A Data-Driven Evaluation of Delays in Criminal Prosecution

Hrafnkell Hjörleifsson , Michelle Ho , Christopher Prince , Achilles Edwin Alfred Saxby , federica B bianco , BetaGov (http://www.betagov.org) , and in collaboration with the Santa Clara District Attorney's office

Abstract

The District Attorney's office of Santa Clara County, California has observed long durations for their prosecution processes. It is interested in assessing the drivers of prosecutorial delays and determining whether there is evidence of disparate treatment of accused individuals in pre-trial detention and criminal charging practices. A recent report from the county's civil grand jury found that only 47% of cases from 2013 were resolved in less than year, far less than the statwide average of 88%. We describe a visualization tool and analytical models to identify factors affecting delays in the prosecutorial process and any characteristics that are associated with disparate treatment of defendants. Using prosecutorial data from January through June of 2014, we find that the time to close the initial phase of prosecution (the entering of a plea), the initial plea entered, the type of court in which a defendant is tried and the main charged offense are important predictors of whether a case will extend beyond one year. Durations for prosecution are found not significantly different for different racial and ethnic population, and do not appear as important features in our modeling to predict case durations longer than one year. Further, we find that, in this data, 81% of felony cases were resolved in less than one year, far greater than the value reported by the civil grand jury.

1 Introduction

In the United States, criminal cases are settled through an adversarial system between the prosecuting attorney who represents the public and the defense attorney who represents the accused. The responsibility of the District Attorney (DA) who prosecutes the case is to bring charges against the accused defendant and prove guilt beyond a reasonable doubt. DA performance is frequently measured by the rate of convictions, plea bargains, or diversions Nugent-Barakove, Budzilowicz, and Rainville (2007). This capstone project focuses on how long it takes for felony cases to be resolved by a District Attorney's office.

This time metric is important to consider because delays in felony case resolutions, or dispositions, places a burden on government resources, leaves defendants uncertain about their futures, and prolongs the wait for closure for victims. Making the criminal justice more efficient while maintaining fairness and due process is beneficial for all parties involved. Association (2006)

The Center for Urban Science and Progress (CUSP) is working with Santa Clara County (SCC) in California and our project sponsor BetaGov at NYU's Marron Institute to investigate the duration and outcomes of SCC's felony cases. In a recent report issued by the SCC Civil Grand Jury, Santa Clara County was found to be the slowest in processing felony cases in California Santa Clara County Civil Grand Jury (2017). The SCC District Attorney's office would like to know how delays and disparities might be explained by case characteristics such as prior convictions, charge enhancements, and defendant characteristics.

The deliverables of this capstone project are two-fold. The first is to provide District Attorneys with an interactive dashboard for exploring and visualizing case progression based on key variables such as the

number of charges, number of defendants, race and age of defendant, and other case characteristics. The second is to provide an in-depth statistical analysis into what variables change the outcome and lengthen the timeline of cases.

2 Literature and Prior Work

2.1 Current condition in Santa Clara County

In a recent report issued by the SCC Civil Grand Jury, it was found that Santa Clara was the slowest of all California counties in resolving its felony cases Santa Clara County Civil Grand Jury (2017). While the rest of the state is able to process 88% of felonies within a year, Santa Clara falls far short. Only 47% of SCC cases in 2013 were resolved within a year. Moreover, according to the same report, SCC has a higher rate of jail incarceration compared to the state. The SCC Civil Grand Jury cited figures from the 2015 Court Statistics report issued by the California Judicial Council Judicial Council of California (2016)

The report found through interviews with officials that a "culture of complacency" that tolerates delays in the county and the DA's approach to charging were contributing factors to case delays, among other reasons. The "culture of complacency" refers to a supposed belief among public officials that everyone in the criminal justice system is already doing their best to move the process forward and that the state Judicial Council's standard of one year for felony case dispositions is unreasonable for complex cases such as gang crimes.

However, it is important to note that the results of the SCC Civil Grand report and the state Judicial Council report are not reproducible. It is unclear what exact data sources were used and whether those sources are publicly available. In addition, the data processing methods that led to the published results and figures are not documented.

2.2 Existing measurements of prosecution performance

In recent years, there have been other data-driven efforts to evaluate and compare court system performance. One such effort is Measures for Justice (2017), an initiative to aggregate and compare the performance of criminal justice systems from arrest to post-conviction for the entire country via an interactive public dashboard. One of the largest challenges is that criminal justice data is neither recorded uniformly across local jurisdictions nor is it publicly available. The solution from Measures for Justice is to reach out individually to jurisdictions to obtain data and then create standardized core measurements for evaluating performance.

Previous studies have examined case processing time as a standardized measurement allowing comparison across jurisdictions Klemm (1986). In order to use case processing time, researches first must subdivide case timelines into appropriate time frames and reduce the scope to time under the control of the court system Neubauer (1983). Early studies have also shown that case complexities such as prior convictions, mandatory minimums, and the number of defendants in specific jurisdictions may contribute to the length of a case Luskin and Luskin (1986); Walsh, Lippert, Edelson, and Jones (2015). These findings align with the expectations of prosecutors at the SCC District Attorney and form the basis for our capstone project.

In addition to parsing and understanding case timelines, another motivation of this capstone is to determine whether the addition of defendant characteristics can explain delays in resolution, which would indicate the presence of disparities. It is widely perceived that racial and ethnic disparities exist in the criminal justice system, and much research has been conducted on biases at the point of arrest and police interaction Ross (2015). However, no previous work was found on the presence of racial disparities criminal case processing times.

2.3 Previous analytical techniques

Machine learning models can be helpful in decision making in the presence of a large amount of data. To be adopted by policy makers, though, they must be easily interpretable and cost-effective. Previous studies on the topic of time to disposition is dominated by linear regression and basic exploratory analysis. The use of machine learning techniques in the field of criminology is just beginning to emerge. Use of tree-based classifiers to model the outcomes of cases Katz, Ii, and Blackman (2017) and advanced techniques in modeling cost-effective treatment regimes to optimize bail decisions Lakkaraju and Rudin (2016) focus on accuracy of prediction and optimization. The employment of advanced models on case processing time could help inform prosecutors in making decisions that both minimize case length and prioritize fair outcomes.

3 Data

3.1 Data sources

The data used in this project were obtained from the DA's office of SCC, which stores its case information in a case management database called CIBERlaw. We recieved data for all felony cases charged by the SCC DA's office between January 1st and June 30th 2014 for adult defendants. The reason for this specific timeframe is twofold. Firstly, on the 5th of November 2014, Proposition 47 was passed in a referendum in California. With Proposition 47 certain non-violent drug and property crimes in the state were reclassified as misdemeanors instead of felonies. It was at the request of SCC DA's office that the time period selected would be one prior to these changes. Secondly, to maximise the proportion of cases concluded at the time of research it is preferrable to examine a not too recent time period. The data arrived as four separate datasets:

- **Case Information**: Case Information has the base information of each case: case ID numbers, defendant ID numbers, the time a case is logged, and other basic information for each case. It also has demographic information for each defendant: race/ethnicity, gender, age and zipcode of residence. The dataset contains 4,794 observations, with the same number of unique defendant ID's and 4,405 unique case ID's.
- **Defendant Charges**: Defendant charges has information on the charges facing defendants relayed in penal codes. It contains 34,421 observations with 15,668 unique case IDs.
- **Charge Enhancements**: Charge enhancements has limited information on prior convictions of defendants as well as enhancements on current charges. When a charge (from Defendant Charges dataset) is enhanced it mandates harsher sentancing. An examle of this is if a driver is charged for driving under the influence of alcohol. Should the alcohol level of his blood be above a certain threshold, or the driver refuses to have his blood tested for alcohol levels, the original charge for driving under the influence is enhanced. The dataset contains 10,831 observations with 3,640 unique case IDs
- **Case Events**: Case events has information on all court events that are related to a case. Each event is timestamped, and falls into one of 155 distinct case event types and one of 469 distict case event results. To construct a simplified timeline of a case, four key events must be recognized and extracted: arraignment, plea, case disposition, and a case last event. In some cases the event is explicitly stated in the categories, in others it must be inferred. Identifying these key events is key to merging the data and gaining insights about prosecution durations. The dataset contains 481,614 observations with 14,983 unique case IDs.

To protect the privacy of defendants both defendant and case ID's are anonymized from their entries in SCC's database. These ID's are then used to merge the four datasets into a single set containing all the relevant information on each case. When merging the Case Information dataset with Defendant Charges 6 observations are lost, taking the total number of observations form 4,794 to 4,788. Merging Charge Enhancements with the resulting dataset from the previous merge has no affect on the number of observations (as only some

cases will carry enhancements the merge is based solely on case ID's and defendant ID's found in the already merged dataset - 'left' merge). Merging Case Events with the outcome of previous merges 278 observations are lost, taking the number of observations from 4,794 to 4,510.

Misdemeanor data were included in the Defendant Charges and Case Events tables which explains why we find much higher numbers of case ID's in those sets of data than in Case Information. All of the misdemeanors get discarded in the merge process. The reason we lose observations in the merging process stems from the fact that some case and defendant ID pairs found in Case Information are missing in Defendant Charges and Case Events. Without direct access to the CIBERlaw system, we cannot know the causes of these discrepancies.

3.2 Construction of timelines

To understand what causes delays in the prosecutorial process, one must first understand the timeline of a case. From the point of view of a prosecutor, a case generally ends at disposition, or resolution. A disposition usually takes the form of either a dismissal, guilty verdict, acquittal, or guilty plea. In the CIBERlaw system, there is no single event that explicitly logs the disposition of a case. Instead there is a number of case event type and results combinations that can represent disposition (the dictionaries that map event categories and subcategories to our event classification are available on the project gitlab repository). By going through the possible combinations, we identified the disposition event for 90% of our cases. The remaining 10% are missing clear disposition dates. This is most likely because the disposition event was logged in a separate database of Santa Clara County courts, or due to the fact that the case has not been concluded yet.

Time to disposition is defined as time from case issuing to the first event having one of the following results: *formal probation granted, credit time served, summary probation granted, sentenced, prison sentenced imposed, defendant deceased, found guilty, found not guilty, defendant released by court, defendant discharged, deferred entry of judgment PC1000, cases consolidated, charges suspended per civil compromise, motion to dismiss interest justice granted*, or *motion to dismiss case granted*.

We are also interested in looking at three other key events for each case: arraignment, plea and last event. The arraignment is identified as the first event for a case of type Arraignment. Plea is identified as the first case event result of one of '*Plead guilty*', '*Plead not guilty*', '*Not guilty plea entered by court*' or '*Plead nolo contendere*'. A plea of *nolo contendere*, or no contest, is a plea where the defendant neither admits nor disputes charges. While it isn't technically a guilty plea it has the same immediate effect. Last event is the very last event logged to a case. 3% of cases have no identifiable arraignment event and 7% of cases have no identifiable plea event.

From these four different events for each case we construct the timelines. The timelines are calculated at the day precision starting from the day a case is issued; *days-to-arraignment*, *days-to-plea*, *days-to-disposition*, and *days-to-last-event*. Out of the 4,510 observations of the merged dataset we find that days-to-arraignment has a negative value for 76 observations, days-to-plea is negative for 69 observations, days-to-disposition is negative for 71 observation and days-to-last is negative for 29 observations. All in all we have 79 negative observations (three arraignments have missing values). These negatives result from the fact that the issuance of a case happens at a later date than might be expected. One specifica and at random example of this is a certain case where disposition happens in September of 2014 and the last event registered to the case is in December of 2015. However the case is issued in April of 2016 rendering all time values negative.

Some of these negative time-lines can be explained with cases being reopened after sentencing. For example, ten of those are due to Proposition 47. Some, though, cannot be easily explained. These 79 observations have been dropped from the dataset, taking observations from 5,510 to 4,431. The resulting timelines can be seen in Table 2.

3.3 Engineered Features

From the attributes of the original sets of data new features were engineered to retain all relevant information we are interested in examining and encode it in a format that enables visualization and modeling. The variables are encoded as either integers (e.g. number of charges for a defendant/case pair), binary (e.g. whether there was a preliminary hearing or not), categorical (e.g. pleas guilty, not guilty, *nolo contendere*), or continuous interval variables (e.g. defendant's age). The features are:

Information on prior convictions ("priors") of defendants turned out to be incomplete so it was not possible to reliably assess whether a defendant had a prior, and more importantly, a strike prior (one which is applied to California's "three strikes" penalties). This information may be important in predicting prosecutorial duration, and we will hopefully be able to incorporate it in a future study.

4 Exploratory Data Analysis

4.1 Time Duration of Cases

Having extracted the time of arraignment, plea, disposition and the last event of a case, timelines for each case can now be constructed. Statistics of the phases of the prosecutorial process for the 4,431 cases issued in January through June of 2014 can be seen in Table 2.

What is immediately interesting from table 2 is the median value of days to disposition: 141.5 days. This directly contradicts the findings of the report issued by the SCC Civil Grand Jury which states that only 47% of cases in SCC are resolved within a year. Furthermore, according to our findings, 81.5% of cases in SCC were resolved within a year. Figure 1 shows us the distribution of case duration and further emphasizes the point that most cases are resolved early in the process. Again, regarding the Civil Grand Jury report, it must be stated that it is not reproducible so direct comparison can not be made. Further discussion on the Civil Grand Jury report can be found in chapter 2.

Event though the picture we get is not as grim as the one depicted in the Grand Jury report a rate of 81.5% case closure within a year is still below the state average of 88% quoted in that report. Furthermore, knowing what drives delays in the prosecutorial process is generally valuable, independent of location and current case closure statistics.

Time to plea is a factor that will play heavily in case duration due to the fact that plea has to take place before disposition, Figure 2. What this means is that if time to plea is long the same will apply to time to disposition. However the opposite is not necessarily true as can be seen from the plot; when time to disposition is long, time to plea isn't necessarily long as well.

To examine what other case factors might be the key drivers of delay we look at case duration for cases with specific characteristics independently. In Figure 3 we look at the distribution of case duration through multiple violin plots. Case duration is measured in days from when a case gets logged in the CIBERlaw system until it is resolved through sentencing or dismissal. The different colors (blue and green) represent case duration for two different subsets of the dataset. The minimum and maximum value of each distribution reflect the shortest and longest case in the dataset. The distributions are normalized and smoothed via kernel density estimate with a Gaussian kernel. We see that cases where the defendent pleads guilty or no contest to charges initially are generally resolved early in the process while cases where the defendant pleads not guilty don't have as obvious most common case duration. This seems in sync with case durations that end in trial and/or have preliminary hearing. The presense of more than one defendant in a case or more than one charge against a defendant do not give an indication of a significant difference in case duration. These last two observations (enhancements and number of defendants) are of specific interest as they had been clearly identified by the SCC DA as possible key contributors to delays in the prosecutorial process.

Type	Feature	Description	Possible values/range
Categorical	Courtroom	in what courtroom did the disposition take place	various
Integer	Courtroom Count	through how many different courtrooms did the case go	1-17
Categorical	Courtroom Type	what is the type of courtroom where disposition happened	case management court, dor tic violence court, drug co general felony court, n county court, south cou court, unknown
Binary	Custody	was the defendant in custody or not at the start of the case	0,1
Categorical	Defender type	type of defender at disposition	public, private, independent ternate, unknown
Binary	Gang enhancements	are gang enhancements present	0,1
Categorical	Initial Plea	did the defendant plea guilty, not guilty or no contest	guilty, not guilty, no contest
Integer	Ncharges	number of charges a defendant is facing	1-47
Integer	NcourtDates*	the number of court dates for a case	1-76
Integer	Ndefendents	number of defendants per case	1-7, 20
Integer	Nenhancements	number of charge enhancements	1-26
Integer	Nfelonies	number of felony charges for a case	1-44
Integer	Nfta	number of times a defendant failed to appear	1-12
Integer	NHS	number of charges due to violation of the Health & Safety code	1-24
Integer	NPC	number of charges due to violation of the general Penal code	1-32
Integer	NpleaDates*	the number of plea dates in a case	1-30
Integer	NVC	number of charges due to violation of the Vehicle code	1-16
Binary	PC12022	are there other critical enhancements connected to the case (the use of a weapon or presence of injury)	0,1
Binary	PC1368	was the defendant deemed incompetent to stand trial at any point	0,1
Binary	Prelim	was there a preliminary hearing or not	0,1
Categorical	Possible Outcome	what was the inferred sentence outcome	prison, probation/jail, unkno
Binary	Public Defender	was the defendant represented by a public defender at any point	0,1
Integer	Time to Plea*	the number of days between when case got created until the defendant's initial plea	0-1222
Binary	Time waived	was there time waived at any point	0,1
Binary	Trial	did the case go to trial or not	0,1
Binary	More than a year	is the time to disposition less than or greater than one year	0,1
Integer	Time to Arraignment*	the number of days between when case got created until the defendant was arraigned	0-1093
Integer	Time to Disposition*	the number of days between when case got created until disposition	0-1181
Integer	Time to Last*	the number of days between when case got created until the last event registered to a case	0-1232

Table 1: Engineered features. *Time to disposition $\geq$ 1 year* is the feature on which the classification, is based. See Section 6.

days to	min	25%	median	75%	max	mean
Arraignment	0	1	5	31	1093	30.5
First plea	0	36	90	180	1222	137.4
Disposition	0	63	141.5	281.8	1181	210.8
Last Court Event	0	178	378	684.5	1232*	455

Table 2: Statistics on the duration of the prosecutorial process in four phases from the day the case was issued for the 4,431 cases issued between January and June of 2014 by SCC with *complete* information (i.e. missing data were removed by row). (*) The last event is the latest event logged, but we have no information to indicate whether future court events are possible or expected.

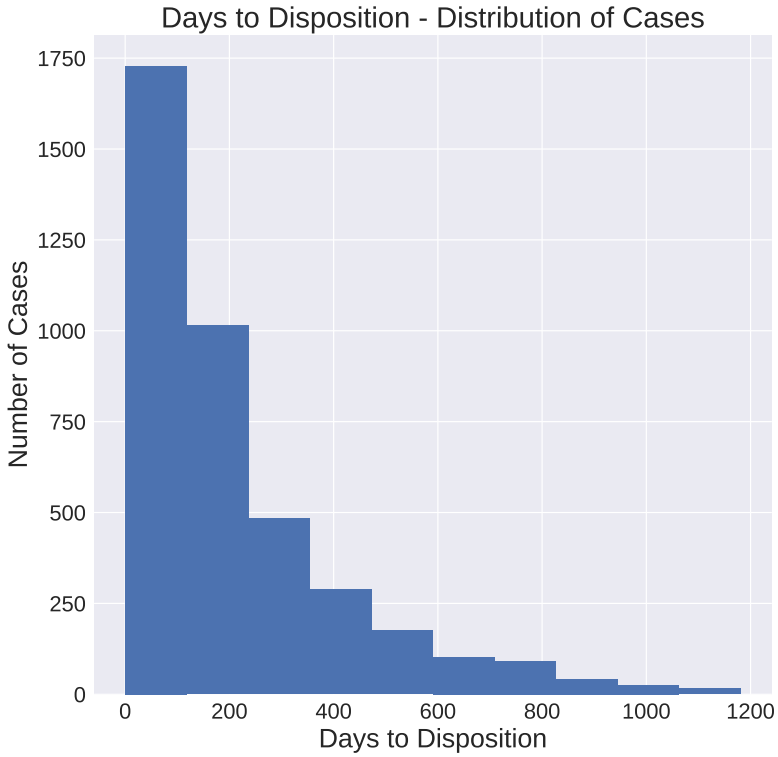

Figure 1: Histogram of the days it takes a case to reach disposition starting from the day the case is issued. 71 cases have negative time to disposition due to the case being re-issued in the course of the prosecutorial process. These cases have been dropped.

We can extend this examination to include other key events of a case. In Figure 4, we see time to arraignment, plea, disposition and the last event for defendants initially in custody against defendants initially not in custody. We see that both arraignment and plea most commonly happens very early in the process for those defendants initially in custody. Based on data from January through June of 2014 the median time to disposition for defendants in custody was 89 days. For those out of custody it was 192.5 days.

In Figure 5 we see the same breakdown for defendants that have at some point during a case been found to be not competent to stand trial plotted against all other defendants. The most common time of arraignment,

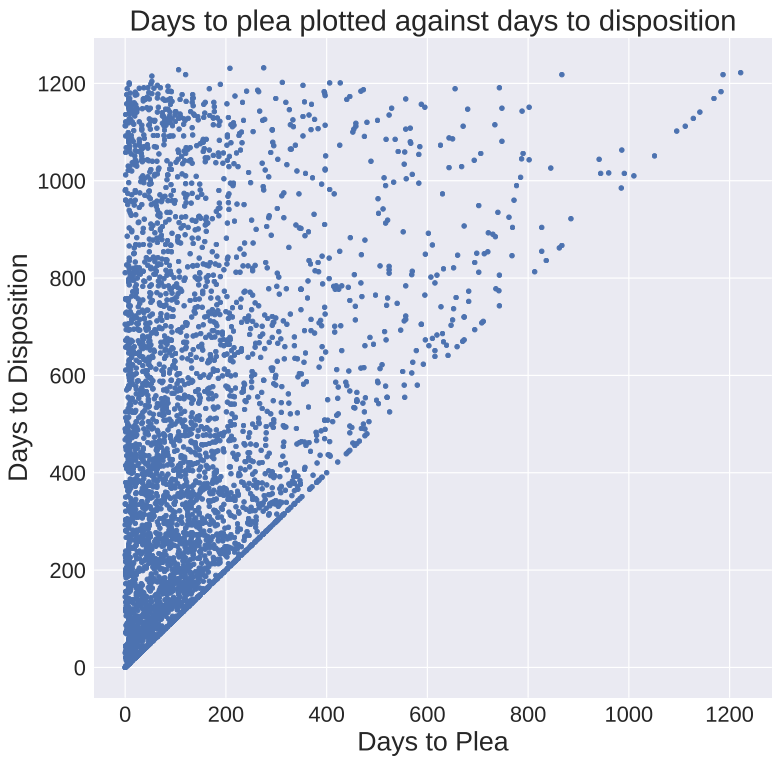

Figure 2: Days to plea plotted against days to disposition for felony cases issued by SCC in January-June 2014. No disposition can happen before the plea, hence the bottom right portion of the plot is empty. The SCC DA indicated the long duration of the prosecutorial process up to plea, which is uncharacteristically long due to peculiarities of the laws that in SCC do not require a defendant to enter a plea early in the case, would drive the long duration of the prosecutorial process to disposition. However, in this plot we see a large fraction of defendant-case pairs at the top left of the plot, with short time to plea, and yet long time-to-disposition, indicating that delays in entering a plea are only partially responsible for delays in the prosecutorial process up to disposition.

plea and disposition for these defendants doesn't seem to be as stark as for the rest of the defendants. Most commonly the last event of these cases happens late in the process, after 1000 days. Based on data from January through June of 2014 the median time to disposition for defendants that are at some point not competent to stand trial was 353 days. For other defendants (excluding aforementioned group) it was 135 days.

Even though it takes more than twice longer to reach disposition for defendants that have at some point been found to be not competent to stand trial, this or any of the other engineered features will not explain delays in the prosecutorial process on their own. The case of being incompetent to stand trial is an exception, applicable to 2% of the defendants in the dataset. In the last section of this paper, we construct models to identify the most prominent drivers of prosecturorial delay.

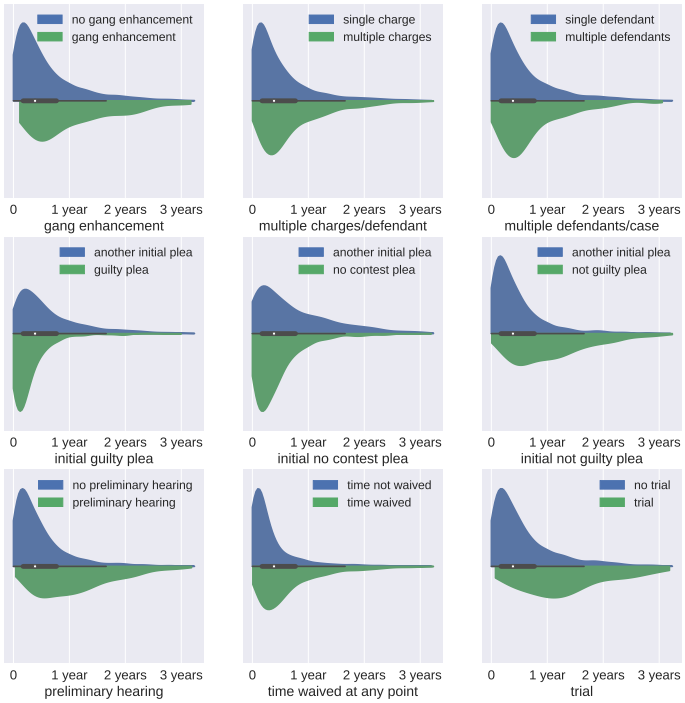

Figure 3: Distributions of duration of the full prosecutorial process, from case issuing to disposition, for felony cases issued between January and June 2014 by the SCC DA. In each plot a distribution is shown as a histogram smoothed with a kernel density estimate for two samples (blue and green) split along the vertical axis for comparison: a so-called *violin* plot. Each violin plot shows the time-to-plea distribution for two subsets of our data. The horizontal bar indicates the IQR (thick bar), full statistical distribution without outliers (thin bar) and median (white dot) for the *top* distribution. We compare time-to-disposition for defendants (from the top left) going *vs* not going to trial, charged of crimes with *vs* without a gang enhancement, which plead guilty *vs* not guilty or *nolo contendere*, *nolo contendere vs* guilty or not guilty, charged with one *vs* more than one charge, who waived *vs* did not waive time (Table 1), who had *vs* did not have a preliminary hearing, charged as a single defendant *vs* with others (often occurring in gang related charges), and that pleas guilty *vs* not guilty or *nolo contendere*

4.2 Demographics

Having information on age and race/ethnicity allows us to explore the demographics of the data. In Figure 6 we see how the defendants's race/ethnicity breakdown compares to that of the population of SCC. There are some disparities between the two with some ethnicities over- or under- represented in the data. The defendants' age decreases steadily (Figure 7) and the majority of defendants are male (Figure 8). Greatest number of defendants come from the zipcodes around San Jose, as well as zipcodes 95037 and 95020 to the south of San Jose, see Figure 9

Using the demographics of the data and the timeline of cases gives us an access to case duration for different demographics. In Figure 10 we see case duration by race/ethnicity. Duration is measured in days between

Figure 4: Duration of the prosecutorial process to disposition of cases for defendants initially in custody (blue) compared to defendants initially not in custody (green). No significant differences are observed. Details of the graphics are as in Figure 3.

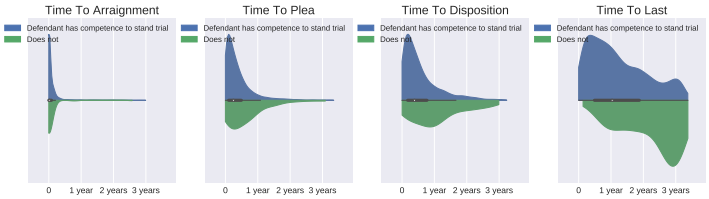

Figure 5: Duration of the prosecutorial process to disposition of cases for defendants initially in custody (blue) compared to defendants initially not in custody (green). Significant differences are observed, especially in the time-to-sentence, the distribution of which peaks later and has more power in the tail, and in the post-sentence duration, with an accumulation of defendant continuing to have court dates scheduled years after the beginning of the case. While these events occur after sentence and do not affect the primary metric we are testing (time-to-disposition and particularly when time-to-disposition extends past a year) it may affect the efficiency of the courts and cause delays in other cases. Details of the graphics are as in Figure 3.

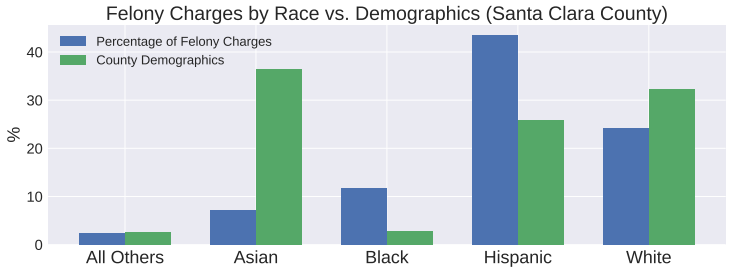

Figure 6: Ethnic breakdown of defendants in SCC felony cases issued between January and June of 2014 (blue) compared to the ethnic breakdown of the population of the county of Santa Clara (green).

the creation of a case until its resolution. From this figure, we can conclude that there is no statistically significant difference in the duration of the process for different races/ethnicities in our case dataset.

In Figure 11 we have isolated the most commonly found charge in the data, violation of Health and Safety Code 11377(a) which is the possession of methamphetamine. The figure shows case duration for different races/ethnicities on the same charge. Again, we conclude that there is no statistically significant difference in the duration of the process for different races/ethnicities.

Lastly, in Figure 12 we have isolated the second most common charge found in the dataset, theft of property (PC 459-460(b). Again there is no statistical difference between the case duration for different ethnicities. However, there are only 312 observations of this kind so statistically robust conclusions cannot be drawn

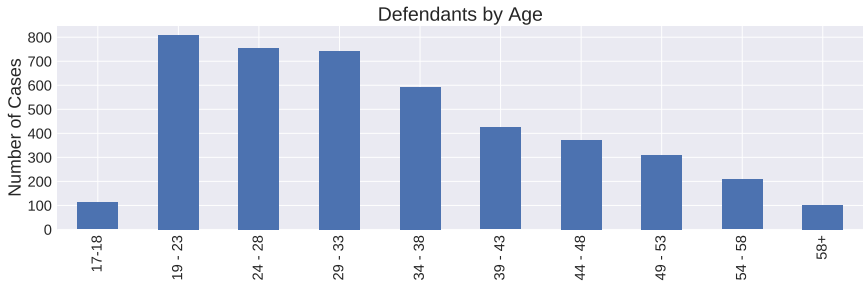

Figure 7: Age of defendants at crime commission in SCC felony cases issued between January and June of 2014 by 5 year age bins. Notice that the data only include defendants tried as adults.

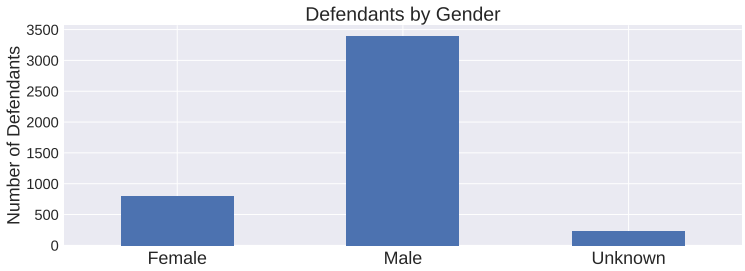

Figure 8: Gender of defendants in SCC felony cases issued between January and June of 2014. 77% of the defendants identified as males, 18% as females, and the gender is unknown for 5%.

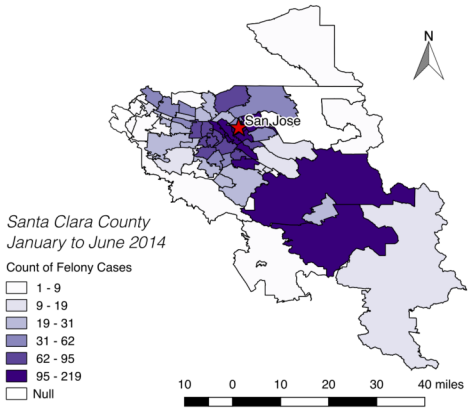

Figure 9: Number of defendants by postal code of residence. Most commonly defendants live in the areas in and around San Jose. Postal codes 95037 and 95020 to the south of San Jose are also prominently featured.

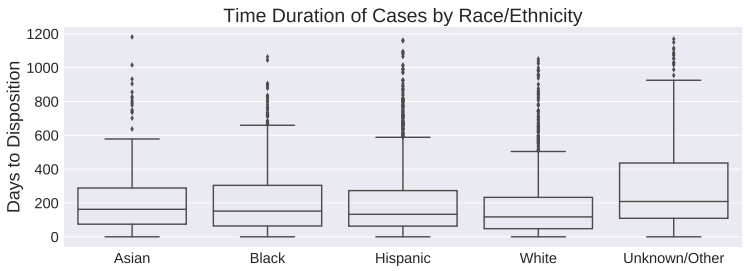

Figure 10: Case duration, measured as days between case issue and disposition, by different race/ethnicitiy. For each ethnic group, the horizontal line within the box represents the median case duration. The box represents the interquartile range (IQR), the "whiskers" represent the full distribution, excluding statistical outliers, which are shown as individual data points. No statistically robust differences appear, as all the medians fall in the 25–75 percentiles of all other groups. Curiously, the distribution for Unknown/Other (missing and uncommon ethnic groups) is only marginally consistent with most of the other distributions. We speculate this may be due to cases issued against defendants that are not in custody and not reachable/fleeing from custody, and wish to test this in the future

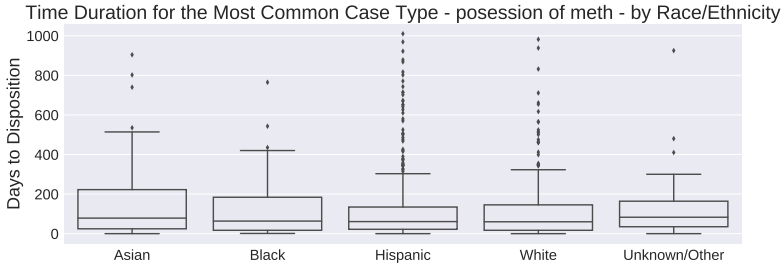

Figure 11: As in Figure 10: prosecutorial process duration by ethnicity/race for the most common charge issued by the SCC's DA in January through June of 2014, HS 11377(a) – possession of methamphetamine – to correct for compounding biases in crime by race (e.g. frequency of crime by ethic group). As for the full charges sample, the distribution of prosecutorial duration is consistent for all ethnic groups.

from this subdivision of the dataset.

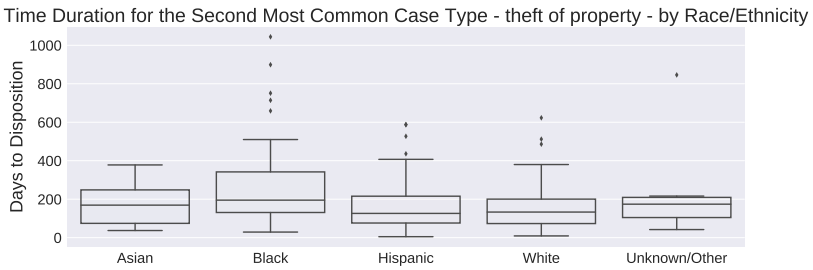

Figure 12: As in Figure 10: prosecutorial process duration by ethnicity/race for the second most common charge issued by the SCC's DA in January through June of 2014, PC 459-460(b) – theft of property. As for the full charges sample, the distribution of prosecutorial duration is consistent for all ethnic groups.

5 Visual analysis tools

5.1 Visual tool to enable data exploration

While the analysis above is informative, it is generated from a typical data science approach: finding, comprehending, merging and sorting data and applying statistical plots and other filters to identify trends in the data. These are not tasks that are suited for a DA's office, which has other important legal tasks to perform. Therefore, it is desirable to automate many of these tasks and provide a way for these prosecutors to interactively engage with their data so that they can identify trends without advanced data skills.

We generated concepts for the visualization using synthetic data sets. These data sets were constructed with a small set of features of various types that we expected would be of interest to the attorneys. This includes the durations of four phases of prosecution, variables for race and gender, and a value for the age of the accused. Although these are important variables to consider in our visualization and modeling activities, we chose these for development purposes so that we could determine the best ways to handle arbitrary variables we may want to display. In particular, we have been able to prototype the ability to filter our data based on binary, categorical, and continuous variables.

The simplest form of this visualization is a stacked horizontal bar plot (Figure 10). Each bar represents a category of comparison that is selected by the user, e.g., individual ethnicities or age ranges. Visual comparisons are made via three information channels for each bar: its location on the x-axis, its width, and its color.

The location of the bar encodes the time for a given phase to commence relative to the start of some other chosen phase. Location attributes are most easily compared by a user when they are placed on the same scale Munzner (2014); Wilkinson (2005). Therefore, we provide the ability to choose which phase to compare against and align the x-axis (time) such that the phase begins at time $t = 0$, and earlier phases are displayed on the negative portion of the scale. For overall case duration comparison, we align to the beginning of the first phase, where the start of each case is displayed at $t = 0$.

The width of the bar encodes the duration of each phase. These values are calculated as the difference of the times from the beginning of each case to the ends of two consecutive phases. Since these times are determined by our own categorization scheme for the case events, the phase durations will be subject to some error depending on how well we can identify the demarcations between the phases in the data and how well the data is entered into the DA's case management system.

The color of the bars encode which of the four phases is being represented. We use four colors drawn widely and uniformly from the viridis color palette van der Walt and Smith (2015). The colormap was developed for the Matplotlib python graphics package and is now its default color palette as of version 2.0. Viridis has two desirable properties: it is perceptually uniform (meaning that the scale is uniformly smooth and does not induce a perception of structure) and robust to common forms of colorblindness. These colors are easily distinguishable.

An additional channel of information is available when hovering the mouse pointer over any aggregated bar, showing a one-dimensional horizontal scatterplot of the underlying data along a time axis. Also displayed is a boxplot of the distribution, as well as the elementary statistics of minimum, maximum and median.

A prototype of this dashboard using synthetic data is available at http://bit.ly/2hbPqrL.

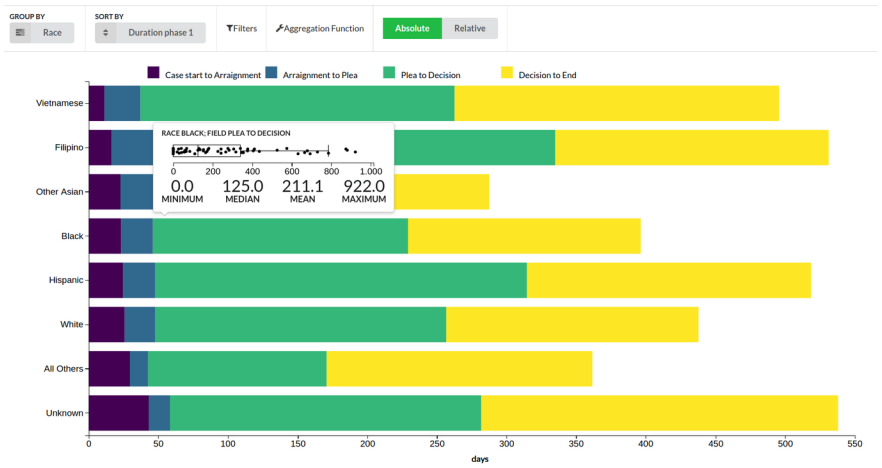

Figure 13: Screenshot of our visualization tool designed to enable exploration of SCC prosecutorial data running in the Chrome web browser. The visualization tool breaks down the prosecutorial process into four phases: case issue-to-arraignment, arraignment-to-plea, plea-to-disposition, disposition-to-last logged event, and enables aggregation, filtering, and sorting on other axes: demographic, court related categories, etc. Here the visualization is using synthetic data, binned and aggregated on age ranges and sorted by the duration of the second phase ("arraignment to plea"). Note that the x-axis (days) is aligned such that the second phase starts at $t = 0$ and the first phase is shown extending in to the negative portion of the domain. Also shown is an example of the distribution information that is displayed when the user hovers over a bar using a pointing device: minimum, maximum, and a box plot showing the entire distribution for that prosecutorial phase (arraignment-to-plea) and the category belonging to that bar (defendant between 21 and 25 years of age).

6 Analytical Models

6.1 Decision Tree Models

We chose to use decision tree-based models to build a classifier that predicts whether a given case would be disposed within a year or not. Decision trees are considered one of the most versatile machine learning methods James, Witten, Hastie, and Tibshirani (2013).

Binary classifiers take the values of the input features x_i and output $y_i \in \{0, 1\}$. Decision trees do this by partitioning the feature space into subspaces, such that the divisions give rise to final regions with learned classifications. The subspaces divided at each step of the construction of the tree represent nodes of the tree, and the final set of subspaces after the desired number of partitions are created are the tree's terminal leaves.

These partitions can be complex, with many splits, leading to many nodes with high accuracy (the so-called "purity" of the leaves, determined by various measures). These trees are "strong" learners, but generally exhibit poor performance on unseen data in high dimensions since they are overfit on training data. Conversely, the partitions can be simple, with few splits (possibly even only one) having nodes with lower purity. These trees are called "weak" learners, but have the advantage of being simple and not overfitting the training data.

Robust against outliers and data transformations, decision trees are fast and their results are interpretable. In isolation, decision trees can perform well and have low bias, but they tend to exhibit high variance as errors in the first node quickly propogate through the children nodes of the tree when applied to data unseen by the model James et al. (2013). In order to reduce this variance, ensemble methods are frequently employed. We attempt to improve the performance of our models using two such techniques: Random Forest and Gradient

Boosted Decision Trees. Both models are implemented in Python using packages scikit-learn Pedregosa et al. (2011) and xgboost Chen and Guestrin (2016).

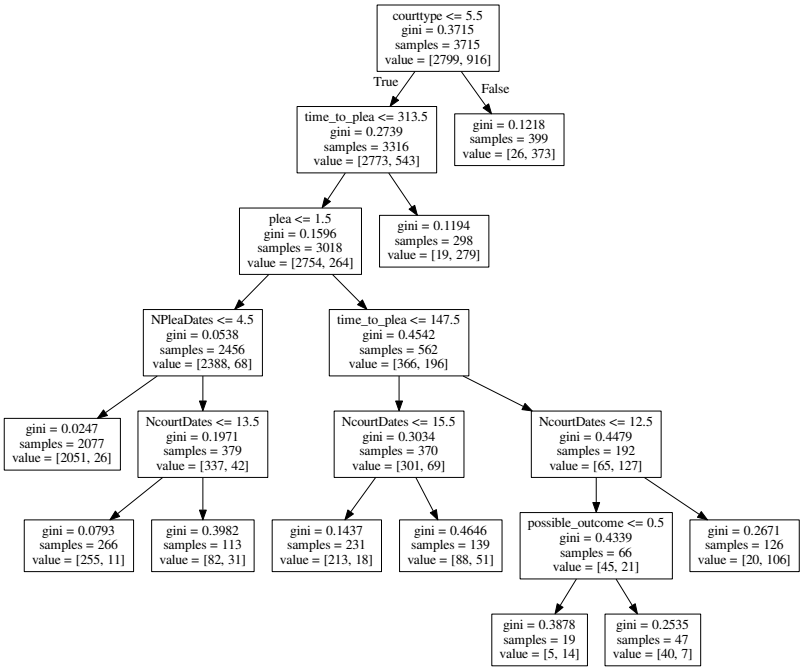

Figure 14: A single decision tree using the training data set, separating the data into the two classes "disposition less than one year" and "disposition greater than one year". At each decision node of the graph, the tree splits on the variable indicated, represented with two arrows drawn below it. The node indicates the boolean test on which the split is performed, the Gini coefficient (representing the purity of the node with respect to the final classification scheme), the number of samples on which the test is performed, and the size of each of the two true classifications. The data is split with the data for which the boolean test is True going to the left child node, and the data for which the boolean test is False going to the right child node. The performance of this tree would be evaluated by measuring how well it classifies a labeled test data set, using the final classifications in the terminal nodes (the "leaves" of the tree), which are assigned to the class having the larger number of observations in the population.

6.2 Leaf purity and feature importances

The Gini impurity is calculated as the sum of the products of the population ratio and the classification error rate over each of N classes,

$$I_{Gini} = \sum_{i=1}^{N} p_i e_i$$

with p_i being the population ratio and e_i being the misclassification rate, both for class i. In the case for $N = 2$, this can be simplified: for a leaf having a members in class 1 and b members in class 2, the impurity can be calculated as

$$I_{Gini} = \frac{2ab}{(a+b)^2}$$

For example, in the tree above, for the rightmost leaf on the bottom level that contains 40 classified in the first class and 7 in second class, the coefficient is calculated as $\frac{2\times 40\times 7}{(40+7)^2} = 0.2535$.

Leaf impurity measures can be used to determine which features of a decision tree model have the most importance to determining the final classifications. The Gini variable importance measure for a variable X_m in a random forest of N trees is given by Louppe, Wehenkel, Sutera, and Geurts (2013)

$$Imp_{Gini}(X_m) = \frac{1}{N}\sum_{T}\sum_{t\in T: v(s_t)=X_m} p(t)(I_{Gini}(t) - p_L I_{Gini}(t_L) - p_R I_{Gini}(t_R))$$

where the summations are over all nodes t in trees T having X_m as the splitting variable, $p(t)$ is the proportion of observations in the forest that are evaluated at node t, and p_L and p_R are the proportions of the population split to the left and right children nodes t_L and t_R, respectively. We use this measure for variable importance throughout the rest of this paper.

In our analysis, the actual performance of the classification is less important than determining the variables that influence the classification. We evaluate the receiver operator characteristic (ROC) plots to validate that the models have some predictive power, but once that is established, the Gini variable importance measures are our primary interest. Weaknesses of this measure include a bias towards (higher reported values for) continuous variables and away (lower reported values for) variables with a small number of categories. It is also possible for a combination of lower importance variables to be jointly predictive, which would not be detected in a simple evaluation of importance rankings Epifanio (2017).

6.3 Treatment of categorical variables

Categorical variables cannot be split at a tree node in a natural way, as a numerical or boolean variable can be. Two techniques are commonly used to transform categorical variables into other types: "one-hot encoding", which produces multiple boolean variables, one for each category; and a simple numerical mapping that assigns integers $i \in \{0, 1, ..., n-1\}$ to each of n classes, such that each category gets a distinct integer label.

There are several weaknesses introduced with this method. For one-hot encoding, the observations within a single category become sparse which might undermine that category's importance. Also, one-hot encoded features of the original feature are dependent on each other. For the second method, by casting categories into integers we are imposing an order relationship to features that may not possess a natural sense of "greater than" or "less than". To counter this, the classification scheme can be permuted to determine if the order changes the outcomes.

To test how the choice of encoding scheme affects the resulting classification, we test a random forest using both methods and compare the resulting top feature importances. Here, one-hot encoding extends the feature space from 26 to 248 covariates. The second method performs the numerical cast on each of the categorical variables as described above, keeping the same number of covariates before and after the transformation. The results using these two schema are shown in Tables 3 and 4. We note that the order of the feature importances is similar between the two runs of the model, after observing that the features are themselves split in the one-hot encoded method. Because both methods (one hot encoding and casting categories into integers) give similar results, we take this to be a indicator of robustness with respect to classification choice, and in the remaining modeling we use only numerical classification.

6.4 Random Forests

Random Forests are an ensemble learning method based on decision trees. The prediction of the Random Forest classifier is determined by majority voting across multiple trees fit on subsamples of the data and subsamples of the features.

We ran the random forest model using four different sets of input variables. For each we optimize the hyperparameters using a grid search routine from scikit-learn.

In our first iteration, we use all of the engineered features. Using hyperparameter grid-searching, we fit 50 trees with a minimum of 10 samples at each leaf node, each tree having a minimum of five features, using a Gini impurity criteria to measure leaf purity.

We then recalibrate and run the model with the demographic features removed. Our motivation for modeling with and without demographic features was to detect disparities. If the predictive power increased with the inclusion of demographic variables, that would be a strong indication that these variables are influential on the model. We chose fit a Random Forest classifier of 50 trees with a minimum of 2 samples at leaf node, again with the Gini criteria. Each of these trees considered a maximum of 20% of the feature space.

Our third iteration of Random Forests is a classifier without timeline related variables. Having timeline variables as features in the trees may be problematic because of their correlation with the target variable. Moreover, we hope to predict the length of cases with information exogenous to the case proceedings; keeping timeline related variables in the classifier is helpful for pointing out where delays may be happening during a case progression, but we would also like to identify which features of our classifier become important when run without this retrospective information.

The fourth model removes both the timeline related variables and the demographic variables, again comparing the performance of each to identify any impact the demographic variables have on the resulting classification.

6.5 Gradient Boosted Decision Trees

Whereas the Random Forest is an ensemble learning method that operates on many decision trees in parallel, the technique known as Gradient Boosted trees is an ensemble method that operates on trees in serial, recursive fashion.

Boosted models are constructed by adding many weak learners into a single model,

$$f_M(x) = \sum_{i=1}^{M} T(x; \Theta_i)$$

where M is the number of learners. Generally, T could be any type of learner, but in a gradient boosted tree $T(x; \Theta_i)$ is the i-th tree of the model defined on the input variables x and whose parameters Θ_i define the structure of the tree. The i+1-th weak learner is generated iteratively by fitting the tree on the residual errors from the model of the first i summed trees. In practice, this is a difficult problem to solve analytically, so numerical methods are substituted to estimate the next optimal tree. In the gradient boosted tree technique, gradient descent is used to find the local minimum of the loss function with respect to the current model. As with other gradient descent learning models, the rate of descent is an additional hyperparameter to tune. The number of trees M may be chosen *a priori* or be allowed to increase until a desired performance is achieved.

Similarly to the random forest models, we run the models four times using the same variable sets as identified above, using the same grid search algorithm to optimize the hyperparameters of the model.

6.6 Results

We found that time to plea is the most important feature for both RF and GBT, regardless of the inclusion of demographic variables (see subplots (a), (b), (c), and (d) in Figure 15). In addition, the number of court dates and type of initial plea appear in the top five of both models, with and without demographic variables. The importance of time-related features is not surprising since they are correlated with time to disposition.

With the addition of demographic variables, we found no change in the predictive power of our models. Although age at offense and zip code appear in the top 10 important features of the GBT, the models are still dominated by the first few time-related features.

After removing time-related features, we find agreement between RF and GBT that possible sentence outcome is among the most important features. The main charge and type of initial plea are the primary importance features for the GBT and RF models, respectively (panels (e) through (h)), however they only appear as at best the fifth-most important feature in the complementary model. It is not clear why the models would predict such different results, although it is possible that there is some covariance between these two features. It is also important to note the relative importances, especially in the GBT model whose importances feature much less variation.

Surprisingly, whether a case went to trial or had a preliminary hearing was not important in determining case duration, even though they appear to have significant differences in the distribution splits in Figure 3 (g) and (i). However, only 100 cases went to trial and only 848 held a preliminary hearing, therefore they may not be greatly important in a classifier. Furthermore, preliminary hearing and trial may correlate with other features, such as court type and initial plea. This correlation may also weaken their importance.

We find that time to plea is the most important feature on predicting disposition time within a year. This corresponds with our expectation based on Figure 2, because a case can not be concluded until a plea has taken place. Splitting the cases on that feature at 365 days will yield a perfectly pure node in the tree.

The presence of a preliminary hearing, the number of court dates, and the type of court where case disposition took place were also important features. The relatively high power of prediction for these two features also aligns with our expectations. The presence of a preliminary hearing in a case indicates that the case has moved beyond early case stages and entered trial preparation stages. A waiver or absence of preliminary hearing would indicate that the case moves directly on to trial or the defendant pleads guilty, both of which shorten the length of the case. We also expect the number of court dates to be positively correlated with the whole length of a case. Court types are associated with specific types of criminal cases such as drug or domestic violence crimes, or they are specific to a different geographic area of the county. Thus, they may be related to the main charge, which is an important feature in the GBT model.

The number of plea dates is also an expected strong indicator of disposition time. Similar to a relatively early time to plea, a low number of plea dates may indicate a defendant has plead guilty early in the process.

Enhancements on a case were very weak features in determining case disposition within a year.

Full details of the feature importances for all eight runs of our models are given in Table 3 in the Appendix.

The feature with the third highest importance is the type of court where case disposition took place. Court types are associated with specific types of criminal cases such as drug or domestic violence crimes, or they are specific to a different geographic area of the county.

Because models determine a binary class, they can be measured using a metric called Receiver Operating Characteristic (ROC), which is based on the probability estimates of the positive class at different thresholds. In Figures 16 and 17, the dashed diagonal lines represent randomly guessing the class. The farther left the model curve falls from the diagonal, the more accurate the model is at predicting the postive class. Our random forest and gradient boosted tree model perform similarly for each set of input variables. Removing the timeline-related data degrades the performance of the model significantly, pulling the curves closer to the 45° line. From this we infer that the timeline-related variables have importance generating predictions of

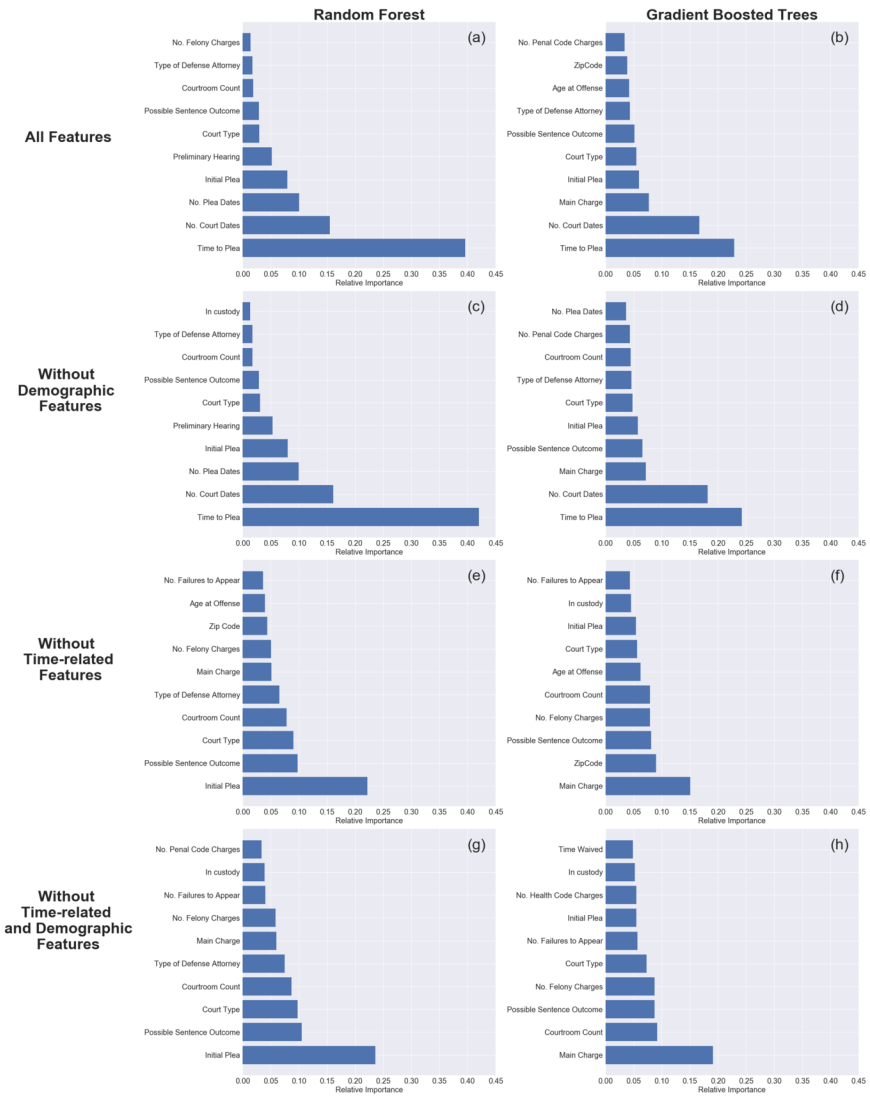

Figure 15: Top ten feature importances for each of the models run for random forest (panels on left, (a), (c), (e) and (g)) and gradient boosted trees (panels on right, (b), (d), (f) and (h)).

long-duration dispositions. By the same analysis, the fact that there is little to distinguish the performance of the models when demographic variables are removed from consideration in either set of models indicates that demographics does not have importance in generating these predictions. This is consistent with the exploratory analysis performed above.

7 Conclusions and future work

Understanding why some criminal cases take a long time to resolve is a complex but important task. We attempted to shed light on this issue by creating a dashboard for exploration of case timelines and by modeling to determine important features of long cases. After the task of parsing and preprocessing case events and characteristics was completed, we were able to begin exploration and modeling.

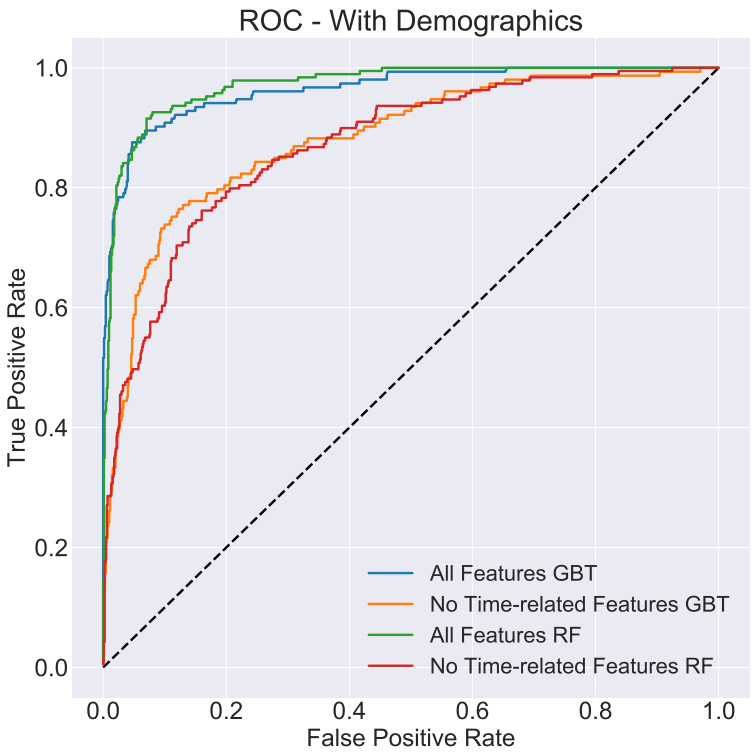

Figure 16: ROC curves for the RF and GBT models, with the addition of Demographic Data. The ROC plots show True Positive (completeness) *vs* False Positive rate (purity) of a classifier. The ROCs change slightly depending on the subset of variables in the model. The two ensemble methods exhibit similar performance. The addition of the Time-Related features improve the accuracy of the predictions and therefore we infer that they are important covariates for the prediction of long duration dispositions.

In this paper, we described the design and construction of a case timeline visualization tool. This tool eases exploratory analysis by representing timelines has horizontal bar charts that can be grouped, filtered, and sorted according to a user's choices.

We also constructed Random Forest and Gradient Boosted decision tree models to isolate important features in determining whether a case is resolved within a year or not. We found that the time to the first plea event is one of the most important features in determining case length. We also found that the type of initial plea and the type of court room for disposition were important features. Because the type of initial plea is highly important in determining duration, the SCC DA could investigate their plea bargaining strategies and change initial plea offers to influence overall case durations.

There is much future potential work that could be done with this dataset. Augmenting this data with external datasets such as sentencing outcomes, bail amounts, concurrent individual court room, attorney case loads, and arrest and incarceration rates in SCC may yield more robust models. A comparison with case data from similar jurisdictions in California would also be valuable.

Our project found no evidence to support the statement from the Santa Clara County Civil Grand Jury report that fewer than half of felony cases are resolved within a year. The exact data and process used by

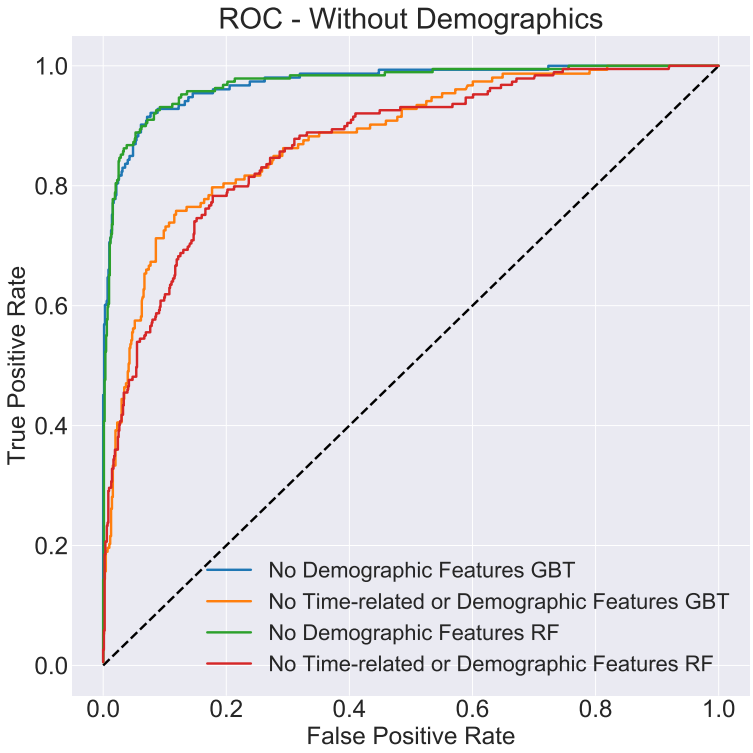

Figure 17: ROC curves for the RF and GBT models, without the addition of Demographic Data. A change in the shape of the plots between this and Fig. 16 (including demographics) would indicate a change in the predictive power; however we do not observe that here. This suggests that there is no evidence of significant disparities with the inclusion of the demographic data.

the Civil Grand Jury and the California Judicial Council to evaluate case duration is not reproducible, but with our 2014 felony dataset and process we found that approximately 74% of our 2014 felony dataset was resolved within a year. While this is still short of the 88% for the rest of California, the problem of felony case durations may not be as dire as claimed by the Civil Grand Jury.

8 Appendix

References

Association, A. B. (2006). *ABA Standards for Criminal Justice*. Washington, D.C..

Chen, T., & Guestrin, C. (2016). XGBoost. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining - KDD '16*. ACM Press. Retrieved from https://doi.org/10.1145%2F2939672.2939785 doi: 10.1145/2939672.2939785

Epifanio, I. (2017, may). Intervention in prediction measure: a new approach to assessing variable importance for random forests. *BMC Bioinformatics*, *18*(1). Retrieved from https://doi.org/10.1186%2Fs12859-017-1650-8 doi: 10.1186/s12859-017-1650-8

	Random Forest				Gradient Boost	
	With time-dependant vars		W/o time-dependant vars		With time-dependant vars	W
	w/ Demogr.	w/o Demog.	w/ Demogr.	w/o Demog.	w/ Demogr.	w/o Demog.	w/
Age at Offense	0.011	–	0.04	–	0.042	–	
Court Type	0.03	0.031	0.09	0.098	0.054	0.048	
Courtroom Count	0.018	0.018	0.078	0.087	0.032	0.045	
Enhancement PC 12022	0.003	0.003	0.012	0.015	0	0	
Enhancement PC 1368	0.001	0	0.004	0.004	0.002	0.002	
Gang Enhancement	0	0	0.002	0.002	0.002	0.002	
Gender	0.002	–	0.006	–	0	–	
In custody	0.013	0.013	0.036	0.039	0.018	0.021	
Initial Plea	0.08	0.08	0.222	0.236	0.059	0.057	
Main Charge	0.013	0.013	0.051	0.06	0.077	0.072	
No. Charges	0.006	0.006	0.021	0.026	0.024	0.032	
No. Court Dates	0.155	0.161	–	–	0.167	0.182	
No. Defendants	0.003	0.003	0.011	0.014	0.005	0.006	
No. Enhancements	0.007	0.007	0.028	0.032	0.005	0.016	
No. Failures to Appear	0.006	0.006	0.036	0.04	0.019	0.014	
No. Felony Charges	0.014	0.012	0.051	0.059	0.022	0.021	
No. Health Code Charges	0.006	0.006	0.017	0.022	0.006	0.011	
No. Penal Code Charges	0.008	0.008	0.029	0.034	0.034	0.043	
No. Plea Dates	0.1	0.1	–	–	0.024	0.037	
No. Vehicle Code Charges	0.001	0.001	0.004	0.005	0.003	0.005	
Possible Sentence Outcome	0.029	0.029	0.097	0.105	0.051	0.065	
Preliminary Hearing	0.052	0.053	–	–	0.006	0.011	
Public Defender	0.002	0.002	0.005	0.006	0.014	0.013	
Race/Ethnicity	0.005	–	0.017	–	0.014	–	
Time Waived	0.006	0.006	0.021	0.025	0.008	0.01	
Time to Plea	0.396	0.42	–	–	0.229	0.242	
Trial	0.004	0.004	0.013	0.016	0	0	
Type of Defense Attorney	0.017	0.017	0.066	0.074	0.043	0.046	
Zip Code	0.013	–	0.043	–	0.038	–	

Table 3: Feature importances found for each of the eight runs of models. "–" indicates the variable was not used in the model.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2013). *An Introduction to Statistical Learning*. Springer New York. Retrieved from https://doi.org/10.1007%2F978-1-4614-7138-7 doi: 10.1007/978-1-4614-7138-7

Judicial Council of California. (2016). *Statewide Caseload Trends, 2005-2006 through 2014-2015* (Tech. Rep.). Retrieved 2017-07-03, from http://www.courts.ca.gov/documents/2016-Court-Statistics-Report.pdf

Katz, D. M., Ii, M. J. B., & Blackman, J. (2017, apr). A general approach for predicting the behavior of the Supreme Court of the United States. *PLOS ONE*, *12*(4), e0174698. Retrieved 2017-07-19, from http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0174698 doi: 10.1371/journal.pone.0174698

Klemm, M. F. (1986, jan). A look at case processing time in five cities. *Journal of Criminal Justice*, *14*(1), 9–23. Retrieved from https://doi.org/10.1016%2F0047-2352%2886%2990023-1 doi: 10.1016/0047-2352(86)90023-1

Lakkaraju, H., & Rudin, C. (2016, nov). Learning Cost-Effective and Interpretable Regimes for Treatment Recommendation. In *arXiv:1611.07663 [stat]*. Barcelona, Spain. Retrieved from http://arxiv.org/abs/1611.07663 (arXiv: 1611.07663)

Louppe, G., Wehenkel, L., Sutera, A., & Geurts, P. (2013). Understanding variable importances in forests of randomized trees. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in neural information processing systems 26* (pp. 431–439). Curran Associates, Inc. Retrieved from http://papers.nips.cc/paper/4928-understanding-variable-importances-in-forests-of-randomized-trees.pdf

Luskin, M. L., & Luskin, R. C. (1986). Why So Fast, Why So Slow: Explaining Case Processing Time. *Journal of Criminal Law and Criminology*, *77*(1). Retrieved 2017-07-03, from http://scholarlycommons.law.northwestern.edu/cgi/viewcontent.cgi?article=6512&context=jclc

Measures for Justice. (2017). *Measures for Justice.* Retrieved 2017-07-03, from https://www.measuresforjustice.org/

Munzner, T. (2014). *Visualization analysis and design.* CRC press.

Neubauer, D. W. (1983). Improving the Analysis and Predentation of Data on Case Processing Time. *The Journal of Criminal Law and Criminology*, *74*(4). Retrieved 2017-07-03, from http://scholarlycommons.law.northwestern.edu/cgi/viewcontent.cgi?article=6410&context=jclc

Nugent-Barakove, M. E., Budzilowicz, L. M., & Rainville, G. (2007, apr). *Performance Measures for Prosecutors: Findings from the Application of Performance Measures in Two Prosecutors' Offices* (Tech. Rep.). National District Attorneys Association, American Prosecutors Research Institute. Retrieved 2017-07-03, from http://www.ndaa.org/pdf/performance_measures_findings_07.pdf

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.

Ross, C. T. (2015, nov). A Multi-Level Bayesian Analysis of Racial Bias in Police Shootings at the County-Level in the United States 2011–2014. *PLOS ONE*, *10*(11), e0141854. Retrieved from https://doi.org/10.1371%2Fjournal.pone.0141854 doi: 10.1371/journal.pone.0141854

Santa Clara County Civil Grand Jury. (2017, jun). *Justice Delayed: Why Does It Take So Long To Resolve Felonies In Santa Clara County?* (Tech. Rep.). County of Santa Clara Superior Court of California. Retrieved 2017-07-14, from http://www.scscourt.org/court_divisions/civil/cgj/2017/Why_Does_It_Take_So_Long.pdf

van der Walt, S., & Smith, N. (2015). *Matplotlib colormaps.* Retrieved 2017-07-03, from http://bids.github.io/colormap/

Walsh, W. A., Lippert, T., Edelson, M. G., & Jones, L. M. (2015, jul). Length of Time to Resolve Criminal Charges of Child Sexual Abuse: A Three-County Case Study. *Behavioral Sciences & the Law*, *33*(4), 528–545. Retrieved from https://doi.org/10.1002%2Fbsl.2187 doi: 10.1002/bsl.2187

Wilkinson, L. (2005). *The Grammar of Graphics*. Springer-Verlag. Retrieved from https://doi.org/10.1007%2F0-387-28695-0 doi: 10.1007/0-387-28695-0