Effective, Explainable and Ethical: AI for Law Enforcement and Community Safety

Campbell Wilson
AiLECS Lab
Monash University
Clayton, Australia
campbell.wilson@monash.edu

Janis Dalins
AiLECS Lab
Australian Federal Police
Canberra, Australia
janis.dalins@afp.gov.au

Gregory Rolan
AiLECS Lab
Monash University
Clayton, Australia
greg.rolan@monash.edu

***Abstract*—We describe the Artificial Intelligence for Law Enforcement and Community Safety (AiLECS) research laboratory, a collaboration between the Australian Federal Police and Monash University. The laboratory was initially motivated by work towards countering online child exploitation material. It now offers a platform for further research and development in AI that will benefit policing and mitigating threats to community wellbeing more broadly. We outline the work the laboratory has undertaken, results to date, and discuss our agenda for scaling up its work into the future.**

***Index Terms*—artificial intelligence, law enforcement**

I. Introduction

Criminal activity is increasingly facilitated by technology; often characterised by the generation, distribution and/or monetisation of illegal material via computer networks. In particular, recent years have seen rampant growth in the production and on-line dissemination of harmful and offensive materials, such as child exploitation material (CEM) and violent media associated with online radicalisation. Furthermore, worrying trends are emerging in the algorithmic generation of realistic abusive material such as "deepfake" imagery. Many of these cybercrimes are organized and transnational. In the course of investigating and prosecuting such offenses, law enforcement agencies deal with significant challenges, both professional and personal. Analysis and classification of such material exposes police and judicial officers to significant psychological harm. This is exacerbated by the increasingly large volumes of data involved in such investigations.

Given their capacity to learn patterns from large datasets and make consequent predictions, *Artificial intelligence (AI)*[1] technologies offer clear advantages for law enforcement in countering these threats while greatly reducing investigator exposure to harm. However, such technologies need to be developed and operationalised in accordance with appropriate frameworks for their legal, ethical, and explainable use, particularly given that preservation of community trust is vital for effective policing.

In this paper we describe a new research initiative: the AI for Law Enforcement and Community Safety (AiLECS) Laboratory [1], a collaboration between the Australian Federal Police (AFP) and Monash University. We envision the lab and its operations as a model and platform for AI research related to law enforcement on an international scale. Such a global orientation is necessary given the cross-jurisdictional nature of the problem domain and complexity of the research issues.

This is a high-stakes endeavour, with potential impact across several of the United Nations Sustainable Development Goals outlined in Resolution A/70/L.1 [1] and adopted by the General Assembly in 2015. All member nations have committed to meeting these goals and their associated targets by 2030. We aim to contribute specifically to Goal 3 Good Health and Well-Being by reducing the incidence and impact of harmful material online, Goal 16 Peace, Justice and Strong Institutions by integrating academia, industry and law enforcement to enhance technological capacity building, and Goal 17 Partnerships for the Goals through establishing sustainable collaboration between researchers and practitioners in information sciences, policing, law, and ethics.

The paper is organised as follows. In section 2, we describe our initial drivers for research in this area, challenges, and related initiatives. In section 3 we discuss the infrastructure underlying the lab in the form of our innovative *data airlock* platform. Section 4 covers work done to date in AI algorithm development. In section 5, we discuss learnings from the partnership and outline the our future plans. We conclude in section 6 with and call for greater collaboration in this domain.

II. Motivations and Challenges

In July 2019, the AiLECS laboratory was officially launched by the AFP and Monash University in Melbourne, Australia. This collaborative initiative represents significant investment by both organizations in building expertise around AI and related technologies for enhancing investigative capabilities. Importantly, it also seeks to move towards limiting exposure of AFP members to traumatizing material. The primary initial motivation behind AiLECS was addressing online child exploitation, building on prior research work the lab investigators.

A. A Motivator: The Scourge of Online Child Exploitation

The disturbing proliferation of online CEM is but one example of the need for urgent investigation of how AI approaches may assist investigators. In 2018 alone, the Australian Center

[1] In this paper, AI refers primarily to machine learning based systems

for Countering Child Exploitation received over 18,000 individual reports of child exploitation, with each one potentially containing hundreds or thousands of abhorrent images and videos. As stated in the introduction, this is a crime that knows no national boundaries and many countries face similar onslaughts. A recent study pointed towards the ”exponential” growth of this material, stating that of the 23.4 million reports of CEM received since 1998 by the US National Center for Missing and Exploited Children (NCMEC), 9.6 million were received in 2017 alone [2].

Combating CEM places an increasingly traumatic drain on the human and digital forensic resources of law enforcement agencies. Development of AI tools to assist in this task is a natural focus of much research and development. When integrated with existing digital investigation workflows, AI techniques do promise much by reducing investigator workload, accelerating investigations, and improving the welfare of those tasked with viewing and labelling such material.

Despite the AI field being long-standing, there have been particularly strong advances in recent years. For example, large strides have been made in improving artificial neural-network based image classification algorithms, thus making semi-automatic triage and categorization of large amounts of seized and previously unseen CEM images and videos more feasible [2], [3]. AI-based natural language processing techniques are also being investigated for assisting with classifying relevant evidentiary textual material [4].

B. AI in Law Enforcement: Challenges

In addition to countering CEM, AI techniques can be deployed to enhance many types of law enforcement activities. They are in a sense an evolution of data analysis capabilities that agencies already possess, and are suited to the increasingly large amounts of data that characterize many contemporary criminal investigations. However, aside from the challenge of research and development into improving the techniques themselves, their use in law enforcement poses a number of challenges:

- **Ethical Application**: Over the past few years there has been a growing discourse concerning the ethics of AI. Particularly in the law enforcement context, we need a clear and practical ethics basis upon which any AI intervention should proceed. Ethics frameworks (typically comprising statements of broad ethics principals) have been proposed by governments, private corporations, professional bodies, standards organisations, academia, and other individual and consortia stakeholders in AI [5]. Understandably, there is an emphasis in this discourse, on justice, fairness, and so on; and some of these can be be addressed through technological measures such as better data wrangling or research into algorithmic bias detection or fairness balancing.
  AI systems based on machine learning algorithms are highly dependent on (possibly voluminous) training data. Human *biases* in the training data are easily reflected in the output of such systems. Training data must be collected in a way that ensures the data is representative, inclusive and accurate. In the law enforcement context, this is particularly important given the potential adverse effects of biased algorithms and consequent loss of community trust in police agencies.
  From a broader perspective, our ethics stance needs to take into account the broader sociological implications of law enforcement work [6]. What are the purposes of law enforcement, and for AI interventions that seek to augment its activities? Along with justice and fairness, how is societal cohesion and interdependence best served?
  For example, the collection of large amounts of data in order to train algorithms in this context must also be mindful of individual rights to privacy. While lawful police activities may be the subject of certain exemptions under relevant legislation governing use of personal data, perception of undue surveillance is again a potential source of erosion of trust in agencies. Further, any degree of automation in law enforcement decision making potentially erodes human agency. How do we automate law enforcement work while retaining interpersonal discretion and *accountability* to those affected by the decision making?
- **Explainability**: Transparent decision making is particularly important in law enforcement. In an AI context, this means that explaining the output of traditionally opaque algorithms (e.g. deep neural networks) is a highly relevant area of research.
  Explainability of the algorithm itself is however not the entire picture—defending an AI enhanced policing approach in a court of law may require clear explanations of the provenance and collection methods of training data, how this data was curated and labelled prior to model training, how results were interpreted, and how predictions made by the model were applied in the investigation.
- **Data Provision and Access**: The use of evidence seized as part of real-world law-enforcement investigations as training data for machine learning algorithms must be carefully considered from legal, ethical and technical perspectives. Relevant legislation in individual jurisdictions will govern how this data can be processed.
  From a purely technical perspective, more training data is better. However the transnational nature of technology facilitated crime poses challenges for the interchange of potentially sensitive data between countries. This is a challenge borne not just of legal restrictions on the export of such data, but also of security and logistical data management concerns.
  Academic researchers outside of law enforcement agencies do not typically have legal access to data held by police. This ostensibly restricts their capability to develop, train, test, and compare machine learning models on such data. While partnerships between academia and law enforcement such as AiLECS do provide a platform for research collaboration, provision and processing of

the training data itself remains fraught from a legal perspective. Additionally, in the case of highly distressing content such as CEM, it is important to be mindful of the psychological harm that may be inflicted on all of those dealing with the data or even with the concepts the data implies.

- **Data Labelling** The labelling of data for training supervised machine learning models in law enforcement poses some particular challenges. In addition to the resourcing challenges of labelling very large volumes of data, potentially across jurisdictions, there are considerations around effecting algorithmic bias through subjective labelling (alluded to earlier). Further, it may be the case that investigators are required to deal with evidential classification schemes that may not be as amenable to machine learning training (e.g. less objective, less distinct classes [3])

The AiLECS lab was created address these and other challenges of AI in law enforcement in mind, and to provide a scalable platform for research and development in this area. Efforts targeting algorithm development, sensitive data management, and national and international collaboration have been core to this mission, in addition to development of goals around ethical use.

Building on its initial motivating remit of countering CEM, the AiLECS lab has prioritised a number of areas of focus, in particular:

- Illicit image and video machine learning classification
- Image localisation (estimating image geolocation by content)
- Scalable near-duplicate image detection
- Natural language processing, in particular of short-text documents
- Explainable algorithms and auditable performance techniques
- The ethics of AI in law enforcement

A key objective of AiLECS is to open source as much of our work as is possible in order to rapidly disseminate research and bolster development in the area.

C. Related Initiatives

Police agencies around the world have investigated, and to varying degrees implemented, AI related technologies. This includes among other areas of application: facial recognition, optimised resource allocation, crime prediction, traffic policing and text/social media analyses.

Much of this has been in conjunction with commercial vendors, while direct collaborations between law enforcement and universities is often ad-hoc. Nevertheless, a number of examples of research collaboration between law enforcement and academia exist. Typically the applications of these initiatives are closely related to the local needs of the relevant jurisdiction. For example:

- The University of Cambridge and Durham Constabulary in the UK have developed a random-forest based reoffending risk model [7]
- Also in the UK, the Turing Institute has worked with the West Midlands Police to study the ethics of data analytics at scale for policing [8]
- The Netherlands Police Lab AI [9] is a collaboration between the Dutch Police, Utrecht University and the University of Amsterdam, and is a close (independently developed) analogue of AiLECS. The lab actively researches both the technical and social/ethical aspects of AI in law enforcement.

III. OPERATIONALISING *AiLECS - Data Airlock Platform*

In order to deal with the challenge of managing sensitive, distressing and legally restricted data, an innovative *Data Airlock* platform [3] is under continued development within AiLECS. This infrastructure is intended to provide controlled and configurable access to large law enforcement datasets for researchers. Such a platform is necessary in order to scale up research in this area, particularly since international collaboration will be vital to further address the large scale technical challenges inherent in combating criminal network activity.

Specifically, the current data airlock platform comprises three zones within its underlying hardware, namely *public*, *sensitive* and *secure*. There is no reliance on any particular underlying hardware or cloud environment, with the platform being highly containerized. The three zones operate under different security and access models.

The raw sensitive data resides in the *secure* zone and runs models on the data in isolated virtual environments ”airlocks”. Layers of encryption and exfiltration controls are in place and configurable, and uploading of data to the *secure* zone is only possible with physical hardware access, with this zone isolated from the broader network.

The *sensitive* zone provides an environment for models submitted to the platform by researchers to be vetted by authorised personnel before they are allowed to execute on the sensitive data in their own airlock in the *secure* zone. This vetting can be either through manual review or via semi-automatic or automatic tools depending on the application.

The *public* zone is a web-enabled environment through which researchers can submit models and view the results of their execution, and any other allowable meta-data as determined by authorised personnel.

The current data airlock infrastructure, given the focus on deep-learning based illicit image classification is based on a combination of secure servers with an NVidia DGX-1 platform providing execution of models within the *secure* zone.

IV. OPERATIONALISING *AiLECS - Algorithm Development*

The AiLECS lab represents a culmination and platform for extension of research into algorithmic approaches to law enforcement already undertaken, some of which is being operationalised by the AFP [10]. We describe below some of this work.

A. Monte-Carlo Filesystem Search

It is often the case that rapid triage of data on seized devices is required in the course of a police investigation. There may be a range of time constraints imposed - by law and/or operational resourcing. Additionally, the search may be undertaken in highly bandwidth constrained environments. Thus, techniques that can decrease the time taken to find potential criminal evidence through prioritised search algorithms are of particular value.

In [11] we describe the *Monte Carlo Filesystem Search* (MCFS), inspired by the success of algorithms that apply *Monte Carlo* methods to searching game trees. The algorithm was designed to be lightweight for in-the-field investigative enviroments and to explicitly incorporate investigative domain knowledge. As with game tree searches, the algorithm leverages file system tree exploration (new branch selection) and exploitation (expansion of visited branch) with the configurable incorporation of bias towards more probable locations of interest. Lightweight machine learning algorithms (e.g. Multinomial Naive Bayes), were used to train the scoring algorithms for filesystem node selection depending on the type of application.

MCFS was evaluated in a realistic investigative setting by training on real case data from police investigations, with speedups of around one third seen in typical file system search scenarios looking for known items of interest compared with uninformed file searching [11]. The extension to this work for web-search and image based similarity search is ongoing.

B. Stonefish Classifier

In [3], we introduce a classifier, "Stonefish" based on deep neural networks to test the capacity of such architectures to reliably identify CEM. Although other such classifiers have been proposed and tested for this task, see for example [12] and [13], a typical bottleneck is access to real world case data. In both of those works, such data was sourced from the Brazilian Federal Police under controlled conditions. In [13] this involved a sandbox approach whereby only feature vectors were exposed to researchers. This is similar in spirit to our data airlock approach, however we provide a more generalised platform to enable research scalability and obviate the need for adhoc virtual environment construction. The Stonefish classifier architecture consists of three levels, namely:

- *Pornography detection:* The first level aims to identify pornographic content with high confidence. This is a mature area of research, so we adapted the pre-trained OpenNSFW [14] model for our environment and deployed it on our test corpus of data sourced from real-world contemporary AFP online child exploitation investigations (with appropriate legal and welfare controls in place [3]). This allowed us to partition the dataset to focus on those images most likely of interest to law enforcement and assess them for possible CEM.
- *Child detection*: We trained a deep neural network model to perform a binary classification task - to label images previously assessed as pornographic as either containing children or not. This is in itself a very difficult task. The estimation of age in images (generally containing faces) is a very active area of research (for a survey of approaches see [15]). In our case, we did not seek to estimate age beyond the very broad categorical assessment of child vs adult. In this layer, our approach utilised a deep convolutional neural network (CNN) pretrained on the ImageNet 1000 class, with the top stack of fully connected layers removed and retrained on images previously labelled according to whether or not they depicted children. Again, training data was provided by the AFP under carefully controlled conditions. While datasets containing illegal and distressing material must not be distributed, it is useful to note also at this point that there is a lack of available standard training sets of completely legal images of children. This is of course understandable given online safety and privacy and ethical concerns.
- *Standardised classification*: The third layer of the Stonefish classifier is tasked with assigning suspected CEM images to the appropriate standardised schema; in our case the CETS scheme, a *defacto* standard in use in Australia and various other jurisdictions [3]. This schema is used as the basis for determining the severity of offending in that it provides labels indicating the nature of the activity depicted. In this layer, a deep CNN architecture is also used to perform this multi-class classification, with training data sourced from real-world AFP investigations. We also tested the model on images downloaded during a random traversal of the Tor network [16]

Results of the Stonefish classifier showed that with this early iteration of the architecture, we could achieve overall accuracies of around 60-70% in identifying real world CEM. This indicates that significant further work is required and is ongoing. However results were certainly adequate for initial triage prior to more in depth examination, thus going some way towards reducing investigator burden.

C. Jurisdictionally Independent CEM Training Schemata

What became particularly striking during our initial work on classifiying CEM was the inadequacy of schemas such as CETS for training machine learning algorithms. There is a lack of standardised terminology in legislation around child exploitation, and schemas such as CETS tend to be overly broad and vague. This is understandable from the point of view of the initial motivation of the schemas - which was to inform categories of sentencing. However this abstract nature is not suitable for machine learning training, where it is valuable to have as little ambiguity as possible (particularly so when large quantities of data and many human labellers are involved). We therefore proposed, also in [3], the *Majura* schema, which provides a manifestly objective and detailed set of categories for labelling CEM that not only covers broad types of activities but also other features such participants, props, subtypes of activities etc. It is our hope that this schema, independent of jurisdictional constraints, will be able to, in addition to the data airlock platform, rapidly accelerate collaborative international development of AI techniques in this area.

V. Discussion and Learnings

The AiLECS lab is scaling up its operations, building on the existing infrastructure and work we have already described. We are broadening the nature of our research and actively seeking collaborators. In doing so, we reflect on our experience so far and on the way forward.

A. The University/Law Enforcement AI partnership

Universities have for many years been working closely with industry on applied research. A variety of models exist for these collaborations, depending to a large extent on the government funding models in place in various countries in addition to the overall strategic vision of the partners. The integration of academic research at scale with law enforcement agencies is relatively new, and in the specific area of AI research, provides great opportunity, albeit subject to a number of important considerations.

Law enforcement agencies are by their nature highly operational and agile, needing to respond rapidly to changing environments to counter threats to community safety. Resource allocation is crucial. It is important that performance indicators relating to collaborative research with universities are formulated such that results of research are as directly and easily applied in practice as possible. For example, modern policing relies heavily on ICT infrastructure tailored to operational requirements, including forensics and intelligence capabilities. From a technical point of view, AI systems developed through collaborative research should be constructed within a mindset of integration with existing systems and workflows.

However, this is not to say there is no place for fundamental longer term research under such arrangements. Indeed this is very much a value-add that is provided through such collaboration given the mission of university research. This forms part of the answer to the question as to why a law enforcement agency would want to partner with a university. In our view, this is further answered through considering the following:

- *Interdisciplinarity:* Collaborations such as AiLECS leverage the natural interdisciplinary capabilities of Universities. In the domain of AI, and in particular when applying AI in the high-stakes environment of policing, computer science is not the only discipline that will inform research. Large research oriented academic institutions bring new perspectives by virtue of their broad topical remit and base of expertise. Naturally, from a technological point of view, computer science plays a crucial role. However, it is clear that to advance AI in law enforcement, research from the social sciences can provide valuable legal, ethical and criminological perspectives on the use of the technology. Further, research in bioinformatics, pharmacology, chemistry, ballistics, and a number of other fields directly relevant to law enforcement operations will likely see tighter integration with data driven AI techniques.
- *Capacity building:* Tighter collaboration between law enforcement and universities through research leads to a cross-pollination of expertise. Not only does this assist agencies to adapt to technological change, but promotes broad understanding of the issues faced by police in the research and higher education sector. This in turn strengthens the community partnership on which law enforcement is best-based.
- *Non-commercial imperatives:* Commercial ICT vendors have been vital partners of law enforcement agencies, and have enhanced police work with a number of tools and systems. However, a mixture of commercial and non-commercial (e.g. government/university) partnerships may avoid risk of vendor lock-in and over-reliance on commercially driven products. This is potentially a particular concern with technologies such as AI which are currently hype-driven in some industry sectors.
- *Research culture:* Universities are typically research oriented and have highly developed infrastructure around the management of data, research student training and supervision and appropriate ethical oversight.

B. The Future: Effective, Explainable and Ethical AI

This is a pivotal time for the application of AI, with research efforts and increases in computational resources driving the technology forward apace. In order to harness AI for application in law enforcement, we believe platforms such as AiLECS will provide a crucial clearing house for research in the area. This can only take place however if the three tenets of *effective* (meeting mission objectives and enhancing cability), *explainable* (transparent and accountable) and *ethical* (respecting the rights, privacy and agency of humans) AI are adhered to.

- *Effective* Our AI development is underpinned by the further development of the data airlock infrastructure for the safe access to real data. We intend to open the data airlock for use by researchers anywhere in the world who wish to test their models against real world data they would not otherwise be able to access. We hope this will also build international collaboration that will further enhance the combating of large scale transnational criminal activity.
- *Explainable* It is important that the entire pipeline of AI application in law enforcement, from methods of data collection and curation, labelling, storage, cleaning, training through to model construction, operation and prediction is as transparent as possible. It is not unforeseeable that the future will see an increase in AI algorithms requiring defence in courts, should they be used in decision support. We are thus working on building frameworks against such transparency, in addition to investigating how explainable AI (XAI) techniques can be applied and improved.
- *Ethical* AI for law enforcement is a focus for AiLECS. We alluded earlier to this issue as a challenge for AI in

law enforcement and, particularly, the danger in ethics frameworks as checklists of overly broad and potentially banal statements such as "do no harm", or "be fair". The more pressing challenge is the mapping of ethics against operational requirements of law enforcement agencies in real environments and building this into tools, data pipelines, and juridical processes. We believe that, in the context of AI, building ethical understandings and implementations that are useful in practice is best achieved with law enforcement practioners working closely with ethicists and researchers, in the context of actual case studies—a trajectory AiLECS will follow in further research in this domain.

VI. Conclusion

We have introduced the AiLECS research laboratory and outlined our work and vision for collaborative research between academia and law enforcement. Squarely in the domain of "AI for good", the work of AiLECS will accelerate progress towards achieving the UN Sustainable Development Goals pertaining to good health and wellbeing, peace, justice and strong institutions as well as those around partnerships for achieving the goals. Our data airlock infrastructure and open source goals are designed to rapidly operationalise research efforts and build international teams to assist in researching and developing ethical AI for community safety applications.

To this end, we actively seek collaboration on a number of fronts. Firstly, with AI researchers and law enforcement agencies, to enhance and scale-up these initiatives as we have described. Secondly, on a cross-disciplinary basis, seeking partnership with domain experts in law, criminology, social-sciences and so on, to ensure ongoing alignment and prioritisation in our problem domain. And finally with ethicists, to ensure the application of these technologies remains consistent with community expectations.

References

[1] "Artificial intelligence for law enforcement and community safety (ailecs) lab." https://www.monash.edu/it/ailecs/, 2019.

[2] E. Bursztein, E. Clarke, M. DeLaune, D. M. Elifff, N. Hsu, L. Olson, J. Shehan, M. Thakur, K. Thomas, and T. Bright, "Rethinking the detection of child sexual abuse imagery on the internet," in *The World Wide Web Conference*, pp. 2601–2607, 2019.

[3] J. Dalins, Y. Tyshetskiy, C. Wilson, M. J. Carman, and D. Boudry, "Laying foundations for effective machine learning in law enforcement. majura–a labelling schema for child exploitation materials," *Digital Investigation*, vol. 26, pp. 40–54, 2018.

[4] M. Ashcroft, L. Kaati, and M. Meyer, "A step towards detecting online grooming–identifying adults pretending to be children," in *2015 European Intelligence and Security Informatics Conference*, pp. 98–104, IEEE, 2015.

[5] A. Jobin, M. Ienca, and E. Vayena, "The global landscape of AI ethics guidelines," *Nature Machine Intelligence*, vol. 1, no. 9, pp. 389–399, 2019.

[6] T. Hagendorff, "The Ethics of AI Ethics: An Evaluation of Guidelines," *Minds and Machines*, vol. 30, pp. 99–120, Mar. 2020.

[7] M. Oswald, J. Grace, S. Urwin, and G. C. Barnes, "Algorithmic risk assessment policing models: lessons from the durham hart model and 'experimental'proportionality," *Information & Communications Technology Law*, vol. 27, no. 2, pp. 223–250, 2018.

[8] Turing Institute and Independent Digital Ethics Panel for Policing, "Ethics advisory report for west midlands police." https://www.turing.ac.uk/research/publications/ethics-advisory-report-west-midlands-police, 2018.

[9] Innovation Center for Artificial Intelligence, "National police lab ai." https://icai.ai/police-lab-ai/, 2018.

[10] Australian Federal Police, "Australian federal police, annual report 2017-2018." https://www.afp.gov.au/annual-report-2017-18, October 2018.

[11] J. Dalins, C. Wilson, and M. Carman, "Monte-carlo filesystem search–a crawl strategy for digital forensics," *Digital Investigation*, vol. 13, pp. 58–71, 2015.

[12] P. Vitorino, S. Avila, M. Perez, and A. Rocha, "Leveraging deep neural networks to fight child pornography in the age of social media," *Journal of Visual Communication and Image Representation*, vol. 50, pp. 303–313, 2018.

[13] M. Islam, A. N. Mahmood, P. Watters, and M. Alazab, "Forensic detection of child exploitation material using deep learning," in *Deep Learning Applications for Cyber Security*, pp. 211–219, Springer, 2019.

[14] J. Mahadeokar, S. Farfade, A. Kamat, and A. Kappeler, "Open nsfw model," *Open NSFW github repository*, 2015.

[15] R. Angulu, J. R. Tapamo, and A. O. Adewumi, "Age estimation via face images: a survey," *EURASIP Journal on Image and Video Processing*, vol. 2018, no. 1, p. 42, 2018.

[16] J. Dalins, C. Wilson, and M. Carman, "Criminal motivation on the dark web: A categorisation model for law enforcement," *Digital Investigation*, vol. 24, pp. 62–71, 2018.