



ARTICLE



<https://doi.org/10.1057/s41599-024-03343-4>

OPEN

Unveiling local patterns of child pornography consumption in France using Tor

Till Koebe^{1✉}, Zinnya del Villar², Brahmani Nutakki¹, Nursulu Sagimbayeva¹ & Ingmar Weber¹ ¹

Child pornography—better known as child sexual abuse material (CSAM)—represents a severe form of exploitation and victimization of children, leaving the victims with emotional and physical trauma. In this study, we aim to analyze local patterns of CSAM consumption across 1341 French communes in 20 metropolitan regions of France between March 16 to May 31, 2019 using fine-grained mobile traffic data of Tor network-related web services. We estimate that approx. 0.08% of Tor mobile download traffic observed in France is linked to the consumption of CSAM by correlating it with local-level temporal porn consumption patterns. This compares to 0.19% of what we conservatively estimate to be the share of CSAM content in global Tor traffic. In line with existing literature on the link between sexual child abuse and the consumption of image-based content thereof, we observe a positive and statistically significant effect of our CSAM consumption estimates on the reported number of victims of sexual violence and vice versa, which validates our findings, after controlling for a set of geographically disaggregated features including socio-demographic characteristics, voting behavior, nearby points of interest and Google Trends queries. While this is a first, exploratory attempt to look at CSAM from a spatial epidemiological angle, we believe this research provides public health officials with valuable information to prioritize target areas for public awareness campaigns as another step to fulfill the global community's pledge to target 16.2 of the sustainable development goals: "end abuse, exploitation, trafficking and all forms of violence and torture against children".

¹Saarland Informatics Campus, Saarland University, Saarbrücken, Germany. ²Data-Pop Alliance, Rennes, France. ✉email: till.koebe@uni-saarland.de

Introduction

“Derrière tout échange d'image ou de vidéo pédopornographique, il y a un agresseur et un mineur agressé.”—Adrien Taquet, 2021.

(Behind any exchange of child pornographic images or videos, there is an attacker and an attacked minor.)

As pointed out by the French Secretary of State for Child Protection Adrien Taquet in 2021, child sexual abuse materials (CSAM) represent both a severe form of exploitation and victimization of children and at the same time a criminal offense (Assemblée Nationale, 2022). Sexual violence leaves affected children with emotional and physical trauma (Pinheiro, 2006). For France, the National Institute of Health and Medical Research (INSERM) estimated in a general population survey conducted between 2020 and 2021 that 1 in 10 French adults, approx. 5.5 million individuals, have been subject to sexual violence in their childhood (Sauvé et al. 2021), with serious health consequences as shown by Brown and Scodellaro (2023). The Independent Commission on Incest and Sexual Violence against Children (CIIVISE) installed by the French president on March 23, 2021, estimates that every year in France alone 160,000 children become victims of sexual violence. Also, research suggests that CSAM consumption is not a rare phenomenon. Seto (2013) estimates that 2% to 4% of all men have consumed CSAM online. Eke et al. (2011) found that 24% of CSAM users from their sample had committed sexual offenses in the past. Similarly, Hall and Hall (2007) reported that 30% to 80% of individuals who viewed CSAM had molested a child. That emphasizes the important link between CSAM consumption and sexual violence against children (Box 1).

Looking at the personal or environmental factors that drive CSAM consumption gives a multi-faceted picture. A study by Price et al. (2015) with 46 CSAM consumers found that the participants were predominantly single or separated/divorced unemployed European males with a pronounced experience of depression and anxiety, loneliness, and childhood abuse. One-third of them had previously engaged in contact sexual offending. A study by Seigfried et al. (2008) where 307 respondents (30 classified as CSAM consumers) completed an online survey revealed that CSAM consumers obtained higher scores on exploitive-manipulative amoral dishonesty traits and lower scores on internal moral choice. Another study by Seigfried-Spellar and Rogers (2010) analyzed responses from 162 female respondents, out of which 10 consumed CSAM. Female CSAM consumers in the study scored lower on neuroticism and higher on moral choice hedonism. In addition, as Fortin and Proulx (2019) point out by referring to a number of studies (cf. Babchishin et al. (2015, 2011); Elliott et al. (2013)) that CSAM consumers and contact sexual offenders have distinct characteristics with the latter being less educated, more often unemployed with more mental health problems, less self-control, more antisocial traits

and more substance abuse. A reason for less clear-cut profiles of CSAM consumers might be that larger-scale studies such as the one by Nurmi et al. (2024) on CSAM consumption behavior are still rare and most of the scientific evidence has been created from studies involving a small number of respondents in individual-level studies. For situations where access to individual-level data is limited (e.g., due to privacy regulations or other data collection challenges), area-level analysis offers an alternative. For example, Chetty et al. (2022) use Facebook friendship ties aggregated to the zip-code level to explain the socio-economic status in these areas. Bruckschien et al. (2019) use area-level aggregates at local levels in Turkey to identify the share of refugees in undeclared employment situations. In another study, Rotondi et al. (2020) use national-level aggregates across 209 countries to find evidence between mobile phone diffusion and health indicators such as contraceptive prevalence. Consequently, we apply a similar strategy in this paper by using CSAM consumption estimates and potentially associated factors both aggregated on the commune-level to complement existing individual-level studies.

When it comes to CSAM detection, various automatic approaches have been proposed. Sae-Bae et al. (2014) developed a classifier with a true positive rate of 83% in detecting explicit-like child images and 96.5% in detecting child faces on a test set of 105 images featuring semi-naked children. Vitorino et al. (2017) utilized convolutional neural networks (CNN) to differentiate regular images from adult pornographic and CSAM content, respectively. Macedo et al. (2018) created a region-based annotated CSAM dataset (RCPD) in collaboration with the Brazilian Federal Police. They combined face-based child detection with a pornography detector and achieved an accuracy of 79.84% on the proposed benchmark. Overall, consistently improving CSAM detection algorithms might prompt illegal content creators and distributors to turn to the so-called “darknet” even more, making it harder for the authorities to assess and prevent CSAM circulation on the web. While the advancement of technology made it easier to moderate and filter abusive and illegal content, it has also provided opportunities for sharing such content with little accountability. CIIVISE states in its interim report that even though France is the fourth-largest online host of CSAM in the world, it only employs 1 cyber-crime investigator per 2.2 million people compared to about 1 investigator per 100,000 people in the Netherlands (CIIVISE, 2021).

With its advanced anonymity and privacy features, the Tor network¹ has been criticized in the past for facilitating illegal activities in the digital space, including the distribution of CSAM (Deutsche Welle, 2019). Gannon et al. (2023) find that child abuse sites are 2000 times more prevalent in the darknet, for which Tor provides the main entry point. But they also find that CSAM communities use both the darknet and the clearnet for content sharing: While live streams of child sexual abuse—predominantly

Box 1 | Descriptions of key terms used in this study

- **Tor:** The Tor network, short for “The Onion Router”, is a privacy-focused network that directs internet traffic through a series of volunteer-operated servers, encrypting it at each step and making it difficult to trace back to the user’s origin, thereby enhancing online anonymity. It is commonly used to access the internet anonymously.
- **Darknet:** The darknet is a part of the internet that is intentionally hidden and not indexed by traditional search engines, often accessible only through specialized software like Tor.
- **Hidden services:** Hidden services, often associated with the Tor network using “.onion” domains, refer to websites and online services that are hosted on servers configured to be accessible only through the Tor network.
- **CSAM:** CSAM, or Child Sexual Abuse Material, refers to explicit media that involves the sexual exploitation or abuse of minors, including images, videos, or other content. As pointed out in Terminology and Semantics Interagency Working Group on Sexual Exploitation of Children (2016), the term CSAM is preferred over the still widely used term *child pornography* as using the word *pornography*—generally referring to “[...] adults engaging in consensual sexual acts distributed (often legally) to the general public for their sexual pleasure”—may trivialize the crime.

taking place in developing countries—are mainly hosted in the clearnet, presumably as the risk of law enforcement agencies being aware of live streams is generally perceived to be low, non-live content is predominantly shared via CSAM forums in the darknet. According to Gannon et al. (2023), CSAM-related hidden services usually showcase archaic layouts and do not use high-security technology. Their main protocol to keep the community safe is to share the sites only with like-minded users, typically by invitation from the site administrators or moderators. Some sites require the user to post similar content before they can access the forums. van der Bruggen et al. (2022) found in a study on a large CSAM forum that while only a fraction of the forum members (0.7%) were responsible for 40% of the content posted, 9 out of 10 forum members tried to download CSAM at least once.

In this work, we present two major contributions to this field of research: First, to the best of our knowledge, this is the first time that consumption patterns of CSAM are estimated at such a high geographic granularity by correlating it with local-level temporal adult porn consumption patterns. Second, we link these fine-granular consumption patterns to both small-area socio-demographic characteristics as well as nearby points of interest and Google Trends² queries. While local patterns of both the consumption as well as production of CSAM are relevant for public health professionals and law enforcement agencies alike, we focus on the *consumption* of CSAM for two reasons: First, we assume that uploads of CSAM are mainly done via fixed internet lines/Wifi rather than via the mobile network. Since we only observe mobile network traffic, we consequently expect download traffic to carry stronger signals related to CSAM-related darknet activities. Second, recalling from above, there is a strong empirical link between the consumption of CSAM and being involved in sexual violence against children. As Insoll et al. (2022) points out: 42% of survey respondents in their study who have viewed CSAM tried to connect with children online afterwards. Therefore, knowledge about local patterns of CSAM consumption in the darknet may also inform about the prevalence of sexual violence against children in the physical world.

The paper is structured as follows: We describe the data used for this study in Section “Data”. In Section “Methodology”, we explain the methodology applied to derive local-level estimates of CSAM consumption and the assumptions used. Commune-level estimates of CSAM consumption for 20 metropolitan regions in France are presented in Section “Results” alongside their links to POIs, Google Trends and other socio-demographic characteristics. Limitations of the study and words of caution are extensively discussed in Section “Discussion”.

Data and methods

In this study, we aim to analyze local patterns of CSAM consumption for 1341 communes across 20 metropolitan regions in France. The communes represent the fourth and thus smallest administrative division in France with considerable political decision-making power on the local level. The population sizes of the communes in the sample range from 80 in Mont-Saint-Martin, Grenoble to 498,596 in Toulouse averaging at 14,802 across all areas (INSEE, 2019).

Data. The data for Tor usage patterns are derived from georeferenced, service-level mobile network traffic data measured by the mobile network operator Orange for 20 major cities in France across 77 consecutive days from March 16 to May 31, 2019, provided on a 100 × 100 m spatial grid, also called *tiles* in the following, through the NetMob 2023 data challenge (Martínez-Durive et al. 2023). The upload and download data obtained from the mobile network operator is normalized by a random value to

conceal the actual traffic of the operator while retaining comparability across web services. Therefore the actual values do *not* have any unit, such as GB, attached to them. It is important to note that mobile network traffic data does not include web traffic generated when connected to fixed-line internet or Wifi. The geographic location of a specific user equipment (UE), e.g., a mobile phone, is captured via the base stations of the mobile network the UE is connected with during a given time interval of 15 minutes. The captured web service-specific network traffic is then distributed within the estimated coverage area of the respective base station. For details on the effect of different coverage area estimation approaches we refer to Koebe (2020). For more details on the data preprocessing performed on the Netmob dataset, we refer to Martínez-Durive et al. (2023).

While data for a variety of web services are provided, we focus on Tor as the main entry point to the darknet. In addition, we consider download traffic from mainly pornographic websites (referred to as “Web Adult” in the following) as a reference for the consumption of pornographic content and download traffic to YouTube as a reference for general mobile video consumption. Both Web Adult and Tor represent multiple web services grouped into a broader category, respectively. However, details on the exact composition of these categories are not available from Martínez-Durive et al. (2023).

In order to investigate spatial relationships of CSAM consumption and local points of interest (POI), we build on the recently released Overture Maps Foundation (OMF) Places dataset that provides information on about 3 million points of interest for France derived from Meta and Microsoft products such as Bing Maps and Facebook pages (Overture Maps Foundation, 2023). Using data from OpenStreetMap (OSM) has also been considered, however, OSM provides comparatively little POI information on local businesses.

Furthermore, we use the reported number of victims of sexual violence as our groundtruth retrieved from the Service Statistique Ministériel de la Sécurité Intérieure (SSMSI) database of the interior ministry of France (Ministère de l’Intérieur et des Outre-Mer) (Ministère de l’Intérieur et des Outre-Mer, 2022). Socio-demographic information provided by the French National Statistical Office INSEE (INSEE, 2019) and voting outcomes from the 2017 French presidential election (Ministère de l’Intérieur et des Outre-Mer, 2017) are used to control for potential confounders when investigating the link between estimated CSAM consumption and sexual violence.

Lastly, we complement our analysis with information on the relative popularity of search terms from Google Trends. Specifically, we consider the following set of partially community-specific keywords inspired by Owens et al. (2022) and complement them with equivalent terms in the French language: *pedoporno*, *porno mineur*, *porno enfant*, *site pedoporno*, *pre-teen hardcore*, *zoo preteen*, *zoo pre-teen*, *pedomom*, *pedodad*, *pthc*, *boylove*, *girllove*, *porno jeune ado*, *video porno ado*, *ado porno*, *porno jeune fille*, *omegle* and *hurtcore*. We extract the relative popularity values of these search terms for each of the 21 regions of France (excluding Corsica, note that Google Trends still uses the regional delineations prior to the 2015 reform) pooled across the years 2017 to 2021 to avoid excessive data sparsity exhibited when using shorter time intervals that more aligned with the time window of the mobile traffic data. We map these values to the departments in our sample. While acknowledging that hidden services cannot be found via Google search queries and that the CSAM community actively exchanges “best practices” to stay anonymous (cf. Gannon et al. (2023)), we expect that these keywords may still be able to capture deviances from these practices. Further details on the variables used in this study can be found in Table 1 of the Appendix.

Table 1 Summary statistics of the potential correction factors.

	Mean	Std	Min	25%	50%	75%	Max
ρ_j	0.1649	0.0750	-0.0493	0.1166	0.1684	0.2139	0.4971
c_j	0.0008	0.0003	0.0000	0.0005	0.0008	0.0010	0.002

Methodology. In order to narrow down from general Tor usage to CSAM consumption via Tor, we follow a simple, yet effective approach: First, we estimate the global share of CSAM-related Tor traffic by combining three interlinked estimates: (i) According to Tor project (2023a), approx. 1.1% of global Tor traffic went to Onion services during our study period (i.e., March 16–May 31, 2019). We believe this number to be a conservative estimate for France as Jardine et al. (2020) report that in “free” countries—as which France classifies according to Freedom House—Tor is used more often to access onion services than in the rest of the world. Specifically, they estimate that approx. 7.8% of Tor users in free countries use Tor to access onion services vis-à-vis ~6.7% on a global level. (ii) Jin et al. (2023) collected 5,437,248 of these .onion-pages during the years 2020–2022 and observed that the category “Pornography” accounted for approx. 41.7% of the collected pages. The authors used the hidden service indexing website Ahmia.fi³ to collect seed addresses for crawling. On the one hand, since Ahmia.fi explicitly blacklists hidden services related to child abuse, we expect that CSAM sites are potentially under-sampled in this dataset (the blacklist contains 40,875 .onion sites as of August 2023). On the other hand, Cloudflare, a major content delivery network and domain name system service provider, allowed Tor browser users from September 2018 onwards to route some of their visits to clearnet websites via one of the ten .onion-addresses of Cloudflare. This could have potentially led to a one-sided increase of onion-traffic that may not have been fully captured by Jin et al. (2023). However, we cannot observe a substantial increase in the share of onion-traffic to overall Tor traffic between 2017 and the end of 2019 (Tor project, 2023a), thus we assume this to have a negligible effect on our approximation. (iii) Al-Nabki et al. (2019) further disaggregated the category ‘Pornography’ in their DUTA dataset and classified 41.5 % of .onion websites in this category to be related to CSAM specifically. Consequently, we conclude that approx. 0.19 % of global Tor download traffic is linked to the consumption of CSAM. However, commune-level CSAM consumption in France most likely deviates from global estimates. Thus, in order to locally adapt the global estimate to the 1341 French communes in our study, we use web service-level mobile traffic information from the Netmob dataset. Specifically, we approximate (ii) with the share of Tor traffic related to pornographic content by correlating the observed activity patterns for Web Adult and Tor for each of the 1341 French communes in our sample on an hourly basis across the whole time window of the study using Pearson’s ρ . The underlying assumption is that the consumption of pornographic content, irrespective of whether adults or children are depicted, follows similar temporal patterns. Thus, locations j with a higher temporal correlation are then assumed to have a larger fraction of their Tor traffic related to pornography in general, with $\rho_j = 1$ corresponding to 100% pornographic content. Figure 1 illustrates the composition of the estimate for global and France, respectively.

The 16.5% for France represents the mean of commune-level correlation coefficients ρ_j . To avoid non-sensible negative estimates of CSAM (in the following abbreviated as *cpc*) due to negative correlation coefficients, we replace them with small




Fig. 1 Composition of CSAM estimates for global and France. The diagram shows how CSAM consumption is estimated from Tor traffic using three analysis steps (i - iii) for both the global level and for France.

positive values near zero, denoting it with ρ'_j . We choose small non-zero replacements to avoid log transformed values going to infinity in later analysis. This affects 14 out of 1341 French communes with negligible effects on the overall distribution. Thus, our commune-level correction factor c_j is defined as $c_j = 0.011 \times 0.415 \times \rho'_j$. Table 1 shows the summary statistics of ρ_j and c_j .

Finally, we define our *cpc* estimates per 1000 inhabitants for all the $J = 1341$ French communes in our sample by

$$cpc_j = \frac{c_j \times Tor_j^{DL}}{pop_j} \times 1000, \tag{1}$$

where c_j denotes the correction factor as described above, Tor_j^{DL} the normalized download traffic related to Tor services and pop_j commune-level population counts. An average c of 0.0008 therefore can be interpreted as an estimated 0.08% of the observed Tor mobile download traffic in our sample of 20 French metropolitan areas being related to CSAM.

We consider this to be a conservative estimate of CSAM consumption via Tor for multiple reasons: First, the 41.5% refers to the share of pornographic .onion-sites that can be linked to CSAM. However, Owen and Savage (2015) found in 2015, that during the 6-month observation period, sites linked to sexual violence against children accounted for only 2 % of the hidden services screened in the study, but 82% of all requests made via Tor. Second, we assume that image-based content (such as CSAM) largely drives traffic. This assumption is backed by the fact that the top 5 web services in terms of download traffic in the Netmob dataset are predominantly image- or video-based (namely Instagram, Facebook, Netflix, YouTube, Facebook Live) (Martinez-Durive et al. 2023). Third, France is the fourth-largest host of online CSAM globally. Assuming a somehow positive relationship between hosting and consuming CSAM, this gives an indication of an overall larger share of CSAM consumption compared to the global average. This assumption is supported by the fact that across the years 2019 to 2022, on average five countries appeared in two Top 10 country lists in the same year, respectively: the “Top 10 countries hosting child sexual abuse URLs” list in the annual reports of the Internet Watch Foundation (cf. Internet Watch Foundation (2022)) and the “Top 10 countries by relay users” list of the Tor Metrics Project (cf. Tor project (2023b)). Lastly and importantly, as Insoll et al. (2021) found in a self-report survey ($N = 3620$) of CSAM users in




Fig. 2 Association of Google Trends search terms related to CSAM with their principal components for French regions pooled across the years 2017-2021. The colour illustrates the strength of the correlation between Google Trends search terms on the x-axis and their three principle components on the y-axis used in further analysis.

the darknet, CSAM is mainly consumed at home (44%), thus handled via Wifi or a fixed internet line. This gives indication that the correction factor for France for these internet connection types to be higher.

While directly validating our estimates with information on the actual commune-level consumption of CSAM in France is not possible due to the lack of ground truth data, we indirectly validate our findings by correlating the cpc estimates with an appropriate proxy indicator, in our case commune-level statistics on the number of victims of sexual violence (both adults and minors) per 1000 inhabitants for communes within our study area. Recalling the link between CSAM consumption and sexual violence against children indicated by Eke et al. (2011), Insoll et al. (2022) and Hall and Hall (2007) in Section “Introduction” and assuming that a non-negligible fraction of victims of sexual violence are minors, we expect our cpc estimates to show stronger correlations with our proxy than general mobile consumption patterns of e.g., YouTube. However, we stress that this proxy most likely just captures the tip of the iceberg of sexual child abuse: First, the indicator includes rape, attempted rape, and sexual assault including sexual harassment. However, somewhat surprisingly, it does not include sexual abuse, where abuse is distinguished from assault per definition as “it is carried out without violence, coercion or surprise” (Ministère de l’Intérieur et des Outre-Mer, 2022). Second, while official numbers report 39,314 victims (minors and adults) of sexual violence in France for the year 2019, CIIVISE (2021) estimates that 160,000 children alone become victims of sexual violence every year in France, as already noted above. Third, the indicator is only reported for those communes with at least five recorded incidences in three consecutive years in total. This statistical disclosure control measure clearly leads to a non-random selection of communes as large communes are more likely to surpass this threshold. Fourth, local variations in reporting behavior, especially in small communes with low overall reported numbers, may impact significantly the observed spatial patterns.

Since simple correlations in complex social settings most likely suffer from confounding factors, we build a hierarchical multi-level regression model in order to single out the influence on the number of reported cases of sexual violence that can be uniquely attributed to our cpc estimates, while controlling for a set of potentially relevant other socio-demographic and spatial features. To the best of our knowledge, this is the first attempt to look at large-scale local-level CSAM consumption from a spatial

epidemiology perspective. We note that this analysis is exploratory and the presented effects do neither imply a causal relationship nor the directionality of any observed relationship. To underline that both directions of influences are possible, we also present analysis results with our cpc estimates as dependent variable.

In addition, we explore points of interest in 0.1% tiles with the highest levels of estimated CSAM consumption. As some of these tiles are located in close proximity to each other, we remove duplicate entries by their unique place identifier. However, we noticed that some places in the OMF Places dataset may still be listed twice, e.g., in two different languages. Thus, duplicate entries may occur, however, we expect these to be negligible. Overture Maps Foundation classifies each POI into categories. We display only those POI categories with $n \geq 3$ in order to limit accidental occurrences on one hand and not to miss out on relevant, but rare categories on the other. To get an estimate for the average download traffic per POI category, we divide the observed download traffic by the number of POIs located for any given tile. In a second step, we average the download traffic across POIs for a given POI category. This leaves us with the average download traffic per POI category. While we acknowledge this to be a crude approximation for the actual traffic generated at a certain POI, we assume that POI categories across the large number of tiles observed are still indicative of existing spatial relationships.

Of the 18 search terms we extract from Google Trends, we discard seven due to complete sparsity. On the remaining 11 search terms, we perform a principal component analysis with varying number of components. We decided to go for three components by balancing the explained variance and the distinctiveness of the components based on visual inspection. Figure 2 shows how the search terms are associated with each of the three components.

Of the three components, we just consider the first (PC1) and the third (PC3) in further analysis as they appear to capture sexual preferences toward children more succinctly.

Results

Estimated CSAM consumption per 1000 inhabitants ranges from 0 in 14 communes to 157,077 in Mondouzil, Toulouse averaging 3703 across all areas between March 16 to May 31, 2019. As noted above, there is no actual unit attached to the traffic volume as it is normalized by the mobile phone operator. For comparison, YouTube download traffic per 1000 inhabitants averages




Fig. 3 Association of Google Trends search terms related to CSAM with their principal components for French regions pooled across the years 2017–2021. The heatmaps show activity patterns for the hour of the day (*x*-axis) by the day of the week (*y*-axis) for download traffic linked to YouTube (a), Web Adult (b), Tor (c) and the CPC estimates for the Top 10 communes (d). The darker the pattern, the higher the download traffic.

3,743,939,828 across all areas during the same time window, thus more than a million times the average Tor download traffic estimated to be related to CSAM. Commune-level results displayed in Fig. 1 in the Appendix.

While more fine-granular estimates, e.g., on the tile-level (100 m or census district (IRIS)-level), are technically possible, the share of census population estimates close to zero grows dramatically for small areas, thus rendering lower-level estimates per 1000 inhabitants increasingly volatile. Therefore, we opt to present commune-level estimates in this study. However, as we observe *mobile* internet traffic only, the locations of (i) the traffic

generation and (ii) the place of residence of the user do not necessarily coincide. Although we account for varying population sizes across communes, we observe that tile-level activity patterns are not necessarily propagated and visible on the commune-level. In other words, highly active tiles do not lead to highly active communes in terms of Tor download traffic, especially if these communes are large. This hints at spatially highly concentrated traffic generation. This argument is also supported when looking at Fig. 3, which shows the normalized download traffic for YouTube, web adult content, and Tor services summarized by weekday and hour across all cities in the sample.

As one might expect, all of the services analyzed show major peak traffic in the evening hours outside of regular business hours, thus hinting at the private entertainment purpose of these services. Download traffic from YouTube and adult content vary smoothly across the hours of the day with additional subtle peaks around 8am and 1pm during weekdays. CPC-related traffic in the 10 communes with the highest CSAM consumption estimates shows a stronger concentration of download activity in the evening hours compared to overall Tor download traffic. However, Tor-based traffic appears more coarse-grained in general. A potential explanation for the pixelated appearance of Tor-based download traffic is that Tor services saw approx. 2.5 million daily visitors globally in 2022 (Tor project, 2023c), while the general internet is used by approx. five billion users per day in 2022 (International Telecommunications Union, 2023). The Tor project estimates 100,537 mean daily Tor users for France during the time window of our study. Thus, it is likely that local-level Tor mobile download traffic via one mobile network operator is driven by a comparatively small subscriber base in our sample, so individual uses have a larger effect on the aggregate.

Validating estimates against official statistics on sexual violence. While direct validation of our methodology is hardly possible due to the lack of statistical data on CSAM consumption habits, we indirectly validate our findings by correlating our cpc estimates with commune-level statistics on the number of victims of sexual violence per 1000 inhabitants as described in Section “Methodology”. Looking back at Eke et al. (2011) and Hall and Hall (2007) in Section “Introduction” that link CSAM consumption and sexual violence, we expect our cpc estimate to indicate a positive association with the reported number of victims of sexual violence than general mobile consumption patterns. Table 2 shows the correlations of the number of victims of sexual violence with download traffic of YouTube, Web Adult, Tor, and cpc estimates, respectively, and whether these correlations are significantly different from zero.

In addition, we perform paired-samples tests for dependent correlation coefficients to check whether the correlation coefficient of our cpc estimates with the reported number of victims of sexual violence differs significantly from the other three web services. We see that the cpc estimates correlate significantly stronger with the number of victims of sexual violence (per 1000 inhabitants) than the other three web services (all three *p*-values < 1e-08). However, relying on correlations to investigate complex social phenomena is prone to confounding influences. Consequently, in further analysis, we link our commune-level cpc estimates to socio-demographic characteristics and other expectedly relevant spatial factors. To do so, we collect demographic data at the levels of communes, intercommunalities, and departments in France from the French statistical office INSEE including data on voting behavior during the 2017 French presidential election, and combine them with the number of certain POIs per 1000 inhabitants and sets of Google Trends

search terms related to CSAM. We chose the POIs based on the argument by Sauvé et al. (2021) that sexual violence against children mostly happens in places where a lot children are, e.g., at home, in schools or in sports clubs. Although child abuse and CSAM consumption may not happen at the same location, it is feasible to assume that offenders are in most cases not strangers to those places and likely live nearby, i.e., in the same commune. As CIIVISE (2021) states: In France, 8 out of 10 victims of child sexual abuse are victims of incest, in most cases committed by the older brother or father. Although both directions of the effect between our cpc estimates and the reported number of victims of sexual violence are plausible and supported by academic literature (cf. Section “Introduction”), we cannot determine the directionality of the relationship in our study design. Thus, we provide results for both directions by fitting one indicator on the other while controlling for a set of potential confounders using an ordinary least squares model with heteroscedasticity-robust standards errors. The results are presented in Table 3.

We observe that both the cpc estimates as well as the sexual violence indicator have a small, but positive and statistically significant impact on the respective outcome. Also, we see that the overall explained variance measured in (adjusted) *R*² is higher for cpc estimates than for the sexual violence indicator. This is expected as we control for download traffic of related web services. Interestingly, the effect of adult porn consumption (*log_Web_Adult_per_1000*) is negative, which hints at a subtle substitution effect: adult porn consumption in the clearnet is to some extent replaced by CSAM consumption in the darknet. Furthermore, we notice little consistency with regard to the direction, significance, and size of the observed effects of the control variables across the two regression setups. Together, this

Table 3 Regression results for our cpc estimates and the number of victims of sexual violence averaged across 2017–2021, by commune.

	Dependent variable (per 1000)	
	log_cpc	Sexual_violence
log_YouTube_per_1000	0.35*** (0.11)	0.04 (0.09)
log_Web_Adult_per_1000	-0.25** (0.10)	-0.09 (0.10)
log_Tor_per_1000	1.13*** (0.04)	-0.10 (0.08)
log_pop_density	0.14*** (0.03)	0.02 (0.08)
Share_of_singles	0.80 (0.66)	3.82** (1.75)
Poverty_rate	-5.21*** (1.17)	2.02 (1.58)
Employment_rate	-8.40*** (2.24)	0.92 (1.39)
Electoral_turnout_2017	-2.27** (0.96)	-1.12*** (0.40)
Share_Le_Pen	0.59 (1.02)	0.73 (0.57)
Share_Macron	2.72* (1.54)	0.05 (0.64)
POI_adult_entertainment	-0.003 (0.01)	0.02 (0.03)
POI_sports_teams	-0.001 (0.01)	0.03 (0.04)
POI_church	-0.01 (0.01)	-0.02 (0.02)
POI_mosque	-0.03** (0.01)	0.003 (0.02)
POI_religious_org	-0.06* (0.04)	0.02 (0.02)
POI_school	-0.02 (0.02)	0.07* (0.04)
PC1	-0.03* (0.02)	0.01 (0.01)
PC3	-0.01 (0.03)	-0.04** (0.02)
Sexual_violence_per_1000	0.14** (0.07)	
log_cpc_per_1000		0.10* (0.06)
Constant	-11.04*** (1.00)	1.26** (0.52)
Observations	630	630
R ²	0.96	0.41
Adjusted R ²	0.96	0.39

Note: Heteroscedasticity-robust SE.
**p* < 0.1.
***p* < 0.05.
****p* < 0.01.

Table 2 Correlation of the reported number of victims of sexual violence with log download traffic, per 1000 inhabitants at the commune-level and by web service.

Indicator	Spearman's ρ	p-Value
log_YouTube_per_1000	0.178	1.33e-06
log_Web_Adult_per_1000	0.161	1.18e-05
log_Tor_per_1000	0.237	7.94e-11
log_cpc_per_1000	0.284	4.99e-15

For *n* = 731 communes for which official numbers are available.

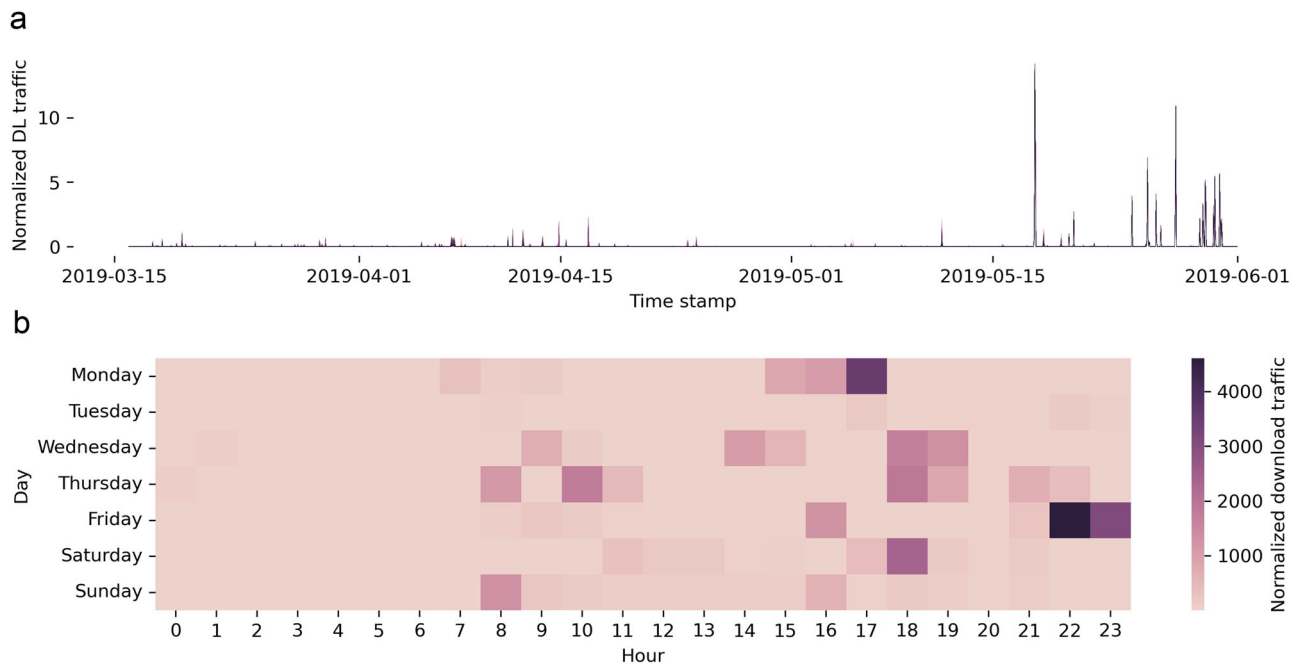


Fig. 4 Commune with the highest cpc estimate. The line plot (a) describes the hourly cpc estimates for each tile within the respective commune. The heatmap (b) shows the cpc estimates by weekday and hour.

hints at the fact that our cpc estimates and the sexual violence capture two distinct behaviors. While this could either support or undermine the validity of our estimate—we are able to single out the signal related to CSAM from the noisy sexual violence indicator vis-à-vis we measure some completely different Tor usage behavior—both the positive and significant association of the two indicators as noted above and the fact that we control for overall Tor download traffic supports the validity of our estimates. To investigate this further, we repeat the analysis for various specifications (see Appendix). We use the reported cases of drug abuse per 1000 inhabitants as a proxy for another presumably popular use of the darknet—ordering drugs. As Table 2 in the Appendix shows, the drug abuse rate does not inform our cpc estimates, giving further indication that we capture (child) porn-related consumption as we do not capture marketplace-related uses of Tor.

Further, we observe that the sexual violence indicator is zero for approx. half of the communes in our sample. Zero inflation may bias our parameter estimates as it hints at unmodelled factors causing the zeros in the first place. In Table 3 in the Appendix, we therefore exclude the communes with no reported cases to check for the impact of a zero-inflated setting. Overall, the significance of the observed effects is reduced which can be to some extent explained by the reduction in sample size, but significant effects do not show a change in sign or size. Lastly, it needs to be pointed out that the level of variation attached to these findings are most likely vastly underestimated, since the uncertainty involved in both the approximation of the correction factor as well as the underreporting of sexual abuse/violence cases is not accounted for, just to name a few. Also, we would like to stress, that this analysis does not, in any way, indicate that people of certain demographics participate in child abuse. Rather, our results should be interpreted as a first step into little-charted territory, namely looking at sexual child abuse via CSAM consumption from a spatial epidemiological perspective.

Investigating spatial relationships of child sexual abuse materials. We further investigate the spatial relationship of estimated

CSAM consumption with the local environment. In Fig. 4, we present the cpc estimates of the commune for which we estimate the highest CSAM consumption per 1000 inhabitants in our sample of 1341 communes. Tile-specific Tor download traffic is multiplied by the respective commune-level correction factor. The correction factor does not vary over time, but has been calculated for the whole time window of our study.

By looking at the timeline of Tor download traffic in this commune, Tor services appear to be used rather irregularly, as already mentioned before. Thus, we not only see a spatially, but also temporarily highly concentrated Tor usage. This would be in line with some common CSAM practices as described by Gannon et al. (2023), where CSAM is usually not streamed on-demand, but downloaded and consumed offline. While this may explain the “front-loaded” cpc download activity apparent in Fig. 3d when compared to adult porn download activity in Fig. 3b and therefore validates our main assumption that porn consumption follows the same temporal pattern, regardless whether adults or children are depicted, it lays open a caveat in it: porn consumption and consumption-related download traffic do not necessarily occur simultaneously, especially in the CSAM community. Based on the visual inspection of Fig. 3b,d, we determine the potential time lag between activity and assumed consumption to be around two hours. Consequently, we lag Tor traffic by 2 h and re-run both the calculation of the correction factor and the subsequent regression analysis. The lagged Tor traffic improves our pairwise correlation with the sexual violence indicator as reported in Table 2 from 0.28 to 0.34 as well as our regression analysis as presented in Table 4 of the Appendix. Even though the patterns observed via the day-of-week by hour-of-day heatmap does not indicate a generalizable usage pattern, one can clearly see that it does not align with regular business hours and therefore indicate private use. This argument is supported by the fact that the most active tiles within the commune displayed here are located in residential or rural neighborhoods as visual inspection of the respective tile locations on Google Earth shows.

By looking not only at the top 10 communes with the highest estimated CSAM consumption, but at the 0.1% of all tiles with the

Table 4 Points of interest in the 0.1 % of tiles with the highest cpc traffic, by web service.

Rank	CPC	Web Adult	YouTube
1	embassy (37)	jazz_and_blues (3)	tennis_court (3)
2	lake (4)	jail_and_prison (3)	attractions_and_activities (17)
3	contract_law (5)	b2b_services (3)	aquarium (3)
4	swimming_pool (9)	roofing (3)	boat_tours (3)
5	eye_care_clinic (8)	agricultural_cooperatives (3)	mining (3)

Number of POIs in parentheses.

highest download traffic ($n = 5259$) for the three different web services in our study, we observe distinct sets of adjacent points of interest (POIs) as shown in Table 4.

Although one could think of plausible explanations for some of the POIs in Table 4 (e.g., concerning the use of web services related to adult pornography around prisons or the use of YouTube at tourist attractions), drawing more general spatial relationships from Table 4 appears challenging, especially for our CPC estimates. For example, it is unclear whether Tor plays an important role in fulfilling diplomatic duties or whether these high levels of Tor mobile download traffic are simply a geographic coincidence. An argument against the latter is that this coincidence is not limited to one larger diplomatic area, but occurs across several cities in France. A detailed look at the POI locations for our CPC estimates reveals that many of the POIs across the mentioned categories are located around Porte de Passy, which surrounding area represents the largest Tor download traffic hotspot in our sample of 20 urban areas in France. However, most of the CPC-related traffic in the corresponding commune is generated throughout the study period at the end or outside of regular office hours.

Noticeable is that many of the identified POIs are located in densely populated areas. One explanation for that is that we look at total traffic on the tile-level as tile-level population statistics are on one hand not readily available and on the other hand potentially misleading, especially in tourist areas. Interestingly, a closer look at the actual POI locations also reveals generally fewer POI locations in the OMF Places dataset vis-à-vis Google Maps.

Importantly, it needs to be stressed here that just because traffic is generated in close proximity to these places, it does not mean that this traffic is generated by the inhabitants, owners, or employees themselves, but by any subscriber near the location. Related to the well-known concept of ecological fallacy, area-level correlations do not necessarily imply individual or POI-level causal relationships. As an example, while prostitution occurs mainly in poorer neighborhoods, the clients are not necessarily the poor locals.

Discussion

In this study, we shed light on a topic usually hidden in the dark from a novel angle: We looked at spatial patterns in the consumption of child sexual abuse material using mobile network data for 1341 small areas across 20 metropolitan areas in France for 77 consecutive days in 2019. To the best of our knowledge, this is the first time that spatial CSAM consumption patterns have been mapped at such a high geographical detail. Validated against the reported numbers of victims of sexual violence at the commune-level, we further explored geographic links to both local socio-demographic characteristics as well as to nearby points of interest and Google search queries. These insights may contribute to a better understanding of the whereabouts of CSAM consumption and thus inform targeting public awareness campaigns such as the one launched in September 2023 by the French government (Le Monde with AFP, 2023). While some of our findings appear to

echo existing literature—for example, we find that higher unemployment levels are associated with higher CSAM consumption—some findings also appear to contradict previous findings, i.e., higher poverty levels are associated with lower CSAM consumption. However, it is important to address the limitations inherent to this study: First, the study analyzes mobile network traffic from one major mobile network provider only; hence, it misses out on web traffic generated both via Wifi or fixed internet connections or via other mobile network operators. Structural differences between mobile-only and overall traffic, especially when it comes down to the consumption of (child) pornography, need to be expected but cannot be further quantified in this study. Second, our estimates build on assumptions as laid out in Section “Methodology”, since detailed information concerning the specific origin of the observed Tor download traffic is not available. While we try to support the assumptions with evidence, they may not hold to a full extent, especially on local levels, as the sample size of actual Tor users generating the observed traffic might be very small. Third, linking consumption patterns with local phenomena such as socio-demographic characteristics or points of interest is subject to additional uncertainty as the mobile traffic is assumed to be generated partly out-of-home, i.e., not exclusively by the inhabitants of that area, but potentially by any visitor. Therefore, relationships observed on the area-level may not hold on the individual-level. Fourth, the sourcing of the POI information is not described in detail by the data provider and, therefore, may be prone to certain selection biases. Especially as residential homes are usually not counted as a point of interest, information on these might be underrepresented or captured indirectly by POIs prevalent in residential areas such as schools. Fifth, our groundtruth indicator, i.e., the reported number of victims of sexual violence, is imperfect in many ways as laid throughout the study, but—to the best of our knowledge—the most suitable proxy for child sexual abuse in France on local levels. Consequently, our CSAM-related consumption estimates need to be considered with caution, especially at the local level. Sixth, while previous studies focused on personality traits such as mental health problems and past experiences to explain CSAM consumption, this data is mostly non-existent on a larger-scale. Thus, we expect that relevant factors explaining CSAM consumption are not sufficiently considered in our regression analysis.

As described in the Netmob dataset description, data collection, processing and aggregation took place in compliance with GDPR under the supervision of the Data Protection Officer of the mobile network operator Orange (Martínez-Durive et al. 2023). Individual-level traffic has been aggregated to 15-minute intervals and spatially distributed across a network coverage grid. Furthermore, the study authors refrain from any detailed depiction of small areas, e.g., presenting geographic coordinates for single tiles that could put people or businesses at risk of being accused of wrongdoing. In addition, we tried to add flags of caution throughout the study to avoid that individual figures or paragraphs can be misinterpreted when taken out of context.

Going ahead, we see multiple ways how this research can be extended: First, the regression could benefit from additional indicators that capture attitudes, behaviors, and opinions in a more nuanced way. This is of particular importance for deriving policy implications from our work. Second, we have not found any major external shock such as take-downs of large CSAM forums in the darknet during the time window of analysis. Re-running the analysis around such an event may provide further insights into the agility and resilience of the community to external interventions. Third, in a related manner, temporal information on forum activities may help to link specific forum activities (e.g., release of a new curated CSAM collection) with traffic patterns. Fourth, extending the analysis to fixed internet connections may allow to capture the full extent of CSAM consumption online and help to quantify the bias induced by observing mobile traffic only. This would also allow to investigate the supply side of the CSAM market more rigorously, namely the upload traffic. Lastly, we hope that the release of the Netmob dataset will strike a precedent for other mobile network operators and internet service providers to provide web service-level network traffic information to researchers in an ethical manner. While the internet has fundamentally transformed the way we behave and communicate, it is still little known how it is actually used in everyday life. Consequently, more such data releases would facilitate research not only on the darknet, but across a wide range of disciplines.

In conclusion, we believe that our study sheds light on the consumption of CSAM from a novel angle using so far little-tapped data source – large-scale web service mobile traffic. In that way, we hope that our study can help in better understanding the spatial relationship between CSAM consumption and child sexual abuse and ultimately help to move forward on target 16.2 of the Sustainable Development Goals: “End abuse, exploitation, trafficking and all forms of violence and torture against children”.

Data availability

• Access to the Netmob dataset can be requested under a licensing agreement via the Netmob Data Challenge website <http://sistemas.inec.cr/pad5/index.php/catalog/113> or by contacting the organizing committee at IMDEA via netmob2023challenge@imdea.org. • Overture Maps Foundation data is openly accessible via <https://overturemaps.org/download/>. Code to query the data and retrieve the POI information used in this study are available from the GitHub repository accompanying this study as described below. • Google Trends data is openly accessible via <https://trends.google.com/trends/>. Code to replicate the queries can be found in the accompanying Github repository. • Socio-demographic data is openly accessible via <https://www.insee.fr/en/outil-interactif/5543645/>. • Voting data is openly accessible via <https://www.data.gouv.fr/fr/datasets/election-presidentielle-des-23-avril-et-7-mai-2017-resultats-definitifs-du-1er-tour-par-communes/>. • Crime data is openly accessible via <https://www.data.gouv.fr/fr/datasets/r/3f51212c-f7d2-4aec-b899-06be6cdd1030>. • The code to reproduce the study is available on GitHub: https://github.com/Societal-Computing/netmob_tor.

Received: 6 March 2024; Accepted: 13 June 2024;

Published online: 21 June 2024

Notes

- <https://www.torproject.org/>.
- <https://trends.google.com/trends/>.
- <https://ahmia.fi/>.

References

- Al-Nabki MW, Fidalgo E, Alegre E, Fernández-Robles L (2019) Torank: identifying the most influential suspicious domains in the tor network. *Expert Syst Appl* 123:212–226
- Assemblée Nationale Amendment no. II-1301. Retrieved on 2023-09-20 (2022)
- Babchishin KM, Hanson RK, VanZuylen H (2015) Online child pornography offenders are different: a meta-analysis of the characteristics of online and offline sex offenders against children. *Arch Sex Behav* 44:45–66
- Babchishin KM, Karl Hanson R, Hermann CA (2011) The characteristics of online sex offenders: a meta-analysis. *Sex Abus* 23:92–123
- Brown E, Scodellaro C (2023) Introduction. les violences envers les populations vulnérables: des réciprocity complexes. *Popul Vulnérables* (9)
- Bruckschen F., Koebe T., Ludolph M., Marino M. F., Schmid T (2019) Refugees in undeclared employment—a case study in turkey. *Guide to Mobile Data Analytics in Refugee Scenarios: TheData for Refugees ChallengeStudy*, 329–346
- Chetty R, Jackson MO, Kuchler T, Stroebel J, Hendren N, Fluegge RB, Gong S, Gonzalez F, Grondin A, Jacob M (2022) Social capital II: determinants of economic connectedness. *Nature* 608:122–134
- CIIVISE (2021) Violences sexuelles: protéger les enfants. conclusions intermédiaires. Technical report, Commission indépendante sur l’inceste et les violences sexuelles faites aux enfants
- Ministère de l’Intérieur et des (2022) Outre-Mer Service statistique ministériel de la sécurité intérieure—base des séries chronologiques. <https://www.data.gouv.fr/fr/datasets/service-statistique-ministeriel-de-la-securite-interieure-base-des-series-chronologiques/>
- Deutsche Welle (2019) German police smash darknet child porn ring. *Deutsche Welle* <https://www.dw.com/en/german-police-smash-child-porn-ring-on-darknet/a-50043833>
- Eke AW, Seto MC, Williams J (2011) Examining the criminal history and future offending of child pornography offenders: an extended prospective follow-up study. *Law Hum Behav* 35:466–478
- Elliott IA, Beech AR, Mandeville-Norden R (2013) The psychological profiles of internet, contact, and mixed internet/contact sex offenders. *Sex Abus* 25:3–20
- Fortin F, Proulx J (2019) Sexual interests of child sexual exploitation material (csem) consumers: four patterns of severity over time. *Int J offender Ther Comp Criminol* 63:55–76
- Gannon C, Blokland, AAJ, Huikuri, S, Babchishin, KM, and Lehmann, RJB Child sexual abuse material on the darknet. *Forensische Psychiatrie, Psychologie, Kriminologie*, pages 1–13. Company: Springer Distributor: Springer Institution: Springer Label: Springer Publisher: Springer Berlin Heidelberg (2023)
- Hall RCW, Hall RCW (2007) A profile of pedophilia: definition, characteristics of offenders, recidivism, treatment outcomes, and forensic issues. *Mayo Clin Proc* 82:457–471
- INSEE Populations légales 2019. Data retrieved on 2023-09-19 (2019)
- Insoll T, Ovaska A, Vaaranen-Valkonen N (2021) csam users in the dark web: protecting children through prevention. *Suojellan Lapsia ry (Protect Children)*. Redirection Survey Report
- Insoll T, Ovaska AK, Nurmi J, Aaltonen M, Vaaranen-Valkonen N (2022) Risk factors for child sexual abuse material users contacting children online: results of an anonymous multilingual survey on the dark web. *J Online Trust Saf* 1(2)
- International Telecommunications Union Internet use (2023) https://public.tableau.com/views/ITUfactsandfigures2022/InternetUse01?:embed=y&:display_count=n&:origin=viz_share_link
- Internet Watch Foundation Annual report (2022)
- Jardine E, Lindner AM, Owenson G (2020) The potential harms of the tor anonymity network cluster disproportionately in free countries. *Proc Natl Acad Sci USA* 117:31716–31721
- Jin Y, Jang E, Cui J, Chung J-W, Lee Y, Shin S (2023) DarkBERT: a language model for the dark side of the Internet. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 7515–7533, Toronto, Canada. Association for Computational Linguistics
- Koebe T (2020) Better coverage, better outcomes? mapping mobile network data to official statistics using satellite imagery and radio propagation modelling. *PloS ONE* 15:e0241981
- Le Monde with AFP France targets incest for first time in national campaign (2023) *Le Monde*. https://www.lemonde.fr/en/france/article/2023/09/12/france-targets-incest-for-first-time-in-national-campaign_6133564_7.html
- Macedo J, Costa F, dos Santos JA (2018) A benchmark methodology for child pornography detection. In 2018 31st SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI), p 455–462
- Martínez-Durive OE, Mishra S, Ziemlicki C, Rubrichi S, Smoreda Z, Fiore M (2023) The netmob23 dataset: a high-resolution multi-region service-level mobile data traffic cartography. Preprint at *arXiv* <https://doi.org/10.48550/arXiv.2305.06933>
- Ministère de l’Intérieur et des Outre-Mer Election présidentielle des 23 avril et 7 mai 2017—résultats définitifs du 1er tour par communes (2017) <https://www.data.gouv.fr/fr/datasets/r/77ed6b2f-c48f-4037-8479-50af74fa5c7a>

- Nurmi J, Paju A, Brumley BB, Insoll T, Ovaska AK, Soloveva V, Vaaranen-Valkonen N, Aaltonen M, Arroyo D (2024) Investigating child sexual abuse material availability, searches, and users on the anonymous tor network for a public health intervention strategy. *Sci Rep* 14:7849
- Overture Maps Foundation Overture places. Release 2023-07-26-alpha.0 (2023)
- Owen G., Savage, N. The tor dark net (2015)
- Owens JN, Clapp K, Craun SW, van der Bruggen M, van Balen I, van Bunningen A, Talens P (2022) Analysis of topic popularity within a child sexual exploitation tor hidden service. *Aggress Violent Behav* 101808
- Pinheiro PS (2006) Violence against children. ATAR Roto Presse SA, Geneva
- Price M, Lambie I, Krynen AM (2015) New zealand adult internet child pornography offenders. *J Crim Psychol* 5:262–278
- Rotondi V, Kashyap R, Pesando LM, Spinelli S, Billari FC (2020) Leveraging mobile phones to attain sustainable development. *Proc Natl Acad Sci USA* 117:13413–13420
- Sae-Bae N, Sun X, Sencar HT, Memon ND (2014) Towards automatic detection of child pornography. In 2014 IEEE International Conference on Image Processing (ICIP), p. 5332–5336
- Sauvé J-M, Atlani-Duault L, Bajos N, Baubert T, Beloucif S, Burguburu J-M, Casagrande A, Cordier A, Damiani C, Devreese A (2021) Les violences sexuelles dans l'église catholique, france 1950-2020. rapport de la commission indépendante sur les abus sexuels dans l'église. Technical report, Commission Indépendante sur les Abus Sexuels dans l'Eglise
- Seigfried KC, Lovely RW, Rogers MK (2008) Self-reported online child pornography behavior: a psychological analysis. *Int J Cyber Criminol* 2(1)
- Seigfried-Spellar KC, Rogers MK (2010) Low neuroticism and high hedonistic traits for female internet child pornography consumers. *Cyberpsychol Behav Soc Netw* 13:629–635
- Seto MC (2013) Internet sex offenders. American Psychological Association
- Terminology and Semantics Interagency Working Group on Sexual Exploitation of Children Terminology guidelines for the protection of children from sexual exploitation and sexual abuse. ECPAT International (2016)
- Tor project Tor metrics (2023a) <https://metrics.torproject.org/hidserv-rend-relayed-cells.html?start=2019-03-16&end=2019-05-31>
- Tor project Tor metrics (2023b) <https://metrics.torproject.org/userstats-relay-table.html?start=2022-01-01&end=2022-12-31>
- Tor project Tor metrics (2023c) <https://metrics.torproject.org/userstats-relay-country.html?start=2022-01-01&end=2022-12-31&country=all&events=off>
- van der Bruggen M, van Balen I, van Bunningen A, Talens P, Owens JN, Clapp K (2022) Even “lurkers” download: The behavior and illegal activities of members on a child sexual exploitation tor hidden service. *Aggress Violent Behav* 67:101793
- Vitorino P, Avila S, Perez M, Rocha A (2017) Leveraging deep neural networks to fight child pornography in the age of social media. *J Vis Commun Image Represent* 50:303–313

Author contributions

BN, IW, NS, TK, and ZD conceived the presented idea and developed the methodology. BN, NS, and TK performed the computations. TK wrote the paper with support of BN, NS, and ZD. IW and ZD helped supervise the project.

Funding

Open Access funding enabled and organized by Projekt DEAL.

Competing interests

The authors declare no competing interests.

Ethical approval

As stated in the original dataset description (Martínez-Durive et al. 2023): “The mobile network traffic dataset we use to generate the dataset was collected, processed and aggregated as described in Section “Methodology” in full compliance with Article 89 of the General Data Protection Regulation (GDPR), under the supervision of the Data Protection Officer (DPO) at Orange. In particular, all data management was performed on a secure platform at the operator’s premises and the raw data was deleted immediately afterward. The resulting service-level time series represent traffic aggregated over all UEs both in space, at eNodeB level, and time, over 15-min intervals. Moreover, the traffic associated to different base stations is further aggregated via the spatial mapping described earlier. The final representation does not allow re-identifying or tracking individual users. Therefore, this article does not contain any studies with human participants performed by any of the authors.

Informed consent

In order to comply with the General Data Protection Regulation in the European Union, the mobile phone operator is required to use data from only those subscribers that opted-in for having their network data analyzed for research purposes. As stated in the original dataset description (Martínez-Durive et al. 2023): “The mobile network traffic dataset we use to generate the dataset was collected, processed and aggregated as described in Section “Methodology” in full compliance with Article 89 of the General Data Protection Regulation (GDPR), under the supervision of the Data Protection Officer (DPO) at Orange. In particular, all data management was performed on a secure platform at the operator’s premises and the raw data was deleted immediately afterward. The resulting service-level time series represent traffic aggregated over all UEs both in space, at eNodeB level, and time, over 15-minute intervals. Moreover, the traffic associated to different base stations is further aggregated via the spatial mapping described earlier. The final representation does not allow re-identifying or tracking individual users.”

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1057/s41599-024-03343-4>.

Correspondence and requests for materials should be addressed to Till Koebe.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024