

Pinning Down Abuse on Google Maps

Danny Yuxing Huang[§], Doug Grundman[‡], Kurt Thomas[‡], Abhishek Kumar[‡], Elie Bursztein[‡],
Kirill Levchenko[§], and Alex C. Snoeren[§]

[§]{dhuang, klevchen, snoeren}@cs.ucsd.edu, [‡]{dgrundman, kurtthomas, abhishekr, elieb}@google.com

[§]University of California, San Diego [‡]Google Inc

ABSTRACT

In this paper, we investigate a new form of blackhat search engine optimization that targets local listing services like Google Maps. Miscreants register abusive business listings in an attempt to siphon search traffic away from legitimate businesses and funnel it to deceptive service industries—such as unaccredited locksmiths—or to traffic-referral scams, often for the restaurant and hotel industry. In order to understand the prevalence and scope of this threat, we obtain access to over a hundred-thousand business listings on Google Maps that were suspended for abuse. We categorize the types of abuse affecting Google Maps; analyze how miscreants circumvented the protections against fraudulent business registration such as postcard mail verification; identify the volume of search queries affected; and ultimately explore how miscreants generated a profit from traffic that necessitates physical proximity to the victim. This physical requirement leads to unique abusive behaviors that are distinct from other online fraud such as pharmaceutical and luxury product scams.

Keywords

abuse; affiliate fraud; local listings; online map

1. INTRODUCTION

Users' online attention is becoming increasingly localized: A recent Google study [1] reports that 4 out of 5 users conduct searches with local intent. A wide variety of local listing services like Apple and Google Maps, Yelp, and Foursquare have emerged to enable users to search for businesses based on physical location. Hence, relevance is no longer sufficient to drive interest; geographical proximity is the coin of the emerging localized-search realm.

In order to bootstrap the process of businesses bridging the physical and digital divide, existing local search services typically allow business owners to create and curate their own listings, often consisting of a company name, address, phone number, and additional metadata. While this crowd sourcing of geospatial information has made millions of legitimate businesses accessible via search, it is

also ripe for abuse. Early forms of attacks included defacement, such as graffiti posted to Google Maps in Pakistan [13]. However, increasing economic incentives are driving an ecosystem of deceptive business practices that exploit localized search, such as illegal locksmiths that extort victims into paying for inferior services [9]. In response, local listing services like Google Maps employ increasingly sophisticated verification mechanisms to try and prevent fraudulent listings from appearing on their services.

In this work, we explore a form of blackhat search engine optimization where miscreants overcome a service's verification steps to register fraudulent localized listings. These listings attempt to siphon organic search traffic away from legitimate businesses and instead funnel it to profit-generating scams. In collaboration with Google, we examine over a hundred-thousand business listings that appeared on Google Maps between June 1, 2014 and September 30, 2015 and were subsequently suspended for abuse. We use this dataset to categorize the types of abuse affecting Google Maps, analyze how miscreants circumvented protections against fraudulent listing registration such as postcard mail verification, identify the volume of search queries that returned abusive listings, and, ultimately, explore how miscreants might have generated a profit from traffic that necessitates physical proximity to the victim. This geographic element distinguishes our work from previous studies of webpage-based blackhat SEO and digital storefronts in the pharmaceutical and luxury product marketplaces [8, 14, 15].

Registering listings on Google Maps requires access to a Google account, a physical address, and a contact phone number in order to satisfy the various verification challenges employed to stem fraud. Despite these requirements, miscreants registered tens of thousands of abusive listings per month during the time period we study, likely spurred in part by short listing lifetimes (a median of 8.6 days between creation and suspension). Abusive business listings that Google was able to detect are concentrated in the United States and India, which combined account for 74% of the addresses of abusive listings. We find at least 40.3% of abusive listings relate to the *on-call* service industry, e.g., locksmiths, plumbers, and electricians. These service providers are typically mobile, and they usually visit customers after being contacted on the phone. In contrast, at least 12.7% of the abusive listings describe *on-premise* businesses, such as hotels and restaurants, where customers visit the service provider.

For fraudulent on-call listings, we find that miscreants primarily acquired access to fresh mailing addresses around the United States by registering post office boxes at UPS stores, in turn re-using the same address to create tens to hundreds of listings. In order to provide a new contact phone number for each listing, miscreants relied on cheap, disposable VoIP numbers provided by Bandwidth.com,

©2017 International World Wide Web Conference Committee (IW3C2), published under Creative Commons CC BY 4.0 License.
WWW 2017, April 3–7, 2017, Perth, Australia.
ACM 978-1-4503-4913-0/17/04.
<http://dx.doi.org/10.1145/3038912.3052590>



Ring Central, Level 3, and others. For on-premise listings, miscreants provided legitimate addresses for restaurants and hotels, but abused the verification process to obtain approval without consent of the business owners. Our findings illustrate the challenge of verifying crowd-sourced locations and ownership where ground truth—even with recent photos of purported storefronts like those available on Google Street Views—is difficult to acquire.

Using these abusive listings, miscreants managed to attract 0.5% of Google Maps’ user impressions during the period of study. Of the user traffic captured by miscreants, some 53.5% of it was forwarded to referral scams for the restaurant and hotel industry, and 3.5% was directed towards deceptive service industries (e.g., unaccredited locksmiths and contractors) operating phone centers to respond to inquiries. We qualitatively assess the organization of each scam and the user harm inflicted. For example, some deceptive on-call services send operatives to a victim’s address in return for exorbitant fees [9]. Due to requirements of physical proximity, these scams are most prevalent in large metropolitan areas like New York, Chicago, Houston, and Los Angeles. In contrast, miscreants operating traffic referral schemes register listings for businesses not yet on Google Maps (or coerce the owners of existing listings) and then forward traffic through affiliate programs to make a profit. While users were likely to ultimately reach the business they intended, we highlight the deceptive practices involved in registering these listings and potential phishing attacks that happened against the business operators.

In summary, we frame our contributions as follows:

- ❖ We present the first systematic analysis of blackhat search engine optimization targeting location-based search.
- ❖ We expose how miscreants circumvented Google’s postcard mail verification, which is similar to those employed by a number of other local listing services.
- ❖ We identify two distinct monetization mechanisms: funneling traffic to deceptive service industries and illegitimate traffic referral portals.
- ❖ We discuss unique constraints that operators and miscreants alike must address in the local-search ecosystem, as opposed to traditional web search.

The remainder of this paper is organized as follows. We start by explaining the life cycle of a typical Maps listing in Section 2 and why Google suspends listings. In Section 3, we describe our dataset and how we transform it to facilitate analysis. We survey the overall landscape of abuse on Maps in Section 4, including per-country and per-category breakdowns of abusive listings. Section 5 focuses on how miscreants created abusive listings, from verifying the listing to connecting with customers. Finally, we quantify the impact of abuse on users in Section 6.

2. BACKGROUND

Before diving into our analysis, we provide an overview of how business owners create or modify Google Maps listings. We then discuss the reasons that Google Maps suspends listings, and consequently, the types of abuse included in our dataset.

2.1 User-generated map listings

Google Maps empowers business owners to create and maintain listings that appear in Google Maps and Search. We outline this process in Figure 1. First, a business owner uses the Google My-Business website to register a new listing [4]. This registration is reflected in a database that is then subject to verification and review, before appearing in the Google Maps database that serves as




Figure 1: Summary of how user-generated content appeared in Google Maps from the My Business and Map Maker front-end websites. The datasets included in our study are highlighted in gray.

a back-end for all location-based searches. This creation path—covered in more detail shortly—is the focus of our study. For completeness, we note that there is a second source of user-generated content in the Google Maps database: wiki-like edits via Google Map Maker that are subject to community voting and review [2]. While the latter path may also contribute to abuse, it is beyond the scope of this paper as abuse requires fundamentally different types of resources (e.g., Sybil accounts for reputation gaming).

Creating a listing: In order for a business owner to create a new listing via Google MyBusiness, they must first control a Google Account. We refer to this account as the *operator*, as it may be the business owner or an authorized third party. To register a new listing, the operator supplies a business name, street address, phone number, an optional website URL, and a business category selected from a dropdown list (e.g., “cafe” or “restaurant”) that is consistent across all countries. If a listing already exists, either from a previous operator or a third-party data source, Google Maps provides a mechanism for a new operator to claim ownership of the listing.

Verifying ownership: Google Maps relies on postcard mail verification to approve all freshly created listings [5]. The process involves Google Maps sending a postcard with a PIN to a new listing’s mailing address. The operator retrieves this code and submits it via a web form to verify their access to the address. The goal of this process is to limit the creation of abusive businesses and ensure the veracity of the data (e.g., address) the operator provided. Google Maps also provides a phone verification option if an operator wishes to claim ownership of an existing listing. This option requires that the said listing must previously exist on Maps but without an operator, and that the phone number associated with the listing came from a source trusted by Google. Phone verification is identical to postcard mail verification, except that Maps delivers the verification code to the operator’s phone number via an automated call. Both verification mechanisms serve as financial and technical hurdles for miscreants, but, as we show in Section 5, neither is insurmountable. Once a listing is verified, it is examined by an automatic or human reviewer before it is published to Maps.

Modifying a listing:

After verification, Google Maps allows a business operator to modify or update a listing to include a new webpage, change photos, or update their self-defined category. In general, such modifications trigger re-verification, except in what Google considers to be low-risk cases. Furthermore, in order to reduce friction on business owners that move within the same city, Google Maps allows operators to update their address without re-verification, so long as the new address is within the same ZIP code.

2.2 Suspending harmful listings

Google Maps periodically scans listings to identify content that violates the site’s Terms of Service around deceptive, misleading, or harmful content [3] and suspends any listings found in viola-

tion of these terms. Upon suspension, a business listing remains in the Google MyBusiness and Google Maps database, but no longer appears in location-based searches and is thus invisible to users.

In our study, we treat any listing that was in a suspended state at any time as *abusive*, while we treat all other currently active listings as legitimate. Creators of abusive listings are henceforth referred to as *abusive operators* or miscreants. This approach mirrors that of previous retroactive studies of abuse in online services [10]. For the purposes of our study, we use the term abusive and suspended interchangeably. We discuss potential biases and limitations related to these labels below.

3. DATASET AND METHODOLOGY

We obtained a snapshot of all business listings registered via Google MyBusiness that appeared on the user-facing Maps service at any time between June 1, 2014 and September 30, 2015. The snapshot was generated on January 7, 2016 and includes all edits post-creation along with whether the listing was *active* or *suspended* as of the snapshot date. As we will show later, this four-month delay provides an ample time window for Google Maps to suspend any abusive listings created at the tail end of September. Our dataset contains over a hundred-thousand suspended listings. Under the conditions of our data sharing agreement, the precise number of active and suspended listings in the snapshot is confidential. For both active and suspended listings, we have the following four categories of data.

(1) Listing metadata: This dataset includes a keyed hash of the operator’s email address, as well as the business listing’s creation timestamp, mailing address, phone number, website, and business type.¹ Using a dropdown menu, an operator can choose from a list of more than 4,000 business types, ranging from generic (e.g. “Restaurant”) to specific (e.g. “Chinese Restaurant”). To facilitate our analysis, we manually cluster similar business types into 32 high-level groups called *business categories*, which, in total, cover 92% of suspended listings in the US, and 84% globally. **Table 1** shows the top 10 business categories along with examples.

(2) Verification method: Whether a listing was verified by mail or by phone. In the case of mail-verified listings, the dataset includes the business address that served as the postcard’s destination. As discussed in Section 2, there are some scenarios where a business owner can change addresses without triggering re-verification. As such, the verification address is not guaranteed to match the listing’s address as displayed on Maps.

(3) Impression count: The number of impressions a listing received from creation up to January 7, 2016. It counts the number of times a listing appears in location-based queries, either on Google Search (where the listing appears as a card) or from searches performed directly on Google Maps. We use this impression count as a metric for understanding the volume of users that encounter listings later determined to be abusive (i.e., spam views), as discussed in Section 6.

(4) Edit history: All modifications made to a listing throughout its history, such as changes in the listing’s website URL, phone number, category, or other content visible to Google Maps users. As discussed in Section 2, such modifications are allowed post-verification, subject to review.

¹Our dataset contains no personally identifiable information. Aside from the hash of the operator’s email address, all information was publicly available on Google Maps until the time of suspension—or remains on Maps, in the case of active listings.




Figure 2: Breakdown of abusive listing creation by month. The abusive listings created in June 2015 account for 13.1% of all abusive listings in our measurement period.

Limitations to our approach: Our study is biased towards abuse caught by the suspension algorithms employed by Google Maps. The main limitation with this approach is that we cannot estimate the number of false negatives, i.e., abusive listings overlooked by Google Maps. (In contrast, we assume there very few false positives as legitimate business owners can appeal suspensions, and had at least four months to do so due to our snapshot methodology.) While in other domains such as email spam or fake accounts it is possible to manually review a sample to estimate the error rate, local listing abuse is far more complicated to verify. For example, if we review Google Street View photos for whether a building exists at the address purported by a freshly created listing, it may be that the photos are out of date. Similarly, if an address refers to a specific suite number, that cannot be verified from street-level photos. Ignoring addresses, if miscreants provide only a phone number for a deceptive locksmith listing, there is no abusive content immediately available for review, unless we manually call the listing’s number and follow through hiring what turns out to be a non-licensed locksmith. These challenges are at the very heart of why local listing abuse is complicated and worth studying; however, it also means we cannot determine whether our analysis uncovers all forms of abuse. Nevertheless, our sample of over a hundred-thousand abusive listings provides one of the first large-scale lenses into how localized search abuse operates.

4. LANDSCAPE OF MAPS ABUSE

We explore how the scale of abusive listings evolved over time, identify the main features of their operation, and expose geographic biases in the locations from which miscreants operated. Our results illustrate that the unique requirements of localized-listing abuse—especially access to physical resources such as mailboxes—yield a distinct abuse strategy compared to email spammers or blackhat SEO for website search engines.

4.1 Volume and duration of abuse

Throughout our analysis period, miscreants registered tens of thousands of new abusive listings each month. **Figure 2** plots the distribution of newly registered abusive listings during our study period. The volume of registration steadily increased until a peak in June of 2015, in which 13.1% of all abusive listings were created. The decline from July onward was a result of Google rolling out a new defense (discussed in more detail in Section 5). If we measure the duration of listings as the time between creation and eventual suspension, we find abusive listings remained active for a median of 8.6 days.

While we cannot disclose the total number of abusive listings registered, it is markedly smaller than the 600,000 accounts scam-

Business Category	Examples	Pop.
Contractors (locksmiths)	Locksmiths	25.7%
Contractors (others)	Plumbers, electricians	14.6%
Food	Restaurants, pizza delivery	7.3%
Hotels	Motels, hotels, bed-and-breakfast	5.4%
Fashion and shopping	Clothing stores, beauty salons	3.8%
Healthcare	Rehab centers, testing services	3.6%
Professionals	Lawyers, consultants, accountants	2.4%
Travel	Limousine, taxi, travel agents	1.9%
Auto	Car repair, towing, dealers	1.7%
Artistic	Photographers, graphic designers	1.5%
Logistics	Movers, packers, shippers	1.5%
Others		30.7%

Table 1: Top-ten business categories associated with abusive listings worldwide. For each category, we include examples of businesses, along with the fraction of abusive listings the category covers, e.g., Locksmiths account for 25.7% of all suspended listings.

mers bulk-registered on Renren and the over 1-million bogus accounts created on Twitter during a similar elapsed time period [10, 16]. We hypothesize this lower rate is an immediate consequence of mail verification, which imposes a higher financial and technical burden compared to phone or email verification. We stress that the abusive listings detailed in this paper actually appeared on Maps before they were suspended. They account for only 15.3% of all abusive registration attempts during the analysis period; the remaining 84.7% of them were suspended even before they reached users and thus are not considered here.

4.2 Abusive business types

We provide a breakdown of the top-ten business categories associated with abusive listings in **Table 1**. For example, we find 25.7% of all abusive listings were categorized as locksmiths at the time they were suspended, followed in popularity by other types of contractors. Combined, the top-ten categories cover 69.3% of all abusive listings. The remaining 30.7% of listings belong to a long tail of business categories. Examples include bail bonds, Internet service providers, real estate agencies, and dating agencies.

Examining the top abusive business types, we can qualitatively divide them into two groups: *on-call* and *on-premise* businesses. An on-call business, such as locksmiths or other general contractors, would typically visit the customer after the customer contacts them over the phone, whereas for an on-premise business, such as restaurants and hotels, customers visit the physical storefront.

We make this qualitative distinction because, as we will show in Section 5, each group exhibits distinct abusive behaviors. For example, abusive on-call businesses are more likely than on-premise businesses to verify multiple business locations with the same street address, change addresses after verification (Section 5.1), or list VoIP phone numbers (Section 5.3).

In addition, the miscreants’ modes of operation differ. An abusive locksmith, for instance, typically places a fake listing over existing locations on Maps (e.g. **Figure 6**). The listing contains what appears to be a local phone number. The locksmith quotes a low price on the phone and, upon job completion, coerces the customer into paying a higher price [9]. By contrast, Google’s internal reports suggest that affiliate fraud is common among abusive restaurants or hotels. The most common approach involves social engineering attacks. First, the miscreants claim the restaurant or hotel listing online, triggering a postcard to be sent to the business. After a few days, the miscreants call up the business, trick the owner or employee into revealing the verification PIN on the postcard, and subsequently take over their Google My Business account. Thereafter, the miscreants replace the original listing with a new one that

Country	Popularity
United States	56.5%
India	17.5%
France	5.0%
United Kingdom	3.1%
Brazil	2.0%
Canada	1.5%
Germany	1.4%
Poland	1.0%
Hungary	0.8%
Turkey	0.7%
Others	10.3%

Table 2: The fraction of abusive listings located in each country.

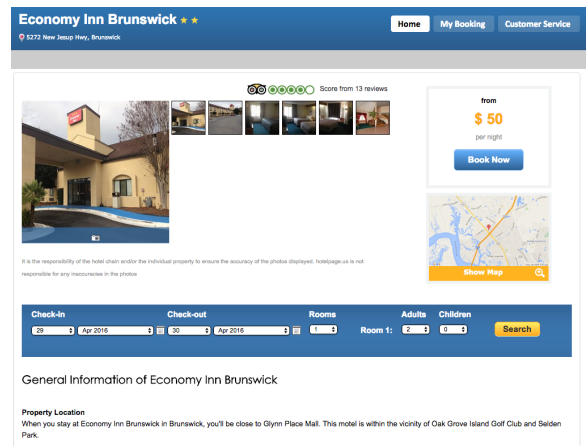


Figure 3: Screenshot of a website linked from a suspended hotel listing that is known to commit affiliate fraud.

links to miscreant-controlled booking/reservation websites. Customers can still order food or make hotel reservations through the new sites, but the miscreants charge a commission per transaction. We show an example in **Figure 3**.

4.3 Global distribution

An abusive listing may require local resources for operation—for example, access to mailboxes at which to verify the listing. As such, there may be a geographic bias in the countries in which miscreants operate. To capture the distribution of abuse globally, we measure both the volume of abuse per country and the top abused business categories per region. **Table 2** provides a ranking of the top-ten countries listed in the addresses of abusive listings. We find 56.5% of abusive listings appear within the United States, followed in popularity by India and France. Combined, these three countries account for 79.0% of all abusive listings, while the top-ten nations account for 89.7% of listings. Zooming in on the United States, we find abusive listings further concentrate their activities in six states: California, New York, Florida, Texas, Illinois, and New Jersey. Combined, these six states contribute 54.0% of suspended listings in the United States, while they account for only 39.9% of the US population.

The types of abusive listings differ drastically across region as illustrated in **Figure 4**. For the United States, locksmiths account for 43.9% of abusive listings while the same abuse is virtually absent from India, Poland, and Hungary. In contrast, the more generic




Figure 4: Per-country breakdown of the top-10 globally abused business categories. Business categories outside the top 10 are labeled as “Others.” Listings with categories that fall outside our business categories (Section 3(1)) are labeled as “Unclustered,” either because the respective categories have fewer than 100 suspended listings, or we do not understand how that business type can be clustered into existing categories.

contractor abuse appears in all top-ten countries. We note, however, that the distribution of categories may be biased toward our understanding of how businesses operate. Our categorization (described in Section 3(1)) is based on what we believe to be businesses with similar operations, which is heavily influenced by our experience in the US. For listings in some foreign countries, we lack knowledge about certain business categories. As a result, the category coverage is relatively lower in non-US countries.

5. REGISTERING ABUSIVE LISTINGS

In order for abusive operators to keep pace with Google Maps’ suspensions, they must continually expend potentially costly resources in the form of fresh Google accounts, physical addresses, websites, and contact phone numbers. Despite these hurdles, the emphasis that Google Maps places on ease of use for new business owners enabled miscreants to ultimately circumvent the intent of the postcard mail and phone verification process.

5.1 Circumventing mail verification

Operators can register a listing via mail verification or phone verification. Together, these techniques account for 63.4% of the abusive listings during our study period. The remaining 36.6% were registered via verification mechanisms that are either outdated or available only upon special request, which we exclude from discussion.

A significant question remains as to how miscreants managed to register the mail-verified listings, which account for 79.8% of all abusive listings verified through mail or phone. In particular, how did they acquire a diverse set of local mail addresses? How did they pick up and respond to the mail verification postcards? To answer these questions, we examine the verification addresses for abusive listings in the United States. This reduction in scope is necessary as our analysis requires language expertise, locale-specific understanding of the addressing system, and knowledge of the businesses that operate within a country.

To start our analysis, we canonicalize all mail verification addresses to strip out non-critical mail routing information. For example, if an abusive operator has access to “123 Park St”, nominal variations such as “123 Park St, Suite 2B” or “123 Park St, Apt




Figure 5: Fraction of listings in the US that are abusive, broken down by the number of listings that share the same mail-verification address. For example, when there are 11–100 listings per verification address, 33.8% of such listings are abusive.

3C” may in fact be fictitious suite or apartment numbers. For the purposes of our study, these addresses are assumed to be identical. More complex address manipulations that target the fault-tolerance built into US postal delivery such as “123 Park St” are beyond the scope of this paper, but may result in an under-counting of the number of abusive listings that all share the same address.

In total, we find 33.0% of mail-verified suspended listings reused the same address at least once compared with 12.0% of active listings. In this section, we investigate these common addresses further. We find dense sets of PO boxes that miscreants used to verify their listings. We also show how miscreants took advantage of what was meant to be an ease-of-use feature on Maps in order to change the address of listings after verification.

Verification hubs

We refer to addresses with ten or more associated listings as *verification hubs*. We find that 25.0% of abusive listings rely on a hub for mail verification, as shown in Table 3a. While a naive defense would be to forbid verification hubs outright, using the same address to verify multiple listings is not exclusively limited to abusive operators. For example, an office building may serve as a verification hub for multiple business professionals like lawyers and accountants; 3.3% of active “professionals” fit into this scenario. In total, 1.0% of active listings also rely on popular, shared addresses.

To quantify whether individual addresses are amenable to black-listing, we calculate the fraction of listings per address later suspended for abuse, broken down into sets of verification addresses that served different numbers of listings each. Figure 5 shows that even when listings are verified at an address that was used to verify between 11–100 listings, 66.2% of such listings are non-abusive. However, listings verified at addresses used to verify over 100 listings are almost always—but not exclusively²—abusive. Our results illustrate that attackers abused the flexibility of Google Maps’ registration system that allowed for multiple listings per address.

If we look at which business types most commonly abuse verification hubs, contractors, auto towing, and logistics (e.g., movers and packers) top the list. In particular, abusive on-call businesses, such as 31.2% of locksmiths and 25.8% of general contractors, used verification hubs. In contrast, on-premise businesses are less likely to use hubs. Only 1.0% of abusive hotel listings and 2.5% of abusive food-related listings rely on verification hubs. As we discuss in Section 4.2, these abusive listings service real hotels and restaurants, but rely on affiliate fraud to benefit from referring traffic to existing hotels and restaurants.

²We cannot confirm whether or not the remaining fraction of a active listings at those addresses are false negatives.

<i>Business Categories</i>	Section 5.1						Section 5.2		Section 5.3	
	(a) Verif Hub		(b) UPS		(c) Moved		(d) Cat changes		(e) VoIP	
	<i>Abusive</i>	<i>Active</i>	<i>Abusive</i>	<i>Active</i>	<i>Abusive</i>	<i>Active</i>	<i>Abusive</i>	<i>Active</i>	<i>Abusive</i>	<i>Active</i>
Contractors (locksmiths)	31.2%	0.8%	17.0%	0.4%	80.4%	9.7%	16.8%	1.6%	90.6%	33.7%
Contractors (others)	25.8%	0.5%	14.0%	0.8%	63.8%	11.9%	9.4%	0.8%	67.4%	16.7%
Food	2.5%	0.3%	0.7%	0.4%	5.6%	3.7%	0.8%	0.4%	4.9%	9.1%
Hotels	1.0%	0.1%	0.2%	0.1%	4.2%	3.1%	0.7%	0.3%	6.4%	6.2%
Fashion and shopping	7.5%	0.7%	3.1%	0.5%	25.6%	7.3%	6.6%	0.9%	15.4%	7.8%
Healthcare	8.2%	1.5%	3.1%	0.4%	24.7%	10.2%	1.0%	1.9%	59.8%	13.0%
Professionals	9.5%	3.3%	1.5%	0.7%	31.4%	12.5%	2.1%	0.9%	55.1%	17.6%
Travel	5.8%	1.0%	2.2%	1.1%	22.8%	12.4%	4.4%	1.3%	29.2%	15.1%
Auto	11.4%	0.2%	0.9%	0.3%	46.3%	7.6%	8.3%	0.5%	37.9%	13.7%
Artistic	2.9%	0.9%	1.0%	0.7%	62.7%	14.8%	3.9%	1.3%	18.5%	17.3%
Logistics	12.2%	0.9%	6.0%	1.6%	29.3%	12.0%	2.5%	1.0%	40.7%	24.2%
<i>Others</i>	13.3%	1.1%	2.6%	0.6%	34.3%	9.4%	4.2%	0.5%	33.4%	9.1%
Overall	25.0%	1.0%	13.0%	0.0%	67.0%	9.0%	10.0%	0.0%	69.0%	11.0%

Table 3: The percentage of listings in the US that exhibited known abusive behaviors: (a) mail-verified at verification hubs, (b) mail-verified at UPS Store addresses, (c) changed postal code after mail-verification, (d) changed into an unrelated business category after verification, or (e) displayed VoIP phone numbers on the listings.

Owners of common addresses

In investigating the most popular verification hubs, we find that miscreants used UPS stores as mailing addresses for 43.5% of all abusive listings that used hubs. Effectively, miscreants created fake listings wherever UPS stores allowed for a PO box number and likely forwarded the mail on to a single retrieval point. In this way, abusive operators removed the requirement of having a physical presence in the location targeted for abuse. For the remaining 56.5% of abusive listings verified through hubs, we discern no obvious patterns.

Even independent of verification hubs, 13.0% of abusive listings that verified through mail did so using a UPS addresses, as shown in **Table 3b**. This behavior was most popular among locksmiths, 17.0% of which were verified at UPS addresses, followed by other contractors and logistics services. For active listings, in contrast, we find very few relied on UPS addresses for verification. For example, strip malls where all tenants share the same address may include a UPS store. Similarly, businesses co-located with UPS stores such as passport services or movers and packers may have legitimate business relationships, as 1.6% of active “logistics” listings were verified at the same canonical addresses as UPS Stores.

For addresses unrelated to PO boxes, we can provide only anecdotal evidence as to how miscreants gained access to the delivered mail. When searching for discussions of how to mail verify listings, we find miscreants posting on how they recruit local residents via advertisements posted to Craigslist for “stay at home jobs,” some of which involved verifying listings via the applicants’ residential addresses. However, absent internal mailing logs like those uncovered for re-shipper scams [6], we cannot measure the prevalence of this approach.

Changing addresses

In order to reduce overhead on business owners, Google Maps allowed businesses to change their listing address without re-verification in one of two situations: (1) the new address was within the same ZIP code (e.g., a shop owner moving across town), or (2) the owner was correcting an address that cannot be parsed by Google Maps, but that nevertheless successfully received a verification postcard. The first case allowed miscreants to use verification hubs like UPS stores or other temporary addresses within a



Figure 6: Street View photo of “700 South State Street, Yadkinville, NC 27055”, the claimed location of an abusive locksmith, which was previously mail-verified at a UPS Store in White Plains, NY.

given ZIP code to serve as an initial listing address that was later updated. The second scenario unintentionally allowed miscreants to move addresses across ZIP code boundaries without triggering re-verification.

Manually reviewing the address history of listings that both moved and were later suspended, we find a common case where miscreants would register a business using an address unparseable by Google Maps, but that the US postal service would nevertheless successfully recognize for delivery due to robust fault-tolerant character recognition, such as “/23 Park St, Anytown, New York” (i.e. a malformatted address) or “123 Park St Suite 7B, Anytown, New York” (i.e. a non-existent suite number). As a backup, Google Maps would request the GPS coordinates from the miscreant for where on the map to display the listing’s pin if the owner successfully received a postcard. While originally intended to improve the accuracy of street addresses, this practice allowed miscreants to provide any GPS coordinates. A miscreant, for instance, could supply a malformatted street address in the state of New York, while providing Google with GPS coordinates in North Carolina. After receiving the postcard in New York, the miscreant would then change the street address to a location in North Carolina that matches the GPS coordinates. In this way, a miscreant could use a fixed physical address to verify a listing anywhere in the country. We show an actual example in **Figure 6**.

In general, 92.1% of listings verified through hubs changed ZIP code. Independent of verification hubs, we provide a breakdown of the frequency with which abusive and active listings change ZIP codes post-registration for each business type in **Table 3c**. We find that 80.4% of abusive locksmiths that were mail-verified changed ZIP codes during their operation, compared with 9.7% of active locksmiths.³ Overall, 67.0% of abusive listings changed ZIP code post-registration, compared with 9.0% of currently active listings. Manually sampling a small fraction of active listings, we find most legitimate address changes involved moves to nearby ZIP codes (e.g., 11000 to 11002), possibly due to editing errors or delivery route changes by the US postal service. Our findings illustrate that any form of ZIP code change should be held to a high degree of scrutiny.

Furthermore, on-call listings, such as locksmiths and general contractors, are more likely to move than on-premise listings, such as hotels and restaurants. In fact, only 5.6% of abusive food listings, along with 4.2% of abusive hotels, ever changed addresses. In contrast, more than 60% of abusive general contractors and locksmiths changed addresses. We speculate that, by changing from a verifiable address to a different, possibly fake, address, the abusive contractors are able to plant listings across a wide area, in an apparent attempt to appear in more user queries and thus attract more customer phone calls. This broad geographic coverage, on the other hand, is not necessary for on-premise listings such as restaurants and hotels.

Mitigation

Given the popularity of these abusive behaviors, Google Maps has been rolling out more stringent checks over the past year. For example, Google limits the rate of verification postcards that can be sent to the same canonical address. Address manipulations, such as adding non-existent suite numbers or spelling addresses with “leet speak”, are no longer a viable attack. While businesses can still relocate, the criteria for relocation without re-verification is further restricted to, for instance, movements within the same ZIP code.

5.2 Post-verification changes

We find evidence of miscreants attempting to evade abuse detection by changing business categories post-verification. In particular, if we look at the business category of a listing at creation time versus suspension time, we find 10.0% of abusive listings changed into unrelated categories (e.g., from Restaurant to Locksmith) as shown in **Table 3d**. By unrelated, we refer to changes from one business category to another, based on what we have constructed in Section 3(1). Changing from Chinese Restaurant to Asian Restaurant, for instance, does not fulfill this criteria, as they both belong to the Food category. Active businesses, on the other hand, rarely experience category changes. For abusive locksmiths, 16.8% changed from an unrelated category to locksmiths, while for abusive hotels and food businesses, this occurred in less than 1% of the listings.

We provide a breakdown of the most popular category transitions in **Table 4**. Of all the suspended listings that changed from one category to another, 74.9% of them changed into locksmiths, and 15.3% changed into other types of contractors. This may be the result of a perception among scammers that registering as a low-risk category reduces the likelihood of scrutiny. Since this is another common abusive behavior, Maps has rolled out extra checks for post-verification changes.

³Our dataset does not indicate whether an address change triggered a re-verification challenge. As such, we cannot detect whether active listings also abused unparsable addresses or used a proper channel. Likewise, we cannot determine which suspended listings abused this security hole.

Terminal Category	Popularity
Contractors (locksmiths)	74.9%
Contractors (others)	15.3%
Auto	1.4%
Travel	0.6%
Healthcare	0.4%
Others	7.4%

Table 4: Top five category transitions for abusive listings in the US, which changed into unrelated business categories after verification. In particular, 74.9% changed from an unrelated category into Locksmiths.




Figure 7: The fraction of listings in the US being abusive, broken down by the number of listings that publish the same phone number. For instance, a listing that publish the same phone number as 99 other listings has a probability of 26.9% of being abusive.

5.3 Communicating with customers

After miscreants successfully create an abusive listing, they require a website or phone number so that victims can contact them. Both of these represent potentially scarce resources that facilitate clustering and blacklisting. For suspended listings in the US, 95.9% of them contain phone numbers, of which 67.7% are unique. Similarly, 65.2% of suspended listings contain URLs, of which 72.8% are distinct. As on-call businesses make up a large portion of suspended listings, this section focuses on the use of phone numbers.

In order to understand whether re-use is a strong signal for abuse, we calculate the fraction of listings per phone number later suspended for abuse, broken down into sets of phone numbers with 1s, 10s, 100s, or 1,000s of listings each. **Figure 7** shows that, unlike verification addresses, popular phone numbers are dominated by legitimate businesses, due largely to thousands of regional or national brands with multiple outlets that use a common phone number.

We examine carrier information tied to each number to identify patterns in how scammers source numbers. As shown in **Table 3e**, we find 69.0% of abusive listings that publish phone numbers rely on cheap, disposable VoIP numbers. This practice is most prevalent among abusive on-call listings, such as 90.6% of locksmiths, 67.4% of general contractors, and 59.8% of healthcare services (e.g., rehab centers). Anecdotal evidence suggests that many of them operate on a referral basis. Customers dial what appears to be local phone numbers but are in fact VoIP numbers. The calls are subsequently routed to call centers, which refer the callers to actual local service providers. In contrast, legitimate on-premise business listings like hotels and restaurants rarely use VoIP phone numbers.

Using a proprietary phone-carrier database obtained on February 2, 2016, we provide a breakdown of the most popularly abused carriers in **Table 5**. These providers match those used by miscreants to bulk register phone-verified email accounts [11]; abusive operators can acquire such numbers for only the cost of a CAPTCHA. In particular, 33.4% of abusive listings with phone numbers used Bandwidth.com as the carrier.

Phone Carrier	Suspended	Active
Bandwidth.com	33.4%	3.0%
Ring Central	13.0%	0.7%
Level 3	9.9%	4.3%
Twilio	6.7%	0.4%
Broadvox	3.2%	0.2%
Google Voice	1.3%	1.9%
Peerless	1.2%	0.4%
Others	31.3%	89.1%

Table 5: Distribution of VoIP carriers for suspended and active listings in the US. For example, 33.4% of abusive listings with phone numbers used the carrier Bandwidth.com. We assign the “Others” label for unknown carriers, or if we are uncertain whether the carrier offers VoIP services.

6. IMPACT ON USERS

The ultimate measure of any form of abuse is the impact it has on users. As Google Maps relies on a ranking algorithm to select which listings to display, the number of abusive listings alone is not an accurate reflection of the state of local-listing abuse. Hence, we consider three additional metrics:

Category Impressions (CI): For each business category, we calculate the volume of impressions that abusive listings receive divided by the total volume of impressions received by all (active and suspended) listings in that category during our period of study. This value estimates the fraction of visitors actually exposed to an abusive listing while searching within a given category.

Aggregate Impressions (AI): For each business category, we calculate the volume of impressions received by abusive listings in that category divided by the number of impressions received by all abusive listings. This metric allows an alternate ranking of categories based on the volume of impressions rather than the number of listings.

Abuse Likelihood (AL): We calculate the number of abusive listings active each day in a particular category divided by the total number of active listings in that category on that day. We define an abusive listing to be active from the time of its creation up until its suspension. We then take the average across all days in our study period. Assuming a uniform query rate, this average approximates the likelihood a user would encounter an abusive listing if Google Maps selected listings uniformly at random rather than based on search quality. Effectively, this metric discounts the (in)effectiveness of any particular listing’s SEO.

We present our analysis for the top-ten abusive business categories in **Table 6**. We restrict our discussion to listings located in the United States, where these categories cover over 84% of abusive listings; coverage in other countries is lower. Overall, fewer than 0.4% of extant listings were abusive during our 16-month study period (based on the Abuse Likelihood metric) and received 0.5% of all impressions on Maps. In some categories, however, this limited impact remains despite a much higher prevalence of abusive listings. For example, abusive locksmiths, the category with the most extreme concentration of abusive listings—42.7% of all listings in the category turn out to be abusive on a day-to-day basis—managed to attract 11.1% of users’ impressions. Even so, such impressions can vary across geographic locations. In particular, users in West Harrison, NY were the most affected—where 83.3% of the search results for locksmiths were abusive. In contrast, 15.6% of search results for locksmiths in New York City were abusive.

The category whose abusive listings had by far and away the most impact on end users, accounting for 47.1% of all impressions

Business Category	CI	AI	AL
Contractors (locksmiths)	11.1%	1.7%	42.7%
Contractors (others)	0.4%	1.8%	0.7%
Food	1.1%	47.1%	0.3%
Hotels	0.5%	6.4%	0.4%
Fashion and shopping	0.1%	2.2%	0.1%
Healthcare	0.3%	1.1%	0.4%
Professionals	0.7%	2.1%	0.2%
Travel	0.9%	0.9%	1.5%
Auto	0.3%	2.6%	0.3%
Artistic	2.1%	2.4%	0.3%
Logistics	0.2%	0.2%	0.5%
Others	0.4%	31.7%	0.1%
Overall	0.5%	100.0%	0.4%

Table 6: Breakdown of user impact metrics for the top ten abusive business categories in the United States: Category Impressions (CI), Aggregate Impressions (AI), and Abuse Likelihood (AL).




Figure 8: Impressions from abusive listings every month in the Food, Hotels and Contractors categories; remaining categories are labeled as “Others.” The y -axis is normalized against the total abusive impressions in the US during our measurement period.

for abusive listings, is Food; Hotels are a distant second, constituting 6.4% of all impressions to abusive listings. As we discussed in Section 4.2, some of these abusive listings profited by referring traffic to legitimate businesses. Hence, from a user’s perspective, there was no evidence of harm: the restaurants and hotels contacted were the same businesses users expected. As such, these impressions likely did not cause harm to any users, though they do represent a financial loss to businesses (in the form of referral fees), or to non-abusive businesses, as miscreants direct user traffic to abusive listings.

We explore how impressions to abusive listings have changed over time in **Figure 8**. We find that each month, the number of impressions toward abusive contractors are significantly smaller than those for abusive restaurants and hotels, despite the fact that thousands of abusive locksmiths were created on a monthly basis. Furthermore, the overall number of abusive impressions declined toward the end of our measurement period, possibly due to a decline in the number of abusive listings.

7. RELATED WORK

Raw materials of abuse: Miscreants that create abusive local listings re-use many of the same raw materials that make up email, social-network, and other online scams. Previous studies have explored how miscreants acquired email address and account credentials via bulk registration [12], VoIP phone numbers by abusing free-tier telephony providers [11], and mailing addresses to

serve as re-shipping hubs by deceiving users into work-from-home scams [6]. Our work focuses on how miscreants combined these components to create fake business listings.

Blackhat search engine optimization: Local listing abuse bears a resemblance to blackhat SEO as both attempt to capitalize on organic search traffic for goods and services. Previous investigations of blackhat SEO found scammers profited by selling illegal pharmaceuticals, counterfeit luxury goods, and dietary supplements via web storefronts [8, 14, 15]. In contrast, the requirements for offline and physical resources, for instance access to mailboxes at scale, yield an entirely different set of monetization strategies that focus on abusive contractors.

8. SUMMARY

Map services constantly need to resolve the tension between security and usability. Excessively stringent security measures may deter abusive listings, but they could be costly to implement and may introduce friction to users. Conversely, lax security leads to more abuse, degrades user experience, and can also incur additional cost to the service provider. The trade-offs in Google Maps' verification strategies also apply to other map services, such as Yelp and Bing Maps, which, likewise, need to verify local listings via phone or mail and which also experience similar abusive issues [7]. The security-versus-usability balancing act is further complicated by geographic variations. In Google's case, mail verification is mostly effective in countries with formal address systems, but for regions that lack formal addresses (e.g. Dubai), verifying the ownership of addresses introduces more challenges—for instance, recruiting human reviewers that understand the language and culture.

In this paper, we examine abuse on Google Maps. In our analysis of suspended maps listings between 2014 and 2015, we show an intricate interplay among the types of abusive listings, the regions targeted, and the verification methods used. Even within a particular locale, the modus operandi of the abuse actors are different, ranging from their choice of verification method, how they circumvented Google's verification, to how they generated revenue. Finally, we develop a number of metrics to measure user impact. While, in general, 0.5% of listings returned by user queries were abusive, certain categories and/or geographic regions were more likely to yield abusive search results, possibly because miscreants were able to monetize local traffic. As such, we may continue seeing focused abuse in these areas.

Acknowledgements

This work was funded in part by the National Science Foundation through grant NSF-1237264. We are grateful to the Google Maps team, and in particular Alex Benton, for their invaluable feedback and access to the dataset.

9. REFERENCES

- [1] Google. Understanding consumers' local search behavior. <https://think.storage.googleapis.com/docs/how-advertisers-can-extend-their-relevance-with-search-research-studies.pdf>, 2014.
- [2] Google. Enrich Google Maps with your local knowledge. <https://www.google.com/mapmaker>, 2016.
- [3] Google. Guidelines for representing your business on Google. <https://support.google.com/business/answer/3038177?hl=en>, 2016.
- [4] Google. Show people you're open for business. <https://www.google.com/business/>, 2016.
- [5] Google. Verify a local business on Google. <https://support.google.com/business/answer/2911778?hl=en>, 2016.
- [6] Shuang Hao, Kevin Borgolte, Nick Nikiforakis, Gianluca Stringhini, Manuel Egele, Michael Eubanks, Brian Krebs, and Giovanni Vigna. Drops for Stuff: An Analysis of Reshipping Mule Scams. In *Proceedings of the Conference on Computer and Communications Security*, 2015.
- [7] Kyle Iboshi. "Pyramid Scheme of Locksmiths" Clogs Portland Market. <http://www.kgw.com/news/investigations/pyramid-scheme-of-locksmiths-clog-portland-market/56421738>, 2016.
- [8] Nektarios Leontiadis, Tyler Moore, and Nicolas Christin. Measuring and Analyzing Search-Redirection Attacks in the Illicit Online Prescription Drug Trade. In *Proceedings of the USENIX Security Symposium*, 2011.
- [9] David Segal. Fake Online Locksmiths May Be Out to Pick Your Pocket, Too. <http://www.nytimes.com/2016/01/31/business/fake-online-locksmiths-may-be-out-to-pick-your-pocket-too.html>, 2016.
- [10] Kurt Thomas, Chris Grier, Vern Paxson, and Dawn Song. Suspended Accounts In Retrospect: An Analysis of Twitter Spam. In *Proceedings of the Internet Measurement Conference*, 2011.
- [11] Kurt Thomas, Dmytro Iatskiv, Elie Bursztein, Tadek Pietraszek, Chris Grier, and Damon McCoy. Dialing Back Abuse on Phone Verified Accounts. In *Proceedings of the Conference on Computer and Communications Security*, 2014.
- [12] Kurt Thomas, Damon McCoy, Chris Grier, Alek Kolcz, and Vern Paxson. Trafficking Fraudulent Accounts: The Role of the Underground Market in Twitter Spam and Abuse. In *Proceedings of the USENIX Security Symposium*, 2013.
- [13] Cadie Thompson. Android bot spotted urinating on Apple in Google Maps. <http://www.cnn.com/2015/04/24/android-bot-spotted-urinating-on-apple-in-google-maps.html>, 2015.
- [14] David Y. Wang, Matthew Der, Mohammad Karami, Lawrence Saul, Damon McCoy, Stefan Savage, and Geoffrey M Voelker. Search + Seizure: The Effectiveness of Interventions on SEO Campaigns. In *Proceedings of the Internet Measurement Conference*, 2014.
- [15] David Y Wang, Stefan Savage, and Geoffrey M Voelker. Cloak and Dagger: Dynamics of Web Search Cloaking. In *Proceedings of the ACM Conference on Computer and Communications Security*, 2011.
- [16] Zhi Yang, Christo Wilson, Xiao Wang, Tingting Gao, Ben Y Zhao, and Yafei Dai. Uncovering Social Network Sybils in the Wild. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 2014.