Mining Case Law to Improve Countries' Accountability To Universal Periodic Review

Soline Aubry[1], Hansdeep Singh[2], Ivan Vlahinic[1], Abhimanyu Ramachandran[1], Sara Fischer[4], Robert O'Callaghan[1], Natalie Shoup[3], Jaspreet Singh[2], David Sangokoya[3], Gabriel Pestre[3], and Carson Martinez[3]

[1]CKM ADVISORS
[2]ICAAD
[3]DATA POP ALLIANCE
[4]GLOBAL INSIGHT

Final Draft for comments

February 2017

TABLE OF CONTENTS

I. BACKGROUND AND RATIONALE

A. Details on the Universal Periodic Review

The United Nations (UN) Universal Periodic Review (UPR) is a process established by the Human Rights Council aiming to monitor and improve the human rights situation in each UN member state. It involves periodic reviews of the human right records of all 193 Member States on a 4.5 year cycle. It comprehensively identifies gaps that undermine the rule of law, since recommendations cover the entire scope of human rights topics. Each review provides an opportunity for the states under review to declare what actions they have taken to improve the human rights situation in their country. The UPR also provides an opportunity for any UN member state to make specific recommendations to another state via a peer-review mechanism, which results in either acceptance, or notation of the recommendation by the state under review. In accepting recommendations, states create a binding obligation to improve upon their current human rights record. Any stakeholder – the state itself, independent national human rights commissions, other UN treaty bodies, and civil society – can then track the implementation of accepted recommendations and provide feedback during the interim period before the next cycle.

The UPR, which aims to create accountability of member states through its recurring, comprehensive, peer-reviewed, and monitored nature, is a unique mechanism with no current equivalent. However, the system also has its weaknesses. One of its major problems is the lack of states' accountability, resulting in poor implementation rates of recommendations. Statistical analysis compiled by UPR-info.org shows that only 18% of 11,527 recommendations were fully implemented, 30% were partially implemented, 48% were not implemented at all, and 4% could not be determined.[1] One of the major reasons for the gap in implementation is the level of specificity of recommendations. Indeed, "general recommendations are difficult to implement because the action is unclear[, and] specific recommendations are also the least likely to be implemented"[2] because their specificity requires the State to take more actionable steps to satisfy the recommendation. Since the review system lacks a comprehensive methodology to monitor each of those steps, the design and collection of indicators on the level of implementation is left to the state under review.[3] Other problems arise from the incapacity to measure and monitor such metrics – since it is often impossible to clearly define what full implementation a given recommendation looks like – and from the lack of proper infrastructure for monitoring implementation.

In an effort to improve states' accountability to the UPR, various initiatives have been developed to strengthen the monitoring processes. For example, the UN Working Group on Human Rights in India devised an implementation tracking tool to review India's first cycle of UPR.[4] The purpose of developing such a tool is two-fold: 1) initiate a cycle of monitoring and reporting; and 2) strengthen

---

[1] UPR Info, *Beyond Promises: The Impact of the UPR on the Ground*, at 17 (2014), http://www.upr-info.org/sites/default/files/general-document/pdf/2014_beyond_promises.pdf
[2] *Id.* at 25.
[3] *Id.*
[4] WGHR, Tracking Implementation: A Monitoring Tool for Recommendations from The United Nations' Universal Periodic Review for India (2012), http://www.upr-info.org/sites/default/files/general-document/pdf/wghr_tracking_implementation_monitoring_tool_2013.pdf.

civil societies' documentation capacity and advocacy efforts.[5] Additionally, the UPR Accountability Initiative, which was founded by International Center for Advocates Against Discrimination (ICAAD), is a public-private partnership aimed at monitoring UPR recommendations regarding access to justice for women and girls in the Pacific Islands. In seeking to remedy the lack of specificity of recommendations, efforts have been made to map structural discrimination for women and minority communities in 12 Pacific Island countries by linking each state's recommendations made during the UPR process to specific governmental policy, legislation, case law, or cultural norms that perpetuate structural discrimination.

However, no systematic methodology has been created to provide metrics to efficiently monitor the level of implementation of recommendations. Additionally, effectively monitoring thousands of recommendations is daunting, and this is where the ability to automate processes for tracking the progress of a given recommendation can potentially be a powerful tool for establishing states' accountability to the UPR system.

B. The Use of Data Science to Monitor UPR accountability

In this study, we hypothesize that leveraging text mining and machine-learning algorithms is a viable strategy for monitoring gender discrimination in sentencing practices of Fiji's judiciary system, which has been the object of recommendations from Norway and Belgium in the UPR cycles of 2010 and 2015, respectively. When focusing on Violence Against Women and Girls (VAWG) in Fiji, two types of offenses are of specific interest: sexual assault (SA) and domestic violence (DV). Sexual assault cases include all sexual crimes, e.g. rape, indecent assault, defilement, statutory rape, incest; domestic violence cases include family violence and intimate partner violence. As we will further elaborate on in Part II, legal action in cases of sexual assault and domestic violence is governed by several different laws in Fiji, but studies have shown that discriminatory practices in how and when these laws are applied may in some instances undermine their effectiveness.[6] Determining whether or not gender discrimination has a systematic impact on the outcome of these sentences requires extensive analysis of case law archives.

In this project we develop algorithms for analyzing and classifying cases, and propose metrics to help characterize the status and evolution of sentencing practices for sexual assault and domestic violence with regards to these laws. We propose these methodologies as a 'proof of concept' of how a data science approach can help produce concrete indicators of progress towards the implementation of recommendations in the UPR. In the particular case of monitoring sentencing practices, we propose a framework for assessing the impact of gender discrimination in Fiji's judiciary system by increasing transparency of outcomes at various levels of the court system, accountability of magistrates towards the laws that govern sexual assault and domestic violence and the vulnerable populations they are intended to protect, and the consistency of sentencing with regards to victims, perpetrators, and authorities in such cases.

---

[5] *Id.* at 4..

[6] Emily Christie, Hansdeep Singh, and Jaspreet Singh, *An Analysis of Judicial Sentencing Practices in Sexual & Gender-Based Violence Cases in the Pacific Island Region,* ICAAD & DLA PIPER (2016), http://www.icaad.ngo/wp-content/uploads/2015/12/ICAAD-Analysis-of-Judicial-Sentencing-Practices-in-SGBV-Cases.pdf.

Our hope is that the outcomes of this study, designed as a collaborative effort between data scientists and lawyers with known expertise in the UPR process, will encourage to develop more systematic and quantitative methodologies to track the implementation of recommendations, resulting in an increased accountability of countries towards the UPR process.

C. Precedents in the Legal Sector

Text mining and machine learning tools have been gaining increasing popularity across many fields over the past several years. In the legal sector, some scholars have begun to use text-mining techniques to automate research tasks that were previously performed by teams of lawyers. Law firms may also explore a wide range of similar cases through automated text-mining to figure out how to best advise their clients and prepare for court. Additionally, systems such as LexisNexis and Westlaw match cases, approaches, and judgments leveraged by legal professionals to better predict case outcomes and strategies based on algorithmic analysis of large sets of legal cases.[7] As these applications show, quantitative legal analysis can help legal scholars automate certain research tasks and provide tools for law firms to better advise clients and prepare cases. Tools for classifying cases and measuring prevalence of different outcomes in an archive of case law data have therefore been developed to meet these needs in the legal sector.

However, outside the legal realm, only a few groups leverage the use of data-mining on legal documents. Additionally, the purpose of using data-mining for those groups rarely goes beyond making quantitative legal predictions.[8] One example of data-mining legal documents for a purpose other than legal prediction is a 2015 study by Sacks, Sainato, and Ackerma. The study examined bail decisions made by judges and their subsequent outcomes from a sample of 975 cases collected by the New Jersey's Criminal Disposition Commission to explore whether defendants were able to meet financial bail in order to be released from jail. Researchers used Bayesian probability analytics to analyze the trends. The research provides evidence that race is a critical factor in pretrial decisions by the court, and more specifically that race exerts a strong influence on a defendant's ability to post bail prior to trial.[9]

Another example of text-mining legal documents comes from Huridocs, an international NGO that works to help human rights organizations use information technologies for advocacy. Huridocs with the Institute for Human Rights and Development in Africa (IHRDA) developed a tool called "Caselaw Analyser" that makes it easier for people to access and browse over 40,000 processed judgments from the European Convention of Human Rights HUDOC database. In a collaborative project with Teradata Partners DataDive, Huridocs shed light on how judges were ranking cases as important and the trends associated with such rankings. The findings were used to learn more about the frequency of requests and court's assessments of which violations were most urgent. This information enables human rights advocates to gain a better understanding of the enforcement of case judgments.

---

[7] Katz, D. M. "Quantitative Legal Prediction--Or--How I Learned to Stop Worrying and Start Preparing for the Data-Driven Future of Legal Services Industry." Emory Law Journal. Vol 62: 909. 2013. http://datascienceassn.org/sites/default/files/Quantitative%20Legal%20Prediction.pdf

[8] *Id.*.

9 Sacks, M., Sainato,V. A., & Ackerman, A.R. *Sentenced to Pretrial Detention: A Study of Bail Decisions and Outcomes.* AM J CRIM JUST (2015) 40:661–681 DOI 10.1007/s12103-014-9268-0

Outside of the bail study and Huridocs studies described above, there is little evidence that data-mining for legal documents is done for reasons other than predicting case outcomes, measuring attorney quality and performance, or predicting the expected bill. Moreover, it seems to be the case that identifying larger trends in court data are not usually the aim of data-mining projects involving legal documents. This study aims to show the potential of these methodologies for applications to local development contexts.

II. THE CASE OF FIJI: CONTEXT, LAWS, AND LEGAL STRUCTURES

A. Background on Violence Against Women in Fiji and Domestic Violence Decree 2009

Approximately 64% of women in Fiji report that they have experienced some form of physical domestic violence or sexual assault[10] in contrast to approximately 35% worldwide.[11] If you include emotional violence, the number in Fiji jumps to over 72%.[12] Violence against women and girls (VAWG) impacts all sectors of society, from health and safety to economics and education. The Reserve Bank of Australia calculated that the annual cost of VAWG in Fiji is equal to 7% of the country's GDP.[13] By contrast, Fiji only spends ~4% of its GDP on education.[14]

Patriarchal beliefs in Fiji have allowed for gender discrimination within law enforcement and courts, the very structures of society that should be providing avenues for justice, redress, and protection. Courts are allowing perpetrators to escape accountability for their crimes, often by imposing low and inconsistent sentences for sexual and gender-based violence (SGBV).[15] Such beliefs and policies impact over 500,000 women and girls in Fiji.

In September 2014, Fiji had its first elections since a military coup in 2006 suspended democratic government. During this eight-year period, the military government, without parliamentary approval, pushed forth over 300 Decrees that became law. One of these Decrees was the Domestic Violence (DV) Decree 2009, which came into force in February 2010.[16] Although it wasn't passed through a

---

[10] UNICEF, Harmful Connections: Examining the Relationship Between Violence Against Women and Violence Against children in the South Pacific, at 4 (2015), http://www.unicef.org/pacificislands/Harmful_Connections(1).pdf.
[11] World Health Organization, *Global and regional estimates of violence against women: Prevalence and health effects of intimate partner violence and non-partner sexual violence*, 2013, http://www.who.int/reproductivehealth/publications/violence/9789241564625/en/ (The WHO methodology surveyed women and girls between 15-49 accessed 14 October 2015)
[12] (FWCC, 2011)
[13] (Laqeretabua, et al., 2009)
[14] *Id.*
[15] *See generally* Emily Christie, Hansdeep Singh, and Jaspreet Singh, *An Analysis of Judicial Sentencing Practices in Sexual & Gender-Based Violence Cases in the Pacific Island Region,* ICAAD & DLA PIPER (2016), http://www.icaad.ngo/wp-content/uploads/2015/12/ICAAD-Analysis-of-Judicial-Sentencing-Practices-in-SGBV-Cases.pdf.
[16] Fiji Domestic Violence Decree (2009.), http://www.paclii.org/fj/promu/promu_dec/dvd2009191/.

democratic process,[17] the progressive nature of the DV Decree is a reflection of the efforts by Fiji Law Reform, Fiji Women's Crisis Center (FWCC), and other women's rights organizations, which prepared a draft Domestic Violence Bill prior to the coup. Nevertheless, even with a progressive Decree, the largest barrier to overcome remains attitudes regarding domestic violence, which leads to a lack of implementation of the DV Decree.[18]

Two components of the DV Decree are of specific interest for fighting practices of SGBV in Fiji: 1) giving greater latitude and flexibility to magistrates and judges to issue protective or restraining orders[19] and 2) making clear that reconciliation practices are inapplicable in domestic violence[20]. In a randomized analysis of 145 sexual assault (SA) and domestic violence (DV) cases in Fiji, gender stereotypes and reconciliation practices were raised in 83% of cases and led to sentence reductions in 52% of cases.[21] A majority of the victims in the selected cases were girls[22]. In an analysis conducted by Nazhat Shameem, former Fijian judge, "resistance to the Domestic Violence Decree comes from the same attitude that saw an eagerness to promote reconciliation for domestic violence cases."[23] When judges give weight to reconciliation that has taken place, especially without consideration for the coercive forces that lead to reconciliation, they introduce gender discrimination into the ruling and give sentences that are not appropriate given the gravity of the crime. Similarly, magistrates fail to issue restraining orders[24] despite the clear language in the Decree that the well-being and safety of the victim is paramount.[25] By tracking weight given to reconciliation and issuing of restraining and protective orders in domestic violence cases, we can identify patterns in whether or not court rulings are in line with these two important factors of the DV Decree. Similarly, we apply some of the principles from our DV analysis to SA cases. Although the legislation related to SA cases is different, the concerns of discriminatory treatment remain relevant and important for study.

---

[17] FWCC, Submission to the United Nations Periodic Review on Violence Against Women: Fiji (2010)., http://lib.ohchr.org/HRBodies/UPR/Documents/Session7/FJ/FWCC_UPR_FJI_S07_2010_Fiji_WomensCrisisCentre.pdf.

[18] Nazhat Shameem, *Fiji Judiciary Criminal Law Workshop for Judges and Magistrates*, at 8 (June 14, 2002)

[19] Fiji Domestic Violence Decree, sec. 23-24 (2009) (Court can issue a restraining order for the "safety and wellbeing of the person against whom the offence appears to have been committed."), http://www.paclii.org/fj/promu/promu_dec/dvd2009191/

[20] Fiji Domestic Violence Decree, Part 11 (2009) ("Section 163 [promotion of reconciliation in the Criminal Procedure Cap21 statute] does not apply to a charge, which in the circumstances of the case, is a charge for a domestic violence offence."), http://www.paclii.org/fj/promu/promu_dec/dvd2009191/. Furthermore, the Criminal Procedure Decree 2009 also makes clear that reconciliation cannot promoted in domestic violence cases. Fiji Criminal Procedure Decree, sec.154(6) (2009), http://www.paclii.org/fj/promu/promu_dec/cpd2009201/.. ...

[21] Emily Christie, Hansdeep Singh, and Jaspreet Singh, *An Analysis of Judicial Sentencing Practices in Sexual & Gender-Based Violence Cases in the Pacific Island Region,* ICAAD & DLA PIPER, at 49 (2016), http://www.icaad.ngo/wp-content/uploads/2015/12/ICAAD-Analysis-of-Judicial-Sentencing-Practices-in-SGBV-Cases.pdf.

[22] Id. at 48.

[23] Nazhat Shameem, *Fiji Judiciary Criminal Law Workshop for Judges and Magistrates*, at 11 (June 14, 2002)

[24] Nazhat Shameem, *Fiji Judiciary Criminal Law Workshop for Judges and Magistrates*, at 9 (June 14, 2002)

[25] Fiji Domestic Violence Decree, sec. 24, 26, 28 (2009), http://www.paclii.org/fj/promu/promu_dec/dvd2009191/

B. Project scope and specific aims

This study was designed as a collaborative effort between data scientists and lawyers with known expertise in the UPR process. The study leveraged case law analysis by the International Center for Advocates Against Discrimination (ICAAD)[26] on the impact of gender discrimination on SA and DV cases in the Pacific Island region to develop appropriate methodologies and facilitate the manual review of computational findings. More specifically, the issue of gender discrimination often arises when judges, including Fijian judges, use gender stereotypes and reconciliation[27] practices as a justification for sentence reduction. In cases involving gender stereotypes, this means that a judge may reduce the sentence in a rape case because the victim/survivor had previous sexual partners, had a drink with the perpetrator, or wore "inappropriate clothing," to give a few examples[28]. In cases involving reconciliation, it could mean a parent or uncle getting food, mats, or payment as a form of apology from the perpetrator, where the victim/survivor is not the beneficiary or even involved in the reconciliation practice[29] As mentioned above, in a randomly selected set of 145 SA and DV cases in Fiji, gender stereotypes and reconciliation practices were raised in 83% of cases and led to sentence reductions in 52% of all cases[30]. A majority of victims/survivors in these cases were girls[31]. These statistics were reflective of a trend across the region, where 908 randomly selected cases were analyzed[32]. The methodology of this study was developed to track whether the Domestic Violence Decree 2009 of Fiji ("DV Decree") was being effectively implemented and whether gender discrimination would impact sentencing in DV and SA cases.

The basis for focusing our analysis on the DV Decree and gender discrimination in DV and SA cases is that they are the subject of recommendations accepted by Fiji during its 2010 and 2015 UPR reviews. During the first UPR cycle in 2010, Norway[33] recommended that Fiji "[a]dopts, in the near future, the proposed laws on domestic violence and sexual offenses, thereby prohibiting practices that legalize violence against women."[34] In the second UPR cycle in 2015, Belgium and Bangladesh offered similar recommendations: "Take the necessary measures to ensure that the decree on domestic violence be effectively implemented and that the perpetrators of violence against women,

---

[26] Emily Christie, Hansdeep Singh, and Jaspreet Singh, *An Analysis of Judicial Sentencing Practices in Sexual & Gender-Based Violence Cases in the Pacific Island Region,* ICAAD & DLA PIPER, at 49 (2016), http://www.icaad.ngo/wp-content/uploads/2015/12/ICAAD-Analysis-of-Judicial-Sentencing-Practices-in-SGBV-Cases.pdf.

[27] *Id.* at 6.

[28] *Id.* at 30-31.

[29] *Id.* at 27.

[30] *Id.* at 49.

[31] *Id.* at 48.

[32] *Id.* at 6.

[33] Chile, Spain, and Belgium all made similar recommendations regarding the proper implementation of domestic violence legislation and combating VAWG. See generally U.N. Human Rights Council, Report of the Working Group on the Universal Periodic Review: Fiji, A/HRC/14/8, at 17 (Mar. 23, 2010), http://daccess-dds-ny.un.org/doc/UNDOC/GEN/G10/124/83/PDF/G1012483.pdf?OpenElement

[34] U.N. Human Rights Council, Report of the Working Group on the Universal Periodic Review: Fiji, A/HRC/14/8, at 17 (Mar. 23, 2010), http://daccess-dds-ny.un.org/doc/UNDOC/GEN/G10/124/83/PDF/G1012483.pdf?OpenElement

including within the family, be duly prosecuted and punished"[35] and "Take concrete measures to eliminate gender stereotypes and discrimination against women."[36]

With the aim of providing tangible data to support levels of accountability on these key recommendations, we selected a dataset to track the level of implementation of the DV Decree and gender discrimination in DV and SA cases. The dataset was a collection of transcribed cases that was made publically available by the Pacific Islands Legal Information Institute (PacLII)[37]. The project was broken down into three specific aims:

- Develop tools to extract and structure relevant features from a case law dataset;
- Use structured data and develop relevant metrics to provide insights on case prosecution, application of specific legislations, and judgments in SGBV cases;
- Evaluate countries' accountability towards the UPR process and provide additional insights on transparency, magistrate's accountability towards the laws and consistency of Fiji judicial system.

III. ANALYSIS AND APPROACH

A. Description of the dataset

The dataset used in this study is a collection of transcribed courtroom cases written in English and hosted on the Pacific Islands Legal Information Institute (PacLII) website[38] in HTML format. The dataset contained 10,173 cases spanning 14 years, from 2000 to 2014. PacLII came into existence around 1999 and likely accounts for increasing number of available cases starting from 2000. The documents were originally sourced from four Fiji courts – Magistrate Court, High Court, Court of Appeals, and Supreme Court – with offenses ranging from minor (e.g. theft) to severe (e.g. murder.) The database includes final judgments as well as other judicial proceedings, such as Voir Dire and Extempore Ruling on Bail.[39]

The transcriptions contain basic case identifiers, including the date of the trial, name of the offender, name(s) of the judge(s), and name of the court, with trial date and courts present as part of the document title. Beyond the basic identifiers, cases differed significantly in terms of specific information transcribed as well as the level of detail provided. For example, while the 900-word DPP v Veresas (2013, FJMC 72) thoroughly describes the sentence, the charges, the offence,

---

[35] U.N. Human Rights Council, Report of the Working Group on the Universal Periodic Review: Fiji, A/HRC/28/8, at 17 (Dec. 17, 2014), http://daccess-dds-ny.un.org/doc/UNDOC/GEN/G14/245/67/PDF/G1424567.pdf?OpenElement

[36] U.N. Human Rights Council, Report of the Working Group on the Universal Periodic Review: Fiji, A/HRC/28/8, at 16 (Dec. 17, 2014), http://daccess-dds-ny.un.org/doc/UNDOC/GEN/G14/245/67/PDF/G1424567.pdf?OpenElement

[37] Pacific Islands Legal Information Institute (PACLII), http://www.paclii.org/

[38] www.paclii.org

[39] Voir Dire refers to the process of jury selection where prospective jurors are questioned about their background and potential biases before being selected as a juror; Extempore Ruling on Bail refers to a judge's determination of whether bail will be granted to the defendant in a particular case.

reference to the Penal Code, aggravating and mitigating factors, reference to the Sentencing and Penalty Decree and precedent cases, the 300-word State v Raituraki (2013, FJMC 278) captures half of the information and lacks much of the contextual detail.

It is this variation in granularity of information, along with the unstructured nature of the transcripts, which poses the greatest challenge in terms of automated extraction of core information required for monitoring of sentencing practices. To address this challenge, a broad spectrum of tools has been developed, ranging from supervised and unsupervised machine learning algorithms to text mining techniques. Additionally, the small number of cases has resulted into substantial amount of noise in the analyses as illustrated by ciseau-shape graphs. Therefore, three-year rolling averages were often implemented to smoothen the trends.

B. Overview – Case Corpus Available for Analysis, Segmented by Court

Dates and courts were extracted from the title of the document by mean of regular expressions and were analyzed to provide high-level insights on the distribution of the dataset along those two dimensions.

B1. The overall volume of cases grown by 25% annually since 2000

In terms of number of cases available for analysis, the dataset shows a consistent increase in the annual volume from about 300 in 2000 to 1300 in 2013 (Figure 1a). Two exceptions could be observed: a) 2006-2009, where the number of cases dropped significantly, most likely attributed to the country's political instability, and b) 2014, which we believe to be the delay in upload of cases to PacLII website (our data download was finalized on March 26th, 2015).

B2. The proportion of magistrate court cases increases after 2010

The breakdown by court over the year 2000-2014 shows a majority of High Court (FJHC) cases (71% on average) and a significant increase in proportion of Magistrate Court (FJMC) cases starting in 2010, going from an average of 4% before 2010 to 20% after 2010 (Figure 1b). Court of Appeal (FJCA) cases represent about 15% of the cases, with a strong increase to 25-33% in 2006-2007. Supreme Court (FJSC) cases represent a minority, averaging 2.5% over the given period – expectedly small proportion because the right to appeal to FJSC is granted only at the discretion of the judges of the Supreme Court.

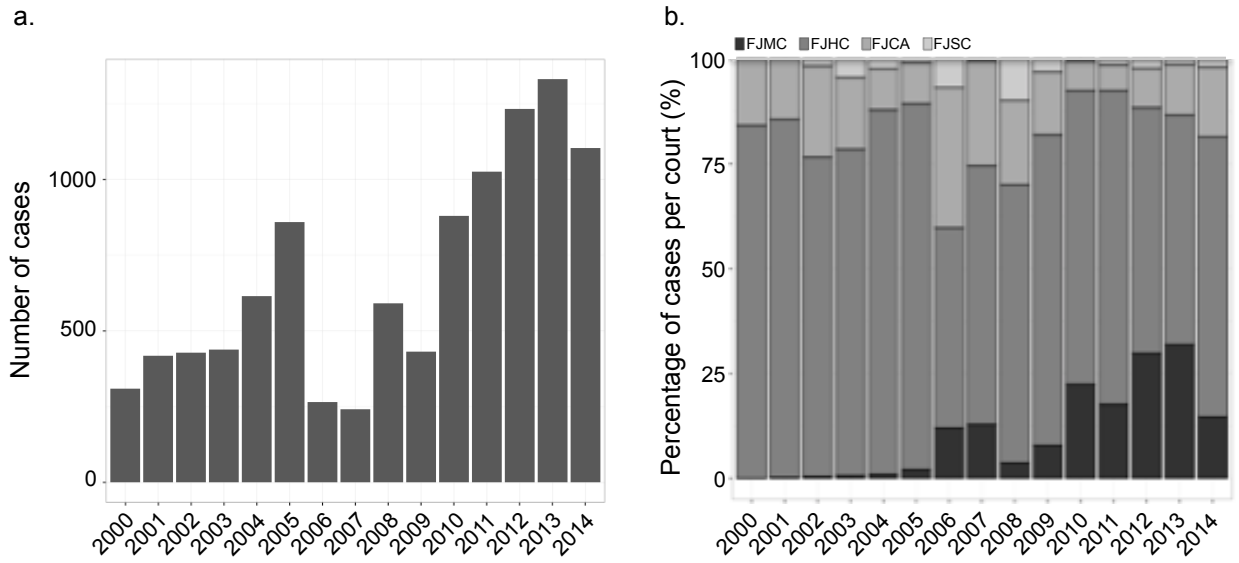

Figure 1. a. Number of cases per year, totaling 10,173 over the period 2000-2014. b. Distribution of cases per court over 2000-2014. FJMC: Magistrate Court, FJHC: High Court, FJCA: Court of Appeal, FJSC: Supreme Court.

C. Trends – Offense Types Within the Overall Case Corpus

Network cluster analysis was leveraged to identify groups of cases within the dataset that are more similar to each other than to those of another group, in much the same way that social groups can be identified based on studies of individuals and their interactions. This was done to provide a high-level overview of types of cases present in the case law dataset over time. The analysis makes no a-priori assumptions as to the total number of groups, an input typically required for classic clustering approaches, e.g. k-means.

C1. Strong increase in prosecution of rape/sexual assault case starting in 2010

We focus the analysis on specific legal language pertinent to this study, i.e. words likely to be used in the context of offenses related to domestic violence (DV) and/or sexual assault (SA) cases, which are of interest when studying Violence Against Women and Girls (VAWG). To this end, a total of 44 such words (mono-grams and bi-grams) were independently aggregated by ICAAD and their use compared across the case dataset. The list includes specific charges, e.g. *indecent exposure,* indicators of severity of offense, e.g. *under 18*, and other language specific to Fiji, e.g. *bulubulu* (for a complete list, see Methods).

We identified 4,503 cases (44% of total) that contained at least one target word. Within this subset, five large groups of cases could be identified[40], broadly representing the following categories:

[40] For completeness, we note that network cluster analysis identified two additional groups, with total population of 25 cases (0.24% of the total). Based on manual review, these were identified as out-of-context cases and excluded from analysis

assault/bodily harm (1,354 cases), rape/sexual assault (1,160 cases), abuse (1,122 cases), murder/manslaughter (681 cases), and restraining order/harassment (161 cases). Plotting the four largest groups by annual volume (Figure 2a) and percent annual volume (Figure 2b), we find: 1) a strong increase in the total volume across all groups starting in 2010 and 2) a proportionally strong increase in prosecution of rape/sexual assault cases also starting in 2010.

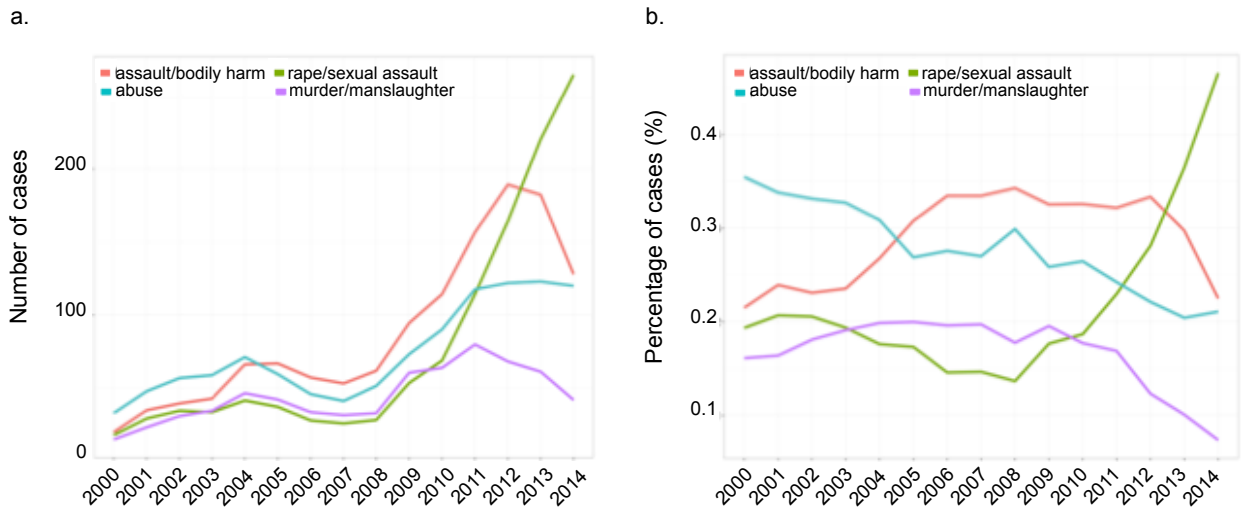

Figure 2. a. Number of cases per cluster over time. b. Distribution of cases per cluster over time. Data reflects three-year rolling average.

It is interesting to note that while sexual assault cases are largely grouped within a single cluster, there isn't a coherent cluster characterizing domestic violence (DV) offenses. This stems from the fact that DV can fall into any one of the existing clusters – with distinguishing feature being not the type of offense, but rather the relationship to the victim (e.g. one's spouse, family member or close relative neighbor). As these distinguishing features were not explicitly included in the target words, the clustering analysis was not able to tease out domestic violence cases into a separate category of crime.

D. Trends – Domestic Violence and Sexual Assault cases

Another strategy to retrieve cases of interest uses supervised machine learning. The algorithm is given example inputs and desired output, and learns to map inputs to outputs based on a subset of manually labeled cases. It can then predict with some probability the class of new cases providing availability of the same inputs. The larger the training set, the more accurate the prediction, so this type of algorithms is not only easily scalable but also easily adaptable to language evolution. Using this technique, we labeled 1035 SA cases (10.2%) and 475 DV cases (4.7%). 193 cases (1.9%) were classified as both SA and DV.

Due to the different nature between murder, SA and DV cases, and the impact it has on sentencing, some analyses would be biased by the dual nature of some of the cases. Therefore, murder cases were segregated out and cases labeled as both SA and DV were classified as SA. A nomenclature using "prime" sign is adopted to distinguish between SA/DV cases that have been labeled by the

algorithms and are not mutually exclusive and SA'/DV' cases that exclude murders and are mutually exclusive. This new classification results in adjusted counts of 973 SA' and 255 DV' cases.

D1. Steady increase in prosecution of domestic violence since 2000; confirmed strong increase in prosecution of sexual assault starting in 2010

The overall volume of DV and SA case shows a strong increase starting in 2010. This increases since 2010 mimic the overall trend across all case types, as identified in section III.C.

A closer look at DV cases shows a steady increase since 2000 to date in proportion of cases being prosecuted – rising at around 0.5% per year and reaching 7% of total in 2013 (Figure 3). On the other hand, after following the same steady increase from 2000 to 2010, a strong relative rise in prosecutions of SA cases is seen through 2014, reaching 20% of the overall case volume in the same year (Figure 3). The latter is also independently corroborated in section III.C.

Despite progress in bringing SA/DV cases to court, as evidenced by proportional (and overall) increase in cases being prosecuted, it's only a start in the fight for improvement of women's right in Fiji. With 64% of women and girls (pop. 290,000) admitting to SA or DV encounters in their lifetime[41], the number of prosecutions is still a relatively small fraction of the overall volume of cases expected to be coming through the court systems – likely due to underreporting and/or lack of willingness to prosecute.

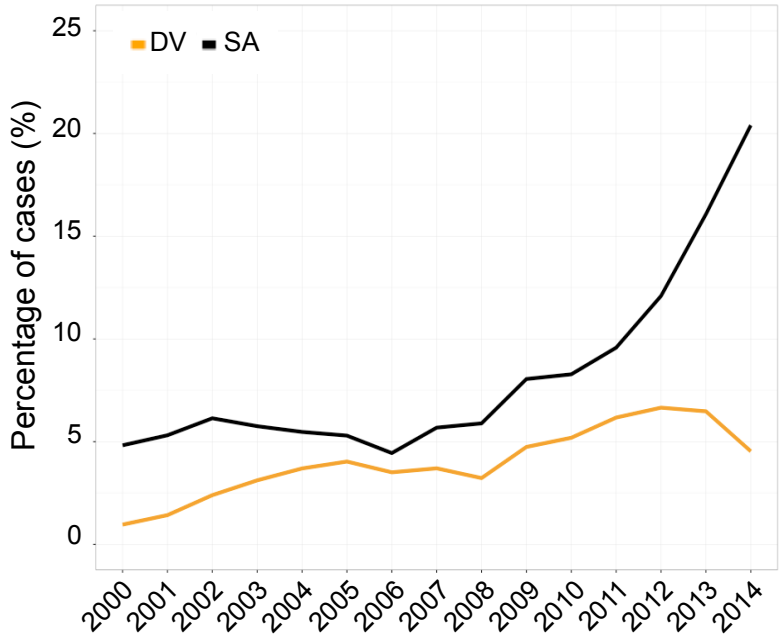

[41] UNICEF, Harmful Connections: Examining the Relationship Between Violence Against Women and Violence Against children in the South Pacific, at 4 (2015), http://www.unicef.org/pacificislands/Harmful_Connections(1).pdf.

Figure 3. Proportion of Domestic Violence (DV) and Sexual Assault (SA) cases over time. Data reflects three-year rolling average.

E. Trends – Granular Breakdown of Charges in Sexual Assault (SA) and Domestic Violence (DV) cases

Using combinations of regular expressions, the top ten most frequent charges encountered were extracted from the set of DV and SA cases: rape, attempted rape, indecent assault, sexual assault, defilement, incest, indecent exposure, assault causing actual bodily harm, unnatural offense, and murder. Those regular expressions, based on manual review of a random sample of cases, captured nuance relative to variations in language encountered across the case corpus. Note that a single prosecuted case may be associated with multiple charges – as a result, percent total of charges across all cases will not add to 100%.

E1. Rape is the most common charge in SA and DV cases

Averaged across years 2000-2014, we find the proportion of charges per year to be steady for DV and SA offenses, with notable exception for the sexual assault charge that shows a significant growth of ~0.6% per year in average in both SA and DV cases[42] (Figures 4a and 4b). Rape charge represents the most common charge in both case groups, with an average of 61% for SA cases and 32% for DV cases over the period 2000-2014. Indecent assault, defilement, and attempted rape are the next three most common charges for SA cases (respectively 16%, 8%, and 6%), and assault causing bodily harm, indecent assault, and murder the three most common for DV cases (respectively 27%, 13%, and 12%). The high prevalence of rape charges within DV cases shows that DV and SA cases are highly intertwined. Indeed, the difference between SA and DV cases does not necessarily rely on the charge being prosecuted, but rather on the relationship between the offender and the victim.

[42] Average trends inferred via linear regression over years 2000-2014

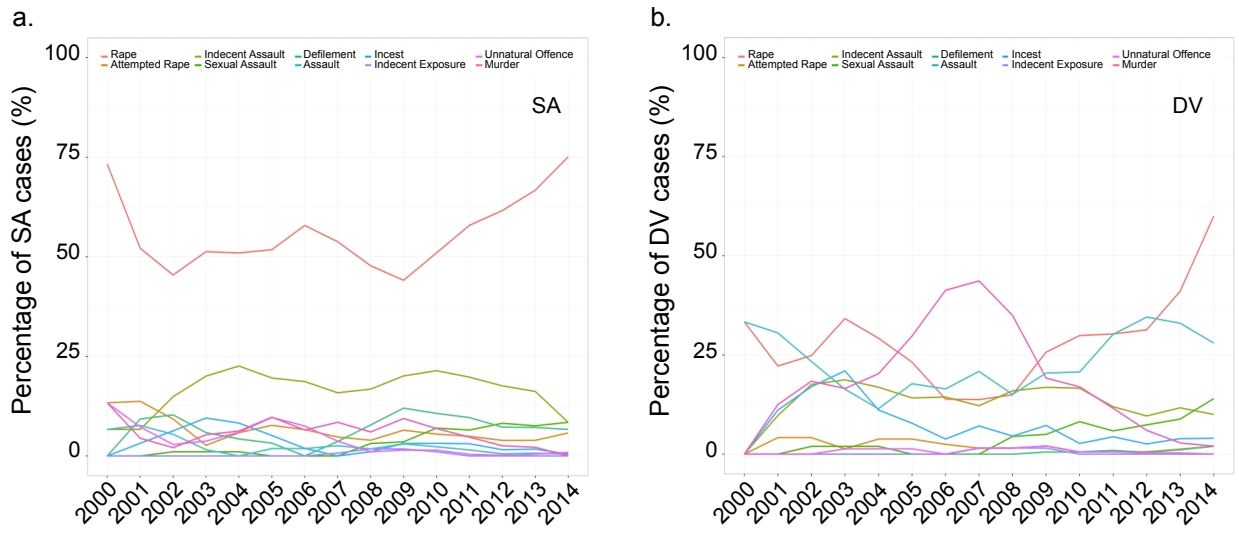

Figure 4. a. Breakdown of SA cases per charge. b. Breakdown of DV cases per charge. Data reflects three-year rolling average.

F. Trends – Adoption of decrees on the front lines of justice system

Using the same technique as the one used in the previous section, i.e. combinations of regular expressions, references to legislations were identified in the case corpus. These include: Penal Code, Crimes Decree, Criminal Procedure Decree, Sentencing and Penalties Decree, and Domestic Violence Decree (DV Decree).

In order to get additional insights on the use of the DV Decree within a judgment, we additionally extracted references to "restraining order" and "reconciliation", and developed a metric to quantify the "degree of emphasis" placed on the decree by the judge. Emphasis was defined using a topic modeling approach, which relies on the intuition that similar topics make use of similar combination of words and phrases. Emphasis, in effect, measures the degree to which a judge relies on the DV Decree in passing a verdict.

Due to the dual nature of some cases and the impact it may have on the use of DV Decree, the following analyses were performed on mutually exclusive groups of SA' and DV' cases[43].

F1. The Decrees enacted in 2009 are displacing the Penal Code over time

A number of decrees were enacted in 2009 to complement and strengthen the Penal Code. Figure 5 shows that there is little gap in adopting the newly enacted legislations, with more than 10% of all cases referencing the new decrees in year 1 of adoption. References also show a steady rise to 2014,

[43] Definition of SA'/DV': SA'/DV' are mutually exclusive groups of cases where all cases containing both SA and DV features are pulled into the SA' category and murder cases are segregated out

the latest year for which the data is available. Rise in new laws are accompanied with concurrent decrease in references to Penal Code.

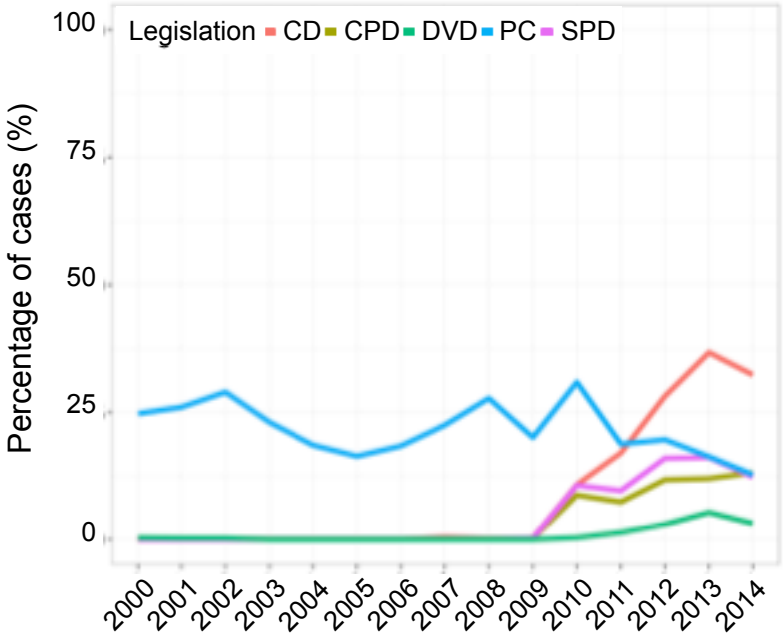

Figure 5. Percentage of all cases referring to legislations over time. CD: Crime Decree, CPD: Criminal Procedure Decree, DVD: Domestic Violence Decree, PC: Penal Code, SPD: Sentencing and Penalties Decree.

F2. The DV Decree is increasingly used in DV cases and tends to favor victims' protection

In order to more specifically evaluate the adoption of the DV Decree by the judicial system, we looked at citations of DV Decree within 255 DV' cases[44]. The findings show a steady increase in citations between 2010 and 2014, reaching 93% in 2014 (Figure 6a) and suggesting that judges have increasingly integrated discussion about the decree into their sentencing decisions. An analysis of the different variations of "reconciliation" and "restraining orders" terms within those cases indicates an increase in usage for the purpose of issuing restraining orders and a decrease for ruling on "reconciliation" (Figure 6b).

[44] Definition of SA'/DV': SA'/DV' are mutually exclusive groups of cases where all cases containing both SA and DV features are pulled into the SA' category and murder cases are segregated out

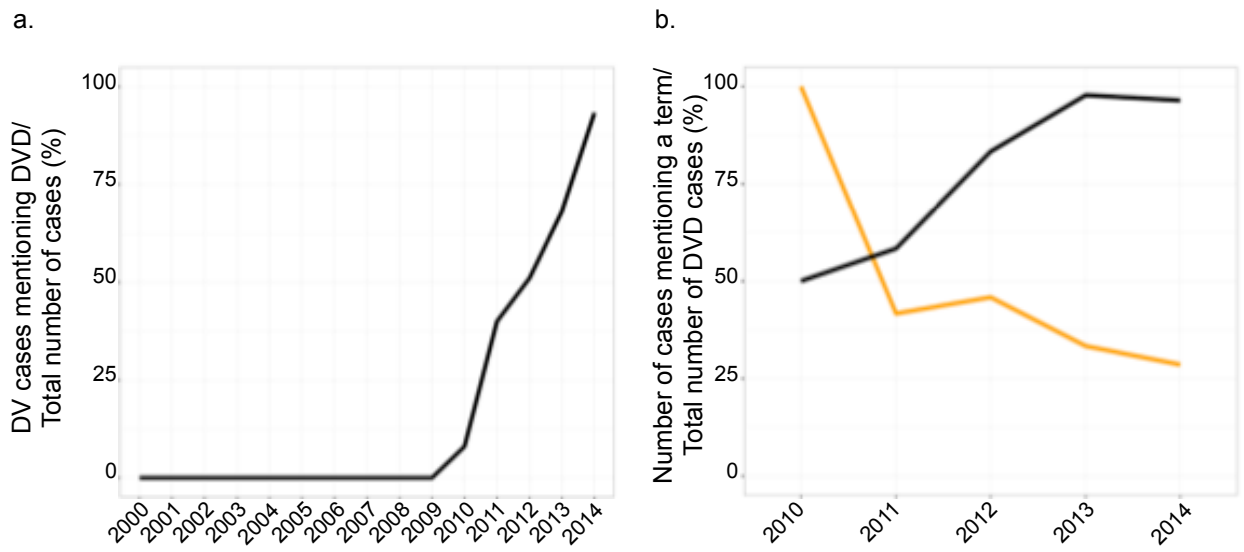

Figure 6. a. Percentage of DV' cases referring to DV Decree over time. b. Percentage of DV' cases referring to DVD and mentioning 'reconciliation' or 'restraining order' over 2010-2014.

This is very encouraging trends in the court from the perspective of advancement of women's rights in Fiji. The newly enacted laws are showing increased use on the front lines of the judiciary system, and are accompanied by increased prosecution of most severe crimes against girls and women (see Sections C1, D1, and E1). Additionally, the increased issuance of restraining orders show that judges acknowledge the necessity to better protect victims and need less and less to discuss situations of reconciliation during the judgment.

F3. References to DV decree via topic modeling

Figure 7a displays a measure of emphasis for each DV case, using a topic modeling methodology, broken down by cases referring, or not, to the DV Decree (in yellow and black respectively). An increase in score within documents that are known to refer to the DV Decree is observed, therefore supporting the use of this metric as a measure of emphasis. In the subset of cases referring to the DV Decree, a consistent decrease from 0.51 to 0.17 is observed (yellow 2010-2014), while the remaining DV cases stay at 0.02 in average (black 2010-2014). The breakdown of cases referring to DVD by court highlights a constant emphasis in the Magistrate Court over the year while the emphasis seems to decrease in the High Court (Figure 7b). It is interesting to see that despite representing 5% of the DV cases, only one case in the Court of Appeal refers to the DV Decree. A decreasing emphasis at the appellate level of the DV Decree may just reflect the lower courts increasing proficiency at applying the DV Decree. Indeed, if a lower court is correctly emphasizing the DV Decree in its decisions, then the High Court would no longer receive appeals based solely on the grounds of failure to adhere to the DV Decree.

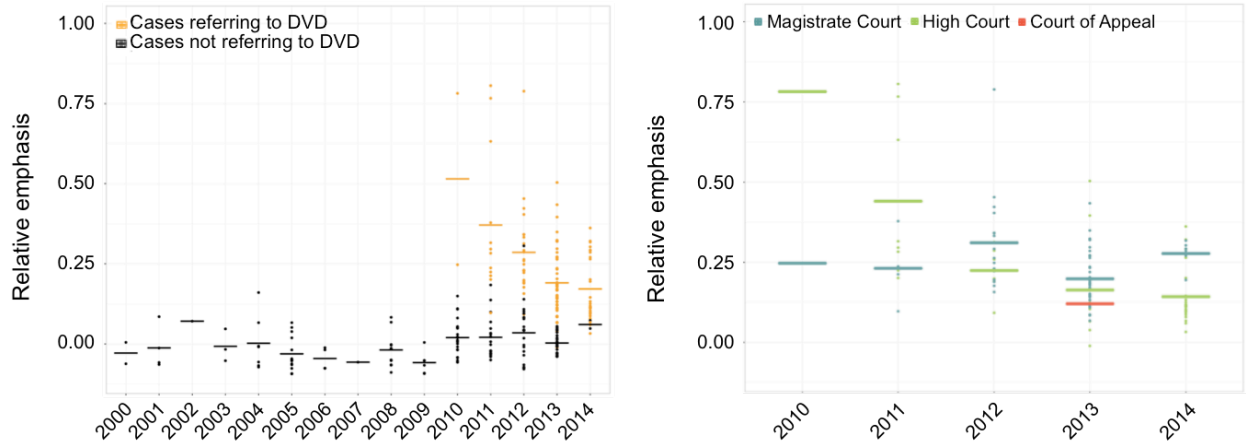

Figure 7. a. Emphasis of DV Decree within 255 DV' cases over time, broken down by cases that have been identified as referring to the DV Decree (yellow) or not (black). b. Emphasis of DV Decree within 111 DV' cases referring to the DV Decree over time, broken down by court. Means were added for each population. FJMC: Magistrate Court, FJHC: High Court, FJCA: Court of Appeal.

G. Trends – impact of age and charge on the sentence

Combinations of regular expressions were used to identify if the victim was under or over 18. Such strategy could not be reliably applied to extract the final sentence due to the lack of consistency in logic and document structure around this information. Therefore, final sentences for 231 SA/DV cases were manually extracted by lawyers. Cases in which the offender was acquitted and cases leading to life sentence (murder) were excluded from the analysis, resulting in selected 220 cases with sentence details.

Due to the dual nature of some cases and the impact it may have on sentencing, the following analyses were performed on mutually exclusive groups of SA' and DV' cases[45].

G1. Victims under 18 years old are the subject of 23% SA'/DV' cases across 2000-2014

In 23% of 1228 SA'/DV' cases, the victim was identified as less than 18 years old. No change in trend was observed over time, apart for the years 2005-2008, which showed a significant decrease in cases involving minors (Figure 8). The distribution of age across the different categories of cases was analyzed and showed that minors are more often victims in SA' cases than in DV' cases (proportion of cases involving minors in SA': 26% vs. DV': 11%). SA' cases are more severe crimes and therefore have a higher chance of being prosecuted. We would therefore assume that the 26% figure for SA' cases is closer to reality than the 11% for DV' cases, which are likely being under-prosecuted.

[45] Definition of SA'/DV': SA'/DV' are mutually exclusive groups of cases where all cases containing both SA and DV features are pulled into the SA' category and murder cases are segregated out

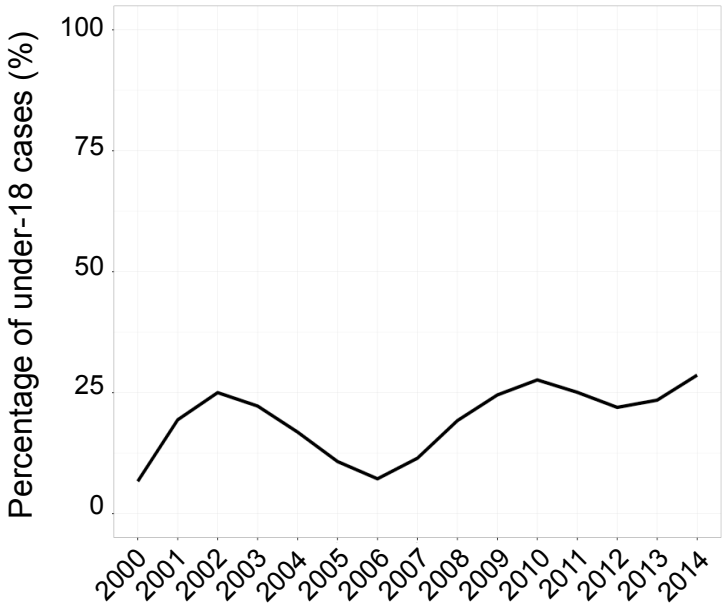

Figure 8. Percentage of 1228 SA'/DV' cases in which the victim is under 18 years old. Data reflects three-year rolling average.

G2. SA cases lead to heavier sentences than DV cases

Using the 220 SA'/DV' cases we found that SA' cases lead to higher mean value of sentences than DV' cases (SA': 7 years vs. DV': 1.42 years). This result confirms previous findings by ICAAD[46]. Additionally, a slight increase in mean final sentences was observed over time (Figure 9), suggesting increased awareness of the severity of these crimes.

[46] An Analysis of Judicial Sentencing Practices in Sexual & Gender-Based Violence Cases in the Pacific Island Region INTERNATIONAL CENTER FOR ADVOCATES AGAINST DISCRIMINATION & DLA PIPER (2016).

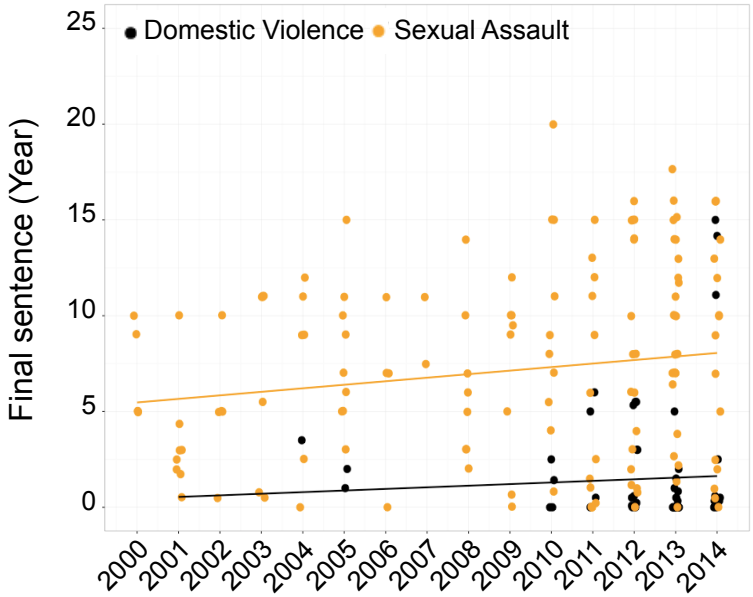

Figure 9. Final sentence over time for SA' and DV' cases. Plots show linear regressions for each category.

G3. Final sentence for rape crime shows high variability

In order to study the final sentence for specific crimes, we decided to focus on 100 rape cases for which the final sentence was extracted. Plotting the final sentence over time shows high variability, suggesting that contextual information is largely taken into account by judges in their judgments (Figure 10). For example, in 2014, the range of sentences for rape varies from 1 to 16 years.

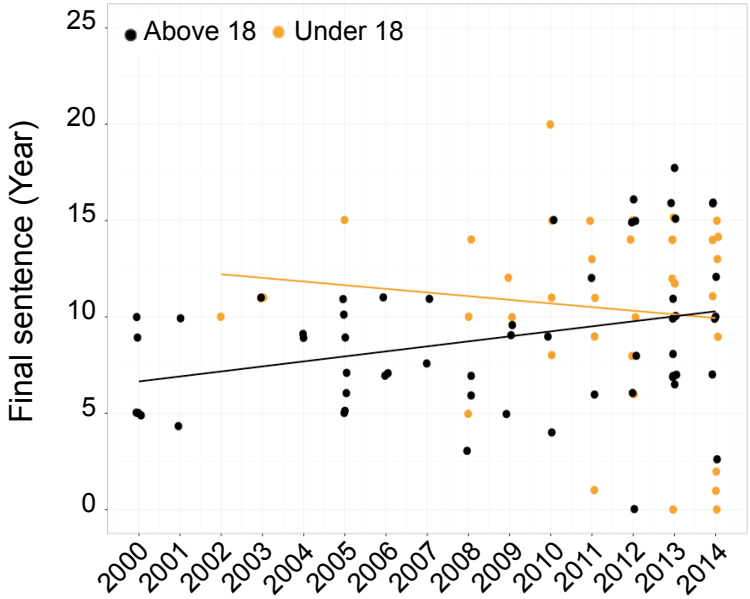

Figure 10. Final sentence for rape crimes over time, broken down by age of the victim. Plots show linear regressions for each category.

G4. Perpetrators who rape minors receive longer sentences but the gap decreases over time

The impact of the victim's age on the sentence was then analyzed. On average, over the years 2000-2014, perpetrators who rape minors receive a longer sentence when compared to those with adult victims, averaging a difference of 1.4 years (under 18: 10.0yr vs. above 18: 8.6yr). However, the difference in sentencing decreases over time (Figure 10). The impact of VAWG in Fiji has gained recognition from the highest levels of government and the judiciary. Increased sentences for minors are consistent with sentencing guidelines that treat criminal acts against minors as an aggravating factor that would justify an increased sentence. However, the gradual decrease in sentencing for perpetrators under 18 over time is a surprising trend amidst judges' recognition of the systemic nature of VAWG throughout society.

IV. IMPACT OF THE STUDY AND DISCUSSION ON THE USE OF DATA ANALYTICS TECHNIQUES ON A CASE LAW DATASET

A. Significance of the results

The use of data analytics techniques on case law datasets can help produce concrete indicators of progress towards the implementation of certain types of recommendations in the UPR, and provide valuable insights on the situation of human rights in states under review. In this project, we provided several metrics aiming at measuring Fiji progress on SGBV in Fiji. In order to assess Fiji's compliance with UPR recommendations it accepted from Norway in 2010 and Belgium and Bangladesh in 2015, we propose a framework that uses the metrics described in this study to measure progress in the following areas.

A1. Transparency of outcomes at various levels of Fiji's court system

The increase in the volume of cases and the proportional increase in magistrate court cases show a greater effort by the Fijian justice system to become more transparent. Indeed, lower court cases are not prioritized in many Pacific Island jurisdictions and PacLII only has appellate court cases for many of them. Therefore, by making accessible a greater number of cases and proportionally greater number of magistrate court cases, the data reveals a positive trend towards greater transparency.

Availability of case data is crucial to enabling external actors (such as NGOs, recommending states, etc.) to track progress towards UPR recommendations. Demonstrating transparency in the judiciary system is therefore an important step for Fiji to make progress towards meeting the goals of its UPR recommendations.

A2. Accountability of magistrates towards the laws that govern sexual assault and domestic violence

The proportional increase in SA/DV cases, and especially SA cases starting 2010, shows Fiji's effort to prosecute more SGBV cases. Additionally, we observed a constant increase in the citation of DV Decree in DV cases from 2010 to 2014, reaching 93% citation rate in 2014, suggesting that the Decree is well adopted by the judiciary system. The next step was to search for occurrences of terms indicating mentions of restraining order or reconciliation in order to gain some insights into the reasons for applying the DV Decree. We found an increase in use of the Decree for enacting restraining orders and a decrease in its use for ruling on reconciliation, which aligns with Norway's recommendation to "prohibit practices that legalize violence against women". Finally, the average final sentences for both SA and DV cases tend to increase over time, showing an increased awareness of the severity of SGBV crimes and making sure that "perpetrators of violence against women, including within the family, [are] duly prosecuted and punished".

Nonetheless, there are trends that raise concern, specifically regarding violence against minors. The percentage of cases involving minors remains constant, showing that no campaign or policy has been successful in improving children's protection. Additionally, data shows a steady decrease of prison time for perpetrators who have raped minors since 2002.

A3. Consistency of sentencing in sexual assault and domestic violence cases

The final sentence for rape cases show high variability suggesting that contextual information is largely taken into account into the judgments and that judges might rule differently on similar cases. Lack of consistency in judgments may undermine court's authority as provides grounds for contesting judges' decisions. Therefore, Fiji should aim at understanding which parameters have an impact on ruling in order to reduce variability.

B. Benefits and limitations of the methodologies used

This project enabled us to explore various computational approaches and evaluate their efficiency. The case classification was performed using supervised machine learning. This technique appears to be very efficient in the classification of SA/DV cases (F-score for SA: 96% and DV: 81%) but showed some limitations in the distinction between relevant and non-relevant cases (see Methods for definition). Indeed, respectively 29% and 38% of true positive DV and SA cases were labeled as non-relevant. Increasing the size of the training set or combining other techniques (e.g. text-mining) with supervised machine learning algorithms could help refine the classification.

A topic modeling approach was also developed to provide a measure of "emphasis" of DV Decree within cases. The limitation of this technique relies on the interpretation of the cosine similarity score, or measure of relative emphasis, and requires close collaboration with legal experts to provide plausible interpretations.

We also used a cluster analysis technique to provide an unbiased overview of the dataset's trends over time with minimal input. This approach provides a quick assessment of the situation, which was subsequently confirmed with more refined analyses.

Most of the remaining features were extracted using more or less elaborated text mining techniques. Text mining remains an efficient approach to detect occurrences of words within documents and, providing a good understanding of lexical and grammatical variations for each feature, turns out to be an efficient technique for extracting features. However, wording/language evolution requires regular update of the code and a manual verification step is required to estimate the errors made for each of those features.

These findings were obtained solely through the analysis of a case law dataset, which emerges as a valuable source of quantitative information when analyzed with the appropriate tools. Rate of application of legislation, specific usage, trend and distribution of number of cases and offenses, variability in judges' ruling, and average sentence per charge are examples of data that can be extracted from such dataset and help inform decision makers or advocates on the adherence to human rights principles within a country.

C. Lessons learned

C1. Leakage along the reporting pipeline negatively impacts data representativeness

The lack of transparency in the accuracy, timeliness, and completeness of case upload to the PacLII database prevents us from having a clear understanding of the representativeness of this dataset (Figure 11). For example, data on the percentage of offenses that are being prosecuted or on the percentage of hearings that are transcribed is missing. There is also no clarity on the identity of the person in charge of the case write up. Consequently, variability in writing styles may have an impact on text mining or topic modeling approaches, which are sensitive to the language used. There is also no clarity on the selection of cases that are being uploaded to the database. Therefore the dataset may suffer from language, social class, tribunal, court, or geographic biases and extrapolation of the findings to the country level would require clarification on the overall pipeline.

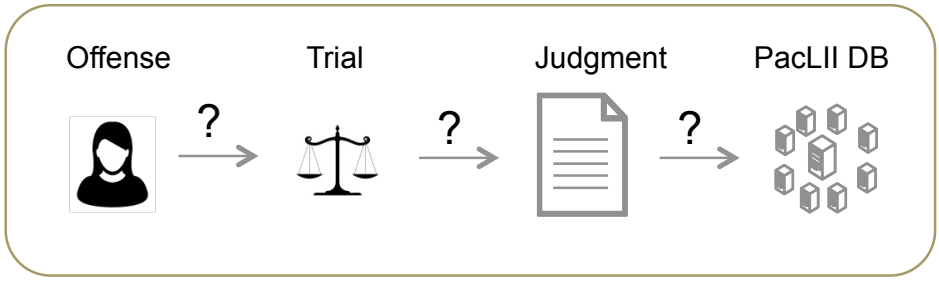

Figure 11. Life cycle of a case from the offense to PacLII database.

C2. Encouraging big and consistently formatted data

Some Pacific Island countries have limited number of publicly available cases. In this situation, the use of text mining and data analytics techniques have limited impact, and manually reviewing such datasets is more accurate and time effective.

Important variability across cases was observed with regard to the structure of the document, the level of details, the formatting of specific information (e.g. charges, sentences, legislations, reference to precedent cases,), and the spelling of names. Many case attributes were there, but unstructured

(e.g., charge, victim's age) or incomplete (e.g., final sentencing, victim's age). This limited the performance of the case classifiers and text mining techniques, and motivated the use of manual intervention to extract the final sentence.

C3. Resources required

Technical capacities as well as country-specific legal expertise are key to the implementation of this study. The team should include one or two data scientist(s) (3-4 months FTE) with expertise in machine learning and text mining techniques as well as one legal expert who can provide contextual knowledge, build dictionaries of key words and manually review the output of the algorithms.

This project can be completed in three to four months and can be divided into four equal phases: 1) acquire contextual knowledge through case reading and interviews with legal experts, 2) design the analyses, build the dictionaries, write and run the code, 3) manually review the results and make some adjustment to the code, 4) interpret the results and provide some insights on some key indicators

METHODS

Software and packages used

R[47] version 3.1.0 including ggplot2[48] and reshape[49] packages. Python version 2.7 including nltk, pattern.en, re, string, pandas, os, io, numpy, shutil, random, math, scipy, glob, itertools, MySQLdb,

Dataset

10,173 judicial cases in the Fiji Islands spanning from 2000 – 2014 were downloaded from the PACLII website (http://www.paclii.org) in HTML format and then converted to text for the purpose of our analyses. The download was completed on March 26th, 2015.

Cleaning of the documents

All documents were processed following a common methodology: a decode/encode step ensuring the readable transcription of all characters, the exclusion of English stop words and punctuations, and the conversion of all word numbers to digits to facilitate age extraction. An additional step of lemmatization was added for the case classifier and topic modeling algorithms. For the topic model, all words were converted to lower case and additional domain stop words were excluded ('case', 'actual', 'appellant', 'fact', 'mindful', 'court', 'occasioning', 'occasion', 'hi', 'particular', 'wa', 'complainant', 'month', 'year', 'learned', 'would', 'r', 'due', 'tell', 'told', 'said', 'say'). Depending on the algorithm used, the output of the cleaning process was a string of words or a list of sentences.

Cluster Analysis

Cluster analysis was used to profile the full set of case transcriptions, in order to identify and track broad categories of cases over time. The algorithm used was the hierarchical Louvain, implemented as part of the python-louvain package. Input to the algorithm was a weighted network of cases, built using networkx python package. Detailed steps performed as part of this procedure are defined below:

1. All cases were first converted to tf-idf vectors using python's scikit-learn package. Tf-idf score, in contrast to a word count-based score, helps adjust for the fact that commonly used words are also less likely to provide strong insight within a given dataset. Tf-idf vectors were limited to 44 words of interest (WOI) identified by ICAAD as most relevant for this study. These are itemized in Table 1 and include a mix of monograms and bi-grams. Prior to analysis, both case transcriptions and WOI were reduced to their root using a standard Porter stemmer via a python nltk package.

2. Each case was compared to every other case in the dataset, and assigned a similarity metric, via the tf-idf score. To define pair-wise case similarity, Jaccard metric was used, which ranges between

---

[47] R Core Team (2014). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL http://www.R-project.org/.
[48] H. Wickham. ggplot2: elegant graphics for data analysis. Springer New York, 2009.
[49] H. Wickham. Reshaping data with the reshape package. Journal of Statistical Software, 21(12), 2007.

zero and one. Zero means that two cases share no WOI in common with one another. One means that two cases not only share all the same WOI but that they share the same degree of emphasis for each WOI.

3. A network was created using a networkx package in python. In a network, each node is associated with a single case, while a Jaccard metric defines the weight of a connection between any two nodes. In total, 5,870 orphan cases were found, defined as cases that either contain no WOI or share no WOI with any other case in the dataset. In technical terms, these are nodes whose weights to all other nodes sum to zero. Clustering analysis was performed on the remainder of the network

Table 1: Words of interest (WOI) used in cluster analysis

Monograms	*abduct, abuse, anal, anus, assault, bulubulu, defile, dvro, enslave, grope, harass, incest, indecently, intercourse, intimidate, manslaughter, murder, penetrate, penis, rape, reconciliation, sex, sexual, slave, vagina, vaginal*
Bi-grams	*attempted rape, bodily harm, carnal knowledge, domestic relationship, domestic violence, grievous harm, gross indecency, indecent assault, indecent exposure, order of protection, restraining order, sexual violence, under 13, under 18, under age, unnatural offense, violence decree, young person*

Table 2. Top 5 words for each cluster and cluster interpretation

Cluster name	Size (#cases)	Top 5 words	Interpretation of the cluster
assault/bodily harm	1354	*Assault, bodily harm, intimidation, incident, reconciliation*	Cases are mostly related to physical assaults, excluding sexual assaults
rape/sexual assault	1160	*Rape, sexual, vagina, intercourse, assault*	Most of the sexual assault cases will be included in this cluster
abuse	1122	*Abuse, under 18, harass, assault, murder*	Cases in this cluster show a large range of type of aggressions but all associated with the notion of abuse
murder/manslaughter	681	*Murder, manslaughter, assault, bodily harm, abuse*	Cases in this cluster are related to more severe aggressions, such as manslaughter or murder
restraining order/harassment	161	*Restraining order, harass, under 13, under 18, order of protection*	This cluster mostly includes cases on minors where protective orders were issued

Case classifier

A supervised machine learning approach was used to identify SA and DV cases. A subset of 1618 cases were labeled by ICAAD law experts, resulting in 232 (14%) relevant SA cases and 154 (9.5%) relevant DV cases. 58 cases were labeled both SA and DV. Cases were deemed relevant when a guilty verdict was issued and when the judge provided rationale for the length of sentence the perpetrator would be receiving. All cases referring to DV or SA but not providing information on the judgment itself (e.g. bail ruling, jury's instructions) were considered irrelevant for the scope of

our analysis. Two different models were developed in order to label documents as SA/non-SA and DV/non-DV. The same method was used to select the best model in each case: firstly, the set of 1618 labeled cases was randomly split into two subsets of documents following a 80/20 distribution for training and testing purpose and respecting the initial distribution of labeled cases (14% for SA and 9.5% for DV); secondly, a term documents matrix was generated using the tf-idf Vectorizer from scikit-learn and using 1,2,3-grams; thirdly, the top 1% most relevant features were selected using the SelectPecentile function and chi2 test from sklearn.feature_selection; fourthly, the different logistic regression models were evaluated via a grid search approach. The best model was selected using a StratifiedShuffleSplit cross validation and F-score (SA: LogisticRegression(C=100, class_weight=None, dual=False, fit_intercept=True, intercept_scaling=1, penalty='l1', random_state=None, tol=0.0001; DV: LogisticRegression(C=1000, class_weight=None, dual=False, fit_intercept=True, intercept_scaling=1, penalty='l1', random_state=None, tot=0.0001).

In order to label the remaining 8555 documents, the term document matrix was generated using the same vectorizer and each logistic regression model was trained on the full set of 1618 labeled cases. Out of the 8555 unlabeled documents, 803 cases were labeled SA and 321 cases were labeled DV. 135 were labeled both.

Topic model

A topic model approach was developed as a measure of "emphasis" of the DV Decree within cases. In order to favor the generation of a topic vector related to the DV Decree, we applied the tf-idf Vectorizer to a subset of 58 DV cases, using 1,2,3-grams and a maximum of 500 features. A non-negative matrix factorization algorithm (scikit-learn) was then applied to generate a 10x500 topic matrix (10 topics). One of the topic vector's top 20 features were closely related to the DV Decree (violence, domestic, domestic violence, victim, decree, domestic violence decree, violence decree, order, restrain, restrain order, magistrate, violence offense, domestic violence offense, 2010, protection, learn magistrate, 2011, provision, make) and was therefore defined as the DV Decree topic vector. The remaining documents were then vectorized following the same procedure and a cosine similarity score was computed between each document vector and the topic vector using scikit-learn's cosine_similarity function. A sample of documents with various scores was given to law experts for interpretation:

- Documents with score above 0.485 show a high density of words. In this case, there are many variations of the DV decree or DV restraining orders throughout the stated decision, indicating that the final decision includes any form of the legislation.
- Documents with scores ranging from 0.351 to 0.485 show a significant density of words. In this case, there are substantive discussions around the DV decree or DV restraining orders but they also include other statutes, such as crime.
- Documents with scores ranging from 0.306 to 0.351 show a low density of words. In this case, there is a low variation of DV and DV restraining orders.
- Documents with score below 0.306 show no evidence of DV.

The "relative emphasis" was defined as: (score – baseline)/(1.00 – baseline). The baseline is the average score of cases that do not refer to the DV Decree. The baseline is equal to 0.17.

Regular expressions

Remaining features were extracted using regular expressions: year, court, references to decrees, penal code, reconciliation, restraining order, charges and age of the victim.

Year and court were extracted from the title of the documents. The legislations and reconciliation features were extracted using a dictionary of expressions including the different decrees, the penal code and specific features of those legislations. For instance, the dictionary developed for the DV Decree includes 'domestic violence decree', 'domestic violence restraining order', 'non contact order', 'non molestation order' and 'DVRO'.

For charges, age of the victim, and restraining orders, a combination of regular expressions within the same sentence was developed. The first step consists of extracting strings of pre-defined size that contain expressions and synonyms of interest, such as 'rape'. Then a check for negation on every topic is run, as well as a check on derived expression, such as 'attempted rape'. Finally, the script makes an attempt to extract the age of the victim in the vicinity of offenses statement and victim qualifiers.

Post processing of the database

The extracted data is later appended to a SQL database. A cross validation check is applied to reclassify some of the DV and SA cases as the DV dimension of SA cases can be undetected by machine learning algorithms. Therefore, all SA cases that were not labeled as DV but were referencing DV decree were automatically additionally classified as DV cases.

Linear Regression analyses

Linear regressions were performed using R.

Results for linear regression for each charge in SA or DV cases over time:

	SA				**DV**			
	Adj. R2	Intercept (%)	Slope (%/yr)	p	Adj. R2	Intercept (%)	Slope (%/yr)	p
Rape	0.040	47.43	0.937	0.23	0.078	17.12	1.265	0.16
Attempted rape	0.118	11.42	-0.585	0.11	0.021	4.19	-0.242	0.28
Defilement	-0.073	6.32	0.085	0.83	0.291	-0.39	0.090	0.02
Assault causing actual bodily harm	0.238	5.88	-0.405	0.04	-0.040	19.29	0.503	0.51
Indecent assault	-0.024	13.35	0.400	0.43	-0.077	13.40	0.003	0.99
Incest	0.019	5.31	-0.257	0.28	0.032	12.37	-0.647	0.25
Sexual assault	0.505	-2.26	0.672	0.002	0.402	-2.30	0.761	0.007

Unnatural offense	0.302	8.76	-0.623	0.02	-0.076	0.66	0.010	0.92
Indecent exposure	-0.055	0.11	0.030	0.61	-0.073	0.17	0.016	0.83
Murder	-0.029	7.67	-0.254	0.45	-0.031	26.35	-0.888	0.46

Manual evaluation of classifiers and features extraction performances

In order to assess the efficiency of the various algorithms developed, random samples of 50 SA/DV cases were provided to legal experts who went through a manual review of a total of 200 documents and corresponding extracted features. This procedure enabled us to estimate the error made on each attrtibute (Figure 12):

Attribute	**TP**	**TN**	**FP**	**FN**	**tp rate = recall**	**fp rate**	**precision**	**accuracy**	**F score**
DV	51	125	17	7	87.93%	12.88%	75.00%	88.0%	80.95%
SA	147	40	9	4	97.35%	20.45%	94.23%	93.5%	95.77%
DVD	14	182	4	0	100.00%	2.20%	77.78%	98.0%	87.50%
DVRO	15	183	2	0	100.00%	1.09%	88.24%	99.0%	93.75%
RO	16	182	2	0	100.00%	1.10%	88.89%	99.0%	94.12%
Rape	91	94	1	14	86.67%	0.93%	98.91%	92.5%	92.39%
Att. Rape	9	186	0	5	64.29%	0.00%	100.00%	97.5%	78.26%
Defilement	19	179	1	1	95.00%	0.56%	95.00%	99.0%	95.00%
Harm	19	180	0	1	95.00%	0.00%	100.00%	99.5%	97.44%
Incest	2	197	0	1	66.67%	0.00%	100.00%	99.5%	80.00%
Indecent	29	171	0	0	100.00%	0.00%	100.00%	100.0%	100.00%
Sexual Assault	13	187	0	0	100.00%	0.00%	100.00%	100.0%	100.00%
Indecent Exposure	0	200	0	0		0.00%		100.0%	
Unnatural Offense	0	200	0	0		0.00%		100.0%	
Murder	10	189	0	1	90.91%	0.00%	100.00%	99.5%	95.24%
Under 18	52	128	2	18	74.29%	1.37%	96.30%	90.0%	83.87%

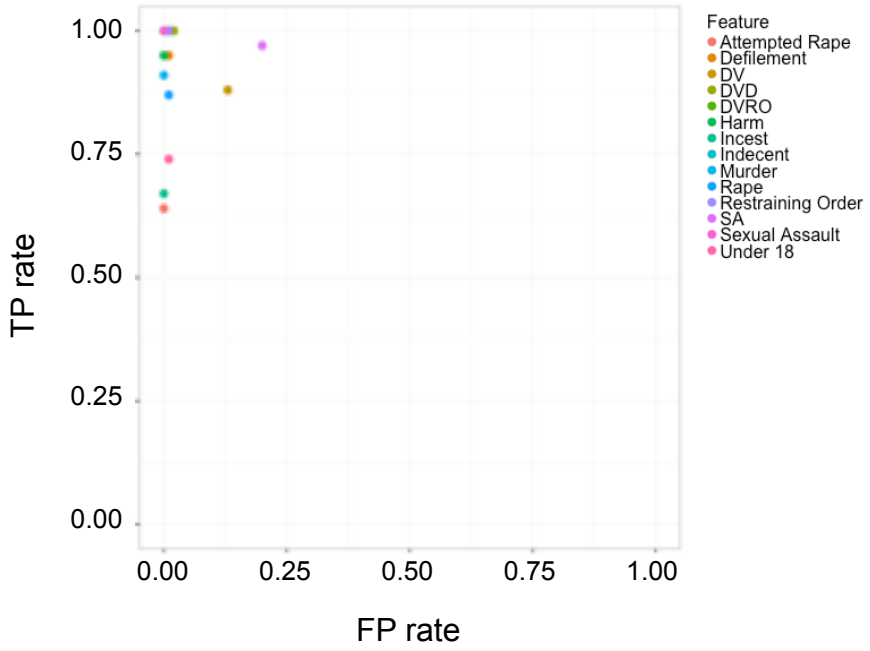

Figure 12. Evaluation of classifiers performance

A review of the scores show that the SA classifier is the most liberal (tp rate = 97% and fp rate = 20%), tending to miss very few SA cases but capturing some false positives. Precision, accuracy and F-score all show good performances. The DV classifier is less liberal and more conservative (tp rate = 88% and fp rate = 13%) – there is a higher chance that the labeled case is a true DV but it also misses more cases. On the other side of the spectrum, the "under 18" feature is very conservative (tp rate = 74% and fp rate = 1%), correctly labeling under 18 cases but missing a high proportion of those.

Manual review of case relevance

During the manual verification process, a subset of documents, labeled by the algorithm as SA or DV, was identified as irrelevant to the scope of the analysis because they do not focus on the sentencing determination of the judge. Those documents are transcripts of specific procedures (pre-trial and during trial) that happen before the judgment and sentencing, such as Summing Up, Voir Dire, Ruling and Extempore Ruling on Bail. Other types of irrelevant cases, although less frequent, are cases where the defendant is acquitted or is a woman, or where a new trial is ordered. Finally, civil cases (divorce, child custody), where sexual assault or domestic violence may be a factor in non-criminal proceedings, are also deemed irrelevant. The proportions of non-relevant cases out of true positive DV and SA cases are 29% and 38% respectively.