

Automate or Assist? The Role of Computational Models in Identifying Gendered Discourse in US Capital Trial Transcripts

**Andrea W Wen-Yi, Kathryn Adamson, Nathalie Greenfield, Rachel Goldberg,
Sandra Babcock, David Mimno, Allison Koenecke**

Cornell University

{andreaswenyi@infosci., ka468, nmg76, rtg67, slb348, mimno, koenecke}@cornell.edu

Abstract

The language used by US courtroom actors in criminal trials has long been studied for biases. However, systematic studies for bias in high-stakes court trials have been difficult, due to the nuanced nature of bias and the legal expertise required. Large language models offer the possibility to automate annotation. But validating the computational approach requires both an understanding of how automated methods fit in existing annotation workflows and what they really offer. We present a case study of adding a computational model to a complex and high-stakes problem: identifying gender-biased language in US capital trials for women defendants. Our team of experienced death-penalty lawyers and NLP technologists pursue a three-phase study: first annotating manually, then training and evaluating computational models, and finally comparing expert annotations to model predictions. Unlike many typical NLP tasks, annotating for gender bias in months-long capital trials is complicated, with many individual judgment calls. Contrary to standard arguments for automation that are based on efficiency and scalability, legal experts find the computational models most useful in providing opportunities to reflect on their own bias in annotation and to build consensus on annotation rules. This experience suggests that seeking to replace experts with computational models for complex annotation is both unrealistic and undesirable. Rather, computational models offer valuable opportunities to assist the legal experts in annotation-based studies.

1 Introduction

Researchers studying legal language have long suggested that gender bias plays a significant role in high-stakes court cases. But building evidence has been challenging because of the subtle, contingent nature of biases and the vast quantity of legal transcripts. New language technologies offer the potential to identify and measure biased language at scale. In this work, we present a case study involving a collaboration between technologists and legal scholars. Our work highlights both the promise and challenges of using powerful new language models to annotate complex concepts. In particular, we focus on the interaction between AI systems and human experts, which leads to both a more developed understanding of gender bias in court transcripts and a more effective computational tool for legal researchers.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This paper studies trial transcripts for women defendants facing capital charges in US courts. The study involves three phases. In the first phase, the legal experts identify and annotate the following four themes of gender-biased language (i.e., language that invokes gender stereotypes) regarding the defendant. The themes are developed upon review of the trial transcripts through Critical Discourse Analysis (CDA):

- describing the defendant's inappropriate expression of emotions, such as joy or lack of grief
- hypersexualizing the defendant
- describing the defendant as evil, deceitful or manipulative
- characterizing the defendant as a bad mother

In the second phase, the NLP experts develop and evaluate computational models to identify the four themes in trial transcripts, following — for this phase — the standard NLP practice of assuming that expert annotations are a “gold-standard.”

In the third phase, we seek a synthesis between expert and model annotations. Three legal experts, who have been closely involved with capital trials, re-read passages with large disagreements between expert and model annotations. They discuss whether to relabel those passages and provide reasons for their decisions. Through participating in the decision-making processes of the legal experts, the full team gathers insights about the desired role of computational tools in the eyes of legal experts in this setting.

The annotation process is time-consuming and difficult — identifying stereotyped language in capital trials requires knowledge about capital trial process, legal practicality and terminology, defendant demographics, tone of speech, as well as deep familiarity of each case. The complexity of the task requires frequent discussions among legal experts to negotiate definitions of themes. This is a task that requires frequent consensus building and cannot be done individually. In standard NLP protocols, labels are intended to be clear and unambiguous. Identifying and defining biased language is, however, a nuanced task that requires a large degree of subjectivity and interpretation. But to carry rhetorical weight, evidence for gender bias in trials must be collected and assessed with as much care to avoid personal bias as possible.

In this work, we ask the question: what is the role of computational models, including large language models

(LLMs), in this highly laborious, complex, and nuanced annotation process? While using a LLM to replicate expert annotations may seem intuitive, we find that completely replacing human judgment is both unrealistic, due to the frequent inconsistencies in annotations, and undesirable, as experts have no interest in ceding ultimate responsibility. Rather, experts want both efficiency gains and the opportunity to be challenged to refine and strengthen their argument. Our concern, therefore, is to study the benefits offered by a finetuned language model in *assisting* the annotation process, as well as its limitations.

The paper is organized as follows: Section 2 outlines the related works in the empirical and theoretical studies of gender and sentencing, and computational approaches in studying gender stereotypes. Section 3 describes transcript collection. Section 4 presents the method and results of the first phase: qualitative annotation. Section 5 discusses the second phase: quantitative modeling. Section 6 discusses the third phase, where legal experts and technologists revisit passages where model disagrees with annotations from first phase. We discuss our findings in Section 7.

Through this process, we identify the challenges of using LLMs for the purpose of quantifying gender bias in large, complicated corpora — we note that there are many nuances that go into replicating the experienced annotations. Furthermore, we derive insights about the benefits of employing imperfect computational models in the process of annotating complex concepts — the legal experts report that they are not looking for models that perfectly replace the need for them to annotate; rather, they look for new perspective that could challenge their own definitions and remind them of their potential blind spots developed over hundreds of hours of annotation. Lastly, we provide design suggestions for the direction of computational model improvements that fit with the interaction patterns desired by the legal experts.

2 Related Work

Gender and Sentencing. Gender disparity in sentencing outcomes is widely studied in the United States. The vast majority of empirical works suggest that women received more favorable sentencing decisions than men in comparable crimes (Crew 1991; Embry and Phillip M. Lyons 2012; Doerner and Demuth 2014; Daly and Bordt 1995; Farnworth and Raymond H. C. Teske 1995). Scholars have also suggested that such disparities in sentencing vary across crimes; for example, the gender effect is inconclusive in more violent crimes (Rodriguez, Curry, and Lee 2006; Nagel and Hagan 1983; Koons-Witt et al. 2014).

Legal and criminal scholars have developed two major theories for the observed gender disparity in sentencing decisions — the *chivalry theory* and *evil woman theory*, both of which are based on the patriarchal point of view on binary gender roles. The *chivalry theory* suggests that women receive preferential treatment in criminal trials because they are stereotyped as a weaker sex, and hence need to be protected by men from suffering (Visher 1983; Rodriguez, Curry, and Lee 2006). On the other hand, the *evil woman theory* posits that the preferential treatment does not extend to women who violate traditional gender roles, and that women

could be punished more harshly than male offenders when committing violent crimes or crimes that are perceived as masculine (Rapaport 1990).

Narrative Analysis in Legal Decision Making. While there are many studies on the relationships between offender gender and sentencing decisions, the majority of these studies are statistical in nature (Streib 1990; Koons-Witt et al. 2014; Doerner and Demuth 2014; Rodriguez, Curry, and Lee 2006). Researchers have looked into narratives and language used in legal decision-making such as asylum claims (Shiff 2021), parole hearings (Greene and Dalke 2021) and rehabilitation treatments (Wyse 2013). However, these studies on language are mostly derived from close readings of a small number of cases (Greene and Dalke 2021; Mogul 2005; Sutton et al. 2022; Dalke and Greene 2023). Studies on courtroom language in capital trials of women are also rare and on only a handful of women (Mogul 2005; Sutton et al. 2022). In this study, we explore the role of computational models in expanding the valuable annotations for systematic investigations of narratives in the courtrooms.

Gendered Discourse in US Capital Trials. Nationwide, one of every fifteen women in prison is serving a sentence of life, virtual life, or death (Nellis 2021). The use of biased language in criminal trials has been identified as a concern, especially in race-based stereotyping (Baldus, Pulaski, and Woodworth 1983; of Appeals for the Fourth Circuit 2016). Murder convictions and death sentences have been overturned due to the use of racial stereotyping language in criminal trials, highlighting the importance of identifying stereotyping narratives in capital trials (Levenson 2022; Liptak 2024a). Gendered discourse in US capital trials has been less studied and only in the cases of a handful of women on death row (Rapaport 1991; Atwell 2007). Yet the stakes are high — a death-sentenced defendant is actively seeking appeal in the Supreme Court due to courtroom languages based on gender stereotypes (Liptak 2024b). Most previous linguistic and legal work on the gendered features of courtroom speech is in non-capital criminal trials (Potts and Weare 2018), civil trials (Conley, O’Barr, and Riner 2019), or cases where the woman was the victim rather than the defendant (Ehrlich 2019, 2021; Potts and Formato 2021). Most research relating to women and capital punishment has focused on the types of crimes for which women have been sentenced to death (Baker 2015; Streib 2002, 2005; Carroll 1996; Rapaport 1991) rather than courtroom language. In this study, we seek to analyze the gendered courtroom language deployed in capital trials of women.

Computational Approaches in Detecting Gender Stereotypes. Detecting language with sexism and gendered stereotypes using Natural Language Processing (NLP) techniques has drawn a large degree of scholarly interest in recent years (Basile et al. 2019), with the majority of applications in online content such as social media (Mozaafari, Farahbakhsh, and Crespi 2020; Anzovino, Fersini, and Rosso 2018; Chiril, Benamara, and Moriceau 2021; Chiril et al. 2020; Vásquez et al. 2022) and news articles (Devinney, Björklund, and Björklund 2020). Unlike trial tran-

scripts, these online texts are very short, usually no more than a few hundred characters, and are meant to be consumed by the general public rather than trained specialists.

Automated analysis of legal language using NLP has also attracted widespread attention in recent years (Bommarito, Katz, and Detterman 2018; Chalkidis et al. 2022), with a particular interest in predicting legal outcomes (Medvedeva and McBride 2023; Vaudaux et al. 2023; Barale et al. 2023; Chalkidis, Androutsopoulos, and Aletras 2019). Scholars have identified that the extremely long sequence length of legal documents, and the legal expertise required to understand the complex terminologies, make automatic processing of legal documents a unique challenge for even the state-of-the-art prompting strategies and LLMs such as GPT-3 and Llama (Thalken et al. 2023; Hakimi Parizi et al. 2023). Our work takes as given these LLM weaknesses, and focuses on how computational models can still benefit the legal community.

3 Data Description and Collection

As of January 2024, there are 47 cis-gender women on US death row. We obtained trial transcripts for all of them. Building this corpus required personal networking with defense attorneys: while transcripts are technically public, acquiring them can cost thousands of dollars. Most of the transcripts are in PDF format. We preprocess court transcripts with the following steps: first, we use Adobe optical character recognition (OCR) to transform PDF transcripts into text files. Second, we lower case and remove all digits in the text files. Lastly, we use Python package NLTK (Bird, Klein, and Loper 2009) to split transcripts into sentences. We only keep sentences with at least three words. While we have 47 trial transcripts, we focus on 8 transcripts with the highest OCR quality in this study.

In the United States, there are two phases in death penalty trials: a guilt phase, where the jury determines whether the defendant is guilty of the crime charged, and a penalty phase, where the jury (or, on rare occasions, a judge or panel of judges) decides what sentence is appropriate. At each of these two stages of the trial, the prosecution and defense can make opening and closing statements to the jury, present witnesses, and cross-examine the witnesses who testify on behalf of the other side. We study the entirety of both stages for the eight trials.

While the absolute number of trials we have may seem small, the total size of the collection is massive. The eight transcripts, documenting months of conversations in court, are each 1,126–3,415 pages long. After preprocessing, the lengths of the eight transcripts range from 178,790 words (roughly the length of Frank Herbert’s *Dune*) to 609,774 words, with an average of 351,000 words per transcript.

4 Qualitative Annotation

In this first phase, legal experts identify gender-biased themes following standard procedures from critical discourse analysis (CDA) (Blommaert and Bulcaen 2000). Broadly speaking, CDA recognizes that discourse both creates and reflects cultural norms. To discern the ways that dis-

course creates social reality, analysis should “examin[e] the linguistic practices through which it is translated into social action” (Conley, O’Barr, and Riner 2019).

Potentially gendered themes of interest are initially developed by a close reading of the prosecution’s opening statement through the guilt phase by two legal practitioners who have been closely involved with capital trials. These initial themes formed the annotation training for a group of graduate assistants (GA) with legal backgrounds. Then, the GA team annotates 8 transcripts in their entirety (including opening statements, witness testimony, and both the prosecution and defense attorney’s closing arguments in both phases of trial). The team uses the qualitative data analysis software, MaxQDA, for annotation.

Each transcript is annotated by one GA with a legal background. While there is no inter-rater reliability to report, the annotation team met weekly to discuss and resolve questions that arose over the annotation process.

The resulting four themes that we explore in this paper are described in Table 1. The annotated sentences are rare in relation to the entire transcript. Furthermore, each sentence could belong to multiple themes.

Qualitative annotation is extremely time-intensive and expensive. In total, it took the GAs more than 300 hours to annotate 8 transcripts, not including the time for training and weekly discussion. Hence, it is infeasible to extend the annotations to all women’s transcripts by manual annotation exclusively. In addition to the investment it required, manual annotations have other challenges. For example, we observe instances that were not coded but should be (see examples in Sec 6). There are also inconsistencies in the unit of annotations — while some GAs assign the theme to sentences, others assign it to an entire conversation, within which not every sentence is meaningful (for example, the short responses such as “yes” and “no” from the witnesses in cross-examinations).

These observations present opportunities and challenges for computational modeling. Hence, our concern is not building a model to replace the manual annotations. In fact, this is not possible because the input data are contestable from a quantitative point of view. Rather, we focus on using a computational model to assist the annotation process, resulting in more consistent annotations.

5 Quantitative Modeling

In the second phase of the study, we develop and evaluate computational methods to annotate gender-biased language. This phase follows standard methodology from NLP, treating human annotations as “gold standard” data for training and evaluation, despite the underlying data concerns outlined above.¹

Since each sentence can belong to more than one theme, we treat each theme as a separate binary classification task. For each theme, we finetune a `LEGAL-BERT` classifier (Chalkidis et al. 2020). Our choice of model is informed by the current literature in the space of legal NLP. Scholars have

¹We release code and appendix at: <https://github.com/andreawwenyi/automate-or-assist>

Theme Name	Code	Descriptions	# sentences	mean % sentences
Emotions	EMOT	Describing the defendant’s emotions, including emotional, emotionless, inappropriate joy or happiness, remorseful, or remorseless.	255	0.136
Hypersexualization	SEX	Describing the defendant as promiscuous, loose, an adulterer; attacking the defendant’s sexual expression, or demonizing their sexual practices.	339	0.174
Betrayal of Gender Norms	NORM	Describing the defendant as manipulative, greedy, evil, deceitful.	793	0.337
Bad Mother	MOM	Describing the defendant as a bad mother.	148	0.064

Table 1: Four themes explored in this paper. Mean % sentences is the average % of annotated sentences across the eight transcripts. On average, each theme represents fewer than 1% of sentences in a trial transcript.

experimented with comprehensive prompting strategies with state-of-the-art LLMs on different kinds of legal tasks. Recent works have found that prompted LLMs perform poorly on various kinds of legal tasks (Hakimi Parizi et al. 2023) and that in-domain `LEGAL-BERT`, albeit smaller, outperforms other larger, newer LLMs on a complex annotation task in the legal space (Thalken et al. 2023).

Including Context in Input and Output The legal experts seldom only rely on a single sentence to determine if there is a gendered strategy at play. The annotation is often judged in context. For example, the question “*What was the defendant’s demeanor like at that time?*” could be related to *Emotions* or not based on the witness’s answer or the other questions that come before or after (see Table 2). Therefore, we’d like the model to learn from context when making predictions, and also be flexible in the output length, adjusting based on the content.

As the language model’s context length is limited, it is infeasible to feed the entire transcript to a model. We thus have to break the transcripts into manageable chunks of text. We use a sliding window approach with window length of 10 and a step of 1. Specifically, an entire transcript is broken down to paragraphs that each contain 10 consecutive sentences and differ from the next paragraph by 1 sentence. This way, each sentence is included in 10 distinct paragraphs.

We predict whether *at least one sentence in each paragraph belongs to a theme of interest* using binary classification. We calculate the score for each sentence by averaging the scores assigned to the 10 paragraphs containing that sentence.

Finetuning and Cross-Validation Because negative labels are much more frequent than positive labels, we under-select negative paragraphs to be three times the number of positive paragraphs. We apply leave-one-out cross-validation to the 8 annotated court transcripts. Specifically, we finetune a model on 7 transcripts, then we apply the trained model on the held-out transcript. This strategy allows us to evaluate how well the model generalizes to unseen transcripts. We finetune `LEGAL-BERT` hosted on Huggingface with a 2×10^{-5} learning rate and 0.01 weight decay.

We also experimented with zero-shot prompting

`FLAN-T5-Large` and `GPT-3.5`. See details in the appendix.

Coreference Resolution Identifying a theme is necessary but not sufficient: trials have many participants, and thematically relevant language does not always refer to the *defendant* (rather, it can refer to another person). For example, the sentence “*You could hear the fear in her voice.*” falls under the theme of “Emotions.” However, it is not necessarily portraying the defendant—it could be describing a witness. To filter to instances where target is not the defendant, we additionally utilize coreference resolution.

Coreference resolution is a popular NLP task that identifies mentions in texts that refer to the same entity. We run the LingMess model from Python package *fastcoref* (Otmazgin, Cattan, and Goldberg 2022) on a 20-sentence passage—the target sentence and its 19 preceding sentences. For each entity mentioned in the target sentence, we examine their coreference clusters. We consider the target sentence to be about the defendant if there is one cluster with direct mentions of the defendant, or if there is one cluster that only contains she/her pronouns.

Quantitative Model Output After we apply the trained `LEGAL-BERT` on unseen transcripts, we obtain a score for each sentence. We group together consecutive sentences that score above 0.5 as one output passage. This helps us to be flexible in output length that presents meaningful chunks of information. Then, we use coreference resolution to identify whether the passages refer to the defendant. Specifically, we run coreference resolution on sentences scored above 0.9, and we only keep passages if they contain at least one sentence scored above 0.9 that mentioned the defendant.

Results The quantitative model performance is not perfect; however, this is expected given the varying quality of annotations and the vastly different contexts across trials.

One metric we use to evaluate the model performance is passage-level precision, calculated as the proportion of passages with at least one sentence annotated positively by experts among passages predicted positive by the model. The mean passage-level precision is 0.151 for *Emotions*, 0.126 for *Hypersexualizations*, 0.272 for *Betrayal of Gender*

Paragraph	EMOT Annotation	Reason
A Prosecutor: What was her demeanor like at that time? Witness: I don't recall.	Negative	A neutral question that does not establish a description of the defendant's emotion.
B Prosecutor: What was her demeanor like at that time? Witness: She looked calm.	Positive	An exchange that describes the defendant's emotion.
C Prosecutor: Did she appear to be grieving? Witness: Possibly. Prosecutor: What was her demeanor like at that time? Witness: I could not see very clearly.	Borderline / Positive	The "grieving" question is setting up for an answer regarding the defendant's emotions. This makes the demeanor question more than a neutral question but leads to a specific answer about the defendant's emotions.

Table 2: The same question “What was her demeanor like at that time?” could result in an annotation of “Emotions” or not, depending on the conversation the question is situated in.

Norms, and 0.063 for *Bad Mother*. On average, the top three passages with the highest scores from each transcript have a precision of 0.417 for *Emotions*, 0.333 for *Hypersexualization*, 0.458 for *Betrayal of Gender Norms* and 0.083 for *Bad mother*. The other evaluation metric we use is sentence-level recall: the proportion of sentences annotated positively by models among sentences annotated positively by humans. The mean sentence-level recall is 0.484 for *Emotions*, 0.341 for *Hypersexualization*, 0.287 for *Betrayal of Gender Norms* and 0.246 for *Bad mother*. Due to the inconsistencies of annotations, passage-level recall is ill-defined and sentence-level precision is a less-than-reliable measure. See details in the appendix.

6 Expert Evaluations on Annotator-Model Disagreements

Neither the CDA nor the NLP phases are sufficient on their own in producing large-scale annotations. Manual annotations are too costly to scale to the full 47 trials, and the influence of individual perspectives raise persistent concerns. In the third phase, we synthesize results from the previous two phases. Legal experts find that the contrast between human and machine annotations was useful and prompted good discussion.

For each theme, three lawyers, who have experience representing defendants in capital trials, examine about 30 passages where the model predictions differ from the graduate assistant (GA) annotations. Specifically, the first author selects the three highest-scored passages (meaning that the model finds them to be highly relevant) that the GA annotated as irrelevant from each trial, along with the 6–8 lowest-scored passages that the GA annotated as relevant. For each passage, three coauthors who are lawyers jointly decide whether at least one sentence within the passage is relevant to the theme. They select from “Positive (✓),” “Negative (✗),” or “Undecided (?)” (if they are unable to reach a consensus). They also provide reasons for their decisions. Before the lawyers decide on a passage, they are not told the GA or model’s annotation of the passage—they only know that the GA and model annotations disagree. This is to avoid

potential biased preferences towards the GA or model’s decision. The first author moderated and documented the discussion processes.

In meetings across three days, the lawyers reviewed passages from the four themes in the following order: *Emotion*, *Hypersexualizations*, *Betrayal of Gender Norms*, and *Bad Mother*. The process is carried out using an annotation tool called Prodigy.

Results We analyze the model-lawyer agreement across four themes. We define model-lawyer agreement in each theme as the number of passages where the lawyers agree with the model predictions, divided by the total number of passages that the lawyers read. Additionally, we analyze the time taken for the lawyers to review passages for each theme. We use *False Positives* (FP) to refer to passages annotated negatively by GA (from the first phase) but positively by the model (from the second phase), and *False Negatives* (FN) to refer to passages annotated positively by GA but negatively by the model. The main metrics are shown in Table 3.

Emotions Model-lawyer agreement is the highest in *Emotions*, at 57.14%. Among FP passages, the lawyers agree with 72.7% of model decisions. This shows that the model has a good understanding of this theme, and that employing the computational model improves the quality of annotations. The lawyers have the least disagreements with each other in *Emotions*, with only 1 *undecided* passage. Words of emotions, such as *calm*, *mad*, *agitated*, *upset*, *remorse* are often cited as justifications for decisions, per Table 5.

However, challenges for computational models arise when one sentence contains references to multiple people. For example,

[The defendant] said that [the victim] looked really mad, really upset.

Additionally, when the passage describes the *lack of* emotions, the model often makes false negative predictions. For example,

Q. Do you recall what was said?

A. [The defendant] was talking about how she had a

Theme	# minutes	# passages (# FP, # FN)	# sentences	# positive, # negative, # undecided in FP	# positive, # negative, # undecided in FN
EMOT	40	28 (22, 6)	359	16 , 5, 1	6, 0 , 0
SEX	65	30 (24, 6)	378	0 , 21, 3	3, 3 , 0
NORM	35	32 (24, 8)	402	13 , 7, 4	6, 1 , 1
MOM	42	26 (18, 8)	479	2 , 12, 4	6, 1 , 1

Table 3: Quantitative results of the lawyers’ evaluations on passages where GA annotator and language model disagree. Numbers where the lawyers agree with the model decision are bolded. FP (False Positives) refers to passages annotated negatively by the GA but positively by the model; FN (False Negatives) refers to passages annotated positively by the GA but negatively by the model.

Reasons for Lawyers’ Disagreement with Model Decisions	EMOT	SEX	NORM	MOM
Unrelated to theme	1	13	1	7
Related to theme but not describing the defendant.	3	3	3	4
Neutral information gathering or factual statements of cases	1	5	2	4
Need longer context	1	3 (+5*)	4	0
Defense’ counter argument to the theme	0	0	1	1

Table 4: We categorize why lawyers disagree with model decisions into five categories. Usually, they highlight one main reason, so the categories are mutually exclusive. In the SEX theme, there are five cases where *longer context* is cited as a secondary reasons, we marked them with a star (*) for emphasis.

lot of food at the house that had been brought in by friends and family members and that she should have brought that to the sheriff’s office because she was sure we were hungry and would want something to eat.

Q. Okay. Anything else you remember she said?

A. There was really nothing else, I was so taken back by that, that after being arrested for the capital murder of her daughter, that she was worried about food.

For more quotes and decisions, see Table 8 in the appendix.

Hypersexualization Model-lawyer agreement is lowest in *Hypersexualization*, with lawyers agreeing with only 10% of model decisions. The lawyers also spent the longest time discussing passages of *Hypersexualization*. Compared to the other three themes, it took the lawyers 1.5–2 times longer per sentence.

Hypersexualization presents challenges for both the lawyers and the model because of the extensive context required (Table 4). The lawyers frequently mention that they need to know information such as the relationships between people in the case and the significance of events being discussed. For example, while reading the following exchange between the prosecutor and a witness, the lawyers mentioned they need to know the relationship between the witness and the defendant, and whether *the brother’s wife* referred to the defendant.

Q: How long were you having sex with your brother’s wife?

A: From about five years.

Q: So you had sex with your brother’s wife for five years; how many times?

A: Not so much.

Q: What does, “not so much” mean; 2, 20, 100?

The ambiguity of passages led to prolonged discussion time — the lawyers spent considerable time proposing hypotheses about why a particular topic was discussed in trials.

The computational model incorrectly picks up passages describing relationships but not necessarily related to sexualization. For example,

[The defendant] had a job for a period of time as a greeter at [a supermarket]. This was the end of May, into the beginning of July. So about a little more than a month. So that is where she met [her second husband]. She took up with him, and they were married in December.

Additionally, the model also picks up passages that the legal experts deem as neutral information gathering or description of crime facts that are not necessarily a manifestation of gender bias. For example,

The defendant returned to the residence [and] discovered her ex-husband’s body just inside the main door leading from the garage. The defendant called 911, and feigned hysteria. The defendant, in her letters to [her lover], had discussed how she would fake grief upon discovering that her ex-husband had been killed.

For more quotes and decisions, see Table 10 in the appendix.

Betrayal of Gender Norms Overall model-lawyer agreement is at 43.75%, showing good but slightly lower alignment than *Emotions*. Among the FP passages, the lawyers agree with the model’s positive predictions 54.17% of the time, showing that the model is helpful in picking up passages that the GA missed annotating in the first phase.

Similar to *Emotions*, direct descriptions of defendants using words such as *greedy, manipulative, evil, calculated* are

often cited by the lawyers as annotation justifications of relevance. For example,

[The defendant], the woman who took every opportunity to line her pockets, a heartless schemer who manipulated and lured men to their peril. She deceived her husband, children, family and friends. She has earned the title premeditated murderer, queen of greed and evil.

Meanwhile, courtroom actors are likely to describe events or use other indirect descriptions to set up an image of the defendant, resulting in disagreements between lawyers and the need for additional context. For example, while the following exchange between the prosecutor and a witness does not seem relevant in isolation, some lawyers suggest that it could be setting up an image of the defendant being performative.

Q. Now, was there anything characteristic about [the defendant] that would catch your attention?

A. Whenever she would drive her car, she would always have her window down and she would always be smoking her cigarette. And just the mannerism that she would hold her cigarette to the side.

Q. How did she do that?

A. Just by flicking her cigarette.

Q. Any particular style that you refer to it as?

A. Well, as I said before, the kids told me it's like a Hollywood style.

For more quotes and decisions, see Table 11 in the appendix.

Bad Mother Model-lawyer agreement of *Bad Mother* is 11.5%, similarly low as *Hypersexualization*. The lawyers struggled to reach consensus with each other: 19.2% of passages are *undecided*, the highest among the four themes. The passages identified positively by the computational model sparked discussions among the lawyers about decoupling factual statements of crime and value judgements towards the defendants, specifically in trials where the victim was the defendant's child. To illustrate, we present two passages that were annotated as positive by the GA annotator but as negative by the model. For the first passage, some lawyers considered it to be *factual* in phase three, leading to be classified as "undecided":

This woman had been trying to get [a man] to kill her daughter for at least months. It got to the point that he went to this child's father and said [the defendant] won't leave me alone about killing [her daughter].

For the second passage, lawyers all agreed that it involves a value judgement that a mother should care about her daughter more than herself, leading to a "positive" decision in phase three:

There has never been a truer statement, by this mom (the defendant). That's the way she felt: [My daughter] is ruining my life. She is ruining my life. Because this was about [the defendant].

Not about [the defendant's daughter].

This is about [the defendant].

Some lawyers suggest that it is fair to discuss crime facts in trials. Therefore, while it is intrinsic to the fact that a woman killing her own children is a "bad mother", such discussion in trial should not be identified as a manifestation of gender bias. Reading several similar passages, the lawyers went back and forth negotiating differences between factual statements and value judgements. For more quotes and decisions, see Table 12 in the appendix.

7 Discussion

Why Language Models Struggle To Replicate Annotations by Legal Experts Identifying gender-biased language in trial transcripts is a hard task even for legal experts. Even engaging in joint discussions, there are still multiple passages where experts could not reach consensus. We find that it is a hard task for language models because legal expertise is required extensively in the annotation process. We share three instances where computational models lack the necessary legal expertise.

First, the model often makes false positive annotations on passages that the lawyers deem as neutral information gathering or factual statements of crime. Distinguishing the differences between factual statements and value judgements is often not an easy task, even for legal experts. Second, discerning subtle tonal differences is another instance that requires legal experience. For example, the lawyers mention the sentence "*she now goes and has another baby*" implies value judgement, but "*She has another baby*" does not. Similarly, "*So you didn't see what her demeanor was at that point in time.*" is a neutral question, compared to the prompting version: "*Did you notice if she was particularly emotional at the time?*" Third, extensive background knowledge of individual cases is often necessary to comprehend who and what is being discussed, and to determine the importance of the discussion. This is especially pronounced when the theme involves discussions of relationships, such as in *Hypersexualizations*.

Benefits and Challenges of Employing Computational Models in a Complex Legal Annotation Pipeline There are multiple benefits in employing a computational model in the annotation process. First, it accelerates the annotation process. Trial transcripts are long but the majority of the transcripts are irrelevant. Therefore, using a computational model in the process to filter out the irrelevant parts of the transcripts greatly improves the efficiency of annotation and enables having multiple legal experts read the same passages for inter-rater reliability. Second, it increases the accuracy of annotation by identifying passages that annotators might overlook. It also helps the legal experts to adjust annotation instructions. For example, upon reviewing the *Emotions* passages, the legal experts found that the GA annotators frequently missed passages describing "lack of remorse." Hence, they decided to add a new subcategory, "remorseless", under the *Emotions* code to assist future annotators. Third, the false positive predictions that the model makes provide opportunities for the legal experts to refine annotation rules. For example, upon reviewing *Bad Mother*, they identified the need to clarify the difference between fac-

Theme	Passage	Decisions	Lawyers' reasons
EMOT	It is beyond comprehension to me that this woman could do what she did, and I submit to you that if you watched her during the course of this trial she has shown no remorse.	GA = ✗ Model = ✓ Lawyers = ✓	Describing that the defendant has not shown any remorse.
EMOT	Q. First of all, Mr. [witness], let's talk about the defendant's attitude and emotional state and what was going on when she came back to your house after [the victim] had left with her. She is accusing [the victim] of taking her money. Is she a little bit more agitated at this point? [...] She tore your bedroom apart, didn't she? A. Yes, she did. [...] Q. So, this calm demeanor that [the defendant] usually had was not present? A. It wasn't there.	GA = ✗ Model = ✓ Lawyers = ✓	Portraying the defendant as agitated, not calm.
SEX	Q. Do you recall [the defendant] speaking to you about [her husband] in any other negative ways? A. Mostly, she just complained he didn't make enough money. It was never enough. [...] Q. How often would she complain to you about [her husband] not making enough money? A. I don't know exactly how often, but more often than not. [...] Whenever it just wasn't enough for her, she would complain.	GA = ✓ Model = ✗ Lawyer = ✗	This is describing a greedy woman. It discusses the relationships between the defendant and her husband, but does not sexualize the defendant.
NORM	[The defendant], the woman who took every opportunity to line her pockets, a heartless schemer who manipulated and lured men to their peril. She deceived her husband, children, family and friends. She has earned the title premeditated murderer, queen of greed and evil.	GA = ✗ Model = ✓ Lawyers = ✓	Portraying the defendant as a person who manipulates, lies, is greedy and evil, etc.
NORM	Q. Did you see any rings? A. I saw one or two small ones, yes. Q. What about earrings, i am pointing to my earrings; do you see them? A. Yes, i do know earrings, yes. Q. Did you see any earrings? A. I didn't see earrings. Q. What else did you see other than these documents, credit cards, the jewelry, three cell phones, and some purses and coins?	GA = ✓ Model = ✗ Lawyers = ✗	Do not see the relevance of this conversation — need more context and background knowledge about this case to know why this was originally coded as related to describing the defendant being "greedy".
MOM	[The defendant] will pose no danger to children if she's punished to life in prison. [The defendant] wanted to be a mom, and she had four kids. Being a mom, that was her dream, but she was never given the necessary tools to do this right. All four of her kids were taken by CPS the same day [the victim] died because the home was so disgusting and unsafe to live in.	GA = ✗ Model = ✓ Lawyers = ✓	Portraying the defendant being an unfit mother as "the home was so disgusting and unsafe to live in".

Table 5: Example of passages where the lawyers agreed with model predictions, and the reasons why they agreed.

tual statements and value judgements. They decided to create a separate coding category to highlight factual descriptions of crimes where the victim was a child of the defendant.

The legal experts also identified challenges when annotating with model-generated annotations. Specifically, they identify the need for more context: “*Machine annotation is done in isolation, (we) lost the context for decision making.*” — it could be hard for them to judge a passage in isolation, especially for themes that often involve multiple people or rely on event descriptions to build up a story. Another challenge is the lack of flexibility. The legal experts mention that when they are doing qualitative annotations on the whole transcript on MaxQDA, they are able to go back and forth to revise previous annotation decisions when they see a running theme. However, reading isolated passages identified by the computational model on Prodigy, they have to make decisions as-is, based on the passage shown. These serve as a good pointer for the design of an interactive annotation pipeline.

Desired Role of a Computational Model In this highly complex annotation task, we find that, first, the legal experts desire a computational model that prioritizes *recall* over *precision*. Second, they appreciate a model that surfaces relevant segments rather than perfect predictions. We discuss these two points in detail.

First, expert annotators would rather see more non-relevant results than potentially miss a relevant result. Therefore, it is better to have a model that retrieves the majority of relevant outcomes, even if it means more irrelevant outcomes are also retrieved. In standard machine learning terms, they prioritize *recall* over *precision*. We caution that this preference for recall may be related to our evaluation setup — it is easy to check whether segments already picked out by human annotators are relevant, but challenging to determine whether non-annotated segments are truly irrelevant. Based on the characteristics of these potentially flawed annotations, recall is a more reliable measure than precision.

Second, the flaws of manual annotation are precisely where the legal experts could use a computational model surfacing relevant segments. As one lawyer coauthor mentioned in phase three: “*The model adds another layer of perspectives. I see [the disagreement of model from the previous annotator] as a benefit — it helps remind me that I may have my own biases. Reading these tricky passages helps me identify my blind spots.*” That is to say, the benefit of the computational model lies in the imperfect predictions it makes. These imperfect predictions prompt the legal experts to refine, concretize, and build consensus for the annotations. And as pointed out by several lawyers, “[the model output] is a great educational tool when we onboard new coders.”.

8 Conclusion

Identifying examples of themes by annotating passages is central to social science research and a significant source of effort and cost. Computational methods, like the language model used in this study, are a promising way to in-

crease scalability and improve annotation consistency. Standard NLP protocols often seek to build systems that *replace* human annotators. In contrast, this study focuses on legal experts, who neither intend nor desire to be removed from the annotation process. In the standard NLP protocol, the primary goal is accuracy, and it is assumed that labels are stable, well-defined, and consistently applicable. Identifying gender-biased narratives in courtroom, on the other hand, is a nuanced task involving subjective decision-making. We identify the value of an *imperfect* computational model in promoting the consistency of annotations in this complex setting. Rather than seeking to *replace* experts, we suggest using computational models to *assist* experts. One of the main concerns of experts in generating their own annotations is the influence of their personal bias, leading to biased and inconsistent annotations within a multi-annotator team. The benefits of a computational model, in the eyes of these experts, is to provide valuable opportunities for them to reflect on their own biases, and a space — grounded in specific examples — for them to engage in discussions with each other to build consensus on the definitions of annotations.

Ethical Statement

Ethical considerations We view our work as a dialogue between computational tools and close qualitative readings, where close reading is an essential part of understanding the gendered discourse in capital trials. We recognize that, though court documents are in the public record, sharing the narratives deployed in a trial could negatively impact the subjects studied in this paper as well as other personnel related to the cases. Topics discussed in a capital trial could involve information regarding deeply private information such as family history, childhood experiences, and relationship history about the defendant as well as witnesses. The fact that we share quotes as examples of gendered discourse gives an additional spotlight to the very kind of language we hope to mitigate. To ameliorate this, we anonymize cases by sharing only background information relevant to the analysis (e.g., excluding names and information such as court location.) In doing so, we aim to minimize the potential harm of linking quotes to defendants. Lastly, due to the ethical and privacy concerns of the defendants being studied, we do not release the raw data for this paper, and only release code for reproducibility.

Positionality The research question in this paper is primarily concerned with the treatment of women in the criminal legal process. The majority of the research team is comprised of women researchers. While our research team does not consist of individuals who have been defendants in criminal trials, the team does consist of legal practitioners who have been closely involved with capital trials for decades. Our team additionally consists of computer scientists who have previously conducted research on fairness topics.

References

Anzovino, M. E.; Fersini, E.; and Rosso, P. 2018. Automatic Identification and Classification of Misogynistic Language

on Twitter. In *International Conference on Applications of Natural Language to Data Bases*.

Atwell, M. W. 2007. *Wretched sisters: Examining gender and capital punishment*, volume 19. Peter Lang.

Baker, D. V. 2015. *Women and capital punishment in the United States: An analytical history*. McFarland.

Baldus, D. C.; Pulaski, C.; and Woodworth, G. 1983. Comparative review of death sentences: An empirical study of the Georgia experience. *The Journal of Criminal Law and Criminology* (1973-), 74(3): 661–753.

Barale, C.; Klaisoongnoen, M.; Minervini, P.; Rovatsos, M.; and Bhuta, N. 2023. AsyLex: A Dataset for Legal Language Processing of Refugee Claims. In Preoțiuc-Pietro, D.; Goanta, C.; Chalkidis, I.; Barrett, L.; Spanakis, G. J.; and Aletras, N., eds., *Proceedings of the Natural Legal Language Processing Workshop 2023*, 244–257. Singapore: Association for Computational Linguistics.

Basile, V.; Bosco, C.; Fersini, E.; Nozza, D.; Patti, V.; Rangel Pardo, F. M.; Rosso, P.; and Sanguinetti, M. 2019. SemEval-2019 Task 5: Multilingual Detection of Hate Speech Against Immigrants and Women in Twitter. In May, J.; Shutova, E.; Herbelot, A.; Zhu, X.; Apidianaki, M.; and Mohammad, S. M., eds., *Proceedings of the 13th International Workshop on Semantic Evaluation*, 54–63. Minneapolis, Minnesota, USA: Association for Computational Linguistics.

Bird, S.; Klein, E.; and Loper, E. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. ”O'Reilly Media, Inc.”.

Blommaert, J.; and Bulcaen, C. 2000. Critical discourse analysis. *Annual review of Anthropology*, 29(1): 447–466.

Bommarito, M. J.; Katz, D. M.; and Detterman, E. M. 2018. LexNLP: Natural Language Processing and Information Extraction For Legal and Regulatory Texts. *InfoSciRN: Legal Informatics (Topic)*.

Carroll, J. E. 1996. Images of women and capital sentencing among female offenders: Exploring the outer limits of the eight amendment and articulated theories of justice. *Tex. L. Rev.*, 75: 1413.

Chalkidis, I.; Androutsopoulos, I.; and Aletras, N. 2019. Neural Legal Judgment Prediction in English. In Korhonen, A.; Traum, D.; and Márquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4317–4323. Florence, Italy: Association for Computational Linguistics.

Chalkidis, I.; Fergadiotis, M.; Malakasiotis, P.; Aletras, N.; and Androutsopoulos, I. 2020. LEGAL-BERT: The Muppets straight out of Law School. In Cohn, T.; He, Y.; and Liu, Y., eds., *Findings of the Association for Computational Linguistics: EMNLP 2020*, 2898–2904. Online: Association for Computational Linguistics.

Chalkidis, I.; Jana, A.; Hartung, D.; Bommarito, M.; Androutsopoulos, I.; Katz, D.; and Aletras, N. 2022. LexGLUE: A Benchmark Dataset for Legal Language Understanding in English. In Muresan, S.; Nakov, P.; and Villavicencio, A., eds., *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 4310–4330. Dublin, Ireland: Association for Computational Linguistics.

Chiril, P.; Benamara, F.; and Moriceau, V. 2021. ”Be nice to your wife! The restaurants are closed”: Can Gender Stereotype Detection Improve Sexism Classification? In *Conference on Empirical Methods in Natural Language Processing*.

Chiril, P.; Moriceau, V.; Benamara, F.; Mari, A.; Origgi, G.; and Coulomb-Gully, M. 2020. He said “who’s gonna take care of your children when you are at ACL?”: Reported Sexist Acts are Not Sexist. In Jurafsky, D.; Chai, J.; Schluter, N.; and Tetreault, J., eds., *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4055–4066. Online: Association for Computational Linguistics.

Conley, J. M.; O’Barr, W. M.; and Riner, R. C. 2019. *Just words: Law, language, and power*. University of Chicago Press.

Crew, B. K. 1991. Sex differences in criminal sentencing: Chivalry or patriarchy? *Justice Quarterly*, 8(1): 59–83.

Dalke, I.; and Greene, J. 2023. Prerequisites and pathways: How social categorization helps administrators determine moral worth. *Theory and Society*, 1–26.

Daly, K.; and Bordt, R. L. 1995. Sex effects and sentencing: An analysis of the statistical literature. *Justice Quarterly*, 12(1): 141–175.

Devinney, H.; Björklund, J.; and Björklund, H. 2020. Semi-Supervised Topic Modeling for Gender Bias Discovery in English and Swedish. In Costa-jussà, M. R.; Hardmeier, C.; Radford, W.; and Webster, K., eds., *Proceedings of the Second Workshop on Gender Bias in Natural Language Processing*, 79–92. Barcelona, Spain (Online): Association for Computational Linguistics.

Doerner, J. K.; and Demuth, S. 2014. Gender and Sentencing in the Federal Courts: Are Women Treated More Leniently? *Criminal Justice Policy Review*, 25(2): 242–269.

Ehrlich, S. 2019. ’Well, I saw the picture’: semiotic ideologies and the unsettling of normative conceptions of female sexuality in the Steubenville rape trial. *Gender & Language*, 13(2).

Ehrlich, S. 2021. Intersections of race and gender in sexual assault trials. *Gender & Language*, 15(3).

Embry, R.; and Phillip M. Lyons, J. 2012. Sex-Based Sentencing: Sentencing Discrepancies Between Male and Female Sex Offenders. *Feminist Criminology*, 7(2): 146–162.

Farnworth, M.; and Raymond H. C. Teske, J. 1995. Gender Differences in Felony Court Processing:. *Women & Criminal Justice*, 6(2): 23–44.

Greene, J.; and Dalke, I. 2021. ”You’re still an angry man”: Parole boards and logics of criminalized masculinity. *Theoretical Criminology*, 25(4): 639–662.

Hakimi Parizi, A.; Liu, Y.; Nokku, P.; Gholamian, S.; and Emerson, D. 2023. A Comparative Study of Prompting Strategies for Legal Text Classification. In Preoțiuc-Pietro, D.; Goanta, C.; Chalkidis, I.; Barrett, L.; Spanakis, G. J.;

and Aletras, N., eds., *Proceedings of the Natural Legal Language Processing Workshop 2023*, 258–265. Singapore: Association for Computational Linguistics.

Koons-Witt, B. A.; Sevigny, E. L.; Burrow, J. D.; and Hester, R. 2014. Gender and Sentencing Outcomes in South Carolina: Examining the Interactions With Race, Age, and Offense Type. *Criminal Justice Policy Review*, 25(3): 299–324.

Levenson, M. 2022. Judge Overturns Murder Convictions, Citing Use of Rap Lyrics At Trial. *New York Times*.

Liptak, A. 2024a. Citing Racist Testimony, Justices Call for New Sentencing in Texas Death Penalty Case. <https://www.nytimes.com/2017/02/22/us/politics/duane-buck-texas-death-penalty-case-supreme-court.html>. [Online; accessed 5-May-2024].

Liptak, A. 2024b. Did Prosecutors' Sex Shaming Help Send Brenda Andrew to Death Row? <https://www.nytimes.com/2024/04/01/us/supreme-court-death-penalty-sex-shaming.html>. [Online; accessed 29-Apr-2024].

Medvedeva, M.; and McBride, P. 2023. Legal Judgment Prediction: If You Are Going to Do It, Do It Right. In Preotiuc-Pietro, D.; Goanta, C.; Chalkidis, I.; Barrett, L.; Spanakis, G. J.; and Aletras, N., eds., *Proceedings of the Natural Legal Language Processing Workshop 2023*, 73–84. Singapore: Association for Computational Linguistics.

Mogul, J. L. 2005. The Dykier, the Butcher, the Better: The State's Use of Homophobia and Sexism to Execute Women in the United States Symposium to Honor the Work of Professor Ruthann Robson: Equality. *New York City Law Review*, 8(2): 473–494.

Mozafari, M.; Farahbakhsh, R.; and Crespi, N. 2020. Hate speech detection and racial bias mitigation in social media based on BERT model. *PloS one*, 15(8): e0237861.

Nagel, I. H.; and Hagan, J. 1983. Gender and Crime: Offense Patterns and Criminal Court Sanctions. *Crime and Justice*, 4: 91–144.

Nellis, A. 2021. In the extreme: Women serving life without parole and death sentences in the United States.

of Appeals for the Fourth Circuit, U. S. C. 2016. Johnny Bennett v. Bryan Stirling. *Administrative Office of the United States Courts*.

Otmazgin, S.; Cattan, A.; and Goldberg, Y. 2022. F-coref: Fast, Accurate and Easy to Use Coreference Resolution. In Buntine, W.; and Liakata, M., eds., *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing: System Demonstrations*, 48–56. Taipei, Taiwan: Association for Computational Linguistics.

Potts, A.; and Formato, F. 2021. Women Victims of Men Who Murder: XML Mark-up for Nomination, Collocation, and Frequency Analysis of Language of the Law. In *The Routledge Handbook of Language, Gender, and Sexuality*. Routledge. ISBN 978-1-315-51485-7.

Potts, A.; and Weare, S. 2018. Mother, Monster, Mrs. I: A Critical Evaluation of Gendered Naming Strategies in English Sentencing Remarks of Women Who Kill. *International Journal for the Semiotics of Law - Revue internationale de Sémiotique juridique*, 31(1): 21–52.

Rapaport, E. 1990. Some questions about gender and the death penalty. *Golden Gate UL Rev.*, 20: 501.

Rapaport, E. 1991. The Death Penalty and Gender Discrimination. *Law & Society Review*, 25(2): 367–383.

Rodriguez, S. F.; Curry, T. R.; and Lee, G. 2006. Gender Differences in Criminal Sentencing: Do Effects Vary Across Violent, Property, and Drug Offenses? *Social Science Quarterly*, 87(2): 318–339.

Shiff, T. 2021. A sociology of discordance: Negotiating schemas of deservingness and codified law in US asylum status determinations. *American Journal of Sociology*, 127(2): 337–375.

Streib, V. L. 1990. Death Penalty for Female Offenders. *University of Cincinnati Law Review*, 58(3): 845–880.

Streib, V. L. 2002. Gendering the Death Penalty: Countering Sex Bias in a Masculine Sanctuary Symposium: Addressing Capital Punishment through Statutory Reform. *Ohio State Law Journal*, 63(1): 433–474.

Streib, V. L. 2005. Rare and Inconsistent: The Death Penalty for Women Special Feature: Women as Perpetrators of Crime. *Fordham Urban Law Journal*, 33(2): 609–636.

Sutton, J.; Mills, J. R.; Merrigan, J.; and Swain, K. 2022. Death by Dehumanization: Prosecutorial Narratives of Death-Sentenced Women and LGBTQ Prisoners.

Thalken, R.; Stiglitz, E.; Mimno, D.; and Wilkens, M. 2023. Modeling Legal Reasoning: LM Annotation at the Edge of Human Agreement. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 9252–9265. Singapore: Association for Computational Linguistics.

Vásquez, J.; Bel-Enguix, G.; Andersen, S. T.; and Ojeda-Trueba, S.-L. 2022. HeteroCorpus: A Corpus for Heteronormative Language Detection. In Hardmeier, C.; Basta, C.; Costa-jussà, M. R.; Stanovsky, G.; and Gonen, H., eds., *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeNLP)*, 225–234. Seattle, Washington: Association for Computational Linguistics.

Vaudaux, O.; Bazzoli, C.; Coavoux, M.; Vial, G.; and Vergès, É. 2023. Pretrained Language Models v. Court Ruling Predictions: A Case Study on a Small Dataset of French Court of Appeal Rulings. In Preotiuc-Pietro, D.; Goanta, C.; Chalkidis, I.; Barrett, L.; Spanakis, G. J.; and Aletras, N., eds., *Proceedings of the Natural Legal Language Processing Workshop 2023*, 38–43. Singapore: Association for Computational Linguistics.

Visher, C. A. 1983. Gender, Police Arrest Decisions, and Notions of Chivalry. *Criminology*, 21(1): 5–28.

Wyse, J. J. B. 2013. Rehabilitating Criminal Selves: Gendered Strategies in Community Corrections. *Gender & Society*, 27(2): 231–255.