

# Using Machine Learning To Understand Suicide: A New Approach To Classifying Australian Coroner's Court Decisions

Ravi Iyer (✉ [raviiyer@swin.edu.au](mailto:raviiyer@swin.edu.au))

Swinburne University of Technology

Elizabeth Seabrook

Swinburne University of Technology

Suku Sukunesan

Swinburne University of Technology

Maja Nedeljkovic

Swinburne University of Technology

Denny Meyer

Swinburne University of Technology

---

## Research Article

**Keywords:** machine learning, Coroner's, NCIS, NLP

**Posted Date:** June 25th, 2021

**DOI:** <https://doi.org/10.21203/rs.3.rs-640308/v1>

**License:**  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

---

# Using machine learning to understand suicide: A new approach to classifying Australian Coroner's Court decisions

Ravi Iyer<sup>1,+\*</sup>, Dr Elizabeth Seabrook<sup>1,+</sup>, Dr Suku Sukunesan<sup>1</sup>, Associate Professor Maja Nedeljkovic<sup>1</sup>, and Professor Denny Meyer<sup>1</sup>

<sup>1</sup>Centre for Mental Health, Swinburne University, Melbourne, Australia

\*raviyer@swin.edu.au

+These authors contributed equally to this work

## ABSTRACT

We aimed to demonstrate how a large collection of publicly accessible Australian Coroner's Court case files ( $n=4459$ ) (2009-2019) can be automatically classified for determination of death by suicide, presence of mental health disorder and sex of deceased via Natural Language Processing (NLP) methods - supervised machine learning and unsupervised dictionary-based and string search based approaches.

We achieved superior levels of accuracy in the machine learning classification (Gradient Boosting vs. Random Forest baseline) of deaths by suicide of 83.3% (sensitivity = 85.1%, Specificity = 79.1%) and an accuracy of 98.3% for the dictionary-based classification of mental health disorder, as defined by the OCD-10 (sensitivity = 99.0%, specificity = 97.9%). Our machine learning approach automatically classified 24.2% (1078/4459) of the case files as referring to deaths by suicide while 63.7% (2940/4459) were classified as exhibiting a mental health disorder<sup>1</sup>.

We employed a two-stage machine learning approach involving feature engineering, followed by predictive modelling in the second. Feature engineering involved several steps including removal of low value text, parts of speech analysis, term document weighting and topic clustering. Predictive classification involved extensive hyperparameter tuning to yield the most accurate model.

We validated our models against a manually pre-coded subsample of case files, and also via binary logistic regression to test the contribution of each classified mental health disorder against determinations of deaths by suicide according to extant literature. This validation step confirmed elevated odds of suicide attributed to diagnoses of Depression, Schizophrenia and Obsessive Compulsive Disorder.

Finally, we offer a short case study to demonstrate the efficacy of our approach in investigating a subset of case findings referring to suicides resulting from family violence. We offer a proof of concept model that demonstrates an objective and scalable approach to the analysis of legal texts. The use of NLP methods in analysing Coroner's Court case findings has important implications for the ongoing development of a real-time surveillance of suicide system in Australia.

## Introduction

Australia has committed to establishing a real time surveillance of suicide system. Such a system would incorporate data from police, hospital emergency departments, the Coroner's Court and other stakeholders<sup>2</sup>. Although in its inception, the real time surveillance of suicide system seeks to make public a range of data and analyses that will inform preventative public policy initiatives.

However, there are considerable methodological challenges involved in predicting suicide events, driven by its often sudden and unpredictable nature, low case numbers, and stigma<sup>3</sup>. In over 50 years of research, a recent meta-analysis revealed both a narrow methodological approach to investigating the risk factors of suicide, and low predictive accuracy of suicide using these identified risk factors, that has not improved over time<sup>4</sup>. Researchers have called for an algorithm-based approach to predicting suicide that utilises machine learning<sup>4</sup>. It is anticipated that this will facilitate a more accurate and efficient detection of suicide.

To facilitate the shift to machine learning methods such as Natural Language Processing (NLP) for predicting suicide, there is a need to leverage complex and comprehensive datasets to utilise the linguistic context within which suicide is discussed. In Australia, evidence collected through the Coroner's process is the richest source of existing collected data relevant to suicide<sup>5</sup>. Coroner's Court data provides a flexible opportunity for retrospective investigation that is suitable for both quantitative and qualitative methods, by reporting in detail the background factors contributing to a reportable fatality (e.g. Milner<sup>6,7</sup>).

The National Coroner Information System (NCIS) was an initiative of the Australian Coroner's Society to instrument the quantification of risk factors for suicide as part of the Court's role in safeguarding public health and safety<sup>8</sup>. Data are routinely

entered into National and State Suicide Registers (e.g., Victorian Suicide Register; NCIS) by Coroner office staff using top-down data dictionaries to code the content of Coroner's Court findings and supporting documentation. Coroner inquests are a medico-legal process that involve the investigation of a death that has occurred in circumstances that are unexpected, unnatural or violent, resulting from accident, injury or medical procedure<sup>9</sup>. Coroner investigations synthesise a range of evidence, typically obtained through autopsy (physical causes), documentation (e.g., text messages, emails, written notes), and witness testimony. However, unlike those case findings released to the public, the NCIS does not contain transcripts of inquests or witness statements<sup>10</sup>.

The Coroner seeks to make findings in relation to: (1) the identity of the deceased, (2) the date and place of a person's death, and (3) the manner (4) and cause of death. Coroner inquests into suicide are detailed, taking up to two years to complete in the majority of cases<sup>11</sup>. The Coroner is at liberty to hold a discretionary inquest in certain circumstances, typically where the findings are of public benefit. Via the publication of discretionary inquests, the Coroner's Court is able to highlight emerging priorities for suicide prevention. These priorities often target important risk factors resulting in policy changes (e.g. netting to prevent jumping from 'known hotspots'; designated officers to monitor prisoners in isolation<sup>10</sup>).

To date, machine learning methods have not been applied to the analysis of NCIS data. Rather, pre-coded categories of the NCIS have been analysed using a range of standard statistical techniques (e.g. regression models<sup>7,12</sup>), while the free text components (e.g. toxicology and autopsy report) have been analysed using conventional qualitative methods (e.g. content analysis<sup>13</sup>). As such, analyses have often focused on single causes (e.g. psychiatric co-morbidity<sup>14</sup>; alcohol use<sup>15</sup>) rather than reflecting the complexity of risk and protective factors that often underpin suicide<sup>16</sup>.

Machine learning methods such as natural language processing (NLP) can provide benefits over traditional forms of statistical analysis by analysing large numbers (i.e. hundreds) of variables simultaneously. Using such NLP methods, the words used in text become the unit of analysis. NLP refers to a suite of methods whose objective is document classification, but can also involve the generation of novel insights via thematic clustering of unstructured text<sup>17</sup>. An NLP workflow typically commences with splitting text into its component parts (tokenisation), word sense disambiguation, attributing grammar to words (parts of speech tagging), weighting of words relative to document occurrence to attenuate rare or commonly used words accordingly, and identifying frequently used combinations of words such as phrases (n-grams)<sup>17</sup>. These pre-processing tasks occur prior to the stages of clustering, classification or dictionary based methods.

The use of supervised machine learning methods to identify the presence or absence of suicide or suicidal behaviour recorded in bodies of text (e.g., Electronic Medical Records (EMRs), social media posts) is an emerging area of investigation that holds promise for the analysis of complex language-based data in Coroner's Court case files<sup>18</sup>. To date, NLP methods have been successfully used to classify EMRs for the presence of suicidal ideation/attempts<sup>19</sup>, and the establishment of suicide risk<sup>20</sup>. Supervised NLP classification algorithms have also been used successfully on complex legal texts in classifying guilty verdicts in the European Court of Human Rights<sup>21</sup>, the United States Supreme Court<sup>22</sup> and other international courts<sup>23-26</sup>. For example, Kaufman and Colleagues<sup>27</sup> utilised a Gradient Boosting Algorithm (GBA) to accurately classify judicial statements made in oral argument to predict majority rulings in US Supreme Court cases with an accuracy of 74.04%<sup>22</sup>.

Recently, Fernandes and colleagues<sup>19</sup> developed two forms of classifier, one employing a top-down dictionary approach to identify suicidal attempt and the other using a Support Vector Machine (SVM) algorithm to identify suicidal ideation, using EMR data. Using a dictionary approach, a lexicon is developed using keywords that best describe the construct under investigation, while SVM is an example of supervised machine learning methods whose aim is to learn from a pre-coded dataset and to minimise classification error when applied to new data. As such, a dictionary approach is useful where a lexicon is known a priori, as in suicide attempts. A supervised machine learning approach is appropriate where the basis of classification is less clear, as is the case with suicidal ideation. The dictionary approach yielded rates of specificity (correctly identified cases that have not involved suicide) and sensitivity (correctly identified cases that have involved suicide) of 87.8% and 91.7% in the prediction of suicidal attempt, while the SVM achieved rates of specificity and sensitivity of 98.2% and 82.8% in the classification of suicidal ideation.

Zheng and colleagues<sup>20</sup> used an alternative form of supervised machine learning model (deep neural networks (DNN)) to estimate the probability of suicide attempt according to four levels of risk (e.g. low, medium, high and very high) within the next year based on a range of EMR information. When compared with a baseline GBA, the DNN achieved accuracies of 77% (vs. 70% in the baseline GBA). There are several reasons why 1 in 4 cases were not correctly classified. Zheng notes that the pre-coded dataset was highly imbalanced (only 0.21% of the study cohort featured suicide attempts) limiting the ability of the machine learning classifier to learn from so few cases. The classifier resulted in a high rate of false positives (57%), but subsequent t-test comparison between actual and falsely predicted suicide attempts on a range of predictors were statistically similar ( $p > 0.35$ ), suggesting that although certain patients fit the profile of suicide attempters (substance use issues, lifestyle problems, history of mental health, depressive and personality disorders), it did not necessarily mean that they would attempt suicide, further supporting Franklin<sup>4</sup>. Finally, it was acknowledged that many suicide attempts may not have been documented in the EMR analysed, either because suicide attempts were not reported, or that the first attempt resulted in death. Zheng<sup>20</sup>

suggest that other more complete databases could be leveraged to remedy these discrepancies.

The work of Fernandez et al.<sup>19</sup> and Zheng et al.<sup>20</sup> have highlighted a number of challenges in applying NLP to the identification of suicide or suicidal behaviours in bodies of text. These include the rarity of such events, especially given that many suicides occur following hospital admission and are therefore not recorded in EMR documentation. The rarity of such events raises an associated issue of imbalanced data, which can pose significant problems for automatic classification algorithms. Finally, EMRs are designed to yield objective clinical utility during time of admission and often do not contain information relevant to when the suicide actually occurs, often following admission.

Given the rare occurrence of suicide, there are suggestions that large publicly available datasets, such as the national death index could be leveraged to provide a richness of input to better optimise machine learning algorithms (29). We propose that the investigative Coronial process into suspected suicide by Coroner's Courts of Australia can provide such richness of data input. Coroner's Court case files confer benefits over other data repositories (e.g. EMRs) by describing, as a matter of procedure, the psychosocial background, as well as the psychological and emotional states of the deceased in close proximity to time of death. As such, Coroner Court data is likely to reflect suicide within the context of the complex interplay of risk and protective factors, including a person's biology, psychology, social environment and life experience that can greatly assist in targeted prevention initiatives<sup>4,28</sup>.

In this study, we seek to extend on previous work<sup>19,20</sup> by applying a comprehensive suite of machine learning methods to the analysis of Coroner's Court case findings. We aim to describe NLP tools used to classify Coroner's Court findings by a determination of death by suicide (using GBA and a Random Forest classification Algorithm (RFA)), according to mental health diagnosis (dictionary-based approach) and sex of deceased (search string based approach). We further aim to evaluate the performance of our approach against manual coding. As additional validation, we employ Binary Logistic Regressions to evaluate the strength of the relationship between deaths by suicide and a number of known psychiatric risk factors (as classified). Finally, we present a case study to illustrate our approach. We classify case findings according to incidence of family violence. We use topic clustering to generate a number of insights into possible predictors of family violence and then classify case findings by additional variables including alcohol and drug use and service utilisation. We then use Binary Logistic Regression to evaluate the strength of relationship between case findings classified by family violence and several risk factors including alcohol and drug use, mental health diagnosis, and service utilisation.

## Results

### Document characteristics and topic clustering

The NSW pre-coded case findings were initially separated into test and validation datasets, using a 70:30 split (test dataset  $n=398$ ; validation dataset  $n=171$ ). The number of words in each case finding ranged between 484 and 258,979 ( $\mu=6,406$ ,  $sd=8,859$ ) (c.f. table 1). Nine topics resulted in the most parsimonious configuration. Table 3 lists the top fifteen terms occurring in each topic cluster. Each topic suggests a clear and distinct thematic cluster, indicating good face validity. For example, Topic 2 refers to investigations of vehicular accidents, while others refer to suicide (Topic 5), maritime deaths (Topic 7) and deaths while incarcerated (topic 8). Only topics 5 and 8 were used by the GBA to classify the corpus documents as determinations of deaths by suicide. These topics were not evident in cluster models with fewer topic configurations. Table 2 describes differences in the numbers of suicides between Australian jurisdictions over a 10 year period. New South Wales recorded the highest number of suicides, followed by Victoria and then Queensland.

### Supervised machine learning algorithm performance

#### GBA

The GBA achieved an accuracy of classification of 83.3% in the validation dataset. This reflected a good balance between the true positive rate (Sensitivity = 85.1%) and the true negative rate (Specificity = 79.1%). The probability that case findings identified by the algorithm as deaths by suicide were actual determinations of suicide was 91.0% in the validation dataset (Positive Predictive Value (PPV)) (c.f. table 5). These metrics of accuracy were stable between datasets (e.g. training accuracy = 81.4%; validation accuracy = 83.3%) suggesting a good fit to the data. Classification of the full corpus with this algorithm identified 24.2% (1078/4459) of publicly available Coroner's findings as determinations of death by suicide.

#### RFA

The RFA achieved an accuracy of 75.4% in the validation dataset. This corresponded to values of Sensitivity of 73.0% and Specificity 81.6%. A PPV of 91.3% in the validation dataset was also achieved, comparable with the GBA approach (c.f. table 6). The sharp decline in sensitivity values (97.0% to 73.0%) between training and validation datasets however suggested overfitting, despite precautions taken via cross-validation and early stopping. Classification of the full corpus with the RFA identified 14.15% (631/4459) of publicly available Coroner's findings as determinations of death by suicide, somewhat lower than our GBA.

## Unsupervised dictionary based algorithm performance

The dictionary based approach achieved an accuracy of 98.3% in the classification of mental health diagnoses. This corresponded to values of sensitivity (99%) and specificity (97.7%). Using this approach, mental health diagnoses were attached to all 4459 Coroner's Court case files where relevant (63.7% = 2940/4459).

## Search string based determination of sex of deceased

Our novel approach to classifying court findings in terms of the sex of the deceased yielded high levels of accuracy. Using this search string approach we obtained levels of accuracy of classification of 97.4% (sensitivity = 100.0, specificity = 90.91%, PPV = 96.4%).

## Error Detection

As part of our auditing processes, it was important to investigate the nature of misclassified suicide case findings in our validation data set. Out of the nine misclassified suicide cases 2/9 (22%) involved sieges and lengthy police negotiation that resulted in suicide of the person of interest; a further 2/9 (22%) involved either mandatory psychiatric evaluation or imprisonment in remand, and were initially assessed by attending personnel as displaying low risk of self-harm; a further 2/9 (22%) involved murder suicides; while the remaining 3/9 (33%) involved death by misadventure: falling from a height, vehicular collision and location of the deceased after being missing. In each of these cases drug intoxication and/or significant mental health impairment were also implicated.

In contrast, where our algorithm incorrectly classified the case finding as constituting a death by suicide, 6/16 (38%) were either accidental (e.g. drowning of an infant; vehicular accident; drug overdose) or misadventure (e.g. combined drug toxicity, absconding from psychiatric detention). In three cases (19%) the cause of death was determined by the Coroner to be natural, but upon closer inspection perhaps warranted further investigation (e.g. "was left alone for 15 minutes... [deceased] later found in the shower without the water running"; "elevated blood levels of prescription drugs"). Three cases (19%) involved possible murders, where the exact cause of death could only be speculated upon, but also where suicide was discussed as a possible consideration. A single misclassified case involved a police shooting, in which the deceased had approached police in a psychotic state.

Only a small number of case findings were misclassified in terms of mental health disorder by our dictionary based algorithm, where several different diagnoses were discussed within the one file. Similarly, only a small number of case findings were misclassified on sex of deceased. In almost all cases, these discussed multiple fatalities.

## Validation by Binary Logistic Regression

We used Latent Semantic Analysis to derive a number of themes from a subset of family violence related case findings. Thirty clusters produced a model with a number of insights. Table 9 summarises the top five words in each of a selected number of clusters. Binary logistic regression analyses were performed to further confirm the validity of our classification methods, by comparing odds ratios against extant literature. A regression model tested the amount of variance explained in classified findings of suicide by mental health diagnosis. Mental health diagnosis (as identified by the unsupervised dictionary-based classifier) was a significant predictor of determinations of death by suicide,  $\chi^2=517.46$ ,  $p<0.001$ , Pseudo  $r^2=0.16$ , AUC=64.8%. The adjusted odds ratios of a determination of death by suicide were significantly elevated when keywords indicating a diagnosis of depression were also present ( $p<0.001$ , OR = 2.86, 95% Confidence Interval (C.I.) = 2.33 - 3.51) and also elevated for diagnoses of Generalised Anxiety Disorder ( $p=0.006$ , OR = 1.39, 95% C.I. = 1.10-1.76), Schizophrenia ( $p<0.001$ , OR = 1.56, 95% C.I. = 1.25 - 1.94) and Obsessive Compulsive Disorder ( $p=0.03$ , OR = 2.18, 95% C.I. = 1.10 - 4.32) (c.f. Table 4).

The adjusted odds of suicide were significantly elevated for males overall,  $p<0.001$ . However, diagnoses of Schizophrenia ( $p<0.001$ , OR = 1.74, 95% C.I. = 1.24 - 2.42) and Post Traumatic Stress Disorder (PTSD) ( $p=0.005$ , OR = 1.69, 95% C.I. = 1.17 - 2.44) proved elevated to a greater extent in female deceased than in males. Table 8 summarises the results of the Binary Logistic Regression analysis.

## Case Study

Two hundred and twenty-nine case findings were identified as referring to family violence. Thirty clusters produced the most parsimonious topic solution. A number of candidate thematic clusters offered novel insights into underlying factors evident in family violence case findings. These are described in table 9. Binary logistic regression analysis was performed to investigate the odds of a number of factors in predicting family violence cases across the corpus as a whole. The combination of predictors including deaths by suicide, mental health diagnosis, sex of deceased, alcohol use, drug use and service utilisation was significantly predictive of family violence cases across the corpus as a whole,  $\chi^2=256.31$ ,  $p<0.001$ , Pseudo  $r^2=0.16$ . The adjusted odds ratios of a classification of family violence were significantly elevated where a death by suicide was also classified ( $p<0.001$ , OR = 2.51, 95% C.I. = 1.87 - 3.37) and were also elevated for alcohol use ( $p<0.001$ , OR = 3.91, 95% C.I. = 2.33 - 4.41), drug use ( $p<0.01$ , OR = 1.55, 95% C.I. = 1.09 - 2.17), ( $p<0.001$ , OR = 2.51, 95% C.I. = 1.87 - 3.37) and service

utilisation ( $p < 0.001$ , OR = 2.49, 95% C.I. = 1.73 - 3.69). The adjusted odds ratios of a classification of family violence were also significantly elevated where several mental health diagnoses were mentioned, including depression ( $p < 0.05$ , OR = 1.46, 95% C.I. = 1.03 - 2.06), Obsessive Compulsive Disorder ( $p < 0.001$ , OR = 9.12, 95% C.I. = 3.35 - 23.97) and Schizophrenia ( $p < 0.001$ , OR = 25.0, 95% C.I. = 9.95 - 54.04). Table 7 describes the proportions of cases of family violence evident across the corpus by year of investigation.

## Discussion

We aimed to demonstrate a proof of concept suite of NLP approaches to classify a large corpus (n=4459) of publicly available Australian Coroner's Court case findings in terms of a determination of death by suicide, presence of mental health disorder and sex of the deceased.

We found that a GBA resulted in high levels of accuracy (83.3%; sensitivity = 85.1% and specificity = 79.1%) when classifying the corpus of documents on determinations of death by suicide, when compared with a subsample of pre-coded case files. This approach proved superior to the baseline RFA (Accuracy = 75.4%; Sensitivity = 73.0%; Specificity = 81.6%). Furthermore, when a dictionary approach was employed, we achieved accuracies of 98.3% (sensitivity = 99%, specificity = 97.7%) in the classification of mental health disorder, when compared to the same pre-coded subsample. Finally, we obtained high rates of accuracy when classifying sex of deceased using a novel search string approach (accuracy = 97.4%, sensitivity = 100%, specificity = 90.91%). To the best of our knowledge, this is the first investigation to use NLP methods to classify Coroner's Court case findings in this way.

We demonstrated superior performance using the GBA rather than the RFA. This is likely due to key differences in how each algorithm classifies known cases to reduce classification error. RFAs randomly sample from the pool of topic clusters and build individual decision trees using this evidence. The developed trees are then averaged to reduce error variance. This strategy is less successful in managing bias due to imbalanced data, such as in the present study where deaths by suicide present in only a minority of cases (24%). In contrast, a GBA grows decision trees by re-weighting in favour of weaknesses in the model, thus adding successive trees until the error of classification is minimised. Where the data is imbalanced, GBA is more likely to achieve a more accurate classification, as we have demonstrated.

The use of NLP in the detection of suicide is new, and as such little comparable data is available. Despite this, our findings compare favourably with the predictive accuracies demonstrated in work classifying EMRs in terms of deaths by suicide. Notably, our supervised machine learning approach using the GBA achieved better precision (positive predictive value) than did Fernandes<sup>19</sup>, who used a hybrid dictionary approach with rule-based post processing (91.0% vs. 82.8%). However, Fernandes achieved higher levels of sensitivity than our GBA (98.2% vs 85.1%). Our approach was more successful in estimating the likelihood of identifying a true determination of death by suicide, but was not as good at detecting non-suicidal findings. Therefore our approach is the more compelling when the aim is the accurate identification of case findings featuring death by suicide.

Our approach is also comparable to the use of other NLP methods in classifying determinations of outcome of appeal reached by international courts of law<sup>25,29</sup>. Virtucio and colleagues<sup>25</sup> classified the outcome of 27,492 appeals made to the Philippines Supreme Court. Virtucio employed a similar methodology to this study: topic clustering followed by RFA but achieved reduced accuracies of classification of 68%. This suggests advantages in using our GBA when classifying binary decisions. The reduction in accuracy can be attributed to the form of sampling used by Virtucio to constitute the reference dataset: one based on string searches for relevant keywords. The high levels of accuracy achieved by our approach suggests that the manual pre-coding of case findings, based on jurisdiction with the highest proportion of determinations of death by suicide, confers advantages by optimising upon the presence of keywords of interest, and establishing a valid ground truth.

There was a degree of equivocation in the language used by each Coroner, with multiple possible causes of death often being considered. Where our algorithm incorrectly classified a case as a determination of death by suicide, there may be benefit in highlighting such cases for further investigation, as these often feature a complex nexus of events, where single causes can often prove difficult to establish *prima facie*.

Case findings differed in the nature and quality of evidence used to reach a determination of cause of death. Coroner's seek to apply the Briginshaw standard, which demands a higher level of evidential proof where an allegation is either serious or unlikely, or where there is a gravity of consequences stemming from a particular finding, as is generally the case in determinations of death by suicide<sup>30</sup>. However, there are differences in how this principle is applied by each Coroner, with some demanding a very high standard of proof, and others arriving at a determination of suicide with relative ease<sup>30</sup>. Our classification algorithm differed at times from the Coroner, where they have relied upon their expert judgement to withhold a determination of suicide, despite textual references to the contrary.

The rates of accuracy for our dictionary based classification of mental health diagnoses are unsurprising, given the consistent way in which psychiatric diagnoses were referenced by Coroners throughout the corpus. It is common when carrying out an investigation into a fatality of unclear cause, that a medical professional is called upon to provide supplementary evidence.

These professionals utilise extant taxonomy (e.g. ICD-10; DSM-V) to reference particular psychiatric profiles, ensuring consistency of language across investigations. Dictionary based classification methods are therefore most appropriate where there is little ambiguity in the language used.

Our high rates of accuracy in the classification of a range of mental health disorders are comparable to other investigations. A recent study by Karystianis and colleagues<sup>31</sup> compiled a dictionary of terms describing a range of mental health disorders using ICD-10 taxonomy to analyse precipitating factors in 492,393 family Violence related police reports in New South Wales across a 10 year time span. Karystianis obtained a similarly high degree of specificity (92.5% among victim and 84.6% among persons of interest) and precision (positive predictive value = 96.1% victim, 99.3% person of interest). Karystianis attributed the lower rates of specificity to the nature of the police reports, being written by non-experts in mental health diagnosis.

Our dictionary approach also exceeded the levels of precision (97.1% vs. 91.7%) and sensitivity (99.0% vs. 87.8%) obtained by Fernandes<sup>19</sup> who also employed a dictionary informed approach. These differences can be explained by the different constructs investigated. It is likely that suicidal ideation may have been referred to inconsistently in the EMRs investigated by Fernandes, and may not have been identified by the dictionary as successfully (e.g. suicidal thoughts, planning etc.). It is likely that the accuracy achieved in our study takes advantage of expert testimony from medical personal as is common to evidence gathering processes in Coroner's Court inquests.

We explored the associations between mental health diagnosis and a determination of suicide within Coroner case findings using a Binary Logistic Regression. This showed well documented associations between psychiatric diagnoses and deaths by suicide. An international systematic review and meta-analysis demonstrated a comparable odds ratio for suicide (OR = 2.20) among patients with more severe depressive symptomatology, that can be expected to also characterise deaths by suicide in the present study<sup>32</sup>. Similarly, a recent systematic review of 50 longitudinal studies found a similar odds ratio of death by suicide attributed to a diagnosis of Schizophrenia of 1.40<sup>32</sup>. Our findings of an elevated odds ratio for suicide (OR=2.18) with a diagnosis of Obsessive Compulsive Disorder were slightly lower than those revealed by a recent systematic review of 63 international studies<sup>33</sup> (Odds Ratios of between 3.02 – 9.08). It is notable that 36.3% of suicide determinations did not involve a mental health diagnosis. It is possible that in specific circumstances, such as for those living in isolation or rural settings, psychiatric history may not be available. This requires further investigation for which supervised NLP may be useful.

To demonstrate our approach, we present a short case study that analyses the odds of several variables in predicting case findings of family violence across the study corpus. Latent Semantic Analysis was used to derive a number of insights using a subset of family violence related case findings. We then used our suite of algorithms to classify cases of deaths by suicide (GBA), mental health diagnoses (dictionary based approach), alcohol and drug use, and service utilisation (string search). We found that the odds associated with several of these variables, including deaths by suicide, alcohol and drug use, diagnoses of depression, Obsessive Compulsive Disorder and Schizophrenia, and service utilisation were significantly elevated for cases of family violence, as classified. Importantly, case findings featuring this mix of variable have been evident within Coroner case findings investigated from at least 2009. Notably, those case findings that were classified as featuring family violence were 150% more likely to have died by suicide, 188% more likely for the deceased to have been female and 150% more likely to have utilised services prior to death.

This illustrates several important points. Firstly that topic clustering can inductively elucidate a range of novel themes evident within a subset of the data. Secondly, these themes can suggest a range of candidate predictors for further investigation. Thirdly, with the use of our suite of classification algorithms, candidate predictors can be unearthed with relative ease. Fourthly, via Binary Logistic Regression, we can deductively estimate the contribution of each identified variable to the prediction of a desired outcome. Finally, by extracting the year in which the inquest was conducted, we have demonstrated that the Coroner's Court has been attempting to publicise the importance of family violence and the complex mix of psychosocial factors that often underpins these decisions, well before the world wide #MeToo movement gained momentum in 2017. Furthermore, family violence related fatalities were more likely to have been involved in services prior to death, suggesting that more could possibly be done to alert services to this nexus of factors.

These findings are not without their limitations. As a proof of concept application of NLP methodology, we chose to analyse a large corpus of publicly available Coroner's Court case findings. The proportions of deaths by suicide as classified are likely to differ if these methods were applied to more representative data repositories (i.e. NCIS). However, witness testimony is not a feature of these repositories, possibly impacting upon the variation in context within which suicide is discussed. Nevertheless, the proportion of deaths by suicide arrived at in the present study should be regarded with caution, given that only a fraction of Coronal inquests are made available publicly. Stratified random sampling across all Australian jurisdictions and by sex of deceased would ensure the training and validation datasets are more representative of the entire corpus. This simple modification might improve classification accuracy in future studies above the 83% achieved.

Certain stages of our analysis workflow were time consuming. Parts of speech tagging took in excess of 30 hours to reach a final outcome on all words in the corpus. Should our approach be applied to larger datasets, this timeframe would likely

increase exponentially. It is unclear at this stage whether the time expense is warranted, and whether this was balanced by a subsequent reduction in time at later stages of analysis (e.g. topic clustering; supervised machine learning).

Similarly, the hyperparameters of both GBA and RFA took considerable time to refine. Perhaps the additional investment in time taken to manually optimise the algorithms via trial and error might best be reserved for real world applications where precision of classification is an imperative.

However, this study also had several strengths. We used the NSW case findings to form the pre-coded dataset to optimise upon references to suicide in text. As noted, the NSW Coroner's Court arrived at determinations of death by suicide in more cases than any other Australian jurisdiction. Our methodological approach was more likely to capture a higher probability of reference<sup>34</sup> to suicide than if we had randomly sampled from the corpus as a whole.

Also of note was our methodological decision to homogenise a range of synonyms under the single term 'suicide'. The reasons for the lack of specificity in the language of Coroners in referencing suicide is unclear, but has been speculated by other authors as relating to persistent community stigma around determinations of death by suicide<sup>30,35</sup>. A commonly agreed upon lexicon would aid in the robustness of any NLP based classification approach.

We also decided to retain many case findings which were *a priori* difficult to classify (e.g. multiple fatalities), thus reducing the levels of classification accuracy that could be obtained. Rather than overly sanitising the data, we deemed it of greater importance to demonstrate classification accuracies with real world case findings.

As a proof of concept, the suite of NLP approaches outlined in this study suggest future directions of enquiry. These approaches could be applied to problems of classification in larger, more comprehensive and representative national and state based suicide related databases. Our approaches could also be used to mine novel insights in extant data repositories such as case findings of the Coroner's Court to inform analyses reported in the real time surveillance of suicide system, as our case study suggests.

It would be useful to compare our use of GBA with other forms of machine learning classification (e.g. Neural Networks) to ascertain whether other approaches can yield higher levels of accuracy. NLP approaches could also be applied to the investigation of other psychosocial predictors of suicide, such as significant life events (e.g. bereavement or relationship breakup). Via our use of topic clustering, novel insights might also be achieved in these areas, helping us to better understand the possible interplay of risk and protective factors in underpinning the decision to suicide.

To the best of our knowledge, this study marks the first proof of concept application of supervised (Gradient Boosting) and unsupervised (Dictionary based; string search based) machine learning methods to Coroner's Court case findings. When compared to a pre-coded sample, our predictive algorithm correctly classified determinations of deaths by suicide in the majority (84%) of Coroner's Court case findings. Furthermore, our dictionary based approach correctly identified instances of mental health disorder in nearly all cases.

The use of text mining has the potential to drill down to the level of individual words, potentially offering a valuable supplement to nationally aggregated suicide data. However, NLP approaches take time to develop and train, with certain steps taking considerable time to ensure adequate precision of prediction. Once developed however, this suite of approaches could provide valuable resources that supplement human based coding of suicide related databases. Such approaches might also identify complex patterns in the data via an exploratory approach rather than through traditional linear hypothesis testing.

Such approaches also involve ethical challenges, especially where NLP approaches reach a determination of death by suicide in the absence of Coroner designation, potentially altering the reputation of the deceased and memories of the living, which may also occur when human coding is employed, such as during pre-coding of case findings.

Our supervised machine learning approach demonstrates that with the aid of a small pre-coded sample of case findings, less clearly defined concepts such as suicide can also be successfully identified from text data. Such an approach, in particular our use of topic clustering, would be well suited to the analysis of free text sections of other databases, with possible e-health applications.

Furthermore, our use of unsupervised dictionary based methods are appropriate where the use of language is clear and unambiguous. Such dictionaries could be defined and refined using pre-coded categories extant in national and state based databases and when applied to new cases, could be used to identify the presence or absence of such information.

To realise the full potential of a real time surveillance of suicide system, we need to leverage data sourced from a range of stakeholders. The Coroner's Court undertakes the most detailed and comprehensive investigations available into deaths of an unexpected nature. Using this rich text based data available in Court case findings, we are well placed to illustrate the complexity of factors that underpin such events as deaths by suicide, family violence and others, once a suite of NLP approaches is leveraged. Furthermore, NLP approaches are scalable and can be applied to any length of document and any size of corpus. Once developed, such approaches provide economically inexpensive forms of analysis of court findings that can easily supplement human based forms of coding and enquiry.

## Methods

### Data collection

4527 publicly accessible Inquest Findings and Findings without Inquest (hereafter: case findings) from the Coroner's Court of each Australian jurisdiction were accessed in September 2019<sup>36-39</sup>. These reflected a reference period of between January 2009 and September 2019. A scraping tool was used to automatically access and download documents from Coroner's Court websites, adapting the methodology of Pina-Sanchez and Colleagues<sup>23</sup>. Optical Character Recognition (OCR)<sup>40</sup> accurately transcribed each .pdf file into .txt format. Twelve files (0.27%) were not able to be transcribed due to read/write protections and were discarded. A further 32 of the extracted case findings (0.71%) were removed from the analysis as they did not contain background information of the deceased relevant to answering the aims of the study. Table 1 shows the number of publicly accessible case findings extracted by Australian jurisdiction.

### Data cleaning

Data cleaning involved removal of low value text, and was made feasible by an analysis of common features common to the structure of each document. Most case findings followed a common pro forma structure: file identification information (title, subject, jurisdiction and identify of the deceased), table of contents; background information or introduction (a summary of the case); sequentially numbered items of evidence detailing the deceased psychosocial background and the circumstances of their death; a determination of the manner and cause of death; recommendations; and concluding remarks (personally directed remarks from the Coroner, often acknowledging the grief and distress of the family members involved).

### Text Content

Data cleaning of case findings removed pro forma sections unrelated to the study aims. This included removing the file identification information, table of contents, recommendations, and concluding remarks. This was performed using a Regex search for different keyword tokens (e.g. "Introduction"; "Background"; "1."), which a random selection of n =25 documents confirmed as corresponding to the beginning of the sequentially numbered summary of evidence.

Commonly used keyword tokens indicating the beginning of the personal remarks were not consistently identified across case findings, as such a Regex search was not feasible for identifying these sections of text. A random sample of documents (n = 25) confirmed however, that the final 250 characters of each file struck a balance between removing information unrelated to the study aims and retaining text referring to cause of death and was decided by the research team as an appropriate cutoff to be eliminated.

### Text Length

Due to the variation in document length across the sample, visual inspection of a subset of case findings was conducted to determine the depth of information contained within shorter documents. A cut off of 4000 characters was set and 10 documents either side of this cut-off point were examined by the research team. Clear qualitative differences were apparent above and below this cut off, with a total of 24 files (0.52%) excluded from analysis that did not contain a sufficient depth of information to meet the aims of the study (e.g. highly redacted cases involving the death of a minor). A final sample of 4459 Coroner's Court case files were included in the final analysis.

### The study design

RStudio<sup>41</sup> was used to develop the data analysis workflow. The full workflow is illustrated in figure 1.

### Sample pre-coding

Four hundred and seventy-two NSW case findings (10.6%) were pre-coded by RI. NSW was selected as the subsample jurisdiction as it experienced the highest number of suicides each year between 2009 and 2018, with a standardised suicide rate of 16.8 per annum per 100,000 population in 2018. A random sample of n=100 of these were chosen for coding by RI and ES. Suicide was coded for where there was clear evidence of behaviours enacted with the intention of ending one's life, irrespective of official Coroner's determination. Mental health diagnoses were coded for using descriptions contained in the ICD-10 classification of mental and behavioural disorders 41 (e.g. Generalized Anxiety Disorder (F41.1), Obsessive Compulsive Disorder (OCD) (F42), Post Traumatic Stress Disorder (PTSD) (F43), Bipolar Disorder (F31), Major Depressive Disorder (F32), Schizophrenia/Schizoaffective Disorder (F20/F25), and Borderline Personality Disorder (BPD) (F60)). Kappa = 0.92 indicated good interrater reliability. Ratings of suicide differed where circumstances alluded to suicide, but without adequate prima facie evidence and in a number of murder/suicides.

The pre-coded dataset was split 70:30 between training and validation. This split was performed 500 times, stratified by determination of suicide. Each split was inspected on the metrics of mean, sd, median, skewness and kurtosis of document length to ensure the training and validation datasets were comparable.

### **Feature engineering**

Nouns, verbs, adjectives, adverbs and pronouns were retained using parts of speech tagging. Eliminating other low value words did not reduce accuracy of prediction overall, but was important in reducing overall processing time. N-Grams <442 were included using Rapid Automatic Keyword Extraction (Rake) (udpipe r-package<sup>42</sup>). The keyword 'suicide' was expanded to include other synonyms commonly used by different Coroners (i.e. 'suicidal', 'kill her/himself', 'deliberate', 'intentional'). Term document weighting using square root/residual inverse document frequency weighting was used to balance rare with common words<sup>42</sup> (Lingmatch r-package<sup>43</sup>). Latent Semantic Analysis was then used to derive the desired numbers of topic clusters

### **GBA and RFA**

An XGBoost GBA was used to classify deaths by suicide and was compared against a RFA benchmark. Several model parameters including values of eta, gamma, maximum tree depth, minimum child weight, subsample and column sample by tree were optimised using an automated grid search with 15000 possible combinations. These parameters were then manually adjusted to further optimise upon values in specificity and sensitivity in validation dataset.

### **Dictionary based approach**

Several OCD-10 psychiatric diagnoses were used to develop the dictionary search strings. Once classified using this approach, Keyword in Context was used to investigate each dictionary term within the sentence context of usage on a random sample of classified case findings to ensure the terms identified corresponded with the desired context.

### **Case study**

We used a string search using the keywords 'domestic' to identify those case findings referring to family violence. Keyword in context was used on a random sample of  $n=10$  case findings to ensure good face validity. Latent Semantic Analysis was used to derive a number of thematic clusters using this subset of case findings. Candidate clusters were inspected for face validity. Terms evident within several novel thematic clusters then informed the search string approach to variable development (e.g. drug use theme = string search 'drug'). Thus we employed a string search algorithm to search all corpus case findings for instances of the terms 'alcohol', 'drug', and 'service'. Keyword in Context was also used to ensure that each term corresponded with the constructs alcohol and drug use and service utilisation respectively. A further string search was used to extract the date of findings.

## **References**

1. World Health Organisation. OCD-10 Classification of Mental and Behavioural Disorders: Clinical Descriptions and Diagnostic Guidelines. Retrieved from <https://www.who.int/standards/classifications/classification/protect/discretionary{\char\hyphenchar\font}{}ofprotect\discretionary{\char\hyphenchar\font}{}diseases> (2019).
2. Nous Group. Complete final report to the national mental health commission: Mental health and suicide prevention monitoring and reporting framework. Retrieved from <https://www.mentalhealthcommission.gov.au/getmedia/6dba2c68-b83e-442c-a964-34362bbfd7c/Mental-Health-and-Suicide-Prevention-Monitoring-and-Reporting-Framework> (2017).
3. Carpinello, B. & Pinna, F. The reciprocal relationship between suicidality and stigma. *Front. psychiatry* **8**, 35 (2017).
4. Franklin, J. C. *et al.* Risk factors for suicidal thoughts and behaviors: A meta-analysis of 50 years of research. *Psychol. bulletin* **143**, 187 (2017).
5. Sutherland, G. *et al.* Implementation and evaluation of the victorian suicide register. *Aust. New Zealand journal public health* **42**, 296–302 (2018).
6. Milner, A., Niven, H., Page, K. & LaMontagne, A. Suicide in veterinarians and veterinary nurses in australia: 2001–2012. *Aust. veterinary journal* **93**, 308–310 (2015).
7. Milner, A., Witt, K., Maheen, H. & LaMontagne, A. Access to means of suicide, occupation and the risk of suicide: A national study over 12 years of coronial data. *BMC psychiatry* **17**, 125 (2017).
8. Saar, E., Bugeja, L. & Ranson, D. L. National coronial information system: Epidemiology and the coroner in australia. *Acad. forensic pathology* **7**, 582–590 (2017).
9. Coroners Act (2008)(Vic.). Retrieved from [http://www6.austlii.edu.au/cgi-bin/viewdb/au/legis/vic/consol\\_act/ca2008120/](http://www6.austlii.edu.au/cgi-bin/viewdb/au/legis/vic/consol_act/ca2008120/) (2008).
10. National Coronial Information System. Coronial recommendations: Fatal facts (2019) (2019).

11. Studdert, D. M., Walter, S. J., Kemp, C. & Sutherland, G. Duration of death investigations that proceed to inquest in australia. *Inj. prevention* **22**, 314–320 (2016).
12. Milner, A., Witt, K., Maheen, H. & LaMontagne, A. D. Suicide among emergency and protective service workers: A retrospective mortality study in australia, 2001 to 2012. *Work* **57**, 281–287 (2017).
13. Manuel, J., Crowe, M., Inder, M. & Henaghan, M. Suicide prevention in mental health services: A qualitative analysis of coroners' reports. *Int. journal mental health nursing* **27**, 642–651 (2018).
14. Churruca, K., Draper, B. & Mitchell, R. Varying impact of co-morbid conditions on self-harm resulting in mortality in australia. *Heal. Inf. Manag. J.* **47**, 28–37 (2018).
15. Chong, D., Buckley, N., Schumann, J. & Chitty, K. Acute alcohol use in australian coronial suicide cases, 2010-2015. *Drug alcohol dependence* **66** (2020).
16. Milner, A., Maheen, H., Currier, D. & LaMontagne, A. D. Male suicide among construction workers in australia: a qualitative analysis of the major stressors precipitating death. *BMC public health* **17**, 584 (2017).
17. Cady, F. *The data science handbook* (John Wiley & Sons, 2017).
18. Burke, T. A., Ammerman, B. A. & Jacobucci, R. The use of machine learning in the study of suicidal and non-suicidal self-injurious thoughts and behaviors: A systematic review. *J. affective disorders* **245**, 869–884 (2019).
19. Fernandes, A. C. *et al.* Identifying suicide ideation and suicidal attempts in a psychiatric clinical research database using natural language processing. *Sci. reports* **8**, 1–10 (2018).
20. Zheng, L. *et al.* Development of an early-warning system for high-risk patients for suicide attempt using deep learning and electronic health records. *Transl. psychiatry* **10**, 1–10 (2020).
21. O'Sullivan, C. & Beel, J. Predicting the outcome of judicial decisions made by the european court of human rights. *arXiv preprint arXiv:1912.10819* (2019).
22. Kaufman, A., Kraft, P. & Sen, M. Machine learning text data and supreme court forecasting. *Proj. Rep.* (2017).
23. Pina-Sánchez, J., Grech, D., Brunton-Smith, I. & Sferopoulos, D. Exploring the origin of sentencing disparities in the crown court: Using text mining techniques to differentiate between court and judge disparities. *Soc. science research* **84**, 102343 (2019).
24. Pina-Sánchez, J., Roberts, J. V. & Sferopoulos, D. Does the crown court discriminate against muslim-named offenders? a novel investigation based on text mining techniques. *The Br. J. Criminol.* **59**, 718–736 (2019).
25. *Predicting decisions of the Philippine Supreme Court using natural language processing and machine learning* (IEEE, 2018).
26. Hitchcock, T. & Turkel, W. J. The old bailey proceedings, 1674–1913: text mining for evidence of court behavior. *Law Hist. Rev.* **34**, 929–955 (2016).
27. Kaufman, A. R., Kraft, P. & Sen, M. Improving supreme court forecasting using boosted decision trees. *Polit. Analysis* **27**, 381–387 (2019).
28. Australian Bureau of Statistics. Research paper: Psychosocial risk factors as they relate to coroner-referred deaths in australia 1351.0.55.062. Retrieved from <https://www.abs.gov.au/ausstats/abs@.nsf/Latestproducts/1351.0.55.062Main%20Features32017?opendocument&tabname=Summary&prodno=1351.0.55.062&issue=2017&num=&view=> (2017).
29. Medvedeva, M., Vols, M. & Wieling, M. Using machine learning to predict decisions of the european court of human rights. *Artif. Intell. Law* 1–30 (2019).
30. Jowett, S., Carpenter, B. & Tait, G. Determining a suicide under australian law: A comparative study of coronial practice. *UNSWLJ* **42**, 534 (2019).
31. Karystianis, G. *et al.* Automatic extraction of mental health disorders from domestic violence police narratives: text mining study. *J. medical internet research* **20**, e11548 (2018).
32. Hawton, K., i Comabella, C. C., Haw, C. & Saunders, K. Risk factors for suicide in individuals with depression: a systematic review. *J. affective disorders* **147**, 17–28 (2013).
33. Albert, U., De Ronchi, D., Maina, G. & Pompili, M. Suicide risk in obsessive-compulsive disorder and exploration of risk factors: a systematic review. *Curr. neuropsychopharmacology* **17**, 681–696 (2019).
34. Juckett, D. A method for determining the number of documents needed for a gold standard corpus. *J. biomedical informatics* **45**, 460–470 (2012).

35. Batty, G. D. *et al.* Psychosocial characteristics as potential predictors of suicide in adults: an overview of the evidence with new results from prospective cohort studies. *Transl. psychiatry* **8**, 1–15 (2018).
36. Coroner’s Court of New South Wales. Coroner’s findings. Retrieved from <https://www.coroners.nsw.gov.au/coroners-court/coronial-findings-search.html> (2019).
37. Coroner’s Court of Northern Territory. Coroner’s findings. Retrieved from <https://justice.nt.gov.au/attorney-general-and-justice/courts/coroners-findings> (2019).
38. Coroner’s Court of Queensland. Coroner’s findings. Retrieved from <https://www.courts.qld.gov.au/courts/coroners-court/findings> (2019).
39. Coroner’s Court of South Australia. Coroners findings. Retrieved from <http://www.courts.sa.gov.au/CoronersFindings/Pages/default.aspx> (2019).
40. AGB. Ocrkit pro v. 19.10.28. <http://ocrkit.com> (2019).
41. Team, R. Rstudio: Integrated development for r, version 1.4.1106. Retrieved from <https://www.rstudio.com/products/rstudio/download> (2020).
42. Wijffels, J. Package ‘udpipe’, version 0.8.6. Retrieved from <https://cran.r-project.org/web/packages/udpipe/index.html> (2019).
43. Iserman, M. Package ‘Ima\_weight’, version 1.0.2. Retrieved from <https://github.com/miserman/lingmatch> (2020).
44. Iyer, R. Required datasets and accompanying r-code file. *cloudstor* <https://cloudstor.aarnet.edu.au/plus/s/aWhd2YEcuKHq5YA> (2021).

## Ethics Declaration

This investigation has been approved by a nationally accredited Ethics committee (SUHREC) (Ref: 20204438-5593).

## Author contributions statement

R.I. was responsible for conceptual development, data acquisition, data coding, algorithm development, and manuscript writing. E.S. was responsible for data coding, manuscript writing and editing. S.S. and M.N. were responsible for manuscript editing. D.M. was responsible for statistical analysis, interpretation of data and manuscript editing.

## Competing interests

There were no competing interests.

## Availability of data and methods

All data is of a sensitive nature and contains personally identifiable information. Written requests should be sent to the corresponding author. Data and r-code (for reviewing purposes are available at<sup>44</sup>. All data was taken from publicly available databases.

State	Cases		No. of words				
	<i>n</i>	%	Min	Max	Median	Mean	<i>sd</i>
New South Wales	472	10.6	690	76,573	5049	6700	6402
Victoria	1705	38.2	484	89,576	2,753	4,925	8,032
Queensland	753	16.9	574	60,704	6,251	8,068	7,508
South Australia	483	10.8	842	258,979	4,536	8,344	15,045
Tasmania	255	5.7	630	45,850	2,056	4,038	6,054
Western Australia	470	10.5	638	116,380	6,015	8,470	9,180
Northern Australia	170	3.8	822	13,673	4,543	4,963	2,664
Australian Capital Territory	151	3.4	1,161	54,060	4,613	6,916	7,105
Total	4,459	100			4,070	6,406	8,859

**Table 1.** No. of case findings and words by Australian jurisdiction (2009 – 2019)

Jurisdiction	Total suicides 2009 – 2018 <sup>a</sup>	Total predicted suicides in corpus 2009 - 2019	% 2009-2018 Total
New South Wales	7627	138	1.81%
Victoria	5896	307	5.21%
Queensland	6696	107	1.60%
South Australia	2133	227	10.64%
Western Australia	3535	191	5.40%
Tasmania	767	49	6.39%
Northern Territory	456	43	9.43%
Australian Capital Territory	384	16	4.17%
Australia	27500	1078	
		Average	5.58%

Note: <sup>a</sup> Does not include 2019 statistics

**Table 2.** No. of deaths by suicide by Australian jurisdiction compared with the study corpus

1 Misspellings	2 Vehicular collision	3 Hospital procedure	4 Coroner's Court	5 Suicide
aŕf	Vehicle	Practice	Safety	Mental
Cffiah	Road	Risk	Victorian	Suicide
Hlfr	Police	Process	Coroner	Harm
Rift©	Driver	Procedure	Recommendation	Health
Afw	Speed	Training	Coronial	Violence
Alh.ha	Travel	Ensure	Section	Self
Athat	Metre	You	Prevention	Suicidal
Mtco	Car	Need	Include	Community
Afcf	Collision	Document	Investigation	Relationship
©at	Constable	Implement	Act	Family
?(c)	Scene	Responsibility	Provide	Behaviour
©att	Front	Assessment	Public	House
Nafmv	Safety	Guideline	Pursuant	Police
rtfim	incident	We	Role	Abuse
6 Collision Driver at fault Vehicle Road Mental Driver Collision Speed Drive Crash Medication Traffic Travel Drug Behaviour Lane	7 Maritime Deaths Water Boat Vessel Rescue Sea Area Locate Safety Beach Fishing Decease Equipment Drown Wave	8 Deaths while incarcerated Prisoner Prison Custody Correctional Cell Officer Sentence Corrective Offence Imprisonment Transfer Inmate Staff Custodial	9 Deaths during Birth Staff Team Family Emergency Assessment Baby Care Child Plan Call Communication Brief Mental management	

**Table 3.** Latent Semantic Analysis (LSA) topic terms

	All suicide determinations		Female suicide determinations		Male suicide determinations	
	Adj. OR (95% C.I.)	p	Adj. OR (95% C.I.)	p	Adj. OR (95% C.I.)	p
Sex	1.63 (1.31, 2.02)	<0.001				
Depression	2.86 (2.33, 3.51)	<0.001	2.89 (2.14, 3.9)	<0.001	2.86 (2.33, 3.51)	<0.001
Generalised						
Anxiety	1.39 (1.10, 1.76)	0.006	1.03 (0.74, 1.46)	0.85	1.39 (1.1, 1.76)	0.006
Disorder						
PTSD	1.28 (0.99, 1.66)	0.06	1.69 (1.17, 2.44)	0.005	1.28 (0.99, 1.66)	0.06
Schizophrenia	1.56 (1.25, 1.94)	<0.001	1.74 (1.24, 2.42)	0.001	1.56 (1.25, 1.94)	<0.001
BPD	1.14 (0.74, 1.76)	0.56	1.47 (0.96, 2.25)	0.08	1.14 (0.74, 1.76)	0.56
OCD	2.18 (1.10, 4.32)	0.03	1.45 (0.64, 3.29)	0.37	2.18 (1.1, 4.32)	0.03

Note: PTSD = Post Traumatic Stress Disorder; BPD = Borderline Personality Disorder; OCD = Obsessive Compulsive Disorder

**Table 4.** Binary Logistic Regression to Identify Suicide Determination from Mental Health Diagnoses Identified in the Australian Corpus

Hyperparameters	Training best AUC	Cut-off Probability	Dataset	Accuracy	Sensitivity	Specificity	Positive Predictive Value
eta = 0.03 gamma = 2 max. depth = 5 min. child weight = 0.8 subsample = 0.2 column sample by tree = 0.9 no. of trees = 12	95.1%	0.36	Training	81.4%	82.0%	79.8%	91.7%
			Test	83.3%	85.1%	79.1%	91.0%

**Table 5.** Performance of Gradient Boosting Algorithm in classification of determinations of suicide of pre-coded data

Hyperparameters	Training best AUC	Cut-off Probability	Dataset	Accuracy	Sensitivity	Specificity	Positive Predictive Value
mtry = 9 node size = 10 no. of nodes = 2 no. of trees = 18 sample size = 15	83.7%	1.0	Training	82.2%	97.0%	44.0%	81.7%
			Test	75.4%	73.0%	81.6%	91.3%

**Table 6.** Performance of Random Forest algorithm in classification of determinations of suicide in pre-coded data

Year	<i>n</i>	%
2009	7	3.07
2010	16	7.01
2011	10	4.38
2012	12	5.26
2013	18	7.89
2014	21	9.21
2015	28	12.28
2016	28	12.28
2017	19	8.33
2018	28	12.28
2019	7	3.07
Total	194	

**Table 7.** Number of cases of family violence across the study corpus by year of Coroner’s Court investigation

Predictor	Adj. OR (95% C.I.)	p
Suicide	2.51 (1.87, 3.37)	<0.001
Female	2.88 (1.04, 11.95)	0.08
Male	1.73 (0.64, 7.14)	0.36
Alcohol	3.19 (2.33, 4.41)	<0.001
Drugs	1.55 (1.09, 2.17)	<0.01
Psychiric Diagnosis		
Anxiety	1.21 (0.76, 1.87)	0.41
Bipolar	1.63 (0.72, 3.32)	0.21
Borderline Personality Disorder	1.12 (0.49, 2.28)	0.78
Depression	1.50 (1.03, 2.06)	<0.05
Obsessive Compulsive Disorder	9.22 (3.35, 23.97)	<0.001
Schizophrenia	25.00 (0.10, 54.04)	<0.001
Service utilisation	2.49 (1.73, 3.69)	<0.001

**Table 8.** Binary Logistic Regression to identify cases of family violence from suicide status, service utilisation and other variables

5 Diagnoses	7 Family Violence	12 Suicide	23 Service	28 Drug Use
psychologist	perpetrated	depression	reportable	intoxication
medication	family	suicidal	professional	drugs
prescribed	violence	stressors	service	cannabis
anxious	controlling	mental	hospital	substances
therapy	sexual	suicide	doctor	affected

**Table 9.** Top 5 terms evident in select topic clusters derived from cases of family violence

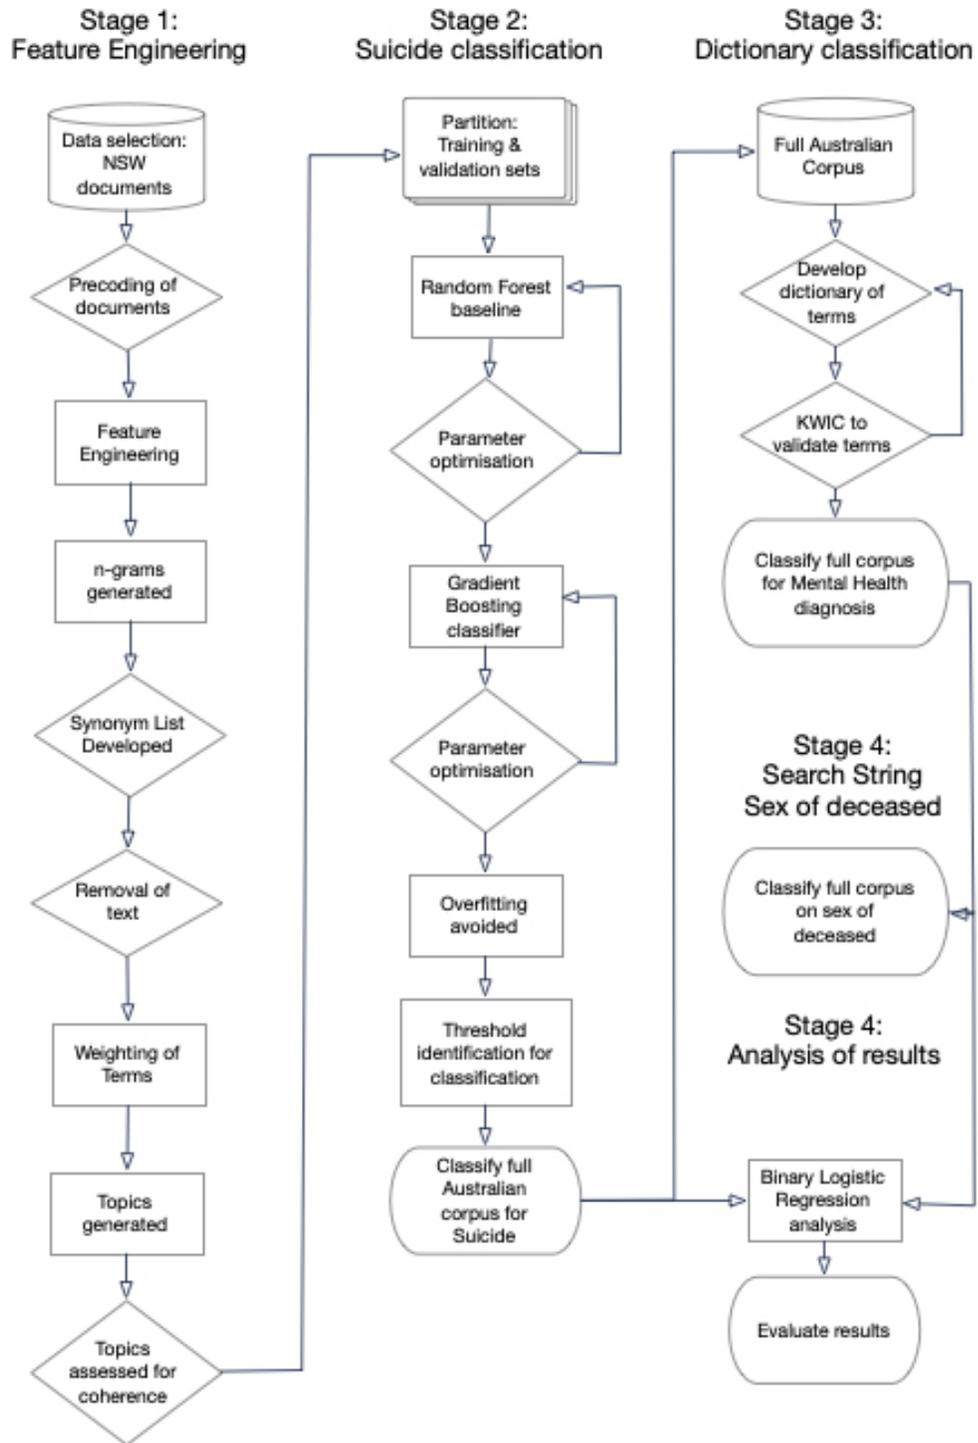


Figure 1. Analysis Flowchart