LOCARD

DELIVERABLE

D4.2 State of the art report on detection of deviant behaviour in Social Networks

Project Acronym:	LOCARD	
Project title:	Lawful evidence collecting and continuity platform development	
Grant Agreement No.	832735	
Website:	http://locard.eu/	
Contact:	info@locard.eu	
Version:	1.0	
Date:	23 October 2019	
Responsible Partner:	ARC	
Contributing Partners:	ARC, KEMEA, TUB	
Reviewers:	Miltiadis Anastasiadis (Motivian) UNIPD Team	
Dissemination Level:	Public	X
	Confidential – only consortium members and European Commission Services	

This project has received financial support from the European Union Horizon 2020 Programme under grant agreement no. 832735.

Revision History

Revision	Date	Author	Organization	Description
0.1	12/07/2019	Nikolaos Lykousas	ARC	Tentative ToC/structure
0.2	07/08/2019	Nikolaos Lykousas	ARC	Version for review
1.0	23/10/2019	Nikolaos Lykousas	ARC	Final version

Every effort has been made to ensure that all statements and information contained herein are accurate, however the LOCARD Project Partners accept no liability for any error or omission in the same.

Table of Contents

1 Executive Summary

In this deliverable, we discuss the state-of-the-art approaches on detection of deviant behaviour in Social Networks, focusing on cybercriminal behaviours related to production/distribution of child sexual abuse media (CSAM) and predatory behaviours targeting children. Moreover, we identify existing gaps in the available data resources and in proposed computational approaches for identifying such online behaviours.

2 Introduction

On-line social networks (OSNs) contain complex forms of inter-linked communities that interact and communicate on different topics. Some communities are characterized by what are usually referred to as deviant behaviours, conducts that are commonly considered inappropriate with respect to the society's norms or moral standards [1]. Cybercrime, eating disorders, drug use, and adult content production/consumption are just a few examples. Many of them are represented, to different extents, on social media. However, since all these topics touch upon different societal taboos, the common-sense assumption is that they are embodied either in niche, isolated social groups, or in communities that might be quite numerous but whose activity runs separately from the mainstream social media life. In line with this belief, research has mostly considered those groups in isolation, focusing predominantly on the patterns of communications among community members [2] or, from a sociological perspective, on the motivations to that make people join such groups [3].

In reality, people who are involved in deviant practices are not segregated outcasts but are part of the fabric of the global society. As such, they can be members of multiple communities and interact with very diverse people, possibly exposing their deviant behaviour to the public. Interestingly enough, these practices are more common that one would expect. For instance, a recent survey organised by the EU Kids Online project showed that children are especially vulnerable to the risks of using the Internet: 9- to 16-year-olds spend 88 minutes a day online on average, with 49% of these adolescents going online in their bedroom — away from adult supervision. Moreover, the study found that 34% of the children who took part in the survey had added people to their social media friends lists they had never met face-to-face, 15% had sent personal information, pictures or videos of themselves to strangers and 9% had agreed to meet in person with someone they had only met online [5]. Based on a study by the European Parliament [6], four key risk factors can be identified that specifically relate to children and the Internet: i) a child can be exposed to harmful and illegal content; ii) social media platforms can be utilised to solicit children, which is often the prelude for child sexual exploitation; iii) children can become engaged in illegal activities, such as disseminating child abuse media or coercing other children into becoming victims; and iv) the over-abundance of virtual communities that enable offenders to remain anonymous or to create false virtual identities facilitates the production and dissemination of child sexual abuse media (CSAM).

3 Background on CSAM

In the context of LOCARD project, we especially focus on the detection of cybercriminal behaviours in the context of social media related to the distribution of CSAM as well as predatory behaviours. In this regard, Colleto et al. [4] aimed at going beyond previous studies that looked at deviant groups in isolation by observing them in context. In particular, they attempted to answer questions relevant to the deviant behaviours related with pornographic material in the social media context, such as i) how much deviant groups are structurally secluded from the rest of the social network, and what are the characteristics of their subgroups who build ties with the external world; ii) how the content produced by a deviant community spreads and what is the entity of the diffusion which reaches users outside the boundaries of the deviant community who voluntarily or inadvertently access the adult content; and iii) what is the demographic composition of producers and consumers of deviant content and what is the potential risk that young boys and girls are exposed to it.

The proliferation of the Internet has transformed child sexual abuse into a crime without geographical boundaries. Child sex offenders turning to the Internet as a means of creating and distributing child pornography has allowed the creation of a network of support groups for child sex offenders, when historically, this was an offense that occurred in isolation [7]. This concern was echoed by Mitchell et al. [8] who recognized that a small percentage of offenders used social networking sites (SNS) to distribute child pornography. While there is scientific debate on whether the online predator is a new type of child sex offender [9] or if those with a predisposition to offend are responding to the opportunities afforded by the new forms of social media [10], empirical evidence points to the problem of Internet based paedophilia as endemic. Recent work shows that nearly half of the offenders who had committed one or more contact offences, i.e., they had directly and physically abused children, had displayed so-called “grooming behaviour” [11],[12]. Grooming refers to the process by which an offender prepares a victim for sexual abusive behaviour:

[Grooming is] a process by which a person prepares a child, significant others, and the environment for the abuse of this child. Specific goals include gaining access to the child, gaining the child’s compliance, and maintaining the child’s secrecy to avoid disclosure. This process serves to strengthen the offender’s abusive pattern, as it may be used as a means of justifying or denying their actions [13]

However, when investigating the possibility to develop automated methods to detect grooming online, researchers are confronted with a number of issues. First, there is only one benchmark dataset available that contains (English) chat conversations written by child sex offenders: t*he PAN 2012 Sexual Predator Identification dataset,* which leverages data from the *Perverted Justice website*[1] (PJ). Concretely, PJ data comprises a single class of chats in the context of PAN 2012 data. Yet, because the victims were actually adult volunteers posing as children, it is likely that these conversations are not entirely representative for online predator-victim communications [14]. Moreover, since the seduction stage often shows similar characteristics with adults’ or teenagers’ flirting, initial studies trying to detect predatory behaviour directly on the user level typically resulted in numerous false positives when they were applied to non-predatory sexually oriented chat conversations in the PAN 2012 dataset [15]. In addition, the increasing amount of child sexual abuse media (CSAM) being shared across borders and with apparent impunity leads to new children being found online every day. Each of these children, often from within the family circle of the offender, is a victim of child sexual abuse [6].

The severity of the problem has already resulted in a number of solutions devoted to monitor such activity. The Child Protection System (CPS) [16] and RoundUp [17], [18], for example, are able to capture data about child sex offender activity and identify CSAM across different peer-to-peer protocols. However, these tools rely on matching the files shared on a network against a hash-value database of known CSAM[2]. As a result, they retrieve thousands of files that have been circulating for several months or even years, but they are not able to identify new CSAM when they are being released on to a network. Nor are they able to detect child sexual abuse media that are not on record.

It is worth to emphasise that one of the main priorities for law enforcement is to identify cases where an offender is actively engaged in the production of new CSAM — they can be indicators of recent or on-going

[1] http://perverted-justice.com/

[2] Such databases are built through post-hoc forensic analysis of seized computers of offenders.

child abuse. Nevertheless, detecting new (or previously unknown) child sexual abuse media is highly challenging, because distributors of CSAM tend to obfuscate the illegal content of their shared files. More specifically, they use specific and trained vocabulary, which contains a variety of keywords, abbreviations, acronyms and even combinations of different languages to avoid (automatic) detection of such files, while making them widely searchable for other offenders. Moreover, this vocabulary proved to be dynamic, i.e. it evolves as existing keywords come to the attention of law enforcement [19]. Finally, detecting new or previously unknown CSAM requires (semi-)automatic analysis of image and video content. However, downloading of all candidate files for automatic image and video analysis is clearly infeasible in, for example, a P2P scenario. Hence, such an approach also requires an intermediate step to reduce the number of candidate files to be downloaded.

4 Methods and alignment with LOCARD Requirements

Although a range of awareness campaigns have already been organized internationally (e.g., the EU Safer Internet Programme[3] and Insafe[4]), only few resources have been employed to investigate novel automated methods to support law enforcement agencies or social network moderators when trying to identify online child sex offenders, in the context of social networks. Additionally, due to both the illicit nature of this topic and the privacy issues that are involved, there is only one website that displays predator-victim chat room conversations, the PJ, which contains over 500 English chat conversations between adult volunteers pretending to be adolescents and as such were approached by an alleged child sex offender.

4.1 Machine Learning approaches and related chat corpora

For machine learning algorithms to be effective in identifying online sexual predators, they need to be trained with both illegal conversations between offenders and their victims and sexually oriented conversations between consenting adults [14]. Since such data are rarely made public, initial studies [14], [20] only experimented with data from PJ. The k-NN classification experiments based on word token n-grams performed by [14] achieved up to 93.4% F-score when identifying the predators from the pseudo-victims. The authors of [21] were the first to include additional corpora in the non-predator class: they included 85 conversations containing adult descriptions of sexual fantasies and 107 general non-offensive chat logs from websites like http://www.fugly.com and http://chatdump.com. When distinguishing between 200 PJ conversations and these additional chat logs, the Naïve Bayes classifier outperformed the Decision Tree and the Regression classifier, which resulted in an F-score of 91.7% for the PJ class. In [22], authors used a corpus of cybersex chat logs and the Naval Postgraduate School (NPS) chat corpus and experimented with new feature types such as emotional markers, emoticons and imperative sentences and computed sex-related lexical chains to automatically detect offenders directly in the PJ dataset. Their Naïve Bayes classifier yielded an accuracy of 92% for PJ predators vs. NPS and 94% for PJ predators vs. cybersex based on their high-level features. However, both [21] and [22] did not filter out any cues that were typical of the social media platforms from which the additional corpora were extracted, which could entail that their models were (to some degree) trained on detecting these cues rather than the grooming content. Moreover, because the high-level features described by [22] were (partially) derived from the PJ dataset itself, these experiments may have resulted in overestimated accuracy when detecting predators from the same dataset.
Recently, the detection of Internet child sex offenders has been extensively investigated in the framework of the PAN 2012 competition, during which efforts have been made to pair the PJ data with a whole range of non-predatory data, including cybersex conversations between adults [15]. Because the PAN 2012

[3] https://ec.europa.eu/digital-single-market/en/content/creating-better-internet-kids-0
[4] https://www.betterinternetforkids.eu/web/portal/policy/insafe-inhope

benchmark dataset was heavily skewed towards the non-predatory class, most participants applied a two-stage classification framework in which they combined information on the conversation level to the user level (e.g. [23]–[26]). Moreover, apart from one submission that used character-gram features, all other studies used (combinations of) lexical (e.g., token unigrams) and “behavioural” features (e.g., the frequency of turn-taking or the number of questions asked). The best results were achieved by [25], who used a Neural Network classifier combined with a binary weighting scheme in a two-stage approach to first identify the suspicious conversations and, secondly, to distinguish between the predator and the victim. Their system achieved an F-score of 87.3%. However, during their study, they assumed that “predators usually apply the same course of conduct pattern when they are approaching a child” [25], which is in contrast with research by [27], which resulted in three different types of predators and, hence, of grooming approaches. Moreover, the PAN2012 dataset was also not cleansed of platform specific cues, which could again have led to overestimated F-scores during the competition. A more detailed overview of the results of the PAN 2012 International Sexual Predator Identification Competition can be found in [15].

4.2 Grooming detection

With regard to the content of predatory chat conversations, the authors of [20] were the first to investigate the possibility to automatically detect different stages in the grooming process. Based on an expanded dictionary of terms they applied a rule-based approach, which categorised a post as belonging to the stage of gaining personal information, grooming (which included lowering inhibitions or re-framing and sexual references), or none. Their rule-based approach outperformed the machine learning algorithms they tested and reached up to 75.1% accuracy when categorising posts from the PJ dataset into one of these stages. A similar approach was used by [28], whose Naïve Bayes classifier achieved a 96% accuracy when categorising predatory PJ posts as belonging to either the gaining access, the deceptive relationship or the sexual affair grooming stage. The second task of the PAN 2012 competition consisted of detecting the specific posts that were most typical of predatory behaviour from the users that were labelled suspicious during the first task. To this end, most participants either created a dictionary-based filter containing suspicious terms [25], [29] or used their post-level predictions from the predator identification task [26], [30]. The best F-score was achieved by [24], who used a dictionary-based filter highlighting the utterances that referred to one of the following grooming stages: sexual stage, re-framing, approach, requests for data, isolation from adult supervision and age- and child-related references. Their approach resulted in a 35.8% precision, a 26.1% recall and a 30.2% F-score. Finally, [31] proposed a method based on Temporal Concept Analysis using Temporal Relational Semantic Systems, conceptual scaling and nested line diagrams to analyse PJ chat conversations. Their transition diagrams of predatory chat conversations seemed to be useful for measuring the level of threat each offender poses to his victim based on the presence of the different grooming stages. Although these studies showed promising results, the issue remains that these methods are applied on a corpus that contains conversations between offenders and pseudo-victims. Hence, the adult volunteers that were posing as children could not accede to requests for “cammin”, sending pictures, etc. As a result, the PJ dataset contains hardly any conversations by groomers, because this type of offender typically does not invest much time in the seduction process and switches to a different victim when his needs are not fulfilled quickly. Moreover, it is highly likely that children would have responded differently to the grooming utterances than the adult volunteers did, which could have influenced the language use of the offenders.

4.3 Detecting Child Sexual Abuse Media

As detecting known CSAM files is relatively straightforward when a hash-value database is available, initial work in this area mainly focused on the ability to disrupt online child exploitation [32], reliability issues regarding mutable identifiers, such as IP addresses and GUIDs [17], [33] and the identification of key sharers [34]. This has already resulted in a number of tools, such as CPS and RoundUp, that can not only monitor such paedophile activities in P2P networks, but also provide additional features such as geolocation capabilities and centralised databases to assist law enforcement in their international struggle against online child exploitation. Moreover, Internet companies such as Google and Microsoft have created software, such as PhotoDNA [35], that enables law enforcement to detect modified versions of known CSAM. Nevertheless, none of these tools offers support for identifying new or previously unknown child abuse media. So far, only few attempts have been made to address this issue, with focus on P2P networks [36]–[41], which were plagued by child pornography sharing behaviours. The authors of [42] demonstrated that collaborative filtering techniques that are typically used in recommender systems, can be successfully applied to identify new media in P2P networks of a certain category (e.g., pornography, piracy software and popular music). Their method is based on the assumption that file-sharing traffic tends to cluster around interest, especially when it involves illegal content, such as CSAM files. Hence, they were able to detect previously unknown examples from these categories without analysing their contents or filenames. Secondly, the MAPAP project [19] specifically targets peer-to-peer file-sharing networks. There, modelling of user activity and identification of CSAM-related keywords is utilised to identify child abuse media. However, the first system was not tested on verified CSAM data and the latter was not evaluated for the scenario of identifying new or previously unseen CSAM files. Finally, to the best of our knowledge, there are only two studies that used language analysis techniques to identify CSAM. As mentioned before, the authors of [19] investigated the feasibility to automatically construct lists of CSAM-related keywords. The second study [43] examined whether techniques used for SMS normalisation could also be used to circumvent the issue of language variation or noise in CSAM filenames.

4.4 Grooming and CSAM Detection in Social Media Interactions

Both police investigators and social media moderators have a limited amount of time and resources. Hence, they would benefit from a system that presents them with a reduced set of possibly suspicious users, which translates into a high-precision system. On the other hand, it is essential that the system does not miss any potential offenders, and hence, that the recall remains high. So far, prior work has attempted to detect predatory behaviour in social networks directly on the user level, which led to a large number of false positives (e.g., online flirting conversations between adults) [15]. Because there are no benchmark datasets available that distinguish grooming from non-grooming messages in online social media, a grooming detection approach required the ability to describe hidden structures from unlabelled data, as described in [44] Specifically in the context of Social Live Streaming Services (SLSSs) which represent a relatively new form of social media, recent work (see [45]) indicates that moderation systems in place are highly ineffective in suspending the accounts of users responsible for producing or distributing pornographic content, which could be potentially linked to CSAM. The authors published two large datasets by crawling the social graphs of SLSSs[5], which were analysed to identify characterizing traits of adult content producers and consumers, as well as highlight network patterns of relationships among them.

4.5 Deviant behaviour detection in LOCARD

Given the scarcity of data resources and tools to battle grooming and CSAM related behaviours in social media, in the context of LOCARD, the consortium members will develop specially crafted crawlers for online services that will collect streamed content, text messages and user interactions, analyse them and identify cases of child exploitation, grooming of underage users and other predatory acts. Unfortunately, the lack of effective measures to prevent such behaviours has led to major problems, such as the exploitation and

[5] https://github.com/nlykousas/asonam2018

grooming of underage users by sexual predators[6]. Moreover, Social Live Streaming Services is a quite new social media domain that has not yet received much attention by the research community. Thus, special attention will be given to Social Live Streaming Services, with the aim to provide end users, LEA investigators, with novel tools that will try to identify such deviant behaviours in online user generated content. Special care in this process will be made to comply with GDPR and other related legal frameworks.

[6] https://www.fox10phoenix.com/news/live-streaming-app-liveme-makes-major-changes-following-award-winning-fox-11-investigation

5 Conclusion

In this deliverable, we present the past research efforts and state of the art approaches to identify deviant behaviours related to production/distribution of CSAM and predatory behaviours targeting children, in the context of social media. Efforts have been focused towards proposing automated text-based methods that can be employed to assist law enforcement in their child protection investigations, as well as the identification CSAM material mostly in P2P file exchange networks. However, research efforts towards the detection and identification of deviant users such as those grooming children in online social media or those distributing new or previously unknown child sexual abuse media that may indicate recent or ongoing child abuse, have been severely limited in scientific literature. Thus, in the context of LOCARD, we will focus our efforts on the design and implementation of efficient approaches and tools for the detection of such criminal behaviours in modern social media platforms.

6 References

[1] M. B. Clinard and R. F. (Robert F. Meier, *Sociology of deviant behaviour*. .

[2] G. Tyson, Y. Elkhatib, N. Sastry, and S. Uhlig, “Are People Really Social on Porn 2.0?”

[3] F. Attwood, “What do people do with porn? Qualitative research into the consumption, use, and experience of pornography and other sexually explicit media,” *Sexuality and Culture*, vol. 9, no. 2. Transaction Publishers, pp. 65–86, 2005.

[4] M. Coletto, L. M. Aiello, C. Lucchese, and F. Silvestri, “Adult content consumption in online social networks,” *Social Network Analysis and Mining*, vol. 7, no. 1, Dec. 2017.

[5] S. Livingstone, L. Haddon, A. Görzig, and K. Ólafsson, “Risks and safety on the internet: the perspective of European children: full findings and policy implications from the EU Kids Online survey of 9-16 year olds and their parents in 25 countries Report Original citation,” 2011.

[6] P. Jeney, “Combatting child sexual abuse online.”

[7] B. G. Westlake and M. Bouchard, “Criminal Careers in Cyberspace: Examining Website Failure within Child Exploitation Networks,” *Justice Quarterly*, vol. 33, no. 7, pp. 1154–1181, Nov. 2016.

[8] K. J. Mitchell, D. Finkelhor, L. M. Jones, and J. Wolak, “Use of Social Networking Sites in Online Sex Crimes Against Minors: An Examination of National Incidence and Means of Utilization,” *Journal of Adolescent Health*.

[9] E. Quayle, G. Holland, C. Linehan, and M. Taylor, “The internet and offending behaviour: A case study,” *Journal of Sexual Aggression*, vol. 6, no. 1–2, pp. 78–96, 2000.

[10] A. L. Cooper, “Sexuality and the internet: Surfing into the new millennium,” *Cyberpsychology and Behavior*, vol. 1, no. 2, pp. 187–193, 1998.

[11] G. M. Winters and E. L. Jeglic, “Stages of Sexual Grooming: Recognizing Potentially Predatory Behaviors of Child Molesters,” *Deviant Behavior*, vol. 38, no. 6, pp. 724–733, Jun. 2017.

[12] P. Zambrano, J. Torres, and P. Flores, “How Does Grooming Fit into Social Engineering?,” in *Advances in Intelligent Systems and Computing*, 2019, vol. 924, pp. 629–639.

[13] S. Craven, S. Brown, and E. Gilchrist, “Sexual grooming of children: Review of literature and theoretical considerations,” *Journal of Sexual Aggression*, vol. 12, no. 3, pp. 287–299, 2006.

[14] N. Pendar, “Toward spotting the pedophile telling victim from predator in text chats,” in *ICSC 2007 International Conference on Semantic Computing*, 2007, pp. 235–241.

[15] G. Inches and F. Crestani, “Overview of the International Sexual Predator Identification Competition at PAN-2012.”

[16] “Child protection system. P2P monitoring software developed at TLO.” [Online]. Available: https://www.tlo.com/. [Accessed: 07-Oct-2019].

[17] M. Liberatore, B. N. Levine, and C. Shields, “Strengthening Forensic Investigations of Child Pornography on P2P Networks,” 2010.

[18] M. Liberatore, R. Erdely, T. Kerle, B. N. Levine, and C. Shields, “Forensic investigation of peer-to-peer file sharing networks,” *Digital Investigation*, vol. 7, pp. S95–S103, Aug. 2010.

[19] M. Latapy, C. Magnien, and R. Fournier, “Quantifying paedophile activity in a large P2P system,” *Information Processing and Management*, vol. 49, no. 1, pp. 248–263, Jan. 2013.

[20] I. McGhee, J. Bayzick, A. Kontostathis, L. Edwards, A. McBride, and E. Jakubowski, “Learning to identify Internet sexual predation,” *International Journal of Electronic Commerce*, vol. 15, no. 3, pp. 103–122, Apr. 2011.

[21] W. Rahmanmiah, J. Yearwood, and S. Kulkarni, “Detection of child exploiting chats from a mixed chat dataset as a text classification task.”

[22] D. Bogdanova, P. Rosso, and T. Solorio, “Exploring high-level features for detecting cyberpedophilia,” *Computer Speech and Language*, vol. 28, no. 1, pp. 108–120, 2014.

[23] E. Villatoro-Tello, A. Juárez-González, H. J. Escalante, M. Montes-Y-Gómez, and L. Villaseñor-Pineda, “A Two-step Approach for Effective Detection of Misbehaving Users in Chats ⋆ Notebook for PAN at CLEF 2012.”

[24] C. Peersman, F. Vaassen, V. van Asch, and W. Daelemans, “Conversation Level Constraints on Pedophile Detection in Chat Rooms.”

[25] C. Morris and G. Hirst, “Identifying Sexual Predators by SVM Classification with Lexical and Behavioral Features Notebook for PAN at CLEF 2012.”

[26] J. María, G. Hidalgo, A. Alfonso, and C. Díaz, “Combining Predation Heuristics and Chat-Like Features in Sexual Predator Identification Notebook for PAN at CLEF 2012.”

[27] P. Gottschalk, “Journal of Computing:: A Dark Side of Computing and Information Sciences: Characteristics of Online Groomers,” vol. 2, no. 9, 2011.
[28] D. Michalopoulos and I. Mavridis, “Utilizing document classification for grooming attack recognition,” in *Proceedings - IEEE Symposium on Computers and Communications*, 2011, pp. 864–869.
[29] J. Parapar, D. E. Losada, and A. Barreiro, “A learning-based approach for the identification of sexual predators in chat logs Notebook for PAN at CLEF 2012.”
[30] A. Kontostathis, W. West, A. Garron, K. Reynolds, and L. Edwards, “Identifying Predators Using ChatCoder 2.0 Notebook for PAN at CLEF 2012.”
[31] P. Elzinga, K. E. Wolff, J. Poelmans, G. Dedene, and S. Viaene, “Analyzing Chat Conversations of Pedophiles with Temporal Relational Semantic Systems.”
[32] K. Joffres, M. Bouchard, R. Frank, and B. Westlake, “Strategies to disrupt online child pornography networks,” in *Proceedings - 2011 European Intelligence and Security Informatics Conference, EISIC 2011*, 2011, pp. 163–170.
[33] L. Dai, “An Ising-based approach for tracking illegal P2P content distributors,” Ames, 2010.
[34] B. G. Westlake, M. Bouchard, R. Frank, B. G. ; Westlake, and M. ; Bouchard, “Finding the Key Players in Online Child Exploitation Networks,” *Policy & Internet*, vol. 3, no. 2, 2011.
[35] ECPAT International, “What is PhotoDNA ?”
[36] D. Hughes, J. Walkerdine, G. Coulsoni, and S. Gibson, “Peer-to-peer: Is deviant behaviourthe norm on P2P file-sharing networks?,” *IEEE Distributed Systems Online*, vol. 7, no. 2. pp. 1–11, Feb-2006.
[37] R. Fournier *et al.*, “Comparing pedophile activity in different P2P systems,” *Social Sciences*, vol. 3, no. 3, pp. 314–325, Sep. 2014.
[38] C. Peersman, C. Schulze, A. Rashid, M. Brennan, and C. Fischer, “ICOP: Automatically identifying new child abuse media in P2P networks,” in *Proceedings - IEEE Symposium on Security and Privacy*, 2014, vol. 2014-January, pp. 124–131.
[39] R. Hurley *et al.*, “Document Title: Measurement and Analysis of Child Pornography Trafficking on P2P Networks, Final Technical Report Measurement and Analysis of Child Pornography Trafficking on P2P Networks.”
[40] J. Wolak, M. Liberatore, and B. N. Levine, “Measuring a year of child pornography trafficking by U.S. computers on a peer-to-peer network,” *Child Abuse & Neglect*, 2013.
[41] C. M. S. Steel, “Child pornography in peer-to-peer networks,” *Child Abuse and Neglect*, vol. 33, no. 8, pp. 560–568, Aug. 2009.
[42] M. Edwards and A. Rashid, “Collaborative Filtering as an Investigative Tool for Peer-to-Peer Filesharing Networks.” 2012.
[43] A. Panchenko, R. Beaufort, and C. Fairon, “Detection of Child Sexual Abuse Media on P2P Networks: Normalization and Classification of Associated Filenames.”
[44] C. Peersman, “Detecting Deceptive Behaviour in the Wild: Text Mining for Online Child Protection in the Presence of Noisy and Adversarial Social Media Communications,” 2018.
[45] N. Lykousas, C. Patsakis, and V. Gomez, “Adult content in social live streaming services: Characterizing deviant users and relationships,” in *Proceedings of the 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2018*, 2017, pp. 375–382.