the early 2000s (see Figure 1), the growing popularity of e-commerce websites like Amazon and Alibaba motivated the emergence of opinion spam (a.k.a. review spam).^{24,27}

According to Liu,²⁷ there are three types of spam reviews: fake reviews, reviews about brands only, and non-reviews. The first type of spam, fake reviews, consists of posting untruthful, or deceptive reviews on online e-commerce platforms, in an attempt to manipulate the public perception (in a positive or negative manner) of specific products or services presented on the affected platform(s). Fake positive reviews can be used to enhance the popularity and positive perception of the product(s) or service(s) the spammer intends to promote, while fake negative reviews can contribute to smear the spammer's competitor(s) and their products/services. Opinion spam of the second type, reviews about brands only, pertains comments on the manufacturer/brand of a product but not on the product itself—albeit genuine, according to Liu²⁷ they are considered spam because they are not targeted at specific products and are often biased. Finally, spam reviews of the third type, non-reviews, are technically not opinion spam as they do not provide any opinion, they only contain generic, unrelated content

(for example, advertisement, or questions, rather than reviews, about a product). Fake reviews are, by far and large, the most common type of opinion spam, and the one that has received more attention in the research community.²⁷ Furthermore, Jindal and Liu²⁴ showed that spam of the second and third type is simple to detect and address.

Unsurprisingly, the practice of opinion spam, and in particular fake reviews, is widely considered as unfair and deceptive, and as such it has been subject of extensive legal scrutiny and court battles. If left unchecked, opinion spam can poison a platform and negatively affect both customers and platform providers (including incurring in financial losses for both parties, as customers may be tricked into purchasing undesirable items and grow frustrated against the platform), at the sole advantage of the spammer (or the entity they represent)—as such, depending on the country's laws, opinion spam may qualify as a form of digital fraud.

Detecting fake reviews is complex for a variety of reasons: for example, spam reviews can be posted by fake or real user accounts. Furthermore, fakes reviews can be posted by individual users or even groups of users.^{27,30} Spammers can deliberately use fake accounts on e-commerce platforms, created only with the scope of posting fake reviews. Fortunately, fake accounts on e-commerce platforms are generally easy to detect, as they engage in intense reviewing activity without any product purchases. An alternative and more complex scenario occurs when fake reviews are posted by real users. This tends to occur under two very different circumstances: compromised accounts (that is, accounts originally owned by legitimate users that have been hacked and sold to spammers) are frequently re-purposed and utilized in opinion spam campaigns;¹¹ and fake review markets became very popular where real users collude in exchange for direct payments to write untruthful reviews for example, without actually purchasing or trying a given product or service. To complicate this matter, researchers showed that fake personas, for example, Facebook profiles,

can be created and associated with such spam accounts.¹⁸ During the late 2000s, many online fake-review markets emerged, whose legality was battled in court by e-commerce giants. Action on both legal and technical fronts has helped mitigating the problem of opinion spam.

From a technical standpoint, a variety of techniques have been proposed to detect review spam. Liu²⁷ identified three main approaches, namely supervised, unsupervised, and group spam detection. In supervised spam detection, the problem of separating fake from genuine (non-fake) reviews is formulated as a classification problem. Jindal and Liu²⁴ pointed out that the main challenge of this task is to work around the shortage of labeled training data. To address this problem, the authors exploited the fact that spammers, to minimize their work, often produce (near-)duplicate reviews, that can be used as examples of fake reviews. Feature engineering and analysis was key to build informative features of genuine and fake reviews, enriched by features of the reviewing users and the reviewed products. Models based on logistic regression have been proven successful in detecting untruthful opinions in large corpora of Amazon reviews.²⁴ Detection algorithms based on support vector machines or naive Bayes models generally perform well (above 98% accuracy) and scale to production systems.²⁹ These pipelines are often enhanced by human-in-the-loop strategies, where annotators recruited through Amazon Mechanical Turk (or similar crowd-sourcing services) manually label subsets of reviews to separate genuine from fake ones, to feed online learning algorithms so to constantly adapt to new strategies and spam techniques.^{11,27}

Unsupervised spam detection was used both to detect spammers as well as for detecting fake reviews. Liu²⁷ reported on methods based on detecting anomalous behavioral patterns typical of spammers. Models of spam behaviors include targeting products, targeting groups (of products or brands), general and early rating deviations.²⁷ Methods based on association rules can capture atypical behaviors of reviewers, detecting anomalies

in reviewers' confidence, divergence from average product scores, entropy (diversity or homogeneity) of attributed scores, or temporal dynamics.³⁹ For what concerns the unsupervised detection of fake reviews, linguistic analysis was proved useful to identify stylistic features of fake reviews, for example, language markers that are over- or underrepresented in fake reviews. Opinion spam to promote products, for example, exhibits on average three times fewer mentions of social words, negative sentiment, and long words (> six letters) than genuine reviews, while containing twice more positive terms and references to self than formal texts.¹¹

Concluding, group spam detection aims at identifying signatures of collusion among spammers.³⁰ Collective behaviors such as spammers' coordination can emerge by using combinations of frequent pattern mining and group anomaly ranking. In the first stage, the algorithm proposed by Mukherjee et al.³⁰ identifies groups of reviewers who all have reviewed a same set of products—such groups are flagged as potentially suspicious. Then, anomaly scores for individual and group behaviors are computed and aggregated, accounting for indicators that measure the group burstiness (that is, writing reviews in short timespan), group reviews similarity, and so on. Groups are finally ranked in terms of their anomaly scores.³⁰

The rise of spam bots. Prior to the early 2000s, most of the spam activity was still coordinated and carried out, at least in significant part, by human

operators: email spam campaigns, Web link farms, and fake reviews, among others, all rely on human intervention and coordination. In other words, these spam operations scale at a (possibly significant) cost. With the rise in popularity of online social network and social media platforms (see Figure 1), new forms of spam started to emerge at scale. One such example is social link farms:¹⁹ similarly to Web link farms, whose goal is to manipulate the perception of popularity of a certain website by artificially creating many pointers (hyperlinks) to it, in social link farming spammers create online personas with many artificial followers. This type of spam operation requires creating thousands (or more) of accounts that will be used to follow a target user in order to boost its apparent influence. Such “disposable accounts” are often referred to as fake followers as their purpose is solely to participate in such link-farming networks. In some platforms, link farming was so pervasive that spammers reportedly controlled millions of fake accounts.¹⁹ Link farming introduced a first level of automation in social media spam, namely the tools to automatically create large swaths of social media accounts.

In the late 2000s, social spam obtained a new potent tool to exploit: bots (short for software robots, a.k.a. social bots). In my 2016 *Communications* article “The Rise of Social Bots,”¹⁶ I noted that “bots have been around since the early days of computers:” examples of bots include chatbots, algo-

rithms designed to hold a conversation with a human, Web bots, to automate the crawling and indexing of the Web, trading bots, to automate stock market transactions, and much more. Although isolated examples exist of such bots being used for nefarious purposes, I am unaware of any reports of systematic abuse carried out by bots in those contexts.

A social bot is a new breed of “computer algorithm that automatically produces content and interacts with humans on the social Web, trying to emulate and possibly alter their behavior.” Since bots can be programmed to carry out arbitrary operations that would otherwise be tedious or time-consuming (thus expensive) for humans, they allowed for scaling spam operations on the social Web to an unprecedented level. Bots, in other words, are the dream spammers have been dreaming of since the early days of the Internet: they allow for personalized, scalable interactions, increasing the cost effectiveness, reach, and plausibility of social spam campaigns, with the added advantage of increased credibility and the ability to escape detection achieved by their human-like disguise. Furthermore, with the democratization and popularization of machine learning and AI technologies, the entry barrier to creating social bots has significantly lowered.¹⁶ Since social bots have been used in a variety of nefarious scenarios (see the sidebar “Social Spam Applications”), from the manipulation of political discussion, to the spread of conspiracy theories and false news, and even by

Figure 2. Video sequence real-time reenactment using AI.³⁴ This proof-of-concept technology could be abused to create AI-fueled multimedia spam.

extremist groups for propaganda and recruitment, the stakes are high in the quest to characterize bot behavior and detect them.^{35,c}

Maybe due to their fascinating morphing and disguising nature, spam bots have attracted the attention of the AI and machine learning research communities: the arms-race between spammers and detection systems yielded technical progress on both the attacker's and the defender's technological fronts. Recent advancements in AI (especially artificial neural networks, or ANNs) fuel bots that can generate human-like natural language and interact with human users in near real time.^{16,35} On the other hand, the cyber-security and machine learning communities came together to develop techniques to detect the signature of artificial activity of bots and social network sybils.^{16,40}

In Ferrara et al.,¹⁶ we fleshed out techniques used to both create spam bots, and detect them. Although the degree of sophistication of such bots, and therefore their functionalities, varies vastly across platforms and application domains, commonalities also emerge. Simple bots can do unsophisticated operations, such as posting content according to a schedule, or interact with others according to pre-determined scripts, whereas complex bots can motivate their reasoning and react to further human scrutiny. Beyond anecdotal evidence, there is no systematic way to survey the state of AI-fueled spam bots and consequently their capabilities—researchers adjust their expectations based on advancements made public in AI technologies (with the assumptions that these will be abused by spammers with the right incentives and technical means), and based on proof-of-concept tools that are often originally created with other non-nefarious purposes in mind (one such example is the so-called DeepFakes, discussed later).

^c It should be noted that bots are not used exclusively for nefarious purposes: for example, some researchers used bots for positive health behavioral interventions.¹⁶ Furthermore, it has been noted the most problematic aspect of nefarious bots is their attempt to deceive and disguise themselves as human users: however, many bots are labeled as such and may provide useful services, like live-news updates.

Social Spam Applications

Political manipulation. In a peer-reviewed study published on Nov. 7, 2016⁶ (the day before the U.S. presidential election), I unveiled a massive-scale spam operation affecting the American political Twitter. With the aid of Botometer, an AI system that leverages over a thousand features to separate bots from humans,³⁵ tens of thousands of bots were identified. By studying the activity signatures of these bots, I noted that they were being retweeted at the same rate than human users, which may have contributed to the spread of political misinformation.³⁶ Since most of these bots aimed at sowing chaos, their presence may have inflamed and further polarized the political conversation, with unknown consequences on the integrity of the democratic process. Since then, dozens of studies corroborated these results; many other studies, before and after mine, showed the perils associated with social spam campaigns in political domains. Most recently, the emerging phenomenon of fake news spreading attracted a lot of attention. Vosoughi et al.³⁶ investigated the role of social media, as well as bots, in the spread of true and false news: the authors showed that humans are more likely to share false stories inspired by fear, disgust, and surprise. This suggests that conditioning and manipulation operations online can affect human behavior.

Public health. Conspiracy and denialism are endemic of social networks. Spam in public health discussions has become commonplace for social media: in a recent study, for example, my team highlighted how bots are used to promote electronic cigarettes as cessation devices with health benefits, a fact not definitively corroborated by science.² The use of bots to carry out anti vaccination campaigns has been the subject of investigation of a DARPA Challenge in 2016.³²

Stock market. Automatic trading algorithms leverage information from social media to predict stock prices. Using bots, spam campaigns have been carried out to give the false impression that certain stocks were spoken positively about on Twitter, successfully tricking trading algorithms into buying them in a pump-and-dump scheme unveiled by the U.S. Securities and Exchange Commission (SEC) in 2015.¹⁵

Data leaks. Social platforms enable the often unwilling disclosure of private user information. A recent study showed that over a third of content shared on Facebook has the default public-visibility privacy settings.²⁸ The amount of content accessible to undesirable users may be even higher when considering privacy settings that allow one's friends to access private information and preferences: Research showed that most users indiscriminately accept friendship connections on Facebook.¹⁸ Spam bots can inject themselves into tightly connected communities, by leveraging the weak-tie structure of online social networks,¹² and obtain private user information on large swaths of users. Phishing is also responsible for data leaks. Attacks based on short-URLs are popular on social media: they can hide the true identity of the spammers and have been proven effective to steal personal data.^{9,19}

In the sidebar “Social Spam Applications,” I highlight some of the domains where bots made the headlines: one such example is the wake to the 2016 U.S. presidential election, during which Twitter and Facebook bots have been used to sow chaos and further polarize the political discussion.⁶ Although it is not always possible for the research community to pinpoint the culprits, the research of my group, among many others, contributed to unveil anomalous communication dynamics that attracted further scrutiny by law enforcement and were ultimately connected to state-sponsored operations (if you wish, a form of social spam aimed at influencing individual behavior). Spam bots operate in other highly controversial conversation domains: in the context of public health, they promote products or spread scientifically unsupported claims;^{2,15} they

have been used to create spam campaigns to manipulate the stock market;¹⁵ finally, bots have also been used to penetrate online social circles to leak personal user information.¹⁸

AI Spam

AI has been advancing at vertiginous speed, revolutionizing many fields including spam. Beyond powering conversational agents such as chatbots, like Siri or Alexa, AI systems can be used, beyond their original scope, to fuel spam operations of different sorts. I will refer to this phenomenon next as spamming *with* AI, hinting to the fact that AI is used as a tool to create new forms of spams. However, given their sophistication, AI systems can themselves be subject of spam attacks. I will refer to this new concept as spamming *into* AI, suggesting that AIs can be manipulated, and even compromised, by spammers (or attackers

Figure 3. Physical-world attacks onto AI visual classifier.¹⁴ Similar techniques could be abused to inject unwanted spam into AI and trigger anomalous behaviors.

in a broader sense) to exhibit anomalous and undesirable behaviors.

Spamming with AI. Advancements in computer vision, augmented and virtual realities are projecting us in an era where the boundary between reality and fiction is increasingly more blurry. Proofs-of-concept of AIs capable to analyze and manipulate video footages, learning patterns of expressions, already exist: Suwajanakorn et al.³³ designed a deep neural network to map any audio into mouth shapes and convincing facial expressions, to impose an arbitrary speech on a video clip of a speaking actor, with results hard to distinguish, to the human eye, from genuine footage. Thies et al.³⁴ showcased a technique for real-time facial reenactment, to convincingly re-render the synthesized target face on top of the corresponding original video stream (see Figure 2). These techniques, and their evolutions,²⁵ have been then exploited to create so-called DeepFakes, face-swaps of celebrities into adult content videos that surfaced on the Internet by the end of 2017. Such techniques have also already been applied to the political domain, creating fictitious video footage re-enacting Obama,^d Trump, and Putin,^e among several world leaders.²⁵ Concerns about the ethical and legal conundrums of these new technologies have been already expressed.⁸

In the future, well-resourced spammers capable of creating AIs pretending to be human may abuse these technologies. Another example: Google recently demonstrated the ability to deploy an AI (Google Duplex) in the real world to act as a virtual assistant, seamlessly interacting with human interlocutors over the phone:^f such technology may likely be repurposed to carry out massive scale spam-call campaigns. Other forms of future spam with AI may use augmented or virtual reality agents, so-called digital humans, to interact with humans in digital and virtual spaces, to promote products/services, and in worse-case scenarios to carry out nefarious campaigns similar to those of today's bots, to manipulate and influence users.

Spamming into AI. AIs based on ANNs are sophisticated systems whose functioning can sometimes be too complex to explain or debug. For such a reason, ANNs can be easy preys of various forms of attacks, including spam, to elicit undesirable, even harmful system's behaviors. An example of spamming into AI can be bias exacerbation: one of the major problems of modern-days AIs (and, in general, of supervised learning approaches based on big data) is that biases learned from training data will propagate into predictions.

The problem of bias,⁵ especially in AI, is under the spotlight and is being tackled by the computing research

community.⁸ One way an AI can be maliciously led to learn biased models is deliberately injecting spam—here intended as unwanted information—into the training data: this may lead the system to learn undesirable patterns and biases, which will affect the AI system's behavior in line with the intentions of the spammers.

An alternative way of spamming into AI is the manipulation of test data. If an attacker has a good understanding of the limits of an AI system, for example, by having access to its training data and thus the ability to learn strength and weakness of the learned models, attacks can be designed to lure the AI into an undesirable state. Figure 3 shows an example of a physical-world attack that affects an AI system's behaviors in anomalous and undesirable ways:¹⁴ in this case, a deep neural network for image classification (which may have been used, for example, to control an autonomous vehicle) is tricked by a “perturbed” stop sign mistakenly interpreted as a speed limit sign—according to the expectation of the attacker. Spam test data may be displayed to a victim AI system to lure it into behaving according to a scripted plot based on weaknesses of the models and/or of its underlying data. The potential applications of such type of spam attacks can be in medical domains (for example, deliberate misreading of scans), autonomous mobility (for example, attacks on the transportation infrastructure or the vehicles), and more. Depending on the pervasiveness of AI-fueled systems in the future, the questions related to spamming into AI may require the immediate attention of the research community.

Recommendations

Four decades have passed since the first case of email spam was reported by 400 ARPANET users (see Figure 1). While some prominent computer scientists (including Bill Gates) thought that spam would quickly be solved and soon remembered as a problem of the past,¹⁰ we have witnessed its evolution in a variety of forms and environments. Spam feeds itself of (economic, political, ideological, among others) incen-

^d See <https://grail.cs.washington.edu/projects/AudioToObama>

^e See <http://niessnerlab.org/projects/thies-2016face.html>

^f <https://bit.ly/2rznYXJ>

^g <https://bit.ly/2lfdtI2>

tives and of new technologies, both of which there is no shortage of, and therefore it is likely to plague our society and our systems for the foreseeable future.

It is therefore the duty of the computing community to enact policies and research programs to keep fighting against the proliferation of current and new forms of spam. I conclude suggesting three maxims that may guide future efforts in this endeavor:

1. *Design technology with abuse in mind.* Evidence seems to suggest that, in the computing world, new powerful technologies are oftentimes abused beyond their original scope. Most modern-days technologies, like the Internet, the Web, email, and social media, have not been designed with built-in protection against attacks or spam. However, we cannot perpetuate a naive view of the world that ignores ill-intentioned attackers: new systems and technologies shall be designed from their inception with abuse in mind.

2. *Don't forget the arms race.* The fight against spam is a constant arms race between attackers and defenders, and as in most adversarial settings, the party with the highest stakes will prevail: since with each new technology comes abuse, researchers shall anticipate the need for countermeasures to avoid being caught unprepared when spammers will abuse their newly designed technologies.

3. *Blockchain technologies.* The ability to carry out massive spam attacks in most systems exists predominantly due to the lack of authentication measures that reliably guarantee the identity of entities and the legitimacy of transactions on the system. The blockchain as a proof-of-work mechanism to authenticate digital personas (including in virtual realities), AIs, and others may prevent several forms of spam and mitigate the scale and impact of others.^h

Spam is here to stay: let's fight it together!

Acknowledgments

The author would like to thank current and former members of the USC Information Sciences Institute's MINDS research group, as well as of the Indiana University's CNetS group, for invaluable research collaborations and discus-

sions on the topics of this work. The author is grateful to his research sponsors including the Air Force Office of Scientific Research (AFOSR), award FA9550-17-1-0327, and the Defense Advanced Research Projects Agency (DARPA), contract W911NF-17-C-0094.

Trademarked products/services mentioned in this article include: WhatsApp, Facebook Messenger, WeChat, Gmail, Microsoft Outlook, Hotmail, Cisco IronPort, Email Security Appliance (ESA), AOL Instant Messenger, Reddit, Twitter, and Google Duplex. **■**

^h It is worth noting that proof-of-work has been proposed to prevent spam email in the past, however its feasibility remains debated, especially in its original non-blockchain-based implementation.²⁶

References

- Adler, B., Alfaro, L.D. and Pye, I. Detecting Wikipedia vandalism using wikitrust. Notebook papers of CLEF 1 (2010), 22–23.
- Allem, J.P., Ferrara, E., Uppu, S.P., Cruz, T.B. and Unger, J.B. E-cigarette surveillance with social media data: social bots, emerging topics, and trends. *JMIR Public Health and Surveillance* 3, 4 (2017).
- Almeida, T.A., Hidalgo, J.M.G. and Yamakami, A. Contributions to the study of SMS spam filtering: new collection and results. In *Proceedings of the 11th ACM Symposium on Document Engineering*. ACM, 2011, 259–262.
- Androustopoulos, I., Koutsias, J., Chandrinis, K.V. and Spyropoulos, C.D. An experimental comparison of naive Bayesian and keyword-based anti-spam filtering with personal e-mail messages. In *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2000, 160–167.
- Baeza-Yates, R. Bias on the Web. *Commun. ACM* 61, 6 (June 2018), 54–61.
- Bessi, A. and Ferrara, E. Social bots distort the 2016 US Presidential election online discussion. *First Monday* 21, 11 (2016).
- Caruana, G. and Li, M. A survey of emerging approaches to spam filtering. *ACM Computing Surveys* 44, 2 (2012), 9.
- Chesney, R. and Citron, D. Deep Fakes: A Looming Crisis for National Security, Democracy and Privacy. *The Lawfare Blog* (2018).
- Chhabra, S., Aggarwal, A., Benevenuto, F. and Kumaraguru, P. Phi.sh/Social: The phishing landscape through short URLs. In *Proceedings of the 8th Annual Collaboration, Electronic Messaging, Anti-Abuse and Spam Conference*. ACM, 2011, 92–101.
- Cranor, L.F. and LaMacchia, B.A. Spam! *Commun. ACM* 41, 8 (Aug. 1998), 74–83.
- Crawford, M., Khoshgoftaar, T.M., Prusa, J.D., Richter, A.D. and Najada, H.A. Survey of review spam detection using machine-learning techniques. *J. Big Data* 2, 1 (2015), 23.
- De Meo, P., Ferrara, E., Fiumara, G. and Provetti, A. On Facebook, most ties are weak. *Commun. ACM* 57, 11 (Nov. 2014), 78–84.
- Drucker, H., Wu, D. and Vapnik, V.N. Support vector machines for spam categorization. *IEEE Trans Neural Networks* 10 (1999).
- Eykholt, K. et al. D. Robust physical-world attacks on deep learning visual classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, 1625–1634.
- Ferrara, E. Manipulation and abuse on social media. *ACM SIGWEB Newsletter* Spring (2015), 4.
- Ferrara, E., Varol, O., Davis, C., Menczer, F. and Flammini, A. The rise of social bots. *Commun. ACM* 59, 7 (July 2016), 96–104.
- Fumera, G., Pillai, I. and Roli, F. Spam filtering based on the analysis of text information embedded into images. *J. Machine Learning Research* 7, (Dec. 2006), 2699–2720.
- Gao, H., Hu, J., Wilson, C., Li, Z., Chen, Y. and Zhao, B.Y. Detecting and characterizing social spam campaigns. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*. ACM, 2010, 35–47.
- Ghosh, S. et al. Understanding and combating link farming in the Twitter social network. In *Proceedings of the 21st International Conference on World Wide Web*. ACM, 2012, 61–70.
- Goodman, J., Cormack, G.V. and Heckerman, D. Spam and the ongoing battle for the inbox. *Commun. ACM* 50, 2 (Feb. 2007), 24–33.
- Gupta, B.B., Tewari, A., Jain, A.K. and Agrawal, D.P. Fighting against phishing attacks: state of the art and future challenges. *Neural Computing and Applications* 28, 12 (2017), 3629–3654.
- Hendler, J., Shadbolt, N., Hall, W., Berners-Lee, T. and Weitzner, D. Web science: An interdisciplinary approach to understanding the Web. *Commun. ACM* 51, 7 (July 2008), 60–69.
- Jagatic, T.N., Johnson, N.A., Jakobsson, M. and Menczer, F. Social phishing. *Commun. ACM* 50, 10 (Oct. 2007), 94–100.
- Jindal, N. and Liu, B. Opinion spam and analysis. In *Proceedings of the 2008 International Conference on Web Search and Data Mining*. ACM, 219–230.
- Kim, H. et al. Deep Video Portraits. arXiv preprint (2018), arXiv:1805.11714.
- Laurie, B. and Clayton, R. Proof-of-work proves not to work; version 0.2. In *Workshop on Economics and Information, Security*, 2004.
- Liu, B. Sentiment analysis and opinion mining. *Synthesis Lectures on Human Language Technologies* 5, 1 (2012), 1–167.
- Liu, Y., Gummadi, K.P., Krishnamurthy, B. and Mislove, A. Analyzing Facebook privacy settings: User expectations vs. reality. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*. ACM, 61–70.
- Mukherjee, A. et al. Spotting opinion spammers using behavioral footprints. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2013, 632–640.
- Mukherjee, A., Liu, B. and Gance, N. Spotting fake reviewer groups in consumer reviews. In *Proceedings of the 21st International Conference on World Wide Web*. ACM, 2012, 191–200.
- Spirin, N. and Han, J. 2012. Survey on Web spam detection: Principles and algorithms. *ACM SIGKDD Explorations Newsletter* 13, 2 (2012), 50–64.
- Subrahmanian, V.S. et al. The DARPA Twitter Bot Challenge. *Computer* 49, 6 (2016), 38–46.
- Suwajanakorn, S., Seitz, S.M. and Kemelmacher-Shlizerman, I. Synthesizing Obama: Learning lip sync from audio. *ACM Trans Graphics* (2017).
- Thies, J., Zollhöfer, M., Stamminger, M., Theobalt, C., and Nießner, M. Face2Face: Real-time face capture and reenactment of RGB videos. In *Proceedings of Computer Vision and Pattern Recognition*. IEEE, 2016.
- Varol, O., Ferrara, E., Davis, C., Menczer, F. and Flammini, A. Online human-bot interactions: Detection, estimation, and characterization. In *Proceedings of International AAAI Conference on Web and Social Media*, 2017.
- Vosoughi, S., Roy, D. and Aral, S. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151.
- Wu, C.H. Behavior-based spam detection using a hybrid method of rule-based techniques and neural networks. *Expert Systems with Applications* 36, 3 (2009), 4321–4330.
- Wu, C.T., Cheng, K.T., Zhu, Q., and Wu, Y.L. Using visual features for anti-spam filtering. In *Proceedings of IEEE International Conference on Image Processing* 3. IEEE, 2005, III–509.
- Xie, S., Wang, G., Lin, S. and Yu, P.S. Review spam detection via temporal pattern discovery. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2012, 823–831.
- Yang, Z., Wilson, C., Wang, X., Gao, T., Zhao, B.Y. and Dai, Y. Uncovering social network Sybils in the wild. *ACM Trans. Knowledge Discovery from Data* 8, 1 (2014), 2.

Emilio Ferrara (emiliofe@usc.edu) is an assistant research professor and associate director of Applied Data Science at the University of Southern California Information Sciences Institute, Marina Del Rey, CA, USA.

Copyright held by author/owner.
Publication rights licensed to ACM.