Time-Zone Geolocation of Crowds in the Dark Web

Massimo La Morgia*, Alessandro Mei*, Simone Raponi*†, and Julinda Stefa*
*Department of Computer Science, Sapienza University of Rome, Italy.
†Hamad Bin Khalifa University.
Email:{lamorgia, mei, stefa}@di.uniroma1.it, sraponi@hbku.edu.qa.

Abstract—**Dark Web platforms like the infamous Silk Road market, or other cyber-criminal or terrorism related forums, are only accessible by using anonymity mechanisms like Tor. In this paper we are concerned with geolocating the crowds accessing Dark Web forums. We do not focus on single users. We aim at uncovering the geographical distribution of groups of visitors into time-zones as a whole. Our approach, to the best of our knowledge, is the first of its kind applied to the Dark Web. The idea is to exploit the time of all posts in the Dark Web forums to build profiles of the visiting crowds. Then, to uncover the geographical origin of the Dark Web crowd by matching the crowd profile to that of users from known regions on regular web platforms. We assess the effectiveness of our methodology on standard web and two Dark Web platforms with users of known origin, and apply it to three controversial anonymous Dark Web forums. We believe that this work helps the community better understand the Dark Web from a sociological point of view and support the investigation of authorities when the security of citizens is at stake.**

I. Introduction

The Dark Web hit the news six years ago with the rise of Silk Road. On Silk Road—a clandestine drug market hidden in the Dark Web—Internet users of Tor [1] could freely buy all sorts of psychedelics with bitcoins in total anonymity. A couple of years later, in 2013, the founder of Silk Road was arrested and the site taken down. At that point it was an estimated $1.2 billion business, and, after it was shut down, countless successors quickly proliferated.

In the popular culture, the Dark Web is associated with criminal activities—drug sale, identity theft, money laundering, computer hacking, botnets, credit card frauds, gun sales, child pornography, and other related cyber-crimes. This is only partly true, anonymity technology like Tor and Bitcoin were designed as a product of debates among technology libertarians in the past decades and Tor and the Dark Web are actually important to support freedom of information and speech in the Internet, especially in countries where the government or other powerful entities try to suppress it. Indeed, an important part of the Dark Web is made of forums where people can debate any matter of interest. Often, these forums are about topics that are controversial or considered questionable by the society. In other cases, they are meeting places where dissidents of authoritarian countries can freely discuss politics without being censored or prosecuted. Examples of political sites in the Dark Web are Strongbox or GlobaLeaks. Strongbox is promoted by the Freedom of the Press Foundation, while GlobaLeaks by The Hermes Center for Transparency and Digital Human Rights. Examples of forums about questionable topics are the CRD Club, a Russian site on computer hacking and technology frauds, or the Dream Market, a forum about the quality of drugs and vendors in the associated marketplace.

In this paper, we consider the problem of uncovering the geographical origin of the crowd of a Dark Web forum. We are not interested in attacking the anonymity of the single forum visitor, we are interested in understanding the geographical distribution of the visitors as a collective property. In these respect, this paper considers the notion of anonymity in the Dark Web with a new angle. Previous work, especially on Tor, has focused on attacking anonymity mostly by using traffic analysis or web browser fingerprinting. In the first case, the adversary controls both the endpoints in the Tor mixing circuit, or even the autonomous systems of the entry and exit points of the circuit, and is able to de-anonymize a single user by correlating the traffic at the two endpoints. In the second case, the adversary controls the local network of the user and is able to understand the destination site of the user browsing session by fingerprinting the traffic and matching the fingerprint against a set of known web sites. In our work, we do not assume any control of the network, we just access the forum and analyze the profile of access to the forum as documented by the postings—information that is available to every member of the forum with no particular privilege.

The fundamental idea is to consider the time of all postings in the Dark Web forum and match it to the profile of Internet activity on standard web forums of crowds from known regions. In this way, we can decompose the profile of posting into components that uncover the geographical origin of the crowd of the forum. In particular, we consider the typical profile of postings during the day of people with verified origin from 14 countries and states, including Germany, Japan, the USA, Brazil, Malaysia, Finland, among others. Surprisingly, the profiles, though with small differences due to culture, are quite consistent, both in the standard web and in the Dark Web. We use these profiles to decompose the profiles of the forum into components centered in the time zones of the world.

We performed experiments in real forums in the Dark Web. The first two, the CRD Club which is in Russian and the Italian DarkNet Community (IDC), validated our findings to be respectively centered in Russia and Italy, as perfectly confirmed by our analysis. Then, we uncovered the crowds of the Dream Market, The Majestic Garden, a site of fans of psychedelic experiences, and the Pedo Support Community, a forum on child abuse. According to our methodology, the first site is mostly European (with an important component from

North America), the second one is mostly North American (with a smaller component from Europe), and the third one, arguably the most controversial, has an important component of the crowd living in Southern Brazil or Paraguay.

We believe that our contribution can be key to better understand the Dark Web and its plethora of forums from a sociological point of view. Not only that, our methodology can give important initial information about the geographical origin of the users of a particular forum and, in case illicit activities are at stake, support the discovery of their real identities by using known de-anonymization techniques in the autonomous systems of the regions where most of them live.

II. Background

A. Tor

Tor [1] is one of the most popular anonymity systems. With over 2 million users, about $7,000$ relays, $3,000$ bridges, and $50,000$ estimated hidden services, it is also one of the largest. Tor can be used to access the Internet anonymously and to use services that are unreachable due to, for example, censorship. The main idea is that the user selects a circuit that typically consists of three relays—an entry, a middle, and an exit node. The user negotiates session keys with all the relays and each packet is encrypted multiple times, first with the key shared with the exit node, then with the key shared with the middle node, and lastly with the key shared with the entry node (also known as the guard). To send a packet to the final destination anonymously, it is first sent to the guard. The guard removes the outer encryption layer and it relays the packet to the middle node. In turn, the middle node removes its encryption layer and relays the packet to the exit node. Lastly, the exit node removes the last layer of encryption and relays the packet to its final destination.

The fundamental idea is that the guard is the only relay that communicates with the user, while it has no information on the final destination. The exit relay is the only one that communicates with the final destination, while it has no information on the user. The middle node relays packets anonymously between the two. Thanks to Tor, the user can get anonymous access to Internet services like standard websites, for example. Not only that, since the user communicates only with the guard and not directly with the service, he can also circumvent local censorship based on the IP of the destination, with the complication that the Tor relay becomes the destination (as far as the local censor is concerned), which can also be IP blocked. Some Tor relays – "bridges" – are not listed in the main Tor directory, to make it more difficult for ISPs or other entities to identify or block access to Tor [2].

B. Hidden services

Tor is also known in the Internet community as one of the core infrastructure to access the Dark Web. The Dark Web is the set of online web resources that are not indexed by common search engines and that can not be explored without using anonymity technologies such as Tor, I2P [3], or Freenet [4]. Technically, the services that run in the Dark Web under Tor technology are called hidden services. Hidden services have their own top level domain which is .onion, and their host name consists of a string of 16 characters derived from the service's public key.

While Tor can be used to access the Internet anonymously, hidden services use the very same Tor infrastructure to expose their services without revealing their real physical identity and location. On a high level, the architecture of the Dark Web under Tor includes the following components: An introduction point, hidden service directories, and a rendezvouz point. The introduction point is a Tor relay chosen by the hidden service. The rendezvous point is a Tor relay chosen by the user. The rendezvous point will be used as the meeting point of the Tor circuit set up by the user and the Tor circuit set up by the hidden service. The hidden service directories are special Tor relays that store all the information useful to allow the client to know the introduction point of the hidden services. This way, both entities are anonymous to each other and no node in the system has complete information about the communication.

To appear in the Dark Web, hidden services need to perform a setup phase. In this phase the hidden service selects the introduction points, open a Tor connection with them, and communicate the descriptor of the service to the responsible hidden service directories. The user willing to connect to the hidden service has to retrieve the descriptor from one of the hidden service directories. Then, the user selects a rendezvous point, sets up a Tor communication to it, and communicates to one of the introduction point of the service the address of the rendezvous point. The introduction point sends the address of the rendezvous point to the hidden service that, in turn, can set up a Tor communication to the rendezvous point. The rendezvous point can now tell the user that a Tor connection is established with the hidden service and that the circuit can be used to communicate anonymously.

This way, the user and the hidden service can communicate with the guarantee that both of them have no information about the real identity and the IP of the counterpart. The Tor infrastructure provides anonymity of the communication with respect to all the entities involved and to external ones. Of course, there are more details in the protocol, especially to improve the security and privacy of the user and of the hidden service. But these details are technical and not fundamental to understand the rest of the paper.

III. The High-Level Description of the Approach

As stated earlier in this work, our goal is to uncover where users post from in Dark Web forums. To reach our goal we start from an observation. Our behavioral patterns, including access to web platforms or Dark Web hidden services, is affected by our everyday life rhythm. In fact, during the day we engage in activities in a systematic way mostly dictated by the local time and daylight—waking up, going to work/school, having lunch breaks, possibly doing afternoon activities, having dinner, resting. Depending on the society we are immersed in, we might get a certain set of habits that are common to people around us but different to others.

E.g., the siesta is common in some cultures, while rare in countries with colder weather. Even so, our everyday life rhythm influences when we do a certain activity during the day, including access to the Internet. This is also shown in [5], [6] where the authors analyzed Facebook and YouTube access patterns. In both services, the requests steadily grow from the early morning to the afternoon with a peak between 17:00 (5pm) and 22:00 (10pm), then the number of requests drops rapidly during the night. In this line, our idea is to use the correlation between the everyday life rhythm (timezone and daylight) and the access or post patterns of users of forums in the Dark Web to uncover where they post from. The first step is to generate access profiles that are common to users of a certain geographical region (e.g. nation). We do so for several regions of the planet. Then, given the access profile of a crowd of users of which we do not know nothing of, we uncover their origin according to the similarity with known profiles.

IV. Building Reliable User and Region Profiles from User Activity Traces

In this section we show how we build profiles of users from a given known population starting from their activity traces. The trace can be of any kind: posts, comments to posts, messages exchanged, access times, or even all the above. We focus on building profiles that describe the level of activity of the population throughout the day. We start by profiling single users. In particular, we try to determine whether a user is or is not typically active at a given hour of the day. For this reason, the profile of a user P_u is represented by an array of 24 elements, one per hour. The element $h \in \{0, \dots, 23\}$ of P_u is the probability that user u is active during hour h of the day on the target platform. Let boolean $a_d(h)$, indicate whether a user has posted in the h^{th} hour of day d. The profile P_u is then defined as follows:

$$P_u = \{P_u[h] | h \in \{0, \dots, 23\}, P_u[h] = \frac{\sum_d a_d(h)}{\sum_{d,h} a_d(h)}\} \quad (1)$$

Intuitively, the profile P_u is the distribution of user u activity throughout the day on the target platform.

To achieve the overall population profile P we aggregate over all the user profiles as follows:

$$P = \{P[h] | h \in \{0, \dots, 23\}, P[h] = \frac{\sum_u P_u[h]}{\sum_{u,h} P_u[h]}\} \quad (2)$$

However, to build reliable region profiles we need to start off from datasets that are rich enough to reflect behavioral patterns of users, and that includes verified information regarding the provenance of the corresponding population. One possibility is the dataset [7] obtained through Twitter livestream APIs representing around 2% of the total Twitter streams in 2016. It includes tweets of $6,058,635$ users from all around the world whose hometown/country is retrievable from their Twitter profile. Using this dataset and the profile methodology described above we have built profiles for 14 countries or regions: Brazil, California, Finland, France, Germany, Illinois, Italy, Japan, Malaysia, New South Wales (Australia), New York, Poland, Turkey, and the United Kingdom. To do so, we have considered daylight saving time for all regions where it is used and we have filtered out periods of particularly low activity, like holidays. In addition, we have also filtered out non active users—users with just a handful of posts, lower than a certain threshold, that do not give enough information to profile their behavior in the long run. In our experiments, we chose the threshold to be 30 posts, as we noticed that it is a reasonable value to get a meaningful profile. Table I shows the regions we have considered along with the number of active users.

TABLE I
Twitter dataset—active users by Country/State.

Country/State	active users (#)
Brazil	$3,763$
California	$2,868$
Finland	73
France	$2,222$
Germany	470
Illinois	794
Italy	734
Japan	$3,745$
Malaysia	$1,714$
New South Wales	151
New York	1417
Poland	375
Turkey	$1,019$
United Kingdom	$3,231$

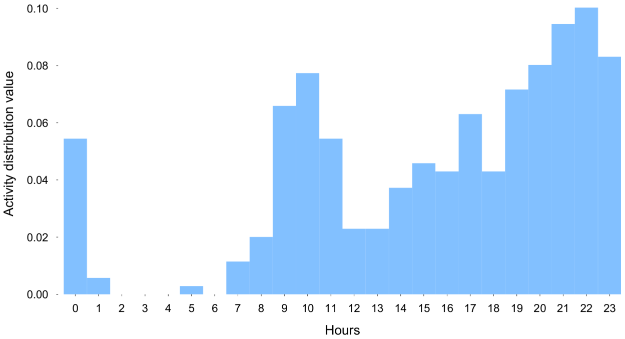

Fig. 1. A German user profile.

As an example, we show in Figures 1 and 2(a) the profiles for a typical German user and the German population, respectively. First, we note that in both the single user profile and the overall population profile we can easily distinguish night hours as those with lower activity (the interval between 1h (1am) and 7h (7am)). In addition, we can observe that the single user activity (see Figure 1) has a first peak in the morning, drops during lunch time, and begins to grow again in the early afternoon till the evening hours, following a typical daily rhythm. Finally, the overall German population profile follows the same pattern found in the Facebook and YouTube [5], [6]. This is true for all the populations of the countries we have considered in Table I: they all have the

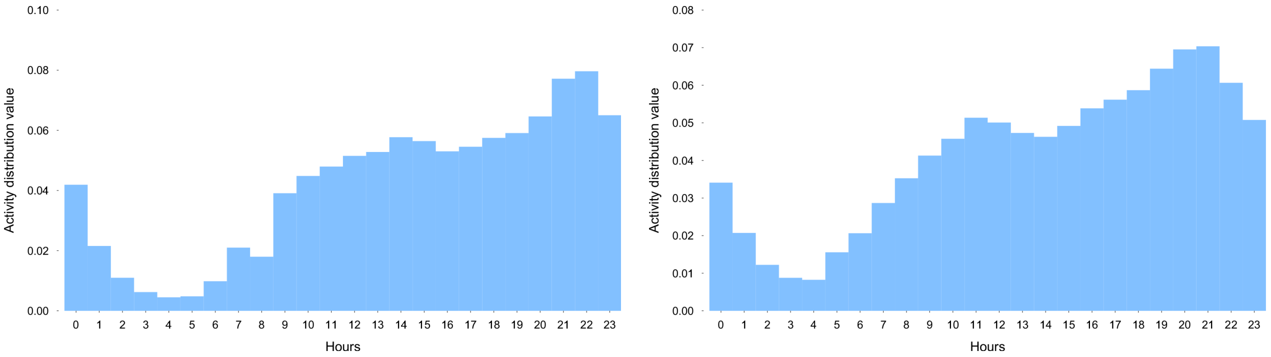

(a) Regional profile built on the German Twitter dataset.

(b) Generic profile built on all the Twitter dataset.

Fig. 2. Crowd profiles on the Twitter Dataset: German ($\mathrm{UTC}+1$) vs Generic profile (UTC).

same trend, with the only difference being the timezone. So, when crowds coming from different timezones are brought to a common one—this can be done by shifting the corresponding crowd profiles to a common timezone, like e.g. UTC—their profiles are almost identical. This is also confirmed by the high Pearson correlation value computed on crowd activity profile distributions after the shift for any two countries of the Table I: It is about 0.9 in average. The trend persists when a generic profile is built upon the activities of all users independently on their region or nationality, after the profiles are properly shifted to a common timezone. As an example, we have plotted the overall Twitter dataset profile aligned to the UTC timezone in Figure 2(b). Note how this profile is very close to the German population one 2(a) aligned to the German local time ($\mathrm{UTC}+1$), with the only difference being a 1 hour shift. Observe, for instance, how the evening activity peak is at $22h$ (10pm) for the German crowd ($\mathrm{UTC}+1$) whereas it is at $21h$ (9pm) for the generic crowd aligned to UTC. Therefore, we can easily build the profile for every region, even those not present in Table I, by just shifting the generic profile according to the time difference between the region's timezone and UTC.

A. *Placing Anonymous Users to Regions*

In this section we describe our methodology to geolocate a crowd of users of unknown origin given their activity profiles. We base our methodology on the following intuition: Users of the same region will have a profile that is very close to that of the corresponding timezone crowd, and further away from crowds of different timezones. So, for every member of an anonymous crowd, we compare its profile with that of all different timezone profiles built with the method described in the previous section. Then, we geolocate that member on the timezone whose activity profile is less distant: The one for which it takes less effort to transform the single user profile into by both shifting and moving probability mass. (Recall that activity profiles are probability distributions). An adequate distance measure in this view is the Wasserstein metric [8], also known as the Earth Mover's Distance (EMD). Given two distributions of earth mass spread on the same space, the EMD measures the least amount of work to move earth around so that the first distribution matches the second.

1) Single-Country Placement: To assess the accuracy of our geolocation methodology, we first apply it to the Twitter dataset crowds, for which we have ground truth information regarding their region. We start off with the German population. For every timezone, we compute the percentage of German population with profiles falling into that timezone according to the EMD. Despite a common nationality, the habits of two different people are not exactly the same. For example, youngsters tend to go to sleep later than older people, parents wake up earlier than teenagers, and so on. This should also be reflected in their online activity profiles. So, while we expect a large number of the German crowd to fall under the timezone of Germany, we also foresee that a portion of this crowd will be placed in neighbor timezones. This is confirmed by Figure 3, which plots the percentage of Germans placed to the 24 timezones according to the EMD. We first observe that there is a peak at $\mathrm{UTC}+1$ timezone, that covers Germany, while the values drop down for timezones further away. Most importantly, we observe that the crowd placement follows a Gaussian distribution, with a standard deviation between the fitted Gaussian and the crowd distribution of 0.013. This is to be expected, considering the slight difference in behavior between people of the same nationality mentioned earlier.

Figures 4 and 5 show the distributions for the populations of France and Malaysia, respectively. Again, we observe that they follow Gaussian distributions centered in the timezone of the corresponding country. The same trend holds for all the other countries in Table I, whose graphs we omit due to space limitations. It is worth mentioning that, after applying curve fitting [9] to the placement distributions, we note that the x axis value corresponding to the peak of the placement matches the mean of the Gaussian distribution. We also found that the average Gaussian standard deviation value for all the countries considered is $\sigma \simeq 2.5$, and that it corresponds to half of the typical hour with lowest activity, between 4am and 5am local time (see e.g. Figure 2(a) for the German population profile).

These observations bring us to the conclusion that, to geolocate a given crowd of people from the same, unknown region, it is enough to build the corresponding activity profiles placement through the EMD distance and curve-fit the result-

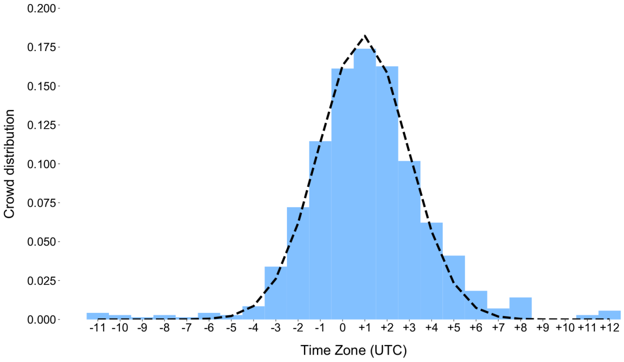

Fig. 3. EMD based placement of the German Twitter crowd.

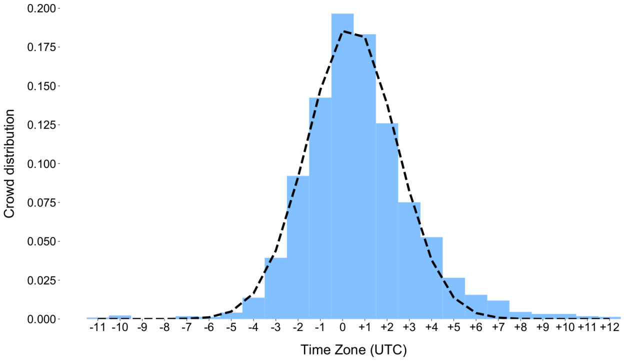

Fig. 4. EMD based placement of the French Twitter crowd.

ing distribution with a Gaussian. The center of the Gaussian will uncover the timezone of the unknown region and thus the geolocation of the crowd.

B. Multiple-Country Placement

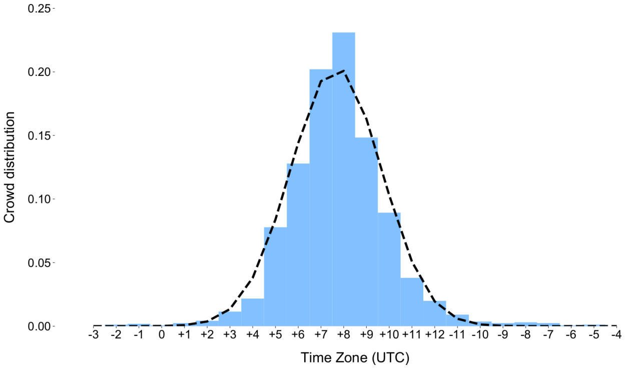

Fig. 5. EMD based placement of the Malaysian Twitter crowd.

TABLE II
GAUSSIAN FITTING METRICS.

Dataset	Average	Standard deviation
Malaysian Twitter	0.009	0.013
German Twitter	0.009	0.009
French Twitter	0.008	0.010
Synthetic dataset (a)	0.011	0.010
Synthetic dataset (b)	0.012	0.010
CRD Club	0.007	0.006
Italian DarkNet Community	0.014	0.016
Dream Market forum	0.011	0.008
The Majestic Garden	0.009	0.011
Pedo support community	0.012	0.010
Baseline	0.081	0.070

Oftentimes, users access a given site from multiple different regions. Since single region crowds follow a Gaussian distribution, we expect that the mixture of multiple region populations exhibits a profile that follows a Gaussian mixture model. Thus, uncovering the Gaussian distributions (i.e. the mean and standard deviation) the model is composed of would allow us to correctly place the members of mixed-country crowds in the corresponding geolocations. However, this is not an easy task. The reason is that we do not know a priori the number of different regions of the crowd. To address this issue we use the Expectation-Maximization (EM) [9] fitting method for Gaussian mixture distributions. EM is an iterative algorithm used to estimate the maximum likelihood parameters of a given model. In our case the model is the Gaussian mixture and the components are the Gaussian curves. To initialize the EM we use the standard deviation $\sigma \simeq 2.5$ observed empirically for the Gaussian fitting curves of single-region placement distributions in the previous section.

We test the effectiveness of the Gaussian Mixture Model (GMM) based geolocation with the Twitter dataset, on which we have ground-truth information regarding the nationality of the users. This time we build two synthetic distributions of multiple-region crowds as follows. The first synthetic distribution is made of a three-way repetition of the Malaysian user activity according to three different timezones: UTC, Californian $(\text{UTC} - 7)$, and the Australian region of New South Wales $(\text{UTC} + 9)$. In the second distribution we simply merge together users from different regions: Illinois $(\text{UTC} - 6)$, Germany $(\text{UTC} + 1)$, and Malaysia $(\text{UTC} + 8)$. The results of the geographical classification for both cases are shown in Figures 6(a) and 6(b). We can notice that we accurately uncover both the different number of regions per crowd given by the number of different Gaussian curves and the corresponding timezones that match the centers of the single Gaussian distributions.

Lastly, in order to quantify how well the fitted Gaussians match the crowd distributions we have computed the average and standard deviation of the point-by-point distance of the two (see Table II for all graphs included in this paper). As benchmark we computed the same metrics for the Malaysian dataset with the corresponding Gaussian fitting shifted of 12 hours (last row of the table). We note that both metrics are very low for both the single-country fitting (first three rows) and the multiple-country fitting (fourth and fifth row of the table). This is particularly true when we compare them to the baseline values, suggesting that the Gaussian curves fit well the crowd distribution.

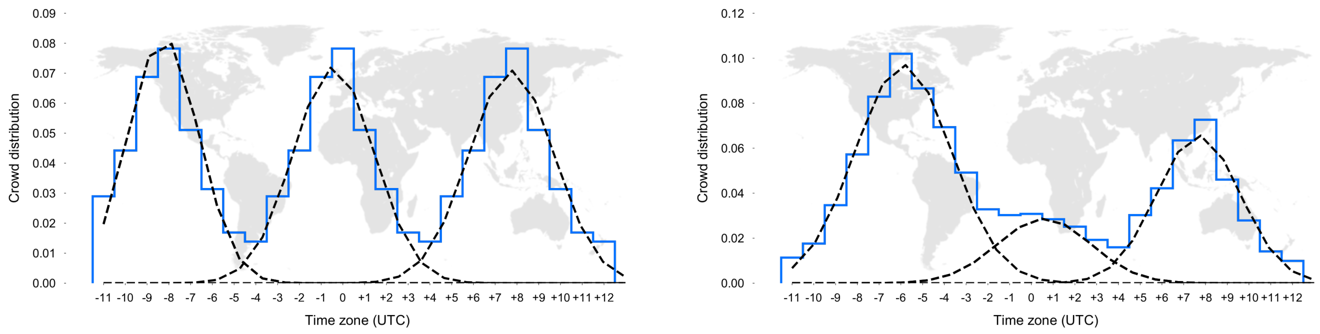

(a) Synthetic dataset modelling the behavior of Malaysian users in three different timezones: UTC, California, and Australia.

(b) Synthetic dataset made of Illinois, German, and Malaysian users.

Fig. 6. Geographical classification of multiple-region crowds.

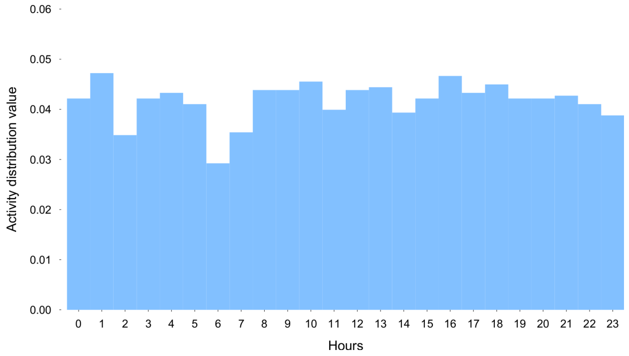

Fig. 7. Example of flat profile.

C. *Polishing the Datasets*

The EMD is also used to filter out users with so called flat profiles: Users whose activity profile are very close to being uniformly distributed over all the hours. See, for example, the profile shown in Figure 7. From an in-depth investigation on the Twitter dataset we saw that these kind of users are typically bots. Rarely, they can be shift workers. At any rate, the flatness of their profile makes so that there is no meaningful information that distinguishes them from a bot machine. In addition, they do not contribute in a meaningful way to the creation of timezone profiles. Thus, we have decided to remove these profiles from the datasets. To do so, we remove all the users whose profiles, according to the EMD, result being closer to an artificial profile created by us where every value is of $1/24$ $(1/(\#\text{daily hours}))$, than to a timezone profile. We apply this procedure in an iterative way to polish all the generic timezone profiles.

V. Results

We used our methodology to geolocate some of the most important Dark Web real forums. First, we collected information from several blogs on Tor and on the Dark Web. The Dark Web system is large and very dynamic, therefore to test our findings we selected five forums amongst the best known and popular ones. Two of these are of known origin: The first, CRD Club, is mostly in Russian, whereas the second one, Italian DarkNet Community (IDC) is the forum of the homonymous Italian marketplace in the Dark Web. We use these two first forums to validate and confirm our methodology, and then apply it to other 3 DarkWeb forums.

The experiments proceed in a similar way for all the forums. First, we sign up in the forum and write a post in the "Welcome" or "Spam" thread to calculate the offset between the server time (the one on the post) and UTC. In some forums the timestamp of the posts is accurate and already in UTC. In some other cases the timestamp does not specify the time zone and we might think that this information alone can uncover the location of the server (but not of the crowd of the forum). Of course, this is not the case since the timestamp can be deliberately shifted. In all cases, once the offset from UTC is known we can collect the timestamps of the posts in a sound and consistent way. Lastly, we also checked that in all of the forums the posts appear with no delay. This has been confirmed for all of the five forums that we have investigated.

A. *CRD Club http://crdclub4wraumez4.onion*

The first case study is a Russian forum called the CRD Club. It is divided in two macro sections, the first one written in Russian (Cyrillic script), while the other one is an international section written in English. On this forum users write about technology, hacking, gambling, online anonymity, credit card frauds and selling. Moreover, there is a subsection where you can find job offers—for example people looking for specialists that can hack a bank account or open a "bank drop" (an account open on fraudulent credentials, often in a fiscal paradise). After our analysis, we can conclude that this forum consists of a technology oriented crowd. Of course, we expect that our methodology locates this crowd in the Russian speaking countries.

We retrieved from the CRD Club 209 active users with $14,809$ posts in Russian. First, we note that the profile of activity of the users of the forum, shown in Figure 8, is very similar to the generic profile based on the whole Twitter dataset (Figure 2(b)). This observation is confirmed by the high Pearson correlation of 0.93 between the two profiles. This result supports the conclusion that the users of the Dark Web have similar access pattern of the users of the standard Web

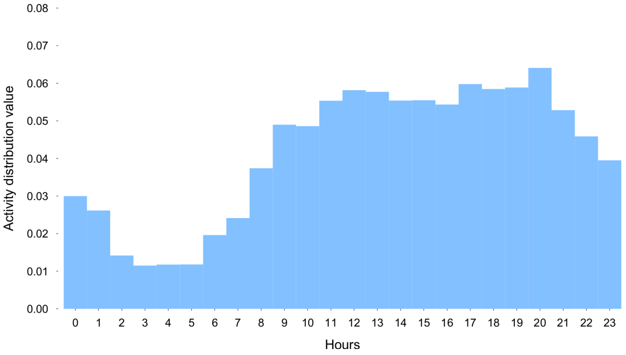

Fig. 8. Regional profile built on CRD Club Forum (UTC + 3)

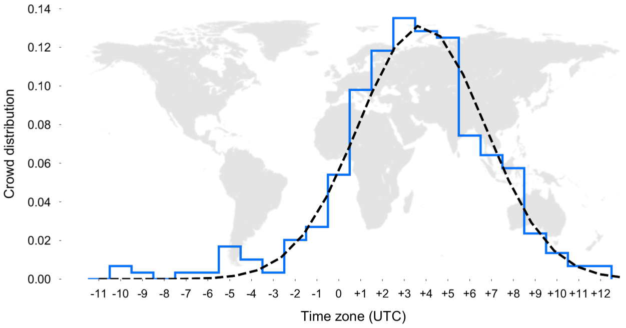

Fig. 9. CRD Club, http://crdclub4wraumez4.onion Russian Forum.

and therefore that the Twitter generic profile can be a good fingerprint for hidden services too.

The results of our geographical classification is shown in Figure 9. As we can see, there is only one Gaussian component, with an average distance of 0.007 and a standard deviation of 0.006. This means that most of the crowd come from a specific geographical area. Moreover, the Gaussian mean falls between the UTC + 3 (Bucharest, Moskow, Minsk) and the UTC + 4 (Abu Dhabi, Tbilisi, Yerevan) time zones. We can note that a very large part of the population of the Russian speaking countries live exactly in these time zones.

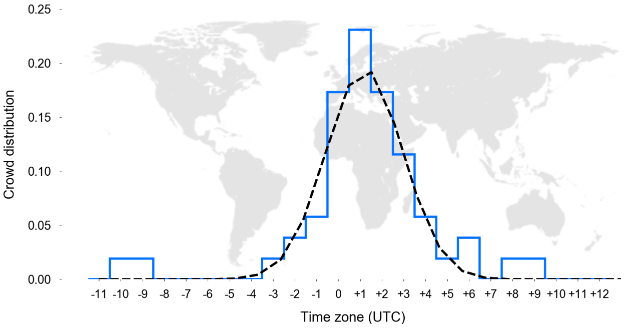

Fig. 10. Italian DarkNet Community, http://idcrldul6umarqwi.onion Italian forum and marketplace

B. *Italian DarkNet Community http://idcrldul6umarqwi.onion*

The Italian DarkNet Community (IDC) is the forum of the homonymous marketplace (http://idcrldiqkb55tjo4.onion). As the name suggests, both the forum and the marketplace are in the Italian language, although inside the forum there are small sections for English and French speakers too. The number of posts per user in these sections is very low, fewer than our threshold of 30 posts in order to get a meaningful profile. Therefore, they are not included in our analysis. The forum covers a large quantity of legal and illegal topics like politics, economy & cryptocurrency, religion, betting, carding, hacking, frauds, drugs, and sex experiences. Users of IDC are classified as regular (can freely join), 'Pro', 'Vendor', and 'Elite'. In order to become a 'Pro', 'Vendor' or 'Elite' user it is necessary to pay a subscription fee that varies in the range $0,007 - 0,3$ Bitcoin. IDC is divided in five macro sections: *Reception*, *Main*, *Bad Stuff*, *Market*, and *Elite*. The first three macro sections—*Reception*, *Main*, and *Bad Stuff*—can be read and written by anyone joining the forum. In addition, most of the posts therein are public. The section *Market* is readable only by 'Pro' users and writable by 'Vendor' users. Finally there is the Elite section that is visible only by 'Elite' members. Here users can get law advice by a lawyer of the community, access to malicious code and advanced tutorials about hacking and card frauds. After the cleaning step we classified 52 users and 1711 posts. In Figure 10 we show the distribution of users for the Italian Dark Net community. The Gaussian Mixture Model identifies a single component centered close to the UTC + 1 and slightly shifted towards UTC + 2. The distribution has a peak in the Italian time zone UTC + 1, a standard deviation of 0.016 and an average distance of 0.014. This result confirms the validity of our methodology being the forum an Italian one.

C. *Dream Market http://tmskhzavkycdupbr.onion*

The Dream Market forum is the official forum of the Dream Market Marketplace (http://lchudifyeqm4ldjj.onion). Most of the discussion is about the quality of goods and vendors in the marketplace, with a separate section to report scam vendors. It is an international forum, where English is the only language allowed. After the data cleaning step we classified 189 users and 14, 499 posts.

In Figure 11 we show the results. As we can see, our methodology discovered two main Gaussian components with an average distance of 0.011 and a standard deviation respect to the crowd distribution of 0.008. The smallest component is centered in the UTC − 6 time zone (Chicago, New Orleans, Mexico City) that is the American Mountain Time Zone. While the largest one is in the UTC + 1 time zone (Berlin, Paris, Rome). We can note that the UTC + 1 time zone, aside from Europe, covers also part of Africa (Namibia, Zimbabwe, Nigeria, etc.), and actually our methodology cannot rule out the fact that part of the crowd is from that part of the time zone. However, the following information should be taken into account: Africa is less developed than Europe from a technological point of view. Rumors [10] suggest that the

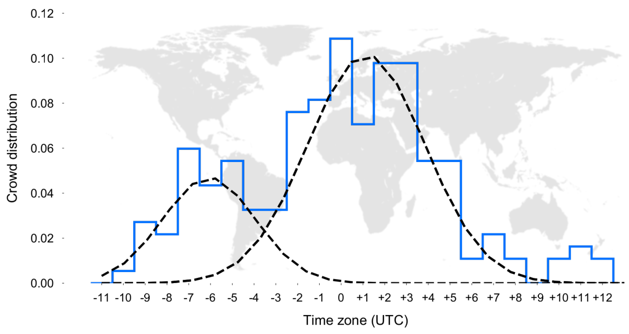

Fig. 11. Dream Market forum, http://tmskhzavkycdupbr.onion Official forum of Dream Market Marketplace.

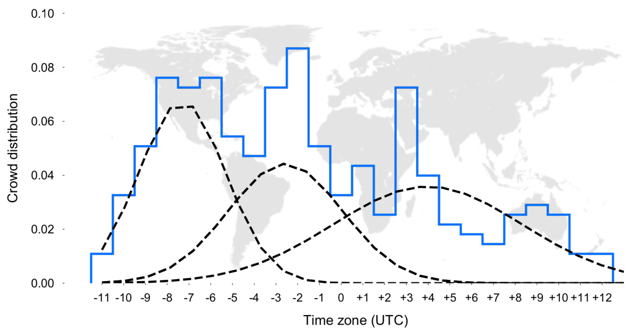

Fig. 13. Pedo support community, http://support26v5pvkg6.onion.

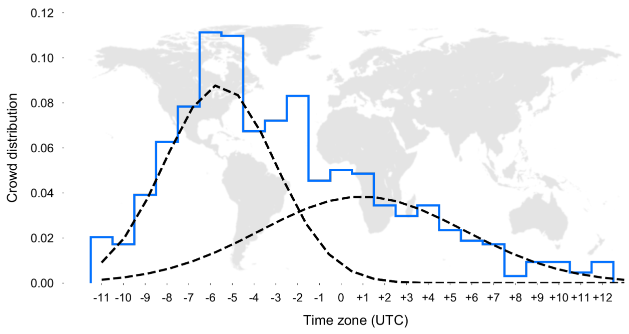

Fig. 12. The Majestic Garden, http://bm26rwk32m7u7rec.onion Psychedelic forum.

hidden service was under law enforcement control by the Dutch police. One of the former administrators, OxyMonster, is French [11]. So, we believe that, we can safely assume that the crowd of the forum classified in the UTC $+$ 1 time zone is mostly from Europe (though we cannot exclude that part of it, or in principle all of it, is from Africa).

D. The Majestic Garden http://bm26rwk32m7u7rec.onion

The Majestic Garden is a meeting point for people who share the passion for psychedelic experiences. It can be thought of as a virtual hippie commune in the Dark Web. In this forum the majority of topics are related to sharing experience after drug assumption, in particular LSD and psychedelic mushrooms. There are also topics related to selling and buying these substances or how-tos that can help you make them at home. In addition, there is a section dedicated to the literature on psychedelic and spiritual experiences.

From this forum we classified 638 active users and $75,875$ posts. In Figure 12 we show the distribution of users for the Majestic Garden. Also in this case we have two main components, an average distance of 0.009 and a standard daviation of 0.011. The largest one is centered on UTC $-$ 6 (Chicago, New Orleans, Mexico City), approximately in the Midwest belt. The mean of the second one falls into UTC $+$ 1 (Paris, Berlin, Rome). This is a mostly American forum.

E. Pedo Support Community http://support26v5pvkg6.onion

As the name suggests, the users of this forum have a common interest in pedophilia. As they say, the forum was born to share their experience far from a "world that does not understand who they really are". They are aware of the immorality and illegality of their interests and behavior—indeed in the forum it is possible to find some ethical discussion about their habits. Moreover, it is forbidden to share pedopornographic material in the forum. Also it is not allowed to disclose the country where the user lives and it is mandatory to use the English language. Lastly, some sections of the forum are hidden and access is allowed only to people that convince the administrators of the forum to be able to contribute to the discussions in a useful way. Of course, we have not done that and therefore we have no data from that part of the forum.

After the cleaning step we classified 290 active users that wrote $44,876$ posts. In Figure 13 we show the distribution of users across the time zones. In this case we have three Gaussian components with a standard deviation of 0.012 and an average distance of 0.01. The highest one is centered between the UTC $-$ 8 and the UTC $-$ 7 (San Francisco, Los Angeles, Las Vegas) time zones. The second important component falls into the UTC $-$ 3 time zone (Rio De Janeiro, Halifax, Sao Paulo). The last one is smaller and centered in the UTC $+$ 4 time zone (Yerevan, Tbilisi, Abu Dhabi).

Differently from the other cases, in this forum we classified a component whose time zone, the UTC $-$ 3 time zone (Rio De Janeiro, Halifax, Sao Paulo), mostly covers countries in the southern hemisphere. The exception is Halifax, Canada. So, intuition would suggest that this part of the crowd lives in South America. To support this intuition, we develop a methodology that we can use to indicate whether this crowd lives in the northern or southern hemisphere of the world. This is described in the next Section.

Lastly, for all five Dark Web forums under investigation we can notice how both the average and standard deviation of the point-to-point distance between the Gaussian curves and the crowd distributions shown in Table II are very low. Even more so when compared to the baseline values—those of the Malaysian distribution and its Gaussian fit shifted of 12 hours—This further supports our findings on these forums.

F. Telling apart the Northern and the Southern Hemisphere

It is well known that daylight saving time consists in advancing clocks during summer. Usually, countries using daylight saving time adjust clocks forward one hour. The idea is to delay sunset during summer at the cost of a delayed sunrise to get more sunlight in the evening and save energy used for lighting. A simple observation is that this is done from (about) March to October in the countries of the northern hemisphere, while it is done from (about) October to February in the southern hemisphere. We can use this simple fact to understand if the people of the crowd lives in the northern or southern hemisphere.

We proceed in this way: If the profile of access to the Dark Web forum of a user in the period from October to March is similar to the profile of the same user in the period March to October adjusted forward one hour, than we rule that the user lives in the northern hemisphere. Conversely, if the profile of access in the period from October to March is similar to the profile of the same user in the period March to October adjusted backward one hour, than we rule that the user lives in the southern hemisphere. If we do not see any particular difference in the two periods, we assign the user to one of the countries that do not use daylight saving time without giving any information on the hemisphere. To assess similarity of access profiles, we use again the Earth Mover's Distance.

To validate this procedure, we classified the five most active users in the datasets of United Kingdom, Germany, Italy, and Brazil. Note that all of these countries use daylight saving time (actually, in the case of Brazil, only the southern part of the country, the most populated, uses it). The 5 users in the dataset of United Kingdom as well the 5 in dataset of Germany and the 5 in the dataset of Italy, all of them, are classified as living in the northern hemisphere. The 5 users in the dataset of Brazil, all of them, are classified as living in the southern hemisphere. Therefore, we believe we can use this methodology with good confidence. We have done it for the Pedo Support Community, due to the controversial nature of the forum and the alleged origin from South America of a good part of the crowd. We limit our analysis to the 5 most active users of the forum, since those are the users for which we have a good number of posts. According to the analysis, 3 out of 5 of the most active users in the Pedo Support Community live in southern hemisphere, while the other 2 in the northern hemisphere. This result confirms our initial intuition that a good part of the crowd of the forum lives in South America. Actually in Southern Brazil or Paraguay, which is the only land in the UTC $-$ 3 time zone that is in the southern hemisphere and that does use daylight saving time.

VI. Related Work

Traffic Correlation Attacks. Many of the attacks to Tor in the literature are traffic correlation attacks to individual users. As Tor is a low-latency network and packet timing and size are not obfuscated, it is well known that an adversary able to observe both endpoints of a Tor circuit can de-anonymize the user [12], [13]. Bouer et al. [14] demonstrate that this kind of attack can be carried out in a quite efficient way. Entry nodes are chosen based on up-time and bandwidth rates reported by nodes, which are not verified by the Tor network. So, malicious nodes can maximize the likelihood to be chosen as entry nodes by reporting incorrect information about their up-time and bandwidth. Then, malicious nodes can also drop all circuits in which either of endpoints is non malicious. Hence the circuit must be rebuilt and there is a new chance to build one with both endpoints under the control of the adversary. Correlation attacks can also be done at autonomous system level. In 2009, Edman et al [15], using BGP (Border Gateway Protocol) historical routing data and simulating the path selection of Tor, show that a significant percentage of paths are vulnerable against an AS-level adversary. Nithyanand et al. [16], in 2013, found out that up to 40% of Tor circuits are vulnerable to the same adversary, 85% to a state-level adversary, and this value rises up to 95% for states like China and Iran. A similar work was done by Murdoch et al. [17] for an IXP-level adversary.

Network Manipulation Attacks. More recently, Sun et al. [18] introduce RAPTOR, a suite of three new attacks to de-anonymize individual users of Tor. In the first attack they show that, instead of monitoring only one direction of the anonymous connections, an AS-level attacker can exploit the asymmetric nature of internet routing. In this way, a malicious observer can increase the chance to observe at least one direction of the connections and use TCP headers in order to correlate them. In the second attack, they show that exploiting both the asymmetric correlation and the BGP churn a long term passive AS-level adversary can increase its surveillance capabilities by up to 50% over a month. The last attack is an active one, where the AS adversary manipulates inter-domain routing by advertising incorrect BGP control messages. In this kind of attack the adversary can observe only one of the two connection endpoints. The adversary launches a BPG hijacking attack against the not controlled endpoint in order to route the traffic into the malicious AS, allowing the adversary to execute an asymmetric traffic correlation attack.

Website Fingerprinting. Other works are based on website fingerprinting [19]–[21]. With this approach the adversary builds a database of network traces of users who visit a set of websites—sequence of packets, size of packets in the sequence, inter-packet intervals. The adversary can do that in several ways, even by deploying his own users. Then, these traces are used to train a classifier. If the adversary can monitor the Tor traffic between a target user and the guard, he can use the classifier to learn what website is being visited by the victim. First results, despite the good accuracy score achieved in a controlled environment, show that in a open-world scenario the success of the attack is significantly lower [22]. Kwon et al. [23] used the same ideas, but this time monitoring only Tor circuits involved in a communication with the hidden services instead the whole Tor traffic. This adjustment greatly reduces the amount of connections to be monitored and makes the fingerprint attack feasible. In 2017, Overdorf et al. [24] repeat previous experiments with a largest dataset of .onion services, achieving an average accuracy score of 80%. Moreover, they show that smaller sites that change the most between visits are the hardest to identify, larger sites instead are more subjects to this kind of attack.

Information Leak. Lastly, other attacks exploit information leaks from applications that were not intended to work over Tor. For example, Biryukov et al. [25] target Bitcoin users. They show that using Bitcoin through Tor exposes the users to a man-in-the-middle attack that can, over multiple transactions, identify the victim. Exploiting the anti-DOS protection of Bitcoin, an attacker can cause the ban of all non-malicious peers that run the Bitcoin protocol over Tor, forcing the victim to use a malicious exit node as peer. Now, the attacker can store a cookie inside the target client in order to identify the victim in future connections. Manils et al. [26] demonstrate that just the use of BitTorrent alone can jeopardize the anonymity of the user. In fact, in some case a BitTorrent tracker responding to a client query can disclose its IP address. So, an attacker that monitor the Exit node is able to unveil the identity of the client. Using a more sophisticated technique, a malicious tracker sends to the client a manipulated list of peers to connect the victim to a malicious peer and retrieve its IP.

In all the above mentioned works the attack targets the anonymity of a specific user. In our work we uncover the geographical origin of a crowd of a Dark Web site, attacking the anonymity of the group instead of directly the anonymity of the single individual. Moreover, to the best of our knowledge, this work is the first one that attacks anonymity by exploiting the collective behavior of a crowd instead of technical weaknesses of the network, the systems, or the protocols used in the Dark Web.

VII. Discussion

No timestamp on posts. Timestamps are always shown in the Dark Web forums under investigation. However, the forum might remove them to protect the time of access of the anonymous user. This is actually not stopping our methodology—it is enough to monitor the forum, see when posts are made and timestamp them ourselves. The process is slightly trickier than just creating a dump of all previous posts, as we have done in this work. One might need to monitor a sufficiently large number of days, depending on the frequency of the posts, in order to collect 30 post per user or more necessary to build meaningful profiles. Nonetheless, the methodology presented in this paper can still successfully be applied.

Forum shows and timestamps posts with random delay. This is possible. But, to be effective, the random delay must be of at least a few hours reducing considerably the forum usability. So, many users could just move to other forums.

What if the crowd coordinates and users deliberately post with a profile of a different region? We assumed that people are not under the control of an adversary. Indeed, coordinating the behavior of hundreds of anonymous users can be very hard. Moreover, if anonymous users are forced to wake up in the night to make a post, most probably they don't, and they either leave the forum or keep behaving normally.

VIII. Ethical Considerations

In this work we analyzed $1,378$ anonymous users of forums in the Dark Web. While doing so, we gathered $151,770$ posts from five different hidden services. The data collected (only author ID and time of posting, without the body of the forum post) was stored for a limited amount of time in our servers in an encrypted form. It was not shared directly nor placed on platforms where it could be downloaded from. Consequently, and accordingly to the policy of our IRB, we did not need any explicit authorization to perform our experiments. Our work is compliant to the Tor research safety guidelines [27]. We believe that the Tor community can benefit from this research, that sheds light on important issues related to the (group) privacy of Dark Web users.

IX. Conclusion

In this paper we focused on geolocating crowds on the Dark Web into the time zones of the World. The approach, that we believe to be unique in its kind, does not use traffic analysis or protocol-related breaches, unlike previous work. The fundamental idea is to build reliable profiles of posting activity on online forums, then, to match Dark Web crowd profiles to those of known regions. Our approach works well with crowds of users coming from a single country and many different countries. We assess the effectiveness of the approach by applying it to three types of datasets: A Twitter dataset that includes information on the geolocation of users, a Russian speaking forum of the Dark Web, and an Italian Dark Web Marketplace that serve as validators of the approach, and three other forums on the Dark Web which we have no information on. Aside from effectively uncovering countries from different time zones, our approach can also be used to discover more fine-grained information on the crowds. An example is that of the most active users of a Pedo Support Forum in the Dark Web. We found out that an important part of the forum crowd comes from a region that covers Southern Brazil and Paraguay.

We believe that the methodology presented in this paper lays down the foundations to shed light on the Dark Web and the multitude of its services from a sociological point of view. At the same time, our geolocation technique can be particularly valuable to authorities performing ongoing investigation and geolocation of users engaged in illicit, cybercriminal, or terrorism related activities in the Dark Web.

Acknowledgment

This work was supported in part by the MIUR under grant “Dipartimenti di eccellenza 2018-2022” of the Department of Computer Science of Sapienza University.

References

[1] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The second generation onion router,” in *Proceedings of the 13th USENIX Security Symposium*, 2004.

[2] Tor Project. Tor: Bridges. Accessed: 2018-05-01. [Online]. Available: https://www.torproject.org/docs/bridges.html.en

[3] zzz and L. Schimmer, “Peer profiling and selection in the i2p anonymous network,” in *PET-CON 2009.1*, 2009.

[4] I. Clarke, S. G. Miller, T. W. Hong, O. Sandberg, and B. Wiley, "Protecting free expression online with freenet," *IEEE Internet Computing*, vol. 6, no. 1, 2002.

[5] M. Kihl, R. Larsson, N. Unnervik, J. Haberkamm, A. Arvidsson, and A. Aurelius, "Analysis of facebook content demand patterns," in *Proceedings of the IEEE International Conference on Smart Communications in Network Technologies (SaCoNeT)*, 2014, pp. 1–6.

[6] A. Arvidsson, M. Du, A. Aurelius, and M. Kihl, "Analysis of user demand patterns and locality for youtube traffic," in *Proceedings of the IEEE 25th International eletraffic Congress (ITC)*, 2013, pp. 1–9.

[7] A. Team. (2016) Archive team: The twitter stream grab. Accessed on 2017-06-12. [Online]. Available: https://archive.org/details/twitterstream.

[8] H. F. K, "The distribution of a product from several sources to numerous localities," *Studies in Applied Mathematics*, vol. 20, no. 1-4, pp. 224–230, 1941.

[9] C. M. Bishop, *Pattern recognition and machine learning*, ser. Information Science and Statistics. Springer-Verlag New York, 2006.

[10] Mashable Inc. (2017) Buyers fear police have secretly seized the last big dark web market standing. Accessed: 2018-03-28. [Online]. Available: https://mashable.com/2017/07/21/dark-web-marketplace-drugs-dream-market

[11] The Guardian. (2017) Trip to world beard competition ends in arrest for alleged dark web drug dealer. Accessed: 2018-03-28. [Online]. Available: https://www.theguardian.com/us-news/2017/sep/28/world-beard-moustache-competition-drug-dealer

[12] P. Syversonl, G. Tsudik, M. Reed, and C. Landwehr, "Towards an analysis of onion routing security," in *Proceedings of the International Workshop on Designing Privacy Enhancing Technologies: Design Issues in Anonymity and Unobservability*. Springer, 2001, pp. 96–114.

[13] V. Shmatikov and M. H. Wang, "Timing analysis in low-latency mix networks: Attacks and defenses," in *Proceedings of the 11th European Symposium on Research in Computer Security (ESORICS)*, 2006, pp. 18–33.

[14] K. Bauer, D. McCoy, D. Grunwald, T. Kohnoi, and D. Sicker, "Low-resource routing attacks against tor," in *Proceedings of the 2007 ACM Workshop on Privacy in Electronic Society (WPES)*. ACM, 2007, pp. 11–20.

[15] M. Edman and P. P. Syverson, "As-awareness in tor path selection," in *Proceedings of the 16th ACM Conference on Computer and Communications Security (CCS)*. ACM, 2009, pp. 380–389.

[16] R. Nithyanand, O. Starov, A. Zair, P. Gill, and M. Schapira, "Measuring and mitigating as-level adversaries against tor," in *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2016.

[17] S. J. Murdoch and P. Zieliński, "Sampled traffic analysis by internet-exchange-level adversaries," in *Proceedings of the 2007 International Workshop on Privacy Enhancing Technologies (PETS)*. Springer, 2007, pp. 167–183.

[18] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal, "Raptor: Routing attacks on privacy in tor," in *Proceedings of the 24th USENIX Conference on Security Symposium (SEC)*, 2015, pp. 271–286.

[19] D. Herrmann, R. Wendolsky, and H. Federrath, "Website fingerprinting: attacking popular privacy enhancing technologies with the multinomial naïve-bayes classifier," in *Proceedings of the 2009 ACM Workshop on Cloud Computing Security (CCSW)*. ACM, 2009, pp. 31–42.

[20] A. Panchenko, L. Niessen, A. Zinnen, and T. Engel, "Website fingerprinting in onion routing based anonymization networks," in *Proceedings of the 10th annual ACM Workshop on Privacy in the Electronic Society (WPES)*. ACM, 2011, pp. 103–114.

[21] T. Wang and I. Goldberg, "Improved website fingerprinting on tor," in *Proceedings of the 12th ACM Workshop on Workshop on Privacy in the Electronic Society (WPES)*. ACM, 2013, pp. 201–212.

[22] M. Juarez, S. Afroz, G. Acar, C. Diaz, and R. Greenstadt, "A critical evaluation of website fingerprinting attacks," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 263–274.

[23] A. Kwon, "Circuit fingerprinting attacks: Passive deanonymization of tor hidden services," Ph.D. dissertation, Massachusetts Institute of Technology, 2015.

[24] R. Overdorf, M. Juarez, G. Acar, R. Greenstadt, and C. Diaz, "How unique is your. onion? an analysis of the fingerprintability of tor onion services," *arXiv preprint arXiv:1708.08475*, 2017.

[25] A. Biryukov and I. Pustogarov, "Bitcoin over tor isn't a good idea," in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, 2015, pp. 122–134.

[26] P. Manils, C. Abdelberri, S. L. Blond, M. A. Kaafar, C. Castelluccia, A. Legout, and W. Dabbous, "Compromising tor anonymity exploiting p2p information leakage," *arXiv preprint arXiv:1004.1461*, 2010.

[27] Tor Project. Tor Research Safety Board. Accessed: 2018-05-01. [Online]. Available: https://research.torproject.org/safetyboard.html