

Using unsupervised machine learning to find profiles of domestic abuse perpetrators

Katerina Hadjimatheou², Alejandro Quiroz Flores³, Ruth Weir^{1,2},, and Taylor Skevington⁴,

¹Violence and Society Centre, School of Policy & Global Affairs, City St George's, University of London, Northampton Square, London EC1V 0HB, UK. E-mail: ruth.weir@city.ac.uk

²Department of Sociology and Criminology, Faculty of Social Sciences, University of Essex, Wivenhoe Park, Colchester CO4 3SQ UK

³Centre for Research in Public Health and Community Care, School of Health and Social Work, University of Hertfordshire, College Lane, Hatfield, AL10 9AB, UK

⁴Department of Sociology and Criminology, Faculty of Social Sciences, University of Essex, Wivenhoe Park, Colchester CO4 3SQ UK

ABSTRACT

In this article we use unsupervised machine learning to discover hidden structures and patterns in a longitudinal police dataset of domestic abuse suspects, to provide a police force with an overarching or 'baseline' picture of how domestic abuse manifests locally. 3 algorithms were used to analyse 12 variables in a longitudinal dataset of over 40,000 suspects, organising them into discreet "clusters" or profiles with common characteristics and highlighting the differences and continuities between these. The quantitative findings, which highlighted clusters of abuse that had not previously been 'on the radar' of domestic abuse services in the specific force area, were then contextualised through qualitative interviews with a range of stakeholders to help identify priorities for intervention and further research. Our study shows how cutting-edge quantitative methods can be applied to improve understanding of prevalence and features of police-recorded abuse; draw attention to previously under-addressed types of abuse; serve as the groundwork for further, more in-depth research; and provide an evidence-base for local decision-making.

INTRODUCTION

Today it is widely accepted that reducing the harms of domestic abuse requires a coherent and evidence-based approach to dealing with perpetrators, though longstanding efforts to safeguard victims also remain vital. In August 2021 the then UK Home Secretary reaffirmed the government's commitment to prioritizing the focus on those responsible for domestic abuse crimes, declaring that:

we must deepen our understanding of who commits them, why they do so, and how it may escalate . . . to better understand key behaviours so we can put a stop to them for good.¹

Nowhere is the need for such an understanding more urgent than in policing. Police have a key role to play in tackling domestic abuse, both by protecting victims and survivors and

by intervening with perpetrators to stop the harm. Domestic abuse makes up almost one in six crimes reported to police, and a third of crimes of violence against the person (ONS, 2023). It is therefore vital that police forces understand the prevalence and nature of abuse within their jurisdictions, so that their limited resources can be targeted effectively in an evidence-based approach.

Police computer systems are sources of rich longitudinal data on the people and types of abuse recorded by officers in their daily work, but police forces do not typically have the capacity, skills, infrastructure, or resources to analyse this data systematically or to standards recognized as rigorous by the scientific community. The SARA (Scanning, Analysis, Response and Assessment) model of policing is an internationally recognized approach to addressing problems of harm and crime in a systematic and preventive way, rather than through reactive responses that only provide short-lived solutions. It involves four stages that can be repeated in an iterative process: scanning through identifying, prioritizing, and selecting problems that need addressing; analysis of the problem; response on the

¹In a speech on 11 August 2021, cited at: <https://www.gov.uk/government/news/police-awarded-113m-for-programmes-to-prevent-domestic-abuse-crimes>

ground; and assessment including evaluation.² Our research helps address the first of these stages, by using cutting-edge data science techniques to help one police force use their data to better understand the problem posed by domestic abuse in their area and make evidence-based decisions about how services and interventions should be resourced, prioritized, and commissioned. Our approach, which to our knowledge is the first of its kind,³ used unsupervised machine learning to discover hidden structures and patterns in a large dataset of domestic abuse suspects, organizing them into discreet ‘clusters’ or profiles with common characteristics and highlighting the differences and continuities between these. This provides police with an overarching picture of how domestic abuse manifests locally, which we contextualize through qualitative interviews with a range of stakeholders to help identify priorities for intervention and further research.

The paper is organized as follows. First, we situate our study in the broader landscape of quantitative and machine learning approaches to analysing large domestic abuse datasets and motivate our interdisciplinary approach. Then we describe the characteristics of the data, the interdisciplinary and collaborative research methods, and the stages of analysis. Next, we present our findings, which identified three distinct clusters and one subcluster of domestic abuse suspects in Essex. The paper closes with proposals and recommendations for commissioning of services and further research.

UNDERSTANDING THE FIELD: USING QUANTITATIVE AND COMPUTATIONAL METHODS TO ANALYSE DOMESTIC ABUSE DATA

We use a broad definition of machine learning (Mitchell, 1997, p. 2): ‘A computer program is said to **learn** from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E ’. Existing approaches to using machine learning for domestic abuse research have largely relied on supervised learning techniques. Supervised learning uses labelled datasets to train algorithms that to classify data or predict outcomes accurately. The aim of most supervised learning in domestic abuse research has been to predict a particular outcome variable—such as the number of crimes committed by an individual—using a set of inputs, such as the age of such an individual. Studies utilizing supervised learning methods have yielded insights into the predictors of escalation and reoffending for risk assessment and forecasting (Adisa *et al.*, 2021; Cunha *et al.*, 2023; Goldstein *et al.*, 2016; Hilton *et al.*, 2004; Kerr *et al.*, 2017; Messing and Thaller, 2013; Turner *et al.*, 2019). However, the vast majority of these focus exclusively on data relating to intimate partner violence (IPV), leaving out other kinds of domestic abuse.

²See the College of Policing’s guidance here: <https://www.college.police.uk/guidance/problem-solving-policing>. Last Accessed 10/07/2024.

³To our knowledge unsupervised machine learning has only been used with police data on domestic abuse by researchers in Australia using text analysis to explore prevalence and gender differences in mental health factors and aspects of victimization (Karytsiannis, 2020, 2022). For a review of machine learning research in domestic abuse see Hui *et al.* (2023).

The few that do consider all domestic abuse incidents tend to treat non-IPV or ‘family violence’ as a relatively homogenous category, rather than a bundle of quite diverse relationships and types of abuse (Goldstein *et al.*, 2016; Holtzworth-Munroe and Stuart, 1994; Johnson *et al.*, 2006; Saunders, 1992). The fact that the legal definition of ‘domestic abuse’ in the UK conflates this already diverse category of ‘family violence’ with IPV is a further challenge when it comes to doing statistical analysis of domestic abuse data.

Our approach differs from supervised learning in a number of important ways. First and foremost, it does not aim to predict anything. Rather, it is exploratory, looking for hidden patterns and therefore meaning in the data. It does this by utilizing unsupervised learning to find groups over *all* features of the data. In unsupervised learning

the researcher feeds unlabelled data to a learning algorithm and allows patterns to emerge, typically based on similarity among observations (within-group homogeneity) and dissimilarity between groupings of observations (between-group heterogeneity) (Wagonner, 2020, p. 1).

In other words, it sorts data—in our case, 12 variables relating to suspects, victims, and the incidents in which they are implicated—into groups or ‘clusters’ whose members are both as similar to each other as possible (internally homogenous) and as different from those in other groups as possible (externally heterogenous). The features of these groups are then analysed against the background of current knowledge about different typologies of abuse and distinct profiles of abusers in order to yield new insights around what kind of people are inflicting what kind of abuse on whom.

Our approach shares some of the features of statistical research methods that seek to distinguish different typologies of perpetration (Holtzworth-Munroe and Stuart, 1994; Johnson *et al.*, 2006). But unlike those, it does not begin with any assumptions or hypotheses about the features of any eventual clusters. Existing approaches tend to begin by investigating the relationship between two or more variables already predetermined as interesting (e.g. age and risk or gender and harm). Or they set out to identify the factors most strongly correlated with, for example serious violence. Instead, our approach starts with 12 variables and lets the data ‘speak for itself’ on the relationship between them, allowing for the emergence of latent profiles, types, and nuances, as well as revealing relative prevalences. Traditional approaches would not be able to look at the relationships between 12 variables simultaneously without taking far longer to complete and producing far less reliable results.⁴ Our approach provides a baseline analysis of force-level data that can serve as an evidence-based justification for devoting resources to investigating specific relationships between variables, such as age and risk. It also provides a rigorous basis for understanding the demand on police

⁴Other traditional quantitative methods could use correlation to organize the same data. However, correlations can only be calculated in pairs, and therefore this traditional analysis would need to explore all possible permutations for all 12 variables included in the analysis. This would result in 132 correlation coefficients (measures

resources in a specific force area and for the commissioning of domestic abuse-related services, such as specialist Independent Domestic Violence Advocates (IDVAs)⁵ and perpetrator-facing services. Finally, our method also acts as a triangulator of hypothesis-driven research, lending rigorous empirical weight to findings from studies that begin with contested assumptions or hypotheses about which kind of people commit what kind of abuse.

METHODS

Our analysis is based on research methods and stages which we developed collaboratively with Essex Police, a regional force in the east of England, who were our main stakeholders for this research and who provided the longitudinal dataset. Essex is one of the largest non-metropolitan police forces in the UK, with a slightly older than the average population (the average age is 41 compared to 40 nationally) and a largely white population at 88.8%, with 4.2% Asian and 3.4% Black, in contrast to its far more diverse neighbour, London. The area has a higher-than-average rate of domestic abuse as evidenced by the number of domestic abuse-related incidents and crimes recorded by the police. In the year ending March 2021, 41,698 incidents and crimes were recorded and 20% of all recorded crimes were classified as domestic abuse-related as compared to 18% nationally (Office for National Statistics, 2021). In the same year, 71% of domestic abuse incidents were subsequently recorded as crimes compared to an average of 58% nationally (Office for National Statistics, 2021). Similarly, the arrest rate in Essex is 38 arrests per 100 domestic abuse-related crimes compared to 32 arrests per 100 nationally (Office for National Statistics, 2021). Essex had a number of high-profile homicides making it the second-highest

of the strength of a relationship between two variables) which would be impossible to make meaning out of. Regression analysis could instead be used to facilitate analysis, but this would require making strict assumptions about a data-generating process, including the definition of a dependent or output variable, and its systematic relationship with independent variables or inputs, and the role and characteristics of a random component that disturbs the relationship between inputs and outputs. For example, one of many potential dependent variables could be the number of crimes committed by a perpetrator. In this very specific regression setting, the number of crimes might be determined by all other 11 independent variables and a random component. This traditional regression setting requires researchers to make assumptions, such as the definition of a probability distribution for the random component, the functional form of an independent variable (e.g. should the suspect's age be measured linearly, or as a logarithm, or as a polynomial?), the presence of interactions with other variables, and so on. Once a researcher has selected one of an infinite number of potential regression models, one of many possible inferences from this model would focus on, for instance, how an additional year in a suspect's age would increase or decrease the number of incidents, along with its confidence interval; this requires an additional assumption about the level of statistical significance and holding all other variables constant at an adequate level. This exercise would then need to be done for all other 10 variables. Moreover, a researcher could choose other dependent variables, so technically this could be done for all other permutations of dependent and independent variables for specific functional forms. This approach can be useful when its many assumptions are met. However, these assumptions are seldom met, and this leads to incorrect inferences about the relationship between variables. We believe that this is a risky approach to the analysis of a sensitive, violent setting where relationships are confounded and simultaneously determined beyond current theoretical frameworks. Our approach is much more general and explorative. We do not make any assumptions about relationships in the features that make domestic abuse, we do not make any claims about the role or distributions of random components, and we do not even claim that one variable precedes another. Instead, we simply let the data organize itself into clusters determined by an independent machine. In other words, traditional analysis can lead to low variance, yet biased point estimates that could be completely wrong, while our approach, while noisy, does not even produce point estimates to begin with, and instead simply organizes the data.

⁵An IDVA is a trained specialist who provides support to victims of domestic violence or abuse, building a trusting relationship to help them navigate the criminal justice system, support from housing and other agencies, and help with personal issues.

rate in the UK in 2019 and leading to a renewed focus on domestic abuse prevention across various agencies and systems, including the launch of a new strategy prioritizing a focus on reducing perpetration (Essex County Council, 2021). Domestic abuse in Essex is addressed through a partnership approach in the form of the Southend, Essex and Thurrock Domestic Abuse Board (SETDAB). SETDAB is a collaboration of services, agencies, and organizations that work to design and implements domestic abuse strategy across the region and offers leadership and guidance around multiagency working, a co-ordinated approach, and how to drive change.

This project was facilitated by the Essex Centre for Data Analytics (ECDA), a partnership for data sharing and analysis across policing, higher education, and the public sector in Essex, for which one of the authors of this paper is the Chief Scientific Adviser.⁶ ECDA responded to the request by Essex police for research that would help to 'gain a general understanding of the situation in Essex', providing 'descriptive analysis to determine granularity and feasibility of further analysis' and to assess whether findings from a literature review on domestic abuse perpetration carried out by ECDA 'matched what is happening in Essex' (Essex Centre for Data Analytics, 2021). The aim was to perform a 'scanning' exercise (the first step in the SARA model of problem-solving policing) to understand the problem posed by domestic abuse in Essex, with a view to informing what was at the time a forthcoming commissioning cycle. A precursor project funded by ECDA attempted to do this analysis by clustering data on five variables linked to domestic abuse suspects, but resource constraints limited the methodological options available and the resulting analysis was insufficiently rigorous to yield reliable insights.⁷ Our approach was co-produced with Essex Police's Head of Analysis and Research and Head of Domestic Abuse, who also provided ongoing input into the analysis. Our interdisciplinary research team involved collaboration between scholars of machine learning, domestic abuse subject matter experts, and qualitative and quantitative criminologists. Figure 1 outlines the steps in our analysis.

At the end of the data cleaning and pre-processing, we produced a usable dataset with 40,488 observations, covering the years 2016–20. The dataset had 59 variables and constructions, many more than we could analyse given time and resource constraints. The challenge we faced was also methodological: machine learning algorithms often fail when they explore large numbers of features. To address this problem, which is popularly known as the 'curse of dimensionality', researchers often implement dimension reduction techniques, such as principal components.⁸ While dimension reduction is a useful technique, it often produces challenges for the interpretation of results. For instance, these techniques may obscure the individual effect of inputs on an output, thus preventing inference. For this reason, we took a different approach and focussed on 12 features of the data only, which are presented in Table 1.

⁶Professor Alejandro Quiroz Flores.

⁷*Understanding Domestic Abuse Perpetrators: Using Unsupervised Machine Learning to Analyse a Longitudinal Dataset of Domestic Abuse Incidents from Essex Police*, cited at https://repository.essex.ac.uk/35676/1/Understanding%20Domestic%20Abuse%20Perpetrators_University%20of%20Essex.pdf

⁸Principal component is a method that reduces or compresses the number of variables or inputs to create a new, smaller set of features while preserving useful information.




Figure 1: Research stages and methods

These are features that enable insight into the different types of domestic abuse and domestic abuse suspects within Essex, allowing us to disaggregate domestic abuse into distinct crime types. For example, it distinguishes IPV from other, less well-understood kinds of domestic abuse, while also highlighting the continuities between them (as the results show, there are both differences and continuities between types of abuse, because while the clusters organized themselves strongly along lines of abuse type, not all cases of abuse in a single cluster were of the same type).⁹

Our data also includes both crimes and recorded ‘non-crime’ incidents, which account for 42% of the recorded incidents. It is worth reflecting on the presence of ‘non-crimes’ in the data. Non-crimes are an undifferentiated category of unspecified behaviours that do not meet the threshold or criteria for criminality according to the police officers who attend a call. National guidance on the recording of non-crime incidents, including in the context of domestic abuse was published in 2011 (Home Office, 2011). When we queried the category of non-crimes in Essex police data with relevant stakeholders, we were told it may include amongst other things low-level threats and harassment, verbal disagreements that did not meet the threshold of crimes, and third-party reports which were not corroborated by the relevant parties and which could not therefore be recorded as crimes. Previous research in other areas has indicated that police under-record domestic abuse offences perceived as ‘low-level’. Though it only analysed snapshot data from a single force, one study found that police were likely under-recording domestic abuse crimes due to a failure to recognize elements of coercive

⁹The algorithms we use are complex and computationally costly. Of course, different computers can complete these tasks more efficiently or in less time than others. For this reason, computer scientists focus on machine-independent measures of computational complexity, including the Big O measure, which is often a function of the size n of a dataset. For reference, a logit model has a time complexity $O(dn)$, where d is the number of variables. The hierarchical algorithms we use here have a time complexity of $O(n^3)$, which is orders of magnitude more complex than in a logistic model. In an iMac with a 3.5 GHz Quad-Core Intel Core i5 processor, a logistic model is estimated in about 15 s while one of our algorithms takes about 48 h to be completed.

control (Myhill and Johnson, 2016). Essex’s track record on the recording of domestic abuse-related incidents is good, with 71% subsequently recorded as crimes compared to an average of 58% across England and Wales in the year ending March 2021 (Office for National Statistics, 2021). In 2019, Essex police crime recording practices were found to be outstanding by the national police ombudsman, including in relation to domestic abuse (HMIC-FRS, 2019). However, the same report also recommended that the force should improve officer and staff knowledge of stalking and of harassment in a domestic abuse context. This suggests that some incidents might still be being recorded as non-crimes when they are in fact elements of a course of conduct crime.

Some police forces do not link non-crimes to an identified suspect or victim, and most datasets for research into suspects or perpetration exclude these kinds of incidents for that reason. Essex police stakeholders explained to us that they record suspect and victim data in relation to non-crimes in part because they recognize that incidents can form part of a series of incidents or instantiations of behaviour that seen as a whole constitutes a crime. Our view is that the inclusion of non-crimes in our study is valuable, because it can shed light on patterns of abuse, the nature of offending, and escalation pathways. Indeed, as our findings show, relative rates of non-crimes differ significantly between clusters.

We then used agglomerative hierarchical clustering (AHC) to find patterns within the data. AHC starts by pairing individual observations based on similarity to each other and then grouping these individuals into larger clusters in a ‘bottom up’ approach (as distinct from divisive clustering, which begins by subdividing a cluster that includes all observations and then subdividing it again). To do this, we started with four different linkage methods or algorithms offered by the cluster R package, namely single link, complete link, average link, and Ward’s link. Each algorithm pairs data according to a different rule and each produces different clusters and different insights. For example, the ‘complete linkage’ method pairs groups of observations based on the maximum distance between them, while the ‘average linkage’ method pairs clusters according to their mean (dis)similarity. The single link method did not produce distinctive clusters and therefore we focus on the results produced by the complete link, average link, and Ward’s link methods. The use of these three methods gave us more robust and rigorous results than using just one, because we could cross-reference the results across the three methods and find strong consistencies between them in a form of triangulation.

Once the clusters for each algorithm were produced, which took about 8 days of computer run time and resulted in 15 distinct clusters, we analysed them to identify clusters that were independently discovered by more than one algorithm, as these are most likely to correspond to reliable features of the data. Once we had selected the clusters and organized them into groups based on the similarity of their features, we further checked the validity of our results by implementing an analysis of association rules. This is another unsupervised learning method that allows us to find clear groups of perpetrators using probability theory. Results from the complementary analysis of association rules confirmed the main characteristics of suspects in the largest cluster group we identified, lending further rigour to our findings.

Table 1: Variables and measurements

Features	Measurement
Suspect is UK National	Yes = 1, no = 0
Suspect is White British (88.8% of Essex residents are White British)	Yes = 1, no = 0
Suspect is also victim	Yes = 1, no = 0
Intimate partner violence	Yes = 1, no = 0
Many crimes above median (5)	Yes = 1, no = 0
Many victims above median (2)	Yes = 1, no = 0
Suspect age above median (32 years in Essex data)	Yes = 1, no = 0
Victim age above median (32)	Yes = 1, no = 0
Suspect gender	Male/female
Victim gender	Male/female
Risk	High/medium/standard
Crime	Violence against the person (46%), non-crime (42.4%) Damage and arson offenses (4.7%), public order offenses (2.2%) Sexual offenses (1.5%) Theft (1.7%) Misc. crimes against society (1%) Possession of weapon (0.04%) Robbery (0.07%) Non-notifiable (0.03%)

Finally, alongside the quantitative analysis, 18 qualitative interviews were undertaken with practitioners working in a range of services and organizations in Essex dealing with suspects of domestic abuse. We asked participants about the different types of suspect or perpetrator they encounter; if they have identified trends in abuse and how; what their priorities are in terms of perpetrator intervention, support, and commissioning; if they felt there were gaps in their knowledge and understanding or in current provision; what challenges they faced in addressing perpetration in their work. Some of these participants were contacted again to help sense-check emerging findings from the clustering analysis and to reflect on their implications for the local context.¹⁰ Implications for local practice and commissioning were also explored in meetings with our police collaborators, the ECDA Directors' Board, and SETDAB. Key themes from the qualitative research have been incorporated into the findings.

DATA LIMITATIONS

We had to delete a significant amount of data during the cleaning process. We deleted observations where key information was missing, repeated, or miscoded, which occurred in some cases with gender, age, country of birth, level of risk, and ethnicity, among others. Our dataset had 59 variables but we could not analyse all of these simultaneously. In close consultation with our police contacts we narrowed these down to those they felt were most relevant to their scanning needs and those were known to

be reliably recorded. We included demographic characteristics of the suspect, including whether they are a UK national, their ethnicity, age, gender, whether they are also a victim, the number of crimes they have committed, and the number of victims they have perpetrated domestic abuse against. For victims, their age and gender were included. The relationship between the victim and perpetrator was recorded as either IPV or familiar violence. To capture the severity of the incident we used both the risk assigned (high, medium, or low) and the Home Office crime group. There were variables we would have liked to include but could not, such as important but regrettably under-recorded factors such as victim self-harm and presence of children.¹¹ Some of the flagged data, such as whether drugs and alcohol had been used had a lot of missing data, so were excluded from the analysis. Readers with prior understanding of domestic abuse behaviours will also notice that the 'crime' variable we used does not record harassment, stalking, and coercive control, all of which are widely acknowledged as common behaviours of perpetrators. The dataset does include a second crime variable that incorporates these and other crimes. Nevertheless, we decided not to add it because doing so would have risked disaggregating the clusters into ever smaller fragments, resulting in a proliferation of sub-clusters that would confound meaningful insight.¹² For the same reason we could not include the more granular data we had about the precise nature of the relationship between suspect and victim. Unfortunately, the data that was provided did

¹⁰SETDAB (Southend and Thurrock Domestic Abuse Board); Multi Agency Public Protection Arrangements; Essex police; Essex Youth Service; Building Better Relationships Programme Treatment Delivery; Probation; Safe Steps; Multi-Agency Risk Assessment Conference; Changing Pathways; The Change Project.

¹¹Data related to links between the domestic abuse incident and children had more than 52,000 missing values; data related to self-harm behaviour in victims had more than 13,000 missing values.

¹²In other words, this crime category would have begun behaving like a continuous variable, which would present technical problems in the context of our algorithm.

not include the geographical location where the incident took place, or the address of the victim and perpetrator.¹³

Machine learning algorithms require that all observations have information for all relevant variables, and therefore it is often the case that some observations will be lost. This increases the quality of the data by removing distortions and bias that would be introduced by a lack of data, and allows the algorithm to be run successfully. However, there is a possibility that this process introduces a measure of ‘selection bias’ into the sample by producing a dataset that does not represent—and therefore could be relied on in ways that *misrepresents*—the population of interest. We are confident that the deletion process did not introduce this kind of selection bias, because our analysis indicated the deleted observations (observations with missing information for a key input) were distributed randomly rather than patterned.¹⁴

Selection biases are often also introduced at the stage of data collection. Research has shown that domestic abuse crimes are not always recorded correctly and that there is a bias in favour of recording incidents involving physical violence (Myhill, 2018; Myhill and Kelly, 2023). Selection biases with respect to gender have also been documented in police identification and recording of suspects and victims. For example, research has shown that police in some forces have assessed as ‘mutually violent’ relationships in which female uses of violence are in fact defensive and retaliatory against a systematically abusive partner, and that officers have taken at face value men’s claims to be subject to violence from a female partner when these are in fact malicious allegations deployed as a tactic of abuse against a woman (Hester, 2013; Robinson and Rowlands, 2006).¹⁵

Selection bias in police data also occurs because most domestic abuse crime is never reported to the police,¹⁶ and those who do report their abuse are not identical in demographic characteristics to those who do not. For example, a recent report by SafeLives (2019) confirms previous research findings that male victims are less likely than female victims to report their abuse (Felson *et al.*, 2002; Mirrlees-Black, 1998; Walby and Allen, 2004). Selection biases arising from reporting and recording practices in the field of domestic abuse are a challenge common to all studies using police-recorded data. While a measure of reassurance can be taken from the excellent recording practices at Essex police, it is inevitable that some selection bias will characterize our data. As it is not possible to remove this bias, it is important to analyse the results in a way that remains alert to

¹³The only location information was the police ward, which is not coterminous with census geographies such as Lower Super Output Areas.

¹⁴In most cases, the observations that had to be deleted seem to be the result of glitches from data management queries, which are random. In other cases, it is impossible to know why some information is missing, such as age or gender—these could be the result of data collection practices happening under difficult circumstances, failure of data terminals, or problematic data linkages. The fact that our findings remain consistent across multiple types of machine learning algorithms gives confidence to our methods.

¹⁵Ethnicity is a further category in which police-recorded data is likely to be unreliable or biased. A recent report by His Majesty’s Inspectorate of Constabulary and Fire and Rescue Services (HMICFRS) found that 12% of cases from a sample of six (unidentified) forces had no information recorded on the ethnicity of the person involved (HMICFRS, 2024). Whilst this remains a problem in historic datasets, HMICFRS have recommended that all forces make changes to their record management systems so that data on self-defined ethnicity is recorded mandatorily by the end of September 2024, which would yield much more reliable datasets for future research (HMICFRS, 2024).

¹⁶According to the Crime Survey for England and Wales, only an estimated 21% of victims report their abuse to the police (Flatley, 2016).

it and cognizant of its potentially distorting effects. We do that in this study by consulting research that uses diverse methods to identify and examine biases in reporting and recording of domestic abuse, by incorporating those findings into our analysis and recommendations, and by avoiding unsubstantiated causal inferences.¹⁷

FINDINGS

Our analysis generated three tables of clusters—one for each algorithm—including a total of 15 clusters. Our interpretation of these findings identified three distinct groups of clusters—or profiles of suspects—which were:

1. Repeat and serial male-to-female IPV (average¹⁸ 45% of suspect data) including a subcluster of younger, more violent suspects.
2. Female-to-male IPV (average 12.5% of suspect data).
3. Repeat and serial familial abuse (average 11.5% of suspect data).

The clusters are represented graphically, in Fig. 2 and then analysed in turn. Detailed quantitative analysis of the clusters, including the cluster tables produced by the algorithm, can be found in the full report to the Home Office.¹⁹

Cluster group 1: Repeat and serial male-to-female IPV

This is the largest group of suspect clusters, containing an average of 45% all suspect data. Its features correspond with what most people would associate with the term ‘domestic abuse’ and what is known to be the most prevalent type of domestic abuse criminality. This cluster group was identified strongly by all three of the algorithms. It consists of mainly white, UK national, male suspects with multiple offenses, multiple female victims, and violent crimes against an intimate partner. Over half of these suspects have a record of more than five domestic abuse (DA) crimes against more than two victims in the 4 years covered by the data. However, as our data was limited to a specific time period of 4 years, we will not have captured the full extent of recorded repeat or serial perpetration which began earlier and/or continued after that period. This means our findings are likely to underestimate the extent of serial and repeat perpetration. Table 2 presents the dominant features of the data in red and uses yellow to show where the total for other features outstrips the dominant feature.

For the largest of these clusters (Table 1, Cluster 2, accounting for 67% of the suspect data in that cluster group) around half of the suspects also appear in the data as victims (Table 3).

¹⁷Bias in police data can also be exacerbated in research using machine learning techniques when the ‘learning’ aims to categorize crimes or people and to predict outcomes and in doing so replicate, and even amplify, selection biases in the data they are fed (Ayre and Craner, 2018). As our analysis does not generate causal inferences or predict where or between whom crime is likely to occur, this specific kind of bias is not a challenge for our study.

¹⁸The average across the clusters in the groups.

¹⁹*Understanding Domestic Abuse Perpetrators: Using Unsupervised Machine Learning to Analyse a Longitudinal Dataset of Domestic Abuse Incidents from Essex Police*, cited at https://repository.essex.ac.uk/35676/1/Understanding%20Domestic%20Abuse%20Perpetrators_University%20of%20Essex.pdf

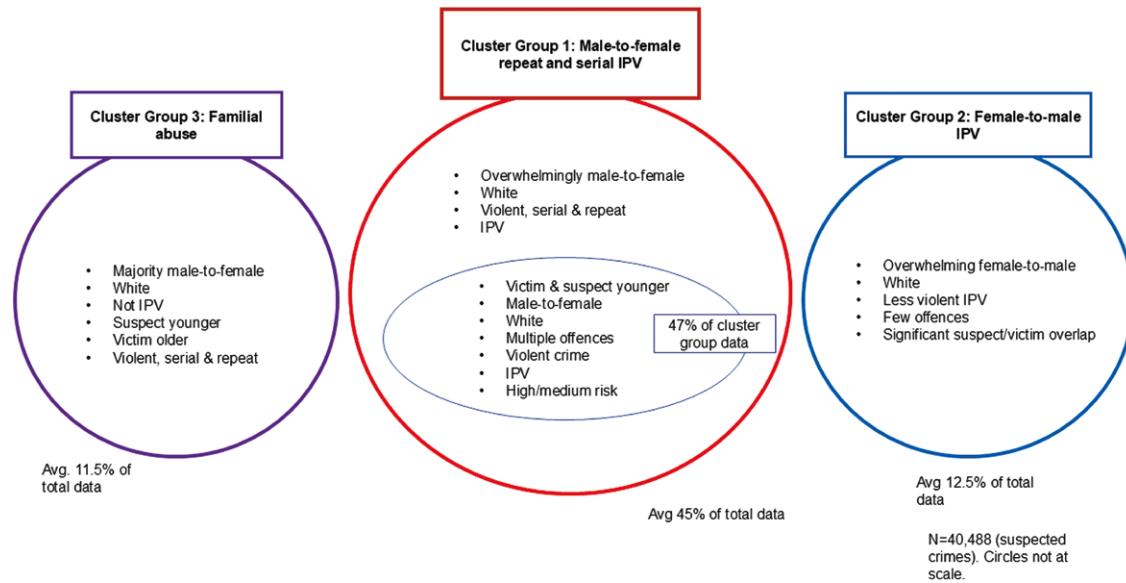


Figure 2: Cluster groups of domestic abuse perpetrators

Table 2: Repeat and serial male-to-female intimate partner violence data

Feature	Table 1, Cluster 2	Table 2, Cluster 2	Table 3, Cluster 2	Table 3, Cluster 3
Sus gender	Male: 27,099 Female: 56	Male: 15,035 Female: 15	Male: 16,499 Female: 114	Male: 13,852 Female: 121
Sus UK National	Yes: 26,951 No: 204	Yes: 12,890 No: 2,160	Yes: 16,479 No: 134	Yes: 11,225 No: 2,748
Sus White British	Yes: 26,074 No: 1,081	Yes: 11,847 No: 3,203	Yes: 16,533 No: 80	Yes: 9,331 No: 4,642
Sus also victim	Yes: 13,888 No: 13,267	Yes: 3,377 No: 11,673	Yes: 9,290 No: 7,323	Yes: 5,677 No: 8,296
Victim gender	Male: 27 Female: 27,028	Male: 26 Female: 15,024	Male: 20 Female: 16,593	Male: 93 Female: 13,880
Crime	Non-crime: 12,011 Violence against the person: 11,889	Non-crime: 6,862 Violence against the person: 6,437	Public order offenses: 7,097 Violence against the person: 7,288	Public order offenses: 6,142 Violence against the person: 6,508
IPV	Yes: 26,416 No: 739	Yes: 14,656 No: 394	Yes: 16,118 No: 495	Yes: 13,332 No: 641
Risk	High: 5,058 Medium: 8,100 Stand: 13,997	High: 3,557 Medium: 4,390 Stand: 7,103	High: 3,935 Medium: 5,087 Stand: 7,591	High: 2,112 Medium: 3,876 Stand: 7,985
Many crimes (>5)	Above: 17,346 Below: 9,809	Above: 9,077 Below: 5,973	Above: 15,451 Below: 1,162	Above: 3,474 Below: 10,499
Many victims (>2)	Above: 13,803 Below: 13,352	Above: 7,103 Below: 7,947	Above: 12,244 Below: 4,369	Above: 3,009 Below: 10,964
Sus age (m = 32)	Above: 14,509 Below: 12,646	Above: 6,588 Below: 8,462	Above: 7,654 Below: 8,959	Above: 9,617 Below: 4,356
Victim age (m = 32)	Above: 12,248 Below: 14,907	Above: 4,663 Below: 10,387	Above: 6,330 Below: 10,283	Above: 7,768 Below: 6,205
Perpetrators	27,155 (67%)	15,050 (37%)	16,613 (41%)	13,973 (35%)

This could indicate a range of things including the presence of mutual abuse, violent resistance on the part of victims, suspects experiencing trauma in other relationships, or a prevalence of malicious counter-allegations. In order to investigate further we

would need to carry out more granular analysis of the cluster to examine more closely any patterns in the incidents between specific victims and perpetrators and any further context provided in the DASH forms.

Table 3: Female-to-male intimate partner violence data

Feature	Table 1, Cluster 4	Table 3, Cluster 4	Table 2, Cluster 5	Table 2, Cluster 7
Sus gender	Male: 349 Female: 4,844	Male: 33 Female: 4,787	Male: 1,680 Female: 2,653	Male: 13 Female: 420
Sus UK National	Yes: 4,756 No: 437	Yes: 4,437 No: 383	Yes: 4,313 No: 20	Yes: 104 No: 329
Sus White British	Yes: 4,600 No: 593	Yes: 4,309 No: 511	Yes: 4,227 No: 106	Yes: 87 No: 346
Sus also victim	Yes: 4,370 No: 823	Yes: 4,300 No: 520	Yes: 3,846 No: 487	Yes: 260 No: 173
Victim gender	Male: 5,077 Female: 116	Male: 4,781 Female: 39	Male: 2,005 Female: 2,328	Male: 399 Female: 34
Crime	Non-crime: 2,213 Violence against the person: 2,586	Public order offenses: 2,078 Violence against the person: 2,388	Non-crime: 1,807 Violence against the person: 2,119	Non-crime: 183 Violence against the person: 236
IPV	Yes: 4,698 No: 495	Yes: 4,630 No: 190	Yes: 2,873 No: 1,460	Yes: 360 No: 73
Risk	High: 198 Medium: 1,083 Stand: 3,912	High: 156 Medium: 973 Stand: 3,691	High: 381 Medium: 947 Stand: 3,005	High: 22 Medium: 92 Stand: 319
Many crimes	Above: 1,789 Below: 3,404	Above: 1,811 Below: 3,009	Above: 2,502 Below: 1,831	Above: 12 Below: 421
Many victims	Above: 1,989 Below: 3,204	Above: 1,924 Below: 2,896	Above: 3,028 Below: 1,305	Above: 45 Below: 388
Sus age	Above: 2,722 Below: 2,471	Above: 2,511 Below: 2,309	Above: 1,796 Below: 2,537	Above: 270 Below: 163
Victim age	Above: 3,184 Below: 2,009	Above: 3,006 Below: 1,814	Above: 983 Below: 3,350	Above: 355 Below: 78
Perpetrators	5,193 (13%)	4,820 (12%)	4,333 (11%)	433 (1%)

A notable subgroup of this cluster group is seen in Table 2, Cluster 2, which accounts for 37% of the data for the group. This cluster is demographically distinct, consisting of younger-than-average males using IPV against younger-than-average females, with most incidents being rated medium or high risk. Here younger-than-average is defined in terms of the dataset rather than population figures. The median age for the dataset is 32 years old, significantly younger than the average age in Essex which is 41 according to the 2021 Census. Cluster 3.2, which also includes a majority of younger-than-average victims, confirms this link between victim age and risk, with 54% of incidents rated high or medium risk.

This cluster had not been previously a focus of police and domestic abuse services in Essex. Young people were not mentioned frequently in the qualitative interviews and were not flagged as a cluster in the precursor study to ours, carried out for the ECDA. There appears to be very little awareness of the heightened risk and violence associated with young people's abuse in Essex. Unlike other force areas in the UK, there is limited provision for young people in Essex,²⁰ no specific IDVA provision for young people who are victims, and no services at all for children and young people who are using abuse in their relationships.

²⁰Since 2019 there has been a female mentoring programme delivered by Goodman and Sisters in Strength in the Southend and Thurrock area, and Break the Cycle is a dedicated CYPVA service for those aged 13–19 years who have witnessed abuse.

Cluster group 2: Female-to-male IPV

This cluster group represents about 12.5% of suspects of DA. It consists mainly of white, UK national, female suspects abusing male victims in the context of an intimate partner relationship. Around half of the incidents are violent, but the other half are recorded as non-crimes or public order offences—a more equal distribution between these two types than for other cluster groups. Risk is more likely to be assessed as standard than medium or high, and more likely to be standard than in other clusters with male suspects. Suspects in the female cluster have fewer recorded crimes and victims than the mean. Notably, between 60 and 70% of suspects in this cluster also appear in the data as victims. These findings confirm, independent of any prior hypothesis about gender differences in domestic abuse perpetration, that female-to-male abuse is distinctive enough across many variables to constitute a self-contained cluster or profile. As discussed in a moment, this has direct implications for the commissioning of perpetrator and victim-facing services, and especially the tendency to repurpose interventions designed for one gender for use with another.

The identification of the female cluster group is particularly significant for this study because our qualitative interviews highlighted the provision for female perpetration as a significant gap in intervention development in the region. There is

Table 4: Repeat and serial familial abuse data

Feature	Table 1, Cluster 1	Table 3, Cluster 1
Sus gender	Male: 2,777 Female: 1,638	Male: 3,546 Female: 1,536
Sus UK National	Yes: 4,364 No: 51	Yes: 4,999 No: 83
Sus White British	Yes: 4,170 No: 245	Yes: 4,810 No: 272
Sus also victim	Yes: 2,842 No: 1,573	Yes: 2,939 No: 2,143
Victim gender	Male: 925 Female: 3,490	Male: 1,216 Female: 3,866
Crime	Non-crime: 1,545 Violence against the person: 2,219	Public order offenses: 1,884 Violence against the person: 2,441
IPV	Yes: 115 No: 4,300	Yes: 478 No: 4,604
Risk	High: 465 Medium: 952 Stand: 2,998	High: 333 Medium: 1,370 Stand: 3,379
Many crimes	Above: 2,595 Below: 1,820	Above: 2,535 Below: 2,547
Many victims	Above: 3,692 Below: 723	Above: 3,553 Below: 1,529
Sus age	Above: 1,478 Below: 2,937	Above: 1,442 Below: 3,640
Victim age	Above: 2,986 Below: 1,429	Above: 3,677 Below: 1,405
Perpetrators	4,415 (11%)	5,082 (12%)

currently no support or rehabilitation provision for female perpetrators from Probation services in Essex.²¹ This was a source of frustration for some Probation officers, who pointed out that the structured programme they offer was adapted from one designed for male perpetrators. This had been developed in-house but was not being delivered due to organizational change. At the same time, one Probation officer complained that in their 4 years of service, they had not received 'a single bit of training on female perpetrators' (P7). This is clearly a gap in Probation knowledge, training, and provision that is unlikely to be unique to Essex. The Change Project does provide a local trauma-informed 1-2-1 female perpetrator intervention service. However, the leaders of that service told us that the lack of research in the field of female perpetration meant they felt their interventions were less evidence-based than those they could offer to men. They also emphasized their view that the differences between male and female abuse typologies spoke strongly in favour of interventions designed specifically for men or women and against the repurposing of interventions across genders,²² a previously anecdotal position that our research supported with evidence.

²¹According to all of our probation participants.

²²In the words of one participant from the Change Project who had previously worked delivering statutory perpetrator programmes: 'In terms of [female clients'] behaviour, it could not be compared to the behaviour of a male perpetrator. Their typology tends to be quite different. I feel really, really passionately that women or perpetrators who identify as female should have equal access to good rehabilitative services so that they can make changes and move forward with their life. And up until now, they just don't.'

Our qualitative interviews revealed that gaps also exist around provision for male victims. Apart from a male IDVA pilot in the Southend area, currently seconding a practitioner to be based in the council offices, there is no specific provision for male victims. Portrayal of domestic abuse services as for females only has been found to be a barrier to seeking help (Huntley *et al.*, 2019), and the three main victim services in Essex are all Women's Aid accredited which means that men cannot attend their buildings. As a result, male victims are currently signposted to Mankind, which is a national rather than a local organization.

Cluster group 3: Repeat and serial familial abuse

The clusters in this group relate to a small cohort containing an average of 11.5% of perpetrator data. About 65% of the suspects are male and 80% of the victims are female—indicating less of a gender divide than in Groups 1 or 4. Crimes are equally likely to involve violence against the person as they are to be public order offenses or non-crimes. But, strikingly, less than 10% of the abuse is IPV. Suspects in this group appear to be prolific and serial offenders, with about 60% committing more than five domestic abuse-related crimes and over 80% abusing more than two victims. However, there is also a significant suspect/victim cross-over with about 60% of suspects also having been victims of DA. Suspects are younger than average and tend to offend against people who are older than average, suggesting that

this data may be capturing child-to-parent and familial abuse (Table 4).

DISCUSSION

Young people and IPV

Nationally, there is very little systematic evidence analysing the prevalence and profiles of victimization and perpetration of domestic abuse amongst young people. There is some small-scale and now-dated research showing that people aged 16–19 are overrepresented in the cohort of victims of domestic abuse,²³ face a higher risk than other age groups, and are less likely to report abuse and receive help.²⁴ And there is growing recognition that people aged around 14–19 can experience domestic abuse (Weir *et al.*, 2024). Police domestic abuse records would not include data for those under 16 as that is the legal age at which domestic abuse is recognized as such. Our analysis lends systematic empirical support to those findings by indicating that young people's perpetration and victimization of IPV in Essex is distinct from that of older adults and therefore merits targeted research and intervention.

Specialist young person's IDVA services are being introduced in other force areas (see, e.g. the use of specialist Children and Young Person's IDVA services in London Boroughs such as Islington, provided by Solace Women's Aid). Our analysis suggests that the potential for further development of services and interventions, especially early prevention through schools and youth hubs, and through trusted relationships with youth workers and youth offending teams for this demographic group should be explored in Essex.

Young people experiencing or perpetrating domestic abuse have distinctive vulnerabilities, dependencies, challenges, and needs, and face a complex transition from childhood to adulthood that impacts on behaviour, decision-making, the way that they understand and respond to abuse as well as the way that they engage with services (Hadjimatheou *et al.*, 2022). They may be less able to recognize relationships as unhealthy or behaviours as abusive than adults and they may also need additional support with practical problems such as relationships with parents, school attainment, moving home, dealing with police and other agencies, and accessing and managing finances. Meeting these needs requires specialist provision, as demonstrated by the recent increase in local authority commissioning of IDVAs working only with children and young people (Hadjimatheou *et al.*, 2022). Adult-facing perpetrator interventions and services may be unsuitable for young men who are vulnerable themselves and for whom labelling as perpetrators could be damaging and counter-productive. Further research looking at precise ages and analysing DASH data and offending/victimization pathways for these would provide more insight into the prevalence and typologies of abuse by young people, which in turn could inform the development of specialist provision in Essex.

Female suspects

The significant crossover between victim and suspect with respect to the female-to-male cluster group merits attention. Existing criminological research identifies two main typologies of female perpetration when there is also victimization. The first is violent resistance, defined by Johnson (2017: p. 152) as 'non-controlling violence exerted in response to intimate terrorism', or 'self-defence' (Babcock *et al.*, 2003). Female victims of abuse are frequently recorded by police as primary suspects when they use violence as resistance against an abusive partner (Johnson 2006; Hester 2013; Voce and Boxall 2018). The second typology, which Johnson (2006) called 'situational' or 'common couple' abuse and claimed was less likely to come to the attention of agencies, involves relatively low-level mutual violence and abuse used in the context of chaotic and toxic heterosexual relationships, often involving alcohol abuse (Hester, 2013). Either or both of these typologies could be characteristic of the suspects in this cluster group, but our analysis does not allow conclusions to be drawn with respect to these hypotheses. Equally, research has shown that police find it difficult to identify the primary perpetrator in domestic abuse incidents (Barlow, 2023) and make erroneous recording decisions in this respect (Christie *et al.*, 2022; Hester, 2013).²⁵ In order to draw conclusions about what combination of explanations accounts for the overlap in this data, and indeed the gender distribution of recorded incidents, more granular analysis of the data is required.

These findings prompted the leadership of Essex Police to invite us to work with them to investigate female suspect data further. At the time of writing, we have secured access to the same dataset but over a longer period of 8 years. We are conducting deeper research using regression analysis and natural language processing of free text DASH data for female suspects, to help disentangle counter-allegations from violent resistance and from situational couple violence and to better understand this cohort. We analyse the presence of risk factors such as mental health, self-harm, and substance abuse (Hester, 2013) and we map the incidence of non-crimes and violent crimes in which the individual is a suspect or a victim, in order to better understand the nature of the offending and any escalation pathways better.

Our follow-on analysis aims to help police distinguish counter-allegations from primary abuse by identifying key indicators of the former, thus better-protecting victims and improving recording practices. It could also improve understandings of the severity, gravity, and nature of the victimization female suspects have experienced, which our Change Project participants said, in a follow-up interview, would be useful in informing the design of targeted female-focussed interventions for them. On the face of it, the suspect/victim crossover by itself confirms an acknowledged need for trauma-informed female-facing interventions, and for specialist screening to ensure that female-recorded suspects who are primary victims are identified correctly and referred to appropriate services.²⁶ But better risk

²³The Crime Survey for England and Wales (CSEW) reports women aged 16–19 years are more likely to be victims of any domestic abuse than women of any other age.

²⁴SafeLives found that this age group is less likely to be referred to support services than older groups, despite demonstrating a higher prevalence of severe abuse (SafeLives, 2017).

²⁵Robinson and Rowlands (2006) have developed guidance for police on distinguishing primary perpetrators from victims.

²⁶See Robinson and Rowlands (2006), for discussion and recommendations on effective screening for primary victimization.

assessment would also be a desirable outcome because violent resistance can kill.

Finally, the research could inform future support and programmes for female perpetrators. There is currently a paucity of services for this group in England and Wales. Investing resources in provision for male victims—or indeed female perpetrators—can be controversial, given the heavily gendered nature of most domestic abuse and the clear need for more resources to prevent and respond to violence against women and girls. However, a 2016 inspection of probation services by the national ombudsman found serious gaps in provision for women, noting that less than one-quarter of responsible officers had received training and guidance around female-specific case management (HMIP, 2016). This was confirmed in our qualitative findings, which found that probation officers and specialist perpetrator services were frustrated with the lack of an evidence base and training in this area.²⁷ Better research and practice in this field are clearly urgently needed.

Repeat and serial familial abuse

Existing research into perpetration typologies of family abuse has focussed mainly on male perpetrators (Holtzworth-Munroe and Stuart, 1994). This means it is of limited relevance to our findings in this cluster group, because 30% of those included as suspects in our cluster are female. Nevertheless, there is some research into child-to-parent abuse, which is likely to be captured in this cluster group. We now consider some aspects of that research which could potentially help explain in part some of the key features and nuances of this group and indicate avenues for further research and provision with this data.

Child-to-parent abuse has been found to involve frequent incidents of verbal and emotional abuse, and notably financial abuse (Ibabe *et al.*, 2014), which may explain the relatively high prevalence of non-crimes and public order offences amongst this cluster group. Child-to-parent abuse has also been shown to be prevalent amongst young people who are involved with the criminal justice system or who present aggressive behaviour in other contexts such as schools (Simmons *et al.*, 2018, p. 35). Our identification of this cluster as a significant profile of domestic abuse in Essex helps to define a research agenda for the region. Specifically, it would be useful to cross-reference our dataset with non-domestic abuse incident data to examine whether and which suspects in this cluster group have been implicated in other crimes. In terms of implications for practice, it may also be worth examining the potential for domestic abuse-related questions to be included in risk assessments already routinely carried out by youth offending and other professionals when children and young people enter the criminal justice system, to identify and open doors to addressing hidden harm.

Research into child-to-parent abuse using data from community samples has found no significant differences in rates of perpetration between females and males (Simmons *et al.*, 2018, p. 33). However, studies using police-recorded data (of which there are many) found that males accounted for 59–87% of suspected perpetrators. This accords with the gender distribution

in our cluster group, in which only 30% of suspects are female. The higher representation of males in police-recorded data such as ours compared to community samples might be explained in part by a range of different factors. Questions used in community samples may not be designed to capture the difference between primary perpetration and self-defence or retaliatory violence, course-of-conduct abuse such as coercive control or to record differences in the severity or frequency of specific incidents of violence and therefore also the gravity of the harm inflicted (Scott-Storey *et al.*, 2023, p. 862-4). Male-perpetrated abuse is also more likely to be severe and therefore more likely to be reported to and recorded by police than that perpetrated by females (Simmons, 2018, p. 33). Research with community samples has also found that males are more likely to under-report the abuse they perpetrate than females (Schmidtgall, 2005).

A systematic review of current research into child-to-parent abuse found evidence across studies of a greater frequency of mental health concerns, in particular depression, among young people who commit abuse compared to those who do not (Simmons *et al.*, 2018, p. 37). Research in community samples suggests that substance use predicts psychological and verbal abuse against both mothers and fathers (Calvete *et al.*, 2015; Pagani *et al.*, 2009). However, research in offender populations suggests that substance use 'is related to an overall pattern of antisocial behaviour rather than child-to-parent abuse specifically' as there are no differences in rates of substance use between young offenders who do and who do not abuse their parents (Simmons *et al.*, 2018; citing Contreras and Cano, 2014; Ibabe *et al.*, 2014). In order to understand better the significance of these risk factors for this cluster group, which in turn would allow police in Essex to assess risk more effectively, it would be useful to analyse DASH data for relevant incidents to identify mental health and substance abuse. To further analyse the profiles or typologies of abuse in this cluster group, more granular analysis of data on the nature of the relations between victims and perpetrators should be undertaken.

CONCLUSION

This paper shows how unsupervised machine learning can be used to provide an overview of suspects of domestic abuse in a specific force area, which improves understanding of the prevalence and features of police-recorded abuse, draws attention to previously under-addressed types of abuse, can serve as the groundwork for further, more in-depth research, and provides an evidence base for local decision-making. For Essex, our study identified two types of domestic abuse that were not 'on the radar' of local services and not previously a focus of commissioning or intervention: male-to-female IPV amongst young people, and female-to-male IPV. It has drawn attention to important gaps in current research and local provision around female suspects and female perpetration; the significance of non-crimes in domestic abuse typologies, escalation, and prediction; and it reiterates the need to investigate the victim/suspect crossover prevalent in specific typologies. The social value of our analysis lies in its provision of a broad but rigorous overview of domestic perpetration in all its forms, supporting police and local commissioning bodies make

²⁷One probation officer we interviewed complained that in their four years of service, they had not received 'a single bit of training on female perpetrators' (Participant 7).

better-informed decisions about where to invest resources and about which kinds of further research to prioritize. Clear takeaways for Essex in that respect are to improve support and interventions for young people involved in or experiencing IPV, and to develop evidence-based interventions and support for females accused of or engaging in domestic abuse against their partner, which they are now doing through investment into more granular, deeper research on female suspects.

ACKNOWLEDGEMENTS

This study was funded by the Home Office Domestic Abuse Perpetrators Research Fund in 2021. Prof. Alejandro Quiroz Flores would like to acknowledge the support of the Business and Local Government Data Research Centre (ES/S007156/1) funded by the Economic and Social Research Council (ESRC) for undertaking this work while he was based at the University of Essex.

REFERENCES

- Adisa, O., Bland, M., and Weir, R. (2021). *Identifying Predictors of Harm Within Black, Asian, and Other Racially Minoritised Communities*. Suffolk: University of Suffolk. <https://www.uos.ac.uk/sites/www.uos.ac.uk/files/FINAL%20Predictors%20of%20Harm%20UOS%20report.pdf> (accessed 3 June 2024).
- Ayre, L. and Craner, J. (2018). 'The Baked-in Bias of Algorithms.' *Collaborative Librarianship* 10(2): 76–78.
- Babcock, J. C., Miller, S. A., and Siard, C. (2003). 'Toward a Typology of Abusive Women: Differences Between Partner-Only and Generally Violent Women in the Use of Violence.' *Psychology of Women Quarterly* 27(2): 153–161.
- Barlow, C., Walklate, S., and Finnegan, E. (2023). *Who is the Victim? Identifying Victims and Perpetrators in Cases of Coercive Control*. <http://138.253.13.50/media/livacuk/sociology-social-policy-and-criminology/2-research/Who-is-the-Victim-Report.pdf> (accessed 3 June 2024).
- Calvete, E., Orue, I., Gamez-Guadix, M., and Bushman, B. J. (2015). 'Predictors of Child-to-Parent Aggression: A 3-Year Longitudinal Study.' *Developmental Psychology* 51(5): 663–676.
- Christie, C., Karavias, I., Bandyopadhyay, S., et al. (2022). 'The CARA (Cautioning and Relationship Abuse) Service Theory of Change, Impact Evaluation and Economic Benefits Study Report [Preprint]; *PsyArXiv*, doi:10.31234/osf.io/jw9uy (accessed 3 June 2024), preprint: not peer reviewed.
- Contreras, L. and Cano, C. (2014). 'Adolescents Who Assault Their Parents: A Different Family Profile of Young Offenders?' *Violence and Victims* 29(3): 393–406.
- Cunha, O., Silva, A., Cruz, A. R. et al. (2023). 'Dropout Among Perpetrators of Intimate Partner Violence Attending an Intervention Program.' *Psychology, Crime & Law* 29(6): 634–652.
- Essex County Council. (2021). *New Domestic Abuse Strategy Launched*. <https://www.essex.gov.uk/news/new-domestic-abuse-strategy-launched> (accessed 28 February 2022).
- Felson, R. B., Messner, S. F., Hoskin, A. W., and Deane, G. (2002). 'Reasons for Reporting and Not Reporting Domestic Violence to the Police.' *Criminology* 40(3): 617–648.
- Flatley, J. (2016). *Intimate Personal Violence and Partner Abuse*. London: Office for National Statistics.
- Goldstein, D. A., Cantos, A. L., Brenner, L. H., Verborg, R. J., and Kosson, D. S. (2016). 'Perpetrator Type Moderates the Relationship Between Severity of Intimate Partner Violence and Recidivism.' *Criminal Justice and Behavior* 43(7): 879–898.
- Hadjimatheou, K., Sleet, A., and Rosbrook-Thompson, J. (2022). *Report: Children and Young Person's Independent Domestic Violence Advocate* — *An Evaluation of Islington Council's Service*. London: MOPAC/Islington Council. <https://repository.essex.ac.uk/35675/> (accessed 28 July 2023).
- Hester, M. (2013). 'Who Does What to Whom? Gender and Domestic Violence Perpetrators in English Police Records.' *European Journal of Criminology* 10(5): 623–637.
- Hilton, N. Z., Harris, G. T., Rice, M. E. et al. (2004). 'A Brief Actuarial Assessment for the Prediction of Wife Assault Recidivism: The Ontario Domestic Assault Risk Assessment.' *Psychological Assessment* 16(3): 267–275.
- HMICFRS. (2019). *Essex Police—Crime Data Integrity Inspection 2019*. <https://www.justiceinspectors.gov.uk/hmicfrs/publications/essex-crime-data-integrity-inspection-2019/> (accessed 27 July 2023).
- HMICFRS. (2024). *Race and Policing: An Inspection of Race Disparity in Police Criminal Justice Decision-making*. <https://hmicfrs.justiceinspectors.gov.uk/hmicfrs/publication-html/inspection-of-race-disparity-in-police-criminal-justice-decision-making> (accessed 15 July 2024).
- HMIP. (2016). *A Thematic Inspection of the Provision and Quality of Services in the Community for Women Who Offend*. Manchester: HMIP. <https://www.justiceinspectors.gov.uk/hmiprobation/wp-content/uploads/sites/5/2016/09/A-thematic-inspection-of-the-provision-and-quality-of-services-in-the-community-for-women-who-offend.pdf> (accessed 15 July 2023).
- Holtzworth-Munroe, A. and Stuart, G. L. (1994). 'Typologies of Male Batters: Three Subtypes and the Differences Among Them.' *Psychological Bulletin* 116(3): 476–497.
- Home Office. (2011). *National Standard for Incident Recording (NSIR) Counting Rules*. London: Home Office. <https://www.gov.uk/government/publications/the-national-standard-for-incident-recording-nsir-counting-rules> (accessed 3 June 2024).
- Hui, V., Constantino, R. E., and Lee, Y. J. (2023). 'Harnessing Machine Learning in Tackling Domestic Violence—An Integrative Review.' *International Journal of Environmental Research and Public Health* 20(6): 4984.
- Huntley, A. L., Potter, L., Williamson, E. et al. (2019). 'Help-Seeking by Male Victims of Domestic Violence and Abuse (DVA): A Systematic Review and Qualitative Evidence Synthesis.' *British Medical Journal* 9(6): e021960.
- Ibabe, I., Arnoso, A., and Elgorriaga, E. (2014). 'Behavioral Problems and Depressive Symptomatology as Predictors of Child-to-Parent Violence.' *The European Journal of Psychology Applied to Legal Context* 6(2): 53–61.
- Johnson, M. P. (2006). *A Typology of Domestic Violence: Intimate Terrorism, Violent Resistance, and Situational Couple Violence*. Boston: Northeastern University Press.
- Johnson, M. P. (2017). 'A Personal Social History of a Typology of Intimate Partner Violence.' *Journal of Family Theory & Review* 9(2): 150–164.
- Johnson, R., Gilchrist, E., Beech, A. R. et al. (2006). 'A Psychometric Typology of U.K. Domestic Violence Offenders.' *Journal of Interpersonal Violence* 21(10): 1270–1285.
- Karystianis, G., Adily, A., Schofield, P. W. et al. (2022). 'Surveillance of Domestic Violence Using Text Mining Outputs from Australian Police Records.' *Frontiers in Psychiatry* 12(1): 1–13. <https://www.frontiersin.org/articles/10.3389/fpsy.2021.787792> (accessed 15 July 2023).
- Karystianis, G., Simpson, A., Adily, A. et al. (2020). 'Prevalence of Mental Illnesses in Domestic Violence Police Records: Text Mining Study.' *Journal of Medical Internet Research* 22(12): e23725.
- Kerr, J., Whyte, C., and Strang, H. (2017). 'Targeting Escalation and Harm in Intimate Partner Violence: Evidence from Northern Territory Police, Australia.' *Cambridge Journal of Evidence-Based Policing* 1(2): 143–159.
- Messing, J. T. and Thaller, J. (2013). 'The Average Predictive Validity of Intimate Partner Violence Risk Assessment Instruments.' *Journal of Interpersonal Violence* 28(7): 1537–1558.
- Mirrlees-Black, C. (1998). *Domestic Violence: Findings from a New British Crime Survey Self-Completion Questionnaire*. Home Office Research Study No. 192. London: Home Office.

- Mitchell, T. M. (1997). *Machine Learning*. New York: McGraw-Hill.
- Myhill, A. (2018). *The Police Response to Domestic Violence: Risk, Discretion, and the Context of Coercive Control*. City, University of London. https://openaccess.city.ac.uk/id/eprint/19905/1/Myhill,%20Andy_Redacted.pdf (accessed 15 July 2023).
- Myhill, A. and Johnson, K. (2016). 'Police Use of Discretion in Response to Domestic Violence.' *Criminology & Criminal Justice* 16(1): 3–20.
- Myhill, A. and Kelly, L. (2023). 'Whose Harm Is it Anyway? Using Police Data to Represent Domestic Abuse Victims' Experiences.' *Policing: A Journal of Policy and Practice* 17: paad013.
- Office for National Statistics. (2021). *Domestic Abuse in England and Wales Data Tool*. <https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/datasets/domesticabuseinenglandandwales-datatool>, <https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/datasets/domesticabuseinenglandandwales-datatool> (accessed 3 June 2024).
- Office for National Statistics. (2023). *Crime in England and Wales: Year Ending September 2023*. <https://www.ons.gov.uk/releases/crimeinenglandandwalesyearendingseptember2023> (accessed 10 July 2023).
- Pagani, L., Tremblay, R. E., Nagin, D. et al. (2009). 'Risk Factor Models for Adolescent Verbal and Physical Aggression Toward Fathers.' *Journal of Family Violence* 24(3): 173–182.
- Robinson, A. L. and Rowlands, J. (2006). The Dyn Project: Supporting Men Experiencing Domestic Abuse Final Evaluation Report. <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=07c-c9a76bc37bba7181cea62255d75824455905> (accessed 3 June 2024).
- Safelives. (2017). *Safe Young Lives: Young People and Domestic Abuse*. <https://safelives.org.uk/sites/default/files/resources/Safe%20Young%20Lives%20web.pdf> (accessed 15 July 2023).
- Safelives. (2019). *Men and Boys' Experience of Domestic Abuse*. <https://safelives.org.uk/sites/default/files/resources/Men%20and%20boys%20experience%20of%20domestic%20abuse.pdf> (accessed 15 July 2023).
- Saunders, D. G. (1992). 'A Typology of Men Who Batter: Three Types Derived from Cluster Analysis.' *The American Journal of Orthopsychiatry* 62(2): 264–275.
- Schmidtgal, K. (2005). *Gender Differences in the Self-reporting of Physical Assault for Domestic Violence Offenders*. Doctoral thesis. <https://www.proquest.com/docview/305021464?pq-origsite=gscholar&fromopenview=true&sourcetype=Dissertations%20&%20Theses> (accessed 08 July 2024).
- Scott-Storey, K., O'Donnell, S., Ford-Gilboe, M. et al. (2023). 'What About the Men? A Critical Review of Men's Experiences of Intimate Partner Violence.' *Trauma, Violence & Abuse* 24(2): 858–872.
- Simmons, M., McEwan, T. E., Purcell, R., and Ogloff, J. R. P. (2018). 'Sixty Years of Child-to-Parent Abuse Research: What We Know and Where to Go.' *Aggression and Violent Behavior* 38: 31–52.
- Turner, E., Medina, J., and Brown, G. (2019). 'Dashing Hopes? The Predictive Accuracy of Domestic Abuse Risk Assessment by Police.' *The British Journal of Criminology* 59(5): 1013–1034.
- Voce, I. and Boxall, H. (2018). 'Who Reports Domestic Violence to Police? A Review of the Evidence', *Trend and Issues in Crime and Criminal Justice*, vol 559. Canberra: Australian Institute of Criminology.
- Waggoner, P. D. (2020). *Unsupervised Machine Learning for Clustering in Political and Social Research*. New York: Cambridge University Press.
- Walby, S. and Allen, J. (2004). *Domestic Violence, Sexual Assault and Stalking: Findings from the British Crime Survey*. London: Home Office. https://eprints.lancs.ac.uk/id/eprint/3515/1/Domesticviolencefindings_2004_5BritishCrimeSurvey276.pdf (accessed 15 July 2023).
- Weir, R., Adisa, O., Blom, N. et al. (2024). 'Adolescent Domestic Abuse and its Consequences: A Systematic Rapid Review.' *Journal of Family Violence* [Forthcoming].