Australian Government
Australian Institute of Criminology

Trends & issues in crime and criminal justice

No. 668

Abstract | This paper demonstrates how biometric features can be extracted from people in child sexual abuse material (CSAM) and examined using social network analysis to reveal important patterns across seized media files. Using an automated software system previously developed by the research team (the Biometric Analyser and Network Extractor), we extract, match and plot multiple biometric attributes (face and voice) from a database of CSAM videos compiled by law enforcement. We apply a series of network measures to illustrate how the biometric match data can be used to rapidly pinpoint key media files associated with an investigation, without the need for an investigator to manually review and catalogue all files. Future directions for this research are also discussed.

Advancing child sexual abuse investigations using biometrics and social network analysis

Russell Brewer, Bryce Westlake, Thomas Swearingen, Stephen Patterson, David Bright, Arun Ross, Katie Logos and Dana Michalski

Trends in the distribution of child sexual abuse material (CSAM) online demonstrate a growing preference by producers and consumers for video and 'on-demand' live streams (Brown, Napier & Smith 2020; Dance & Keller 2020; Maxim et al. 2016). Such preferences are also confirmed by recent reporting to the National Center for Missing and Exploited Children (2022), which in 2021 received more reports for videos than images (44.8 million vs 39.9 million), representing an increase of 41.7 percent over the previous year. These trends highlight a growing need for effective tools for analysing videos in child sexual abuse (CSA) investigations. The proliferation of video files amplifies the challenges investigators face, given that software tools available to process and analyse video lag behind those developed for images (Sanchez et al. 2019). This necessitates painstaking manual review and verification to extract key information about individuals, or significant patterns across multiple videos (eg victims or offenders appearing in several videos, instances of co-offending and co-victimisation).

Child Sexual Abuse Material Reduction Research Program

Manual processing of videos furthers the existing problems investigators face with unmanageable workloads and burnout, as well as significant psychological harms, including secondary traumatic stress disorder, emotional exhaustion, intrusive thoughts, and interpersonal and marital problems (Bourke & Craun 2014; Burns et al. 2008; Powell et al. 2015; Seigfried-Spellar 2018). This has led to an emergence of automated techniques, including those using hash values (ie digital fingerprints) and, more recently, biometric characteristics to ameliorate these challenges and enhance investigations.

Recent work by the present authors has demonstrated the utility of combining multiple biometric modalities (face and voice) from CSA videos using a custom-designed automated software system, entitled the Biometric Analyser and Network Extractor (BANE). Through a series of performance tests, this research showed that multiple biometric cues (ie faces and voices) can be successfully extracted and matched in CSAM (see Westlake et al. 2022 for a detailed overview). These performance results demonstrate the potential utility of this automated software infrastructure to augment investigations undertaken by law enforcement. This software offers a means of automatically grouping victims/offenders by face and/or voice matches, in ways that would be very difficult, if not impossible, to accomplish manually.

Although using multiple biometric attributes can improve accuracy in identifying the same individuals from one video to the next, it is possible to take this further and identify significant patterns pertaining to co-offending or co-victimisation, both within and across investigations (and across time). Establishing such connections is important, as there is considerable evidence to suggest that children are often victimised by multiple offenders, and that it is not uncommon for an offender to appear, visibly or audibly, in CSAM (Canadian Centre for Child Protection 2017; Interpol 2018; Salter & Whitten 2022; Seto et al. 2018). In this paper, we demonstrate how such patterns can be revealed by using social network analysis to examine face and voice biometric match data extracted from CSA videos. To this end, this work describes how this analytical approach can point investigators towards media files that should be prioritised for manual analysis. This also has the potential to dramatically reduce investigator workloads by obviating the need to manually review all media files, as well as the attendant psychological harms associated with viewing such materials.

This paper is presented in three parts. First, we provide a methodological account of this research. Second, using a database of 530 CSA videos compiled by law enforcement, we demonstrate how specific social network metrics (community detection, degree centrality and betweenness centrality, which are explained below) can be used and visualised to rapidly pinpoint key media files associated with an investigation, without the need for an investigator to first manually review and catalogue files. Third, we outline the implications of integrating and improving this analytical approach in future research.

Methodology

Given the nature of the content being studied, BANE was supplied to Australian law enforcement agencies to extract biometric match data from a collection of CSA videos. Match data were extracted from BANE as metadata, to plot a network map for visualisation and analysis. These processes are detailed below.

Compilation of the video database

A collection of 1,308 videos was compiled by Australian law enforcement. These videos represent all videos seized by law enforcement relating to a large and recent CSA investigation (ie all video files from hard drives, mobile phones, tablets, online accounts, messaging applications etc), with each file being verified by investigators as containing CSA, as defined under Australian law. The collection contained numerous duplicate files, which were removed using a script written to identify duplicates (on the basis of a file being of the same size, length and extension), leading to a sample of 553 unique videos. The final sample of videos contained a variety of forms of CSA, involving children of a range of ages (roughly between 3 and 17 years old), as well as adults. Videos also varied in length, with the shortest being five seconds and the longest approximately 48 minutes. At no point did members of the research team view or have access to CSAM.

Data procedures

Law enforcement officers imported the 553 videos into BANE, which successfully processed 530 videos. A small proportion of video files (*n*=23, or 4.2%) were not included in the analysis because either they were not encoded in one of the formats BANE is presently designed to process (ie *.3gp, *.3gpp, *.asf, *.avi, *.divx, *.mkv, *.mp4, *.mpg, *.mpeg, *.wmf, *.wmv, *.vob), they were corrupted, or they could not be decoded using ffmpeg v 4.0 (see https://ffmpeg.org).

This study deployed a processing pipeline similar to the one described in Westlake et al. (2022), with the only change being the use of a different face recognition library. This study used a pipeline consisting of the Dlib face detector (King 2009) and an open-source face recognition system (see https://github.com/ageitgey/face_recognition). The included face recognition pipeline achieves 98.9 percent accuracy on the 'Labeled Faces In The Wild' dataset (Huang et al. 2007), which is in line with state-of-the-art face recognition systems. Facial features were successfully extracted from 258 videos and audio features from 352 videos. Such performance aligned with expectations, given that CSAM does not always contain a clearly visible face or audible voice (Salter & Whitten 2022; Tejeiro et al. 2020). In total, biometric features (a face and/or a voice) were extracted from 445 videos. These extracted features were then matched to all other faces and voices from other videos in the database, producing a match score for each pair.

We used an inductive approach to selecting an appropriate threshold from which matches would be derived for further analysis. To our knowledge, no previous study has sought to match face and voice biometrics on a raw database of CSA videos (Westlake et al. 2022 is the closest, using a small, labelled dataset of CSAM). Consequently, the research team had limited empirical evidence with which to determine appropriate match threshold levels. Discussions with investigators revealed that the database comprised numerous videos containing the primary offender (whose computer, phone, accounts etc were seized) offending against multiple children. In addition, there were various other separate collections featuring victims that the primary offender had not had contact with, which the primary offender had found online or had been sent by other offenders.

Accordingly, we expected that the matching process would return one large cluster of nodes (corresponding with the primary offender), and numerous smaller clusters (varying in size, according to the number of videos containing each individual). We evaluated several thresholds for matching both faces and voices. Faces were matched using a dissimilarity score, meaning that scores closer to 0 are more likely to indicate a genuine match, while voices were evaluated on a similarity score, meaning that scores closer to 1 are more likely to indicate a genuine match (for further detail, see Westlake et al. 2022). Through this process, it was determined that the expected pattern was generated using thresholds of less than 0.47 for face and greater than 0.96 for voice. In making this selection, we acknowledged that presenting false positive matches may slow down or be harmful to an investigation and reduce the practical utility of the software. As such, we sought thresholds corresponding with a low false match rate, even though this will also result in a lower true match rate.

Analytical framework

Once processed by BANE, inter-video match scores for the 445 videos were converted into matrices, to permit analysis of biometric matches using social network analysis (SNA). SNA is an analytical technique that can be used to triangulate multiple file attributes (eg face and/or voice matches), as well as identify and visualise structural patterns across large datasets. With its roots in mathematical graph theory, this analytical approach focuses on a set of actors or ‘nodes’ (in this case, CSA videos) that are tied by one or more types of relations (in this case, biometric matches) (Wasserman & Faust 1994). SNA can be used to examine patterns of relations across different levels of networks, including the complete network of relations, relations within clusters or subgroups, and relations between pairs of nodes. SNA can also be used to identify well-connected or strategically positioned nodes within a network. This approach has the advantage of yielding new information about the structural characteristics of a network, and allows analysts to visualise and identify significant connections between nodes that may not otherwise be revealed (Marin & Wellman 2011). Various research studies have extolled the potential of this approach as a means to better understand and control crime (see Brewer 2017), particularly in the context of CSA (eg Bursztein et al. 2019; Krone 2004; Westlake & Bouchard 2016; Westlake & Frank 2016).

To permit SNA in the present study, the full list of match pairs derived from the dataset by BANE was converted into a machine-readable format (ie an edge list) and imported into a separate SNA software program, Gephi (v 0.97), for further analysis. This made it possible to generate a visual representation of the network and to calculate several measures of connectivity (density) and centrality (degree centrality and betweenness centrality) and to detect communities (using the Girvan–Newman algorithm). This revealed both important patterns and key videos across the network. Ethics approval for this work was granted by the University of Adelaide Human Research Ethics Committee.

Results and discussion

Biometric data (faces and voices) extracted from the 445 distinct videos were matched, combined and modelled as a network map, visually represented in Figure 1. Here, each of the 445 black dots, or nodes, represents a video, while the 222 lines, or ties, linking nodes denote at least one biometric match between people contained within the two connected videos (one video can contain multiple people, who are each matched to other videos). These matches are diverse and represent one of three tie types: a face match (red), a voice match (violet) or a match between *both* face and voice (green). The diversity of match types may be attributable to the fact that faces and voices often do not appear in CSAM. To further assist with visualisation, the Force Atlas algorithm was applied to drive isolates (the nodes not linked to any other nodes) to the centre of the map, leaving clusters on the periphery of the network for closer inspection. This modelling makes it possible to draw links between people (victims or offenders) across videos by constructing more complex networks. This network visualisation moves beyond simply identifying the same victim or the same offender appearing in multiple videos and can also establish co-offending and co-victimisation relationships across a large holding of videos.

Figure 1: Network of videos extracted from a single investigation

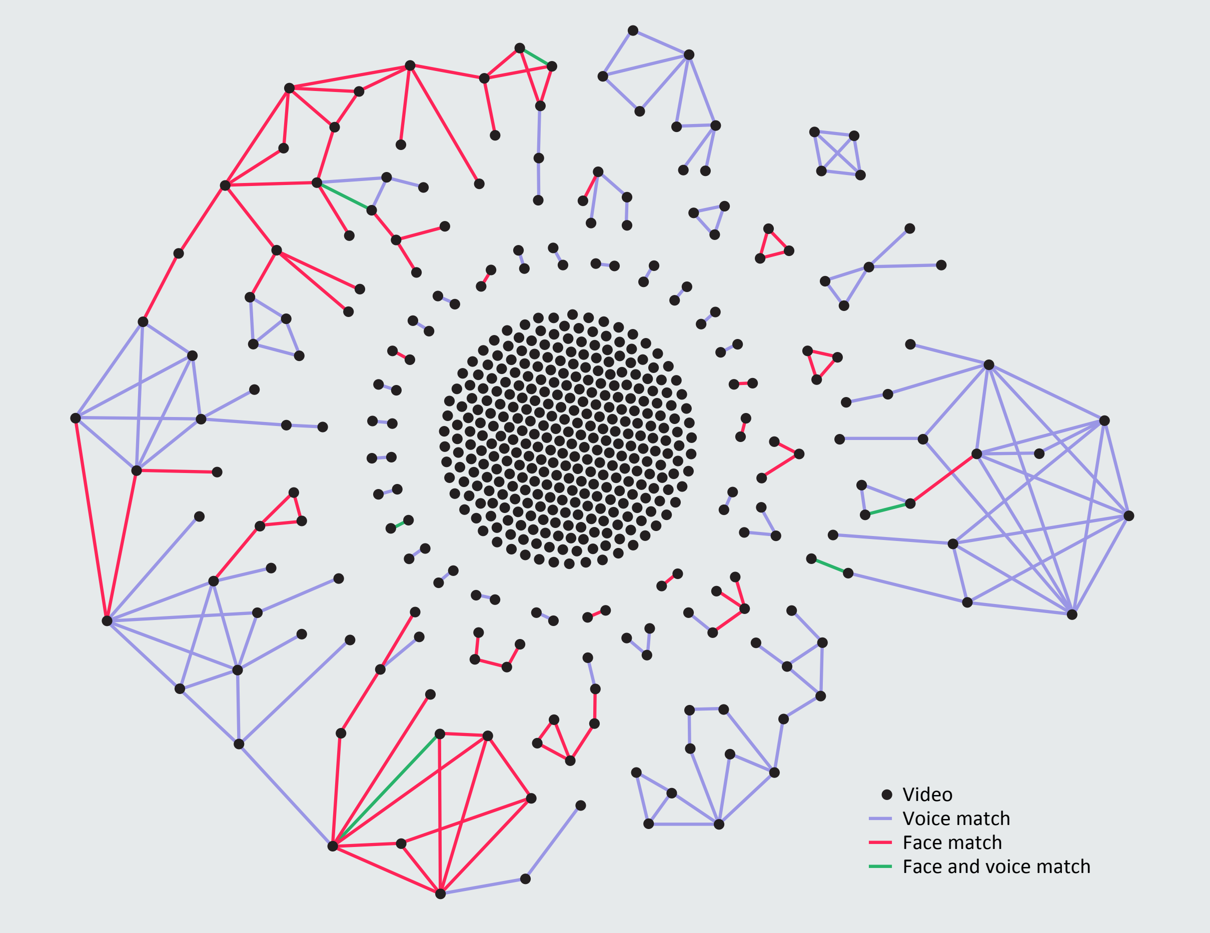

As expected, the network is relatively low in density, with matches being returned across 0.02 percent of the total matches possible between videos (Figure 1). Matches are, nevertheless, largely concentrated in clusters. Just under half of all matches (42%) are concentrated into one large cluster (likely depicting the primary offender and associated victims and co-offenders), located in the left-hand portion of the network. Videos throughout this cluster are indeed connected by face matches, voice matches or both, with the loss of either biometric severely fragmenting the cluster. Elsewhere, 15 other smaller clusters populate the periphery of the network, likely representing victims and/or offenders not directly associated with the primary offender. These smaller clusters vary in size, each containing between three (bottom right) and 17 (far right) nodes. Closer to the centre of the graph are 25 matches that link only two videos (ie 2 nodes and a single tie), while 324 isolates (videos with no matches) reside at the centre of the network map. Further analysis of these connections can reveal key patterns within the data that can direct investigations.

Identifying subgroups of related videos using community detection

While mapping these 16 clusters offers initial macro-level insights into how video files are connected, it does not provide investigators nuanced information about how individual subjects are represented within clusters. For investigators, making such distinctions is crucial, as clusters may contain subgroups of videos depicting multiple victims or offenders. That is, a subgroup within a cluster could represent multiple victims being abused by the primary offender or multiple offenders victimising the same child. To this end, it is possible to detect subgroups that form structurally separate entities, commonly referred to in the literature as 'communities'.

Numerous methods exist for identifying communities in network data. In this study, we applied the commonly used Girvan–Newman algorithm, which identifies discrete communities (subgroups of nodes) by locating structurally important ties, whose removal fragments the network (Borgatti, Everett & Johnson 2018). This method revealed six communities contained within the largest cluster, distinguished by colour in Figure 2: green (10 nodes), blue (14 nodes), pink (7 nodes), yellow (9 nodes), orange (14 nodes) and purple (13 nodes). The communities present within the clusters suggest that there are separate yet interconnected collections of different people within the primary cluster that investigators may wish to focus on. That is, each of these communities may contain a particular victim or offender, or potentially be part of the same series of videos. For this dataset, beyond the primary cluster, no additional communities were found to be contained *within* other clusters or pairings (and were therefore not visualised below). Note that a uniform colour (black) has been applied to all ties (biometric matches) in Figure 2 for clarity purposes. Tie types can still be determined by referring to Figure 1.

Figure 2: Identifying communities of victims/offenders within the primary cluster

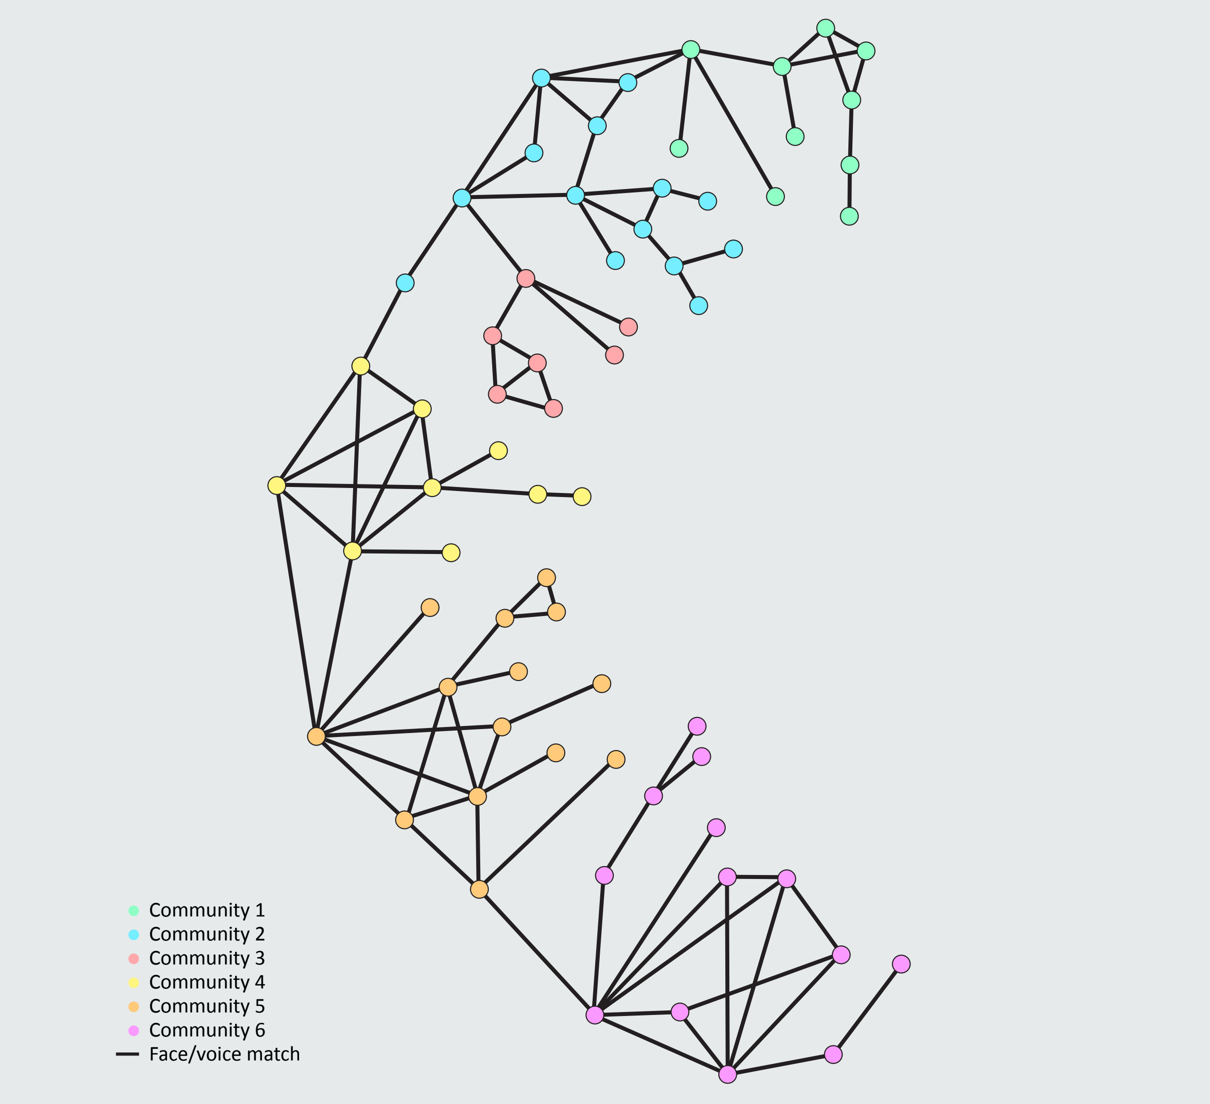

Identifying videos most central to the network

Beyond detecting communities, alternative network measures can be used to identify the most ‘central’ (ie highly connected) nodes within the network. This has the potential to provide investigators important insights about where to focus an investigation, whether within a community or within a particular cluster. In the current application, this means establishing which videos have the largest number of connections to other videos in the network. In practical terms, this could allow investigators to flag key videos that are likely to contain:

- a victim or offender who appears in many videos;
- the highest-quality biometric samples that match with other videos (ie videos matched using face and voice rather that only one biometric feature); or
- a number of different people from which matches can be drawn.

The most commonly used centrality metric is degree centrality, which is simply a sum of the number of ties (in this case, biometric matches) a node has with other nodes (Borgatti, Everett & Johnson 2018). This is illustrated in Figure 3, where the size of nodes is weighted in direct proportion to the number of direct matches they have with other nodes (the degree). That is, the larger the node, the greater the number of videos it matches. Visual inspection of the network makes clear that several nodes have greater importance than others. For example, nodes 252, 278, 446, 121 and 105 attract the most matches, relative to other nodes. See the *Appendix* for a full list of degree centrality scores for this network.

Figure 3: Degree centrality across the primary cluster

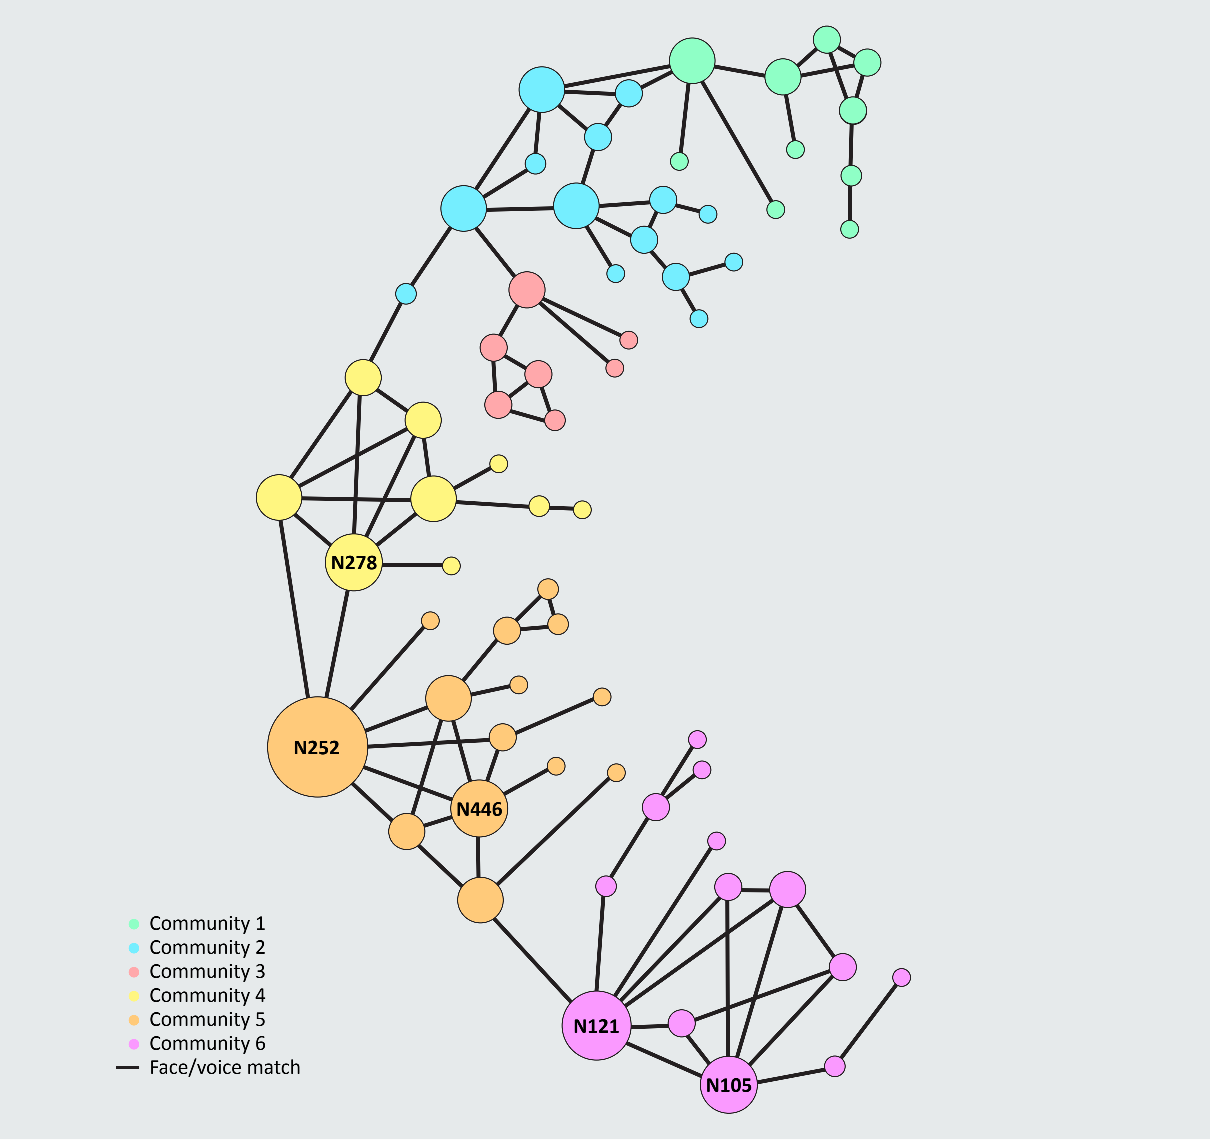

Identifying strategically-placed videos in the network

Beyond degree centrality, there are other measures of centrality that can effectively be used to focus investigations. One other commonly used centrality metric is betweenness centrality, which measures the extent to which any one node lies on the shortest path between all other pairs of nodes in the network. Betweenness has been used to identify strategically placed ‘brokers’ in networks (eg Morselli 2009). Determining the extent to which a video is strategically placed within a network (ie assumes a ‘brokerage’ function) may be particularly useful in investigatory contexts—permitting an investigator to pinpoint a specific video that creates a path between two otherwise disconnected videos. This may benefit an investigation in two ways. First, Figure 4a provides an example where a single video (Node 2) contains multiple biometric attributes of a single victim (Victim A) and thus brokers a connection between videos containing only one biometric attribute of that victim (Node 1 and Node 3). Second, Figure 4b depicts a situation where brokerage can be used to identify connections between discrete subjects contained within videos. For example, Node 2 (video containing Victims A and B) brokers a link between two otherwise disconnected nodes—Node 1 (a video containing Victim A) and Node 3 (a video containing Victim B).

Figure 4: Biometric match brokerage configurations

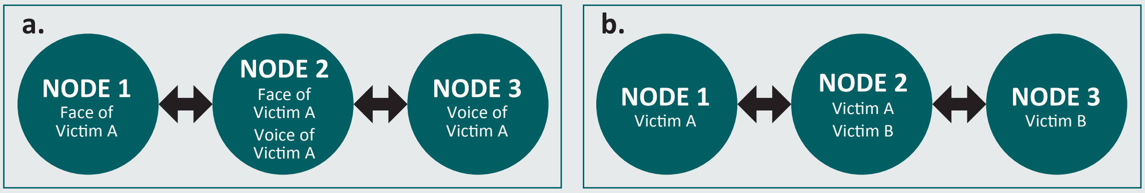

We assessed the strategic position of nodes by calculating betweenness centrality scores for every node in the network (for full results, see the *Appendix*). These scores are also depicted visually in the network map appearing at Figure 5, where node sizes are adjusted according to betweenness centrality (larger nodes have higher betweenness centrality scores). Here, the nodes intersecting communities are most prominent. These nodes—such as those numbered 40, 41, 156 and 252—unite various otherwise disconnected components. Identifying these videos may thus be of immense value to an investigation, as high betweenness scores suggest that these nodes may be the key linking videos that draw connections between previously unmatched victims or offenders contained in different sets of videos, as compared to degree centrality, which likely flags videos that contain individuals appearing in numerous other videos.

Figure 5: Betweenness centrality across the primary cluster

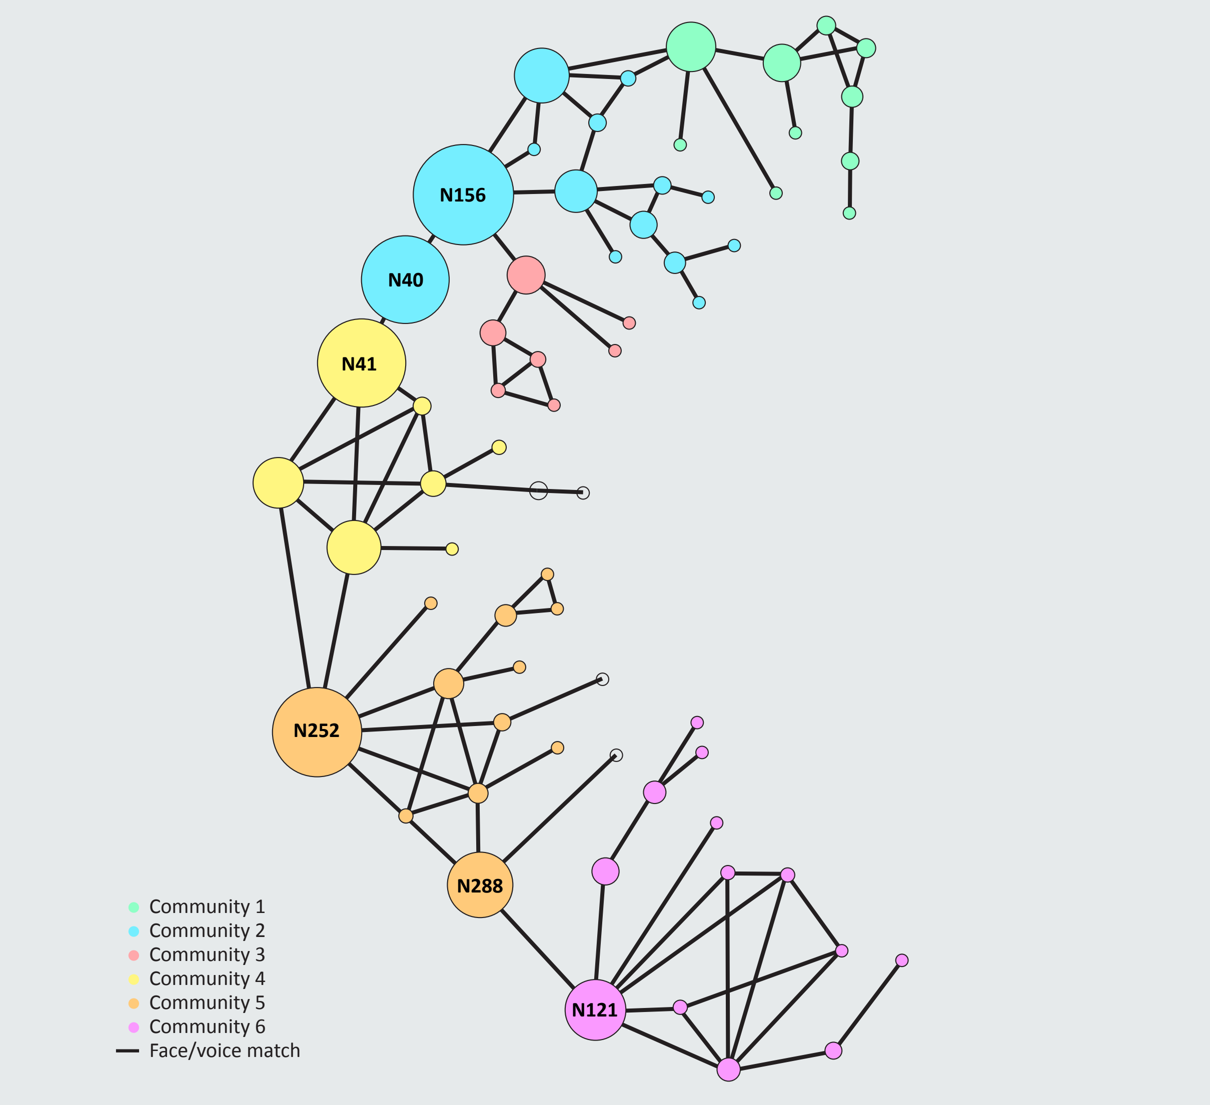

Limitations and directions for future research

Using biometric data extracted from CSA videos as linking attributes, SNA can aid investigators in identifying which videos to prioritise, without having to review and catalogue hundreds of different videos. We demonstrate how various measures can be used to take a macro view of a network and potentially parse out individuals via clusters, or to provide more granular insights by identifying distinct communities. In addition, measures of centrality can be used to identify the videos in a network that match most frequently with other videos (degree centrality), as well as those videos that link groups of videos together, possibly by virtue of unique characteristics being portrayed (betweenness centrality). In undertaking this analysis, however, we recognise that the results are based on some assumptions about the data, and that future research should seek to address these limitations. Three of these are discussed in turn.

First, we acknowledge the practical utility of SNA in assisting investigations relies on the veracity of the underpinning ties—in this case, the accuracy of the face and voice match data. The algorithms used by BANE (see Westlake et al. 2022 for additional details) were developed for other purposes and were not trained using CSAM, which limits their performance. This is a common problem experienced by researchers in this field and can be addressed through the creation of large, labelled datasets of CSAM (in partnership with law enforcement) for training and testing purposes.

Second, and relatedly, while care was taken to ensure that videos in the testing database contained CSA and were reflective of 'real-world' conditions, it was not possible to verify the face and voice matches returned by the software at the selected thresholds. As such, it was not possible to determine the accuracy of the network, according to the corresponding true/false match rates. Future research will need to incorporate labelled data with an established ground truth so that accuracy can be evaluated. Given the graphic nature of the content, and the legal implications of possessing CSAM, such activities will need to be completed in partnership with law enforcement (see Bright, Brewer & Morselli 2021 for further elaboration on how this can be accomplished).

Finally, this paper presented a network derived from face and voice biometric match data. This can be problematic, as a proportion of CSA videos being distributed online contain neither a face nor a voice (as was the case in 85 videos contained in the current dataset). This suggests a need to extend the software's extraction and matching capabilities to include algorithms capturing additional biometrics such as age, gait, gender, hair colour and ethnicity (eg Macedo, Costa & dos Santos 2018; Moser, Rybnicek & Haslinger 2015; Sae-Bae et al. 2014; Yiallourou, Demetriou & Lanitis 2017), as well as objects (eg artwork, food packaging, newspapers and magazines), camera sensors (Bennabhaktula et al. 2020; Timmerman et al. 2021), and file encoding properties (Lyons & Epstein 2021, 2020). Such algorithms can be integrated into future iterations of BANE and may further enhance matching performance as well as the depth and breadth of networks derived.

References

URLs correct as at November 2022

Bennabhaktula GS, Alegre E, Karastoyanova D & Azzopardi G 2020. Device-based image matching with similarity learning by convolutional neural networks that exploit the underlying camera sensor pattern noise. In M De Marsico, GS di Baja & A Fred (eds), *Proceedings of the 9th International Conference on Pattern Recognition Applications and Methods: Volume 1*. SciTePress: 578–584. https://doi.org/10.5220/0009155505780584

Borgatti SP, Everett MG & Johnson JC 2018. *Analyzing social networks*, 2nd ed. Los Angeles: SAGE

Bourke ML & Craun SW 2014. Secondary traumatic stress among Internet Crimes Against Children Task Force personnel. *Sexual Abuse: A Journal of Research and Treatment* 26(6): 586–609. https://doi.org/10.1177/1079063213509411

Brewer R 2017. Controlling crime through networks. In P Drahos (ed), *Regulation, institutions and networks*. ANU Press: 447–464. https://doi.org/10.22459/RT.02.2017.26

Bright D, Brewer R & Morselli C 2021. Using social network analysis to study crime: Navigating the challenges of criminal justice records. *Social Networks* 66: 50–64. https://doi.org/10.1016/j.socnet.2021.01.006

Brown R, Napier S & Smith RG 2020. Australians who view live streaming of child sexual abuse: An analysis of financial transactions. *Trends & issues in crime and criminal justice* no. 589. Canberra: Australian Institute of Criminology. https://doi.org/10.52922/ti04336

Burns CM, Morley J, Bradshaw R & Domene J 2008. The emotional impact on coping strategies employed by police teams investigating internet child exploitation. *Traumatology* 14(2): 20–31. https://doi.org/10.1177/1534765608319082

Bursztein E, Clarke E, DeLaune M, Elifff DM, Hsu N, Olson L, Shehan J, Thakur M, Thomas K & Bright T 2019. *Rethinking the detection of child sexual abuse imagery on the internet*. World Wide Web Conference, 13 May, pp 2601–2607. https://doi.org/10.1145/3308558.3313482

Canadian Centre for Child Protection 2017. International Survivors' Survey. https://protectchildren.ca/en/resources-research/survivors-survey-results/

Dance GJX & Keller MH 2020. Tech companies detect a surge in online videos of child sexual abuse. *New York Times*, 20 February. https://www.nytimes.com/2020/02/07/us/online-child-sexual-abuse.html

Huang GB, Ramesh M, Berg T & Learned-Miller E 2007. *Labeled Faces in the Wild: A database for studying face recognition in unconstrained environments*. Workshop on Faces in 'Real-Life' Images: Detection, Alignment, and Recognition. https://hal.inria.fr/inria-00321923

Interpol 2018. *Towards a global indicator on unidentified victims in child sexual exploitation material: Technical report*. https://ecpat.org/resource/technical-report-towards-a-global-indicator-on-unidentified-victims-in-child-sexual-exploitation-material/

King DE 2009. Dlib-ml: A Machine Learning Toolkit. *Journal of Machine Learning Research* 10: 1755–1758

Krone T 2004. A typology of online child pornography offending. *Trends & issues in crime and criminal justice* no. 279. Canberra: Australian Institute of Criminology. https://www.aic.gov.au/publications/tandi/tandi279

Lyons B & Epstein B 2021. *Source identification of unknown video files*. National Cyber Crime Conference, May, online

Lyons B & Epstein B 2020. *The truth in video files: Introducing a novel approach to video source identification/authentication*. VirtualLEVA Digital Multimedia Evidence Training Symposium, 26 October, online

Macedo J, Costa F & dos Santos JA 2018. *A benchmark methodology for child pornography detection*. 2018 31st SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI), Paraná, Brazil, pp 455–462. https://doi.org/10.1109/SIBGRAPI.2018.00065

Marin A & Wellman B 2011. Social network analysis: An introduction. In P Carrington & J Scott (eds), *The SAGE handbook of social network analysis*. Sage: 11–25. https://doi.org/10.4135/9781446294413.n2

Maxim D, Orlando S, Skinner K & Broadhurst R 2016. *Online child exploitation material: Trends and emerging issues*. Canberra: Australian National University Cybercrime Observatory and Office of the Children's eSafety Commissioner. https://doi.org/10.2139/ssrn.2861644

Morselli C 2009. *Inside criminal networks.* New York: Springer. https://doi.org/10.1007/978-0-387-09526-4

Moser A, Rybnicek M & Haslinger D 2015. *Challenges and limitations concerning automatic child pornography classification*. Proceedings of the 10th International Conference on Computer Vision Theory and Applications, Berlin, pp 492–497. https://doi.org/10.5220/0005344904920497

National Center for Missing & Exploited Children 2022. 2021 our impact. https://www.missingkids.org/content/dam/missingkids/pdfs/2021-Our-Impact.pdf

Powell M, Cassematis P, Benson M, Smallbone S & Wortley R 2015. Police officers' perceptions of their reactions to viewing internet child exploitation material. *Journal of Police and Criminal Psychology* 30(2): 103–111. https://doi.org/10.1007/s11896-014-9148-z

Sae-Bae N, Sun X, Sencar HT & Memon ND 2014. Towards automatic detection of child pornography. 2014 IEEE International Conference on Image Processing, Paris, pp 5332–5336. https://doi.org/10.1109/ICIP.2014.7026079

Salter M & Whitten T 2022. A comparative content analysis of pre-internet and contemporary child sexual abuse material. *Deviant Behavior* 43(9): 1120–1134. https://doi.org/10.1080/01639625.2021.1967707

Sanchez L, Grajeda C, Baggili I & Hall C 2019. A practitioner survey exploring the value of forensic tools, AI, filtering, & safer presentation for investigating child sexual abuse material (CSAM). *Digital Investigation* 29: 124–142. https://doi.org/10.1016/j.diin.2019.04.005

Seigfried-Spellar KC 2018. Assessing the psychological well-being and coping mechanisms of law enforcement investigators vs. digital forensic examiners of child pornography investigations. *Journal of Police and Criminal Psychology* 33(3): 215–226. https://doi.org/10.1007/s11896-017-9248-7

Seto MC, Buckman C, Dwyer RG & Quayle E 2018. *Production and active trading of child sexual exploitation images depicting identified victims*. National Center for Missing & Exploited Children and Thorn. https://www.missingkids.org/theissues/csam

Tejeiro R, Alison L, Hendricks E, Giles S, Long M & Shipley D 2020. Sexual behaviours in indecent images of children: A content analysis. *International Journal of Cyber Criminology* 14(1): 121–138. https://doi.org/10.5281/zenodo.3743390

Timmerman D, Bennabhaktula S, Alegre E & Azzopardi G 2021. Video camera identification from sensor pattern noise with a constrained ConvNet. In M De Marsico, GS di Baja & A Fred (eds), *Proceedings of the 10th International Conference on Pattern Recognition Applications and Methods: Volume 1*. SciTePress: 417–425. https://doi.org/10.5220/0010246804170425

Wasserman S & Faust K 1994. *Social network analysis: Methods and applications*. Cambridge University Press. https://doi.org/10.1017/CBO9780511815478

Westlake BG & Bouchard M 2016. Liking and hyperlinking: Community detection in online child sexual exploitation networks. *Social Science Research* 50: 23–36. https://doi.org/10.1016/j.ssresearch.2016.04.010

Westlake BG, Brewer R, Swearingen T, Ross A, Patterson S, Michalski D, Hole M, Logos K, Frank R, Bright D & Afana E 2022. Developing automated methods to detect and match face and voice biometrics in child sexual abuse videos. *Trends & issues in crime and criminal justice* no. 648. Canberra: Australian Institute of Criminology. https://doi.org/10.52922/ti78566

Westlake BG & Frank R 2016. Seeing the forest through the trees: Identifying key players in online child sexual exploitation distribution networks. In T Holt (ed), *Cybercrime through an interdisciplinary lens*. Routledge: 189–209. https://doi.org/10.4324/9781315618456

Yiallourou E, Demetriou R & Lanitis A 2017. *On the detection of images containing child-pornographic material*. 24th International Conference on Telecommunications, Cyprus. https://doi.org/10.1109/ICT.2017.7998260

Appendix

Table A1: Full centrality results

Node	Degree centrality	Betweenness centrality
7	3	65.3
36	3	65.0
40	2	1080.0
41	4	1085.0
50	1	0.0
52	4	369.0
85	1	0.0
87	3	129.0
95	1	0.0
99	1	0.0
105	6	149.3
118	2	65.0
121	7	687.8
136	2	65.0
142	1	0.0
148	4	20.5
149	1	0.0
151	1	0.0
153	1	0.0

Table A1: Full centrality results

Node	Degree centrality	Betweenness centrality
156	5	1265.7
170	3	128.0
208	1	0.0
219	3	32.0
221	2	0.0
226	5	528.0
238	1	0.0
242	2	189.0
252	8	1111.0
257	1	0.0
260	1	0.0
262	5	587.8
263	3	31.7
265	3	65.0
271	3	32.0
272	1	0.0
278	6	607.2
285	3	93.0
286	2	0.0
288	5	741.0
310	3	129.0
335	5	251.0
356	4	25.0
357	1	0.0
359	1	0.0
370	4	365.0
372	5	542.2
373	5	441.5
375	4	42.7
379	3	20.0
380	3	93.0
389	3	0.3
408	3	128.0
414	3	0.0
419	1	0.0
439	3	189.0
444	1	0.0
446	6	111.0

Table A1: Full centrality results		
Node	**Degree centrality**	**Betweenness centrality**
452	1	0.0
458	1	0.0
459	1	0.0
473	2	0.0
486	1	0.0
489	3	189.0
503	1	0.0
504	2	65.0
532	5	191.0
533	2	0.0

Russell Brewer is an Associate Professor at the University of Adelaide.

Bryce Westlake is an Associate Professor at San Jose State University.

Thomas Swearingen is a doctoral candidate at Michigan State University.

Stephen Patterson is a Detective Sergeant for the Joint Anti-Child Exploitation Team, South Australia Police.

David Bright is a Professor at Deakin University.

Arun Ross is a Professor at Michigan State University.

Katie Logos is a Lecturer at the University of Adelaide.

Dana Michalski is a Defence Scientist at the Defence Science and Technology Group.

This project is part of the AIC's Child Sexual Abuse Material Reduction Research Program, funded under section 298 of the Commonwealth *Proceeds of Crime Act 2002*.

General editor, *Trends & issues in crime and criminal justice* series: Dr Rick Brown, Deputy Director, Australian Institute of Criminology. Note: *Trends & issues in crime and criminal justice* papers are peer reviewed. For a complete list and the full text of the papers in the *Trends & issues in crime and criminal justice* series, visit the AIC website: www.aic.gov.au

ISSN 1836-2206 (Online) ISBN 978 1 922478 94 8 (Online)

https://doi.org/10.52922/ti78948

©Australian Institute of Criminology 2023

GPO Box 1936
Canberra ACT 2601, Australia

Tel: 02 6268 7166

Disclaimer: This research paper does not necessarily reflect the policy position of the Australian Government