

Race in the Record: Measuring Racial Signal and Model Inference in Police Arrest Narratives

Arka Dutta^{◇*}, Ted Forsyth^{♠*}, O. Nicholas Robertson[◇], Ashiqur R. KhudaBukhsh[◇]

[◇]Rochester Institute of Technology [♠]Syracuse University

ad2688@rit.edu, tdforsyt@syr.edu, onrgcj@rit.edu, axkvse@rit.edu

Abstract

We present RPD-UoF, a corpus of 3,091 police use-of-force narratives obtained via Freedom of Information Law (FOIL) requests from the Rochester Police Department, combining structured demographic metadata with officer-authored free-text accounts. We use this dataset to investigate three related questions: whether officer narratives contain recoverable racial signal after demographic redaction; whether large language models amplify racial framing when conditioned on subject race; and whether that signal survives text redaction to re-emerge in multimodal outputs. Our experiments suggest that redacted narratives are statistically separable by subject race; LLMs produce structurally longer narratives for White profiles than Black profiles from identical seed texts; and CLIP classifications of GPT-5-generated images differ significantly depending on the racial origin of the source narrative. Our analyses provide an initial empirical lens on how race appears in police narratives and how modern language models (and vision language models) interact with such sensitive attributes. The dataset and analyses are intended to support future research on fairness, accountability, and transparency in computational studies of policing and criminal justice.

1 Introduction

The killing of George Floyd by a U.S. police officer marked a watershed moment in public discourse on racial justice and policing, catalyzing sustained scrutiny of racial disparities in the use of force against Black Americans. Decades of empirical research demonstrate that Black civilians are disproportionately subject to police stops, searches, arrests, and use of force, including lethal force (Pier-son et al., 2020; Edwards et al., 2019; Fryer Jr,

2019), even after accounting for situational and contextual factors. Nationally, Black Americans are several times more likely than White Americans to experience police use of force (Fryer Jr, 2019), a pattern that has been documented across cities, policing contexts, and methodological approaches. These disparities are not merely statistical artifacts; they reflect longstanding structural inequities embedded within institutional practices of law enforcement (Goldkamp, 1976; Knox et al., 2020).

Recent work in computational social science and NLP has shown that racial disparity in policing is not confined to outcomes alone but is also reflected in language and interaction. Prior studies demonstrate that officers use systematically less respectful language when interacting with Black community members than with White community members during routine traffic stops, even after controlling for officer race, infraction severity, location, and stop outcome (Voigt et al., 2017). Such findings suggest that police-authored text, whether spoken or written, may encode racialized patterns that extend beyond explicit references to race and into subtler forms of linguistic framing.

At the same time, the increasing integration of artificial intelligence into criminal justice systems (Caianiello, 2021; Gebru et al., 2021; Gans-Combe, 2022; Basu et al., 2025) introduces new sociotechnical risks. AI systems are now used or proposed for a range of high-stakes applications, including report drafting, risk assessment, evidentiary summarization, and internal review (Narayanan Venkit et al., 2024). When trained on or prompted with institutional records, these systems may inadvertently learn, amplify, or operationalize racialized signals embedded in administrative data. Even when sensitive attributes such as race are explicitly removed, models may still infer or reconstruct them from correlated textual cues, raising concerns about proxy discrimination, infer-

*Arka Dutta and Ted Forsyth are equal-contribution first authors. Ashiqur R. KhudaBukhsh is the corresponding author.

ential leakage, and downstream harm to already vulnerable communities (Staab et al., 2024).

This paper investigates how race is reflected in police use-of-force documentation and how modern language and multimodal models interact with such records. To this end, we present a novel corpus of 3,091 police use-of-force narratives obtained through Freedom of Information Law (FOIL) requests. Each record combines structured demographic and incident metadata with free-text narratives authored by officers describing the encounter. Leveraging this dataset, we conduct a multi-stage empirical analysis.

First, we examine whether officer-written narratives alone contain sufficient linguistic signal to predict the race of the arrested individual, even when race is not explicitly mentioned. Second, we analyze how large language models behave when tasked with generating arrest narratives conditioned on demographic profiles, asking whether model outputs differ systematically across racial categories. Finally, we study how multimodal generative systems further complicate the sociotechnical landscape by transforming race-redacted narratives into synthetic images and character descriptions, and whether racial information can still be recovered downstream. Of course, our goal is not to enable automated race inference, but to surface and measure the extent to which racialized information is already present, recoverable, and potentially amplified by AI systems operating on policing data. By empirically tracing how racial signal propagates from human-authored narratives to language models and onward to multimodal representations, this work contributes to computational social science by informing audits, governance, and accountability mechanisms for AI deployed in criminal justice settings.

We investigate the following three research questions:

RQ1. *Do officer-written use-of-force narratives encode racial signal that survives demographic redaction?*

RQ2. *Does conditioning a language model on subject race amplify framing disparities beyond what human officers produce?*

RQ3. *Can racial signal survive text-level redaction and re-emerge in images generated from those narratives?*

2 Related Work

Racial disparities in policing outcomes are well-documented. Large-scale studies of traffic stops find persistent disparities in stop and search rates (Pierson et al., 2020; Gelman et al., 2007), and use-of-force records show higher rates of non-lethal force against Black and Hispanic civilians in comparable situations (Fryer Jr, 2019). Administrative datasets can mask this bias, because selective recording and enforcement already reflect structural inequities (Knox et al., 2020). Social psychology provides the perceptual mechanisms: racialized primes shift visual attention and threat judgments in crime-relevant contexts (Eberhardt et al., 2004). Language and interaction are also sites of disparity. Officers use less respectful language with Black community members during routine traffic stops (Voigt et al., 2017), the structure of police-civilian dialogue can be modeled at scale (Prabhakaran et al., 2018), and escalated stops differ linguistically in their earliest moments, with listeners reacting differently to playback of those opening seconds (Rho et al., 2023; Field et al., 2023).

The deployment of AI in criminal justice introduces compounding risks. Policing data can embed the effects of unlawful practice (Angwin et al., 2022; Richardson et al., 2019), and fairness definitions applied without institutional context can fail (Selbst et al., 2019). Language models encode social hierarchies from pretraining text (Caliskan et al., 2017), and benchmark performance does not imply low bias (Nadeem et al., 2021; Blodgett et al., 2020). Our RQ1 result sits within the literature on the failure of de-identification: Staab et al. (2024) show that models infer protected attributes from ostensibly anonymous text at scale, and we extend that failure mode to an institutional record where the inference runs off an officer’s own framing rather than off self-disclosure. Our RQ2 result connects to work measuring demographic disparity in open-ended LLM generation (Shieh et al., 2026; Fang et al., 2024); relative to that work, the disparity we recover in a high-stakes institutional genre is carried less by the content of what models write than by whether they agree to write at all. Hofmann et al. (2024) show that models produce racially disparate decisions in response to dialect cues alone, without any explicit demographic signal, a particularly relevant finding for policing narratives that carry indirect identity markers. For multimodal systems, demographic stereotypes appear at scale

in text-to-image generation (Bianchi et al., 2023), and face analysis systems show disparate performance across skin type and gender (Buolamwini and Gebru, 2018; Raji and Buolamwini, 2019).

3 Data and Dataset Contribution

One of the key contributions in this paper is a novel use-of-force records dataset¹. In what follows, we describe the collection process and subsequent extraction, cleaning, and human validation steps.

3.1 Data Collection via New York FOIL

We obtained the use-of-force records analyzed in this work through the New York State Freedom of Information Law (FOIL), which provides the public a statutory process to request and receive government records. Under FOIL, agencies must grant access, deny access in writing, or acknowledge receipt of a request within five business days; acknowledgments must also provide an approximate date for a determination, typically within twenty business days, or else specify a written reason for delay and a “date certain.” Failures to comply with these timelines constitute a denial that can be appealed administratively (New York State Committee on Open Government, a,b; New York State Department of State).

Following this process, we submitted a FOIL request to the Rochester, NY Police Department (RPD) via the City’s webform on June 12, 2020, requesting all Subject Resistance Reports (SRRs) from 2001–2019. The City assigned the request FOIL number RR20-02509. The City granted the request in part on June 29, 2021. The release included SRRs from 2015–2019 (five years) rather than the full requested window, as older records had been destroyed under the agency’s retention schedule (Forsyth, 2026).

The released corpus comprises 3,453 scanned PDF documents containing 5,297 subject-level SRR records: a single incident generates one record per subject, so document and record counts differ. SRRs are encounter-level records in which one or more officers document force used against one or more subjects for a single incident; records are linked via a crime report number (CR#). SRRs include structured subject descriptors (e.g., sex, race, ethnicity, height, weight) and incident metadata (e.g., location, date, time). Personally identify-

ing information about subjects (e.g., name, date of birth) is redacted by the City, with additional redactions for minors. SRRs also contain structured fields for subject actions and police force options, including the temporal order in which force options were used and an officer-reported effectiveness rating, as well as free-text officer narratives describing the encounter; narratives may be continued via addenda when necessary (Forsyth, 2026). Table 1 shows example excerpts.

3.2 RPD Use-of-Force (RPD-UoF) Corpus

Scale and scope. The raw FOIL release consists of 5,297 Subject Resistance Reports spanning 2015–2019, delivered as 3,453 individually scanned PDFs. A separately obtained 2022 batch contributes 1,698 additional records in monthly PDFs. Table 7 traces how these become the 3,091-record analysis corpus. After race-label matching the corpus spans six annual windows: 2015 ($n=579$), 2016 ($n=632$), 2017 ($n=380$), 2018 ($n=683$), 2019 ($n=641$), and 2022 ($n=176$).²

Each document is a first-person officer account following the police report genre (Coulthard et al., 2007): situating context, triggering event, subject behavior, force sequence, and disposition. This formulaic structure, shared across officers and incidents, makes word-choice and framing differences more visible by holding narrative shape constant.

OCR extraction pipeline. The 2015–2019 documents are scanned PDFs varying in resolution, page orientation, and handwritten content. We evaluate three extraction systems in parallel (Cutting and Cutting-Decelle, 2021): (1) Gemini-2.5-Flash, which passes each page to a multimodal API performing joint layout analysis and text recognition, returning structured JSON with CR number, narrative, race, and gender; (2) pypdf followed by GPT-based token repair for OCR artifacts; and (3) dots.ocr by Rednote (Ouyang et al., 2025) as an independent reference. We measure cross-system agreement via character-level Levenshtein edit distance (Levenshtein, 1966) on a 200-document sample (5.8% of the 2015–2019 docu-

¹Our datasets are available at <https://github.com/Social-Insights-Lab/race-in-the-records>

²The 2017 subset is smaller because of attrition at two stages, not sampling: 2017 accounts for 42 of the 293 OCR failures (14.3% of failures against 13.2% of the raw corpus), consistent with that year’s physical archive being in poorer condition, and 34 of the 414 successfully extracted 2017 records were dropped at the race-label step (20 Hispanic, 9 unknown, 5 Asian). The force-type distribution in 2017 is proportional to other years (χ^2 -tested), so year-stratified analyses are unaffected.

CR#	Race	Gen.	Force	Resistance	Narrative (excerpt)
15-114726	Black	M	OC	Verbal	R/O observed S at [ADDRESS] from a previous harassment incident. S walked aggressively toward my vehicle stating profanities and threats. I informed S of arrest; S was initially compliant. [In vehicle, S began kicking and head-banging; I deployed OC to chest/face, effective. Spit hood applied.]
15-217645	Black	M	Ground Stab.	Verbal	ROs responded to family trouble: S locked in bathroom with knife, reported intoxicated. ROs attempted verbal contact; S blocked the door and appeared to prepare to lunge. RO deployed Taser (one five-second cycle). S secured via two-point landing; transported for MHA evaluation.

Table 1: Example Use-of-Force Narratives (excerpted) with Redacted Identifiers

ments), drawn uniformly at random from the 2015–2019 pool; the resulting per-year breakdown—2015 ($n=40$), 2016 ($n=44$), 2017 ($n=26$), 2018 ($n=48$), 2019 ($n=42$)—mirrors the corpus composition. Gemini and dots.ocr align closely (mean 6.4 ± 4.1 characters), with human inspection slightly favoring Gemini for paragraph structure preservation. We adopt Gemini as the primary extractor; it returned a parseable record for 3,160 of the 3,453 documents (293 failures, concentrated in low-quality scans). The 2022 monthly PDFs are natively digital and are split by matching the header pattern “Incident type: Use of force”, then extracted with pypdf.




Figure 1: Demographic distribution of RPD-UoF records by race and gender.

Human verification of the extraction pipeline.

Extraction errors correlated with document quality could introduce measurement bias if quality itself correlates with subject demographics. We sampled 750 case reports (24.3% of the analysis corpus) for independent annotation by two trained graduate students working from scanned PDFs, under the supervision of an expert social scientist. The pipeline agrees with the adjudicated human label on 87.7% of race labels and 91.1% of gender labels (Figure 4). Disagreements concentrate in

two cases: non-standard race abbreviations (e.g., “B/M” for Black male) and records where poor scan quality causes the human annotators themselves to disagree. Records where the human-override label differs are flagged for downstream users.

Data cleaning and deduplication. The raw release contains duplicates from two sources: multi-subject incidents generate one SRR per subject under the same CR number, and administrative re-filings sometimes replace earlier drafts. We deduplicate by CR number, retaining the longest narrative as the canonical version; combined with successful OCR extraction this reduces the 2015–2019 material to 3,160 records (Table 7). We then apply a three-step quality filter: remove records whose narrative failed to extract; remove records missing race in administrative data (unknown $n=86$); and remove groups with fewer than 50 verified examples (Hispanic $n=151$; Asian $n=8$). The final analysis corpus is 3,091 records: 2,277 Black subjects (73.7%) and 814 White subjects (26.3%).

PII redaction and narrative cleaning. Officer names, precinct identifiers, and location-specific language remain in the FOIL release and can serve as indirect race proxies, because certain officers patrol predominantly Black neighborhoods and certain streets could be demographically segregated (Knox et al., 2020). We apply a two-stage redaction process. First, a rule-based system replaces officer names, badge numbers, exact addresses, and institutional names with generic tokens (e.g., [OFFICER], [ADDRESS], [SCHOOL], etc.), removing 62.3% of identifiable references. Second, Claude Haiku via API masks remaining proper nouns associated with specific communities while preserving syntactic structure. We also standardize subject pronouns to “they/them” using Gemini-2.5-Flash to remove gender signal

before race classification. Post-redaction, a multi-LLM judge panel (Claude-Sonnet-4.5, GPT-5.2, Gemini-2.5-Flash) confirmed that no individual officer name from the RPD roster appeared in more than 1% of test-set narratives, verifying that officer identity was effectively suppressed; the residual under 1% was removed by manual inspection.

4 Experimental Results

4.1 RQ1: Linguistic Bias in Officer-Authored Narratives

We probe for racial signal in redacted narratives using both lexical analysis and classification as an indicator of separable linguistic cues.

4.1.1 Differential Lexical Analysis

We apply the log-odds ratio with a Dirichlet prior (Monroe et al., 2008), which produces a signed z -score per word: positive values indicate over-representation in Black-subject narratives, negative values in White-subject narratives. Two classes of token would otherwise dominate this ranking for reasons unrelated to framing. Explicit race terms (*black*, *white*, *caucasian*, and their variants) appear in 14.0% of narratives and simply restate the label. Officer surnames rank highly because patrol assignments are racially segregated, not because officers write differently. We therefore remove both before estimation, alongside standard stop words and Rochester-specific geographic terms. Person names are removed using the 695-token name vocabulary induced from our redaction pipeline (Section 3.1), which suppresses every RPD surname we audited to zero occurrences; this is a stricter filter than a hand-curated roster list.

Across 3,990 words, 178 are significantly associated with Black subjects ($z > 1.96$) and 392 with White subjects ($z < -1.96$). Words at the top of the Black-associated list include *suspect* ($z = 7.05$), *vehicle* ($z = 7.69$), *gun* ($z = 7.00$), *running* ($z = 6.80$), *waistband* ($z = 6.45$), *handgun* ($z = 6.15$), and *chase* ($z = 5.28$). Similarly, top words in White-associated list include *ambulance* ($z = -10.21$), *gurney* ($z = -9.88$), *suicidal* ($z = -8.59$), *intoxicated* ($z = -7.22$), *bridge* ($z = -6.93$), *hospital* ($z = -6.23$), and *mental* ($z = -5.70$). Hence, encounters involving Black subjects are framed in terms of pursuit, weapon suspicion, and criminal threat, whereas encounters involving White subjects are framed in terms of medical crisis, substance intoxication, and mental

health intervention. This is a consistent pattern. A subject characterized through the vocabulary of threat is processed as a security problem; a subject characterized through the vocabulary of crisis is more likely to receive a welfare response indicating well-established different institutional outcomes (Wacquant, 2020; Muhammad, 2019).

Interestingly, *suspect* ($z = 7.05$) appears in Black-subject narratives at 1.8 times its rate in White-subject narratives (rates normalised by corpus size). This result is consistent with the findings from body-camera transcript analysis (Rho et al., 2023) and social-psychological research on race and threat perception (Eberhardt et al., 2004). *Waistband* ($z = 6.45$), which signals an officer’s attention to possible weapon concealment, appears at 3.0 times its rate in White-subject narratives, consistent with research on weapon-related stops (Gelman et al., 2007).

Framing or incident mix? A natural objection is that this vocabulary contrast reflects a different distribution of incident types across groups rather than different framing of comparable incidents. Using the officer-reported structured `force_used` field as an incident-type proxy, the distributions across the four most common categories (Verbal Commands, Physical Force, OC Spray, Ground Stabilization) do not differ by subject race, $\chi^2(3) = 0.44$, $p = 0.93$. Subject-resistance codes are similarly proportional (Avoiding Custody: 17.3% Black vs. 16.1% White; Verbal Resistance: 7.7% vs. 10.3%). Encounters of each recorded type occur at essentially the same rates for both groups, yet the lexical contrast persists within them. The disparity is therefore in how comparable encounters are written up, not in which encounters occur.

4.1.2 Neural Classification

The lexical analysis shows that certain words show strong association with subject race. We next consider if the narrative texts are statistically distinguishable based on the subject race. We examine this possibility using classification-based separability tests. Separability tests have been used in prior NLP research to examine framing differences across diverse settings, including news coverage of policing (Dutta et al., 2022), LLM-assisted summarization (Vijay et al., 2025), and partisan patterns in courtroom speeches (Fang et al., 2023). The underlying assumption is straightforward: if two sets of samples are drawn from the same underlying

distribution, then, on a balanced test set, a classifier trained to distinguish between them should achieve accuracy close to chance. Conversely, classification accuracy substantially above chance indicates that the two distributions are separable.

We consider four models spanning two architectural families. From encoder-based models, we fine-tune BERT-base (Devlin et al., 2019) and DeBERTa-v3-small (He et al., 2021). From decoder-only model, we fine-tune Mistral-7B (Jiang et al., 2023) and Llama-3-8B (Grattafiori et al., 2024) with Low-Rank Adaptation (Hu et al., 2022). All four models are trained with class-balanced weighting for binary classification. We run three seeds each and report mean balanced accuracy and macro F1.

Model	Balanced Accuracy	F1	White Recall	Black Recall
BERT	0.61 ± 0.03	0.60 ± 0.03	0.50 ± 0.04	0.72 ± 0.12
DeBERTa	0.58 ± 0.02	0.56 ± 0.01	0.39 ± 0.05	0.77 ± 0.08
Mistral	0.54 ± 0.01	0.52 ± 0.02	0.69 ± 0.14	0.38 ± 0.14
Llama3	0.50 ± 0.01	0.50 ± 0.01	0.41 ± 0.07	0.59 ± 0.05

Table 2: Race classification results on the held-out test set (141 Black / 141 White; chance = 50% balanced accuracy). Means over three seeds.

Table 2 summarizes our results. BERT-base achieves the highest balanced accuracy at 60.99% while among the decoder models, Mistral reaches 53.7%. A one-sample *t*-test on the three seed values confirms that BERT-base, DeBERTa, and Mistral each exceed the 50% chance baseline significantly while Llama-3-8B does not.

The recall asymmetry across models reveals a secondary pattern. Both encoder models (BERT and DeBERTa) are better at identifying Black-subject narratives than White-subject narratives: BERT and DeBERTa respectively achieve recall ≈ 0.72 and 0.77 on Black narratives and 0.50 and 0.39 on White narratives. This follows from the lexical results: Black-subject narratives cluster around a narrow pursuit-and-weapon vocabulary, making them more homogeneous and easier to separate as a class. White-subject narratives cover a wider range of scenario types, ranging from bar fights and bridge standoffs to medical crises, and that diversity makes them harder to distinguish. Mistral and Llama-3, along with low accuracy, show inconsistent recall patterns, offering no reliable signal. Overall, this experiment shows that the signal is not confined to explicit demographic fields; it is distributed across the vocabulary of the narrative

itself (Dutta et al., 2022; Prabhakaran et al., 2018).

Does the signal survive redaction? Because 14.0% of narratives contain an explicit race term, we re-ran the two encoder models under an identical protocol on text in which every race term and every token from the 695-item name vocabulary is replaced by [REDACTED]. Both models remain significantly above chance (BERT $0.618 \rightarrow 0.589$, $p = .010$; DeBERTa $0.566 \rightarrow 0.563$, $p = .018$; three seeds, one-sample *t*-test against 50%), and neither drop is itself significant ($p = 0.14$ and $p = 0.83$). Explicit demographic tokens are therefore not what the classifiers are using: the separability reported above survives their removal.

This is further supported by a lexical erasure experiment (Figure 2): removing the 100 most differentially weighted words leaves accuracy essentially unchanged (0.590 against 0.575 at baseline), and even after removing 500 words the classifier holds at 0.520 , reaching chance only once roughly 800 words are stripped. The racial signal is not localized to a small set of marker words; it is distributed across the full narrative vocabulary. The curve does thin faster between 100 and 200 removed words than elsewhere, indicating a moderately weighted middle band rather than a perfectly flat distribution.




Figure 2: Balanced accuracy as the most differential words are erased.

4.2 RQ2: Bias in LLM-Generated Narratives

Recent literature raises the concern that AI-assisted report drafting could introduce systematic distortions into institutional records (Dement and Inglis, 2024; Ferguson, 2025). We now ask whether race-conditioned generations differ from real-record framing patterns.

4.2.1 Controlled Generation Experiment

We create 1,000 synthetic subject profiles (500 Black, 500 White), each specifying race, sex, and

age drawn from corpus distributions. We prompt five models under two conditions: zero-shot and few-shot. In the zero-shot setting, the model receives one unlabeled example narrative and is asked to generate a new narrative for a specified profile. This tests whether models encode racial assumptions even without race-labeled examples. In the few-shot setting, five race-labeled example narratives per group are prepended; this tests whether explicit race-labeled context causes models to reproduce or amplify race-correlated patterns.

We consider five models Gemini-2.5-Flash, Gemma 2 9B, Llama 3.1 8B, Mistral-7B v0.3, and Mixtral 8x22B, generating 8,002 narratives total (3,730 Black, 4,272 White; imbalance from content-filter refusals on Black-profile prompts). Each narrative is scored for word count, force keyword density (47-term lexicon), VADER sentiment (Hutto and Gilbert, 2014). Lexical analysis of the generated texts (Table 3) shows that differential adjective use is most pronounced in the few-shot condition, where models sometimes carry racial labels directly into the generated text. In zero-shot outputs, surface-level differences are smaller.

In addition, following the same logic behind RQ1 exploration, we fine-tune BERT-base-uncased and DeBERTa-v3-small to measure whether generated text perpetuates or introduces any recognizable bias. Balanced accuracy across all models and conditions hovers near chance (range: 0.437–0.524), and paired $\Delta P(B)$ gaps are near zero, with the large majority of comparisons non-significant after FDR correction ($q > 0.05$ for 13 of 16 model-classifier pairs). This indicates that LLM-generated narratives are largely impartial with respect to the racial framing signal encoded in real police reports and the models do not systematically reproduce the situational language patterns that distinguish Black-subject narrative from White-subject narratives in human-authored text. Appendix contains full report on model-wise classification results with significance tests (Table 8).

4.2.2 Counterfactual Pair Analysis

To differentiate the effect of race conditioning from prompt variation, we constructed 200 matched counterfactual pairs. For each pair, a single redacted seed narrative from the real corpus was used to generate two outputs from the same model—one with a Black profile, one with a White profile—holding everything else constant. Any systematic

difference is attributable to race conditioning alone.

The results indicate differential generation. White-profile narratives are considerably longer (153.8 ± 48.7 words) than Black-profile narratives (102.1 ± 27.7 words). White-profile narratives also contain more force-related keywords (mean 1.47) than Black-profile narratives (mean 1.03). Sentiments are opposite (VADER compound): Black-profile narratives average 0.604 versus -0.234 for White-profile. All these results are statistically significant as described in Table 4. Models apply a domestic-crisis or mental-health escalation template to White profiles which are inherently longer because it narrates a sequence of de-escalation attempts; on the other hand it assigns an action-summary template to Black profiles. Template selection, not explicit force vocabulary insertion, is the primary mechanism; the few-shot condition amplifies it by making the demographic-to-template link explicit in context. Per-model force keyword differences do not survive Bonferroni correction (Gemma 2 9B $\Delta = 0.02$, Llama 3.1 8B $\Delta = -0.07$, Mistral-7B $\Delta = 0.01$, Gemini $\Delta = 0.17$), which suggests the signal is structural rather than lexical.


These findings are not in contradiction. The counterfactual pairs show that models select different narrative templates depending on race; the classifier results show that those templates do not carry the lexical patterns that separate real Black-subject narratives from real White-subject narratives. The model’s racial assumptions operate at the level of structure, including scenario type, length, and escalation sequence, rather than word choice. Human-written and LLM-generated records therefore encode race differently.

4.3 RQ3: Inferential Leakage in Multimodal Outputs

In the third analysis, we ask whether racial signal survives text-level redaction to re-emerge in downstream visual outputs, a failure mode we term inferential leakage, where a downstream model may recover, re-introduce and accentuate race-linked information in a new form (Staab et al., 2024; Betley et al., 2025). Text-to-image systems amplify demographic stereotypes at scale (Bianchi et al., 2023), and face analysis systems show disparate performance across skin type and gender (Buolamwini and Gebru, 2018; Raji and Buolamwini, 2019). Our question is whether these phenomena compound when the generation prompt is itself a narrative

Prompting	Model	White-skewed adjectives	Black-skewed adjectives
Zero-shot	mistral-7b	white, haired, quiet, caucasian, loud, disruptive, front, apparent, unpredictable, open	black, senior, takedown, high, precinct, incoherent, outstanding, stable, suspect, evasive
	gemini-flash	alcoholic, flushed, welcome, apparent, intoxicated, strong, unsteady, small, abusive, own	heated, separate, low, drive, other, sustained, double, commercial, upright, calm
	gemma2-9b	unsteady, multiple, intoxicated, limp, busy, disorderly, full, past, unresponsive, disruptive	vague, compliant, unsuccessful, final, erratic, unharmed, momentary, male, reassuring, measured
	llama3-8b	handed, public, possible, past, unharmed, irate, several, secondary, joint, successful	opposite, respectful, single, counterfeit, concerned, restrained, upset, rear, non, violent
Few-shot	gemma2-9b	white, last, empty, intoxicated, past, left, subsequent, unaware, broken, domestic	black, racist, own, pat, tense, senior, tall, electrical, unpredictable, dark
	mistral-7b	white, visible, possible, standard, following, justified, administrative, second, quiet, front	black, african, american, high, first, suspicious, young, old, fearful, subject
	llama3-8b	white, slight, past, sudden, situational, aware, strong, handed, electrical, cooperative	black, demographic, continued, combative, serious, handcuffed, domestic, assertive, suspect, difficult
	gemini-flash	white, unsteady, extreme, many, intoxicated, underage, protective, alcoholic, strong, oncoming	black, dark, limited, ongoing, well, full, momentary, common, entire, unpaid

Table 3: Top Differential Adjectives by Model and Prompting Condition



(a) Generated from a White-origin narrative (CR# 16-057041).



(b) Generated from a Black-origin narrative (CR# 18-214755).



(c) Generated from a White-origin narrative (CR# 19-158468).

Figure 3: Synthetic mugshot-style images generated by GPT-5 from fully redacted narratives. No race label appears in any prompt; visible variation reflects residual lexical signal in the source text.

Comparison	Statistic	Significant?
Word count	$t(199) = -14.65$, $p < .0001$, $d = 1.31$	Yes
Force kws	$t(199) = -6.60$, $p < .0001$, $d = 0.65$	Yes
Sentiment	$t(199) = 15.36$, $p < .0001$, $d = 1.36$	Yes

Table 4: Significance summary for RQ2 for counterfactual generations.

with racial bias cues.

4.3.1 Character Sketch Generation

We prompted GPT-5.2 to produce a character sketch (personality, likely occupation, likely hobby, general characterization) for 987 records (493 Black-origin, 494 White-origin, balanced by year) based solely on the redacted narrative. Sentiment is negative for both groups—73.0% of Black-origin and 77.9% of White-origin sketches score negative on VADER—because any sketch drawn from a use-of-force narrative describes someone in a moment of conflict. The group means differ only

slightly, and in the direction opposite to a stereotyping account (Black -0.425 , White -0.509 ; $t(985) = 2.07$, $p = 0.039$, $d = 0.13$). Inferred occupations are similarly uniform: manual labor and unemployment together account for 75.1% of Black-origin and 69.8% of White-origin sketches, and the six-category distribution does not differ significantly ($\chi^2(5) = 10.49$, $p = 0.062$, $V = 0.10$). This uniformity reflects the socioeconomic profile of use-of-force records generally, not differential racial stereotyping by the model.

4.3.2 Image Generation and Race inference

A note on intent. This subsection measures skin tone in machine-generated images, and that method has an obvious surface resemblance to physiognomy. We state the intent plainly here rather than only in Section 7: we are not inferring character or criminality from appearance, and we make no claim that appearance carries such information. We measure the rendered skin tone of synthetic images purely as a read-out of what a generative system did with a text prompt that contained no race label.

The unit of analysis is the model, not any person. ITA (Del Bino et al., 2006) is used because it is the established colorimetric standard in the dermatology literature for placing a skin sample on a light-to-dark axis, which makes it a stable and independently interpretable baseline against which to check a learned classifier.

We generated synthetic mugshot-style images from race-redacted narratives using GPT-5. From 987 eligible records (493 Black-origin, 494 White-origin), GPT-5 produced 875 images within three retries (446 Black-origin, 429 White-origin); Gemini-2.5-Flash refused 91.0% of prompts (89 of 987 succeeded) under default safety settings and is excluded from quantitative analysis.

We first measured skin tone directly with ITA, computed over a central face patch under a standard YCbCr skin-pixel gate. Images generated from Black-origin narratives are rendered significantly darker than those from White-origin narratives (mean -8.75 ± 15.62 against -6.35 ± 13.90 degrees; $t(873) = -2.40$, $p = 0.017$, $d = -0.16$), where lower ITA indicates darker skin. The effect is small, which is what a purely colorimetric read-out of stylised pencil-sketch renders should be expected to give. We therefore complement it with CLIP-based classification. To establish a baseline, we sampled 1,000 real images each for Black and White subjects from FairFace (Karkkainen and Joo, 2021) and scored them using a continuous skin-tone prompt ranging from 1 (very dark) to 5 (very light). The resulting distributions were statistically separable ($p < 0.001$) along expected racial lines, confirming that the approach is discriminative on real photographs. Appendix contains figures (8, 9) showing clear distinction in CLIP-based racial analysis. We then applied the same CLIP-ViT-base (Radford et al., 2021) pipeline with zero-shot race labels to all 875 generated images.

For White-origin records, CLIP predicts “White” in 59.4% of cases; for Black-origin records, only 44.0%. The 15.4 percentage-point gap is statistically significant (Table 5), and it agrees with the ITA measurement in both direction and magnitude ($\phi = 0.16$ against $d = -0.16$), so two independent read-outs of the same images converge.

No image in this pipeline contains an explicit race label. Yet images derived from Black-origin narratives are classified as depicting non-White individuals at a reliably higher rate. The most plausible mechanism is the vocabulary contrast from

Origin (n)	Black	Hispanic	Unknown	White
Black (446)	16.6%	33.0%	6.5%	44.0%
White (429)	7.2%	26.1%	7.2%	59.4%

Table 5: CLIP race predictions from AI-generated images ($n=875$).

Section 4.1.1: a prompt saturated with *suspect*, *waistband*, and *chase* activates different visual templates than one built from *ambulance*, *suicidal*, and *bridge*. The scenarios are themselves racialized in the model’s learned associations, and the images inherit that signal without any explicit cue. Text-level redaction is not sufficient to prevent racial signal from propagating into downstream visual outputs.

Comparison	Statistic	Significant?
CLIP prediction gap	$\chi^2(1) = 21.0$, $p < .001$, $\phi = 0.16$	Yes
ITA skin tone	$t(873) = -2.40$, $p = 0.017$, $d = -0.16$	Yes
Character sketch sentiment	$t(985) = 2.07$, $p = 0.039$, $d = 0.13$	Marginal
Occupation distribution	$\chi^2(5) = 10.49$, $p = 0.062$, $V = 0.10$	No

Table 6: Significance summary for RQ3.

5 Conclusion

We presented RPD-UoF, a corpus of 3,091 officer-authored use-of-force narratives with verified race labels, and used it to trace racial signals across three stages of an AI pipeline. We observe that (1) redacted narratives are separable by subject race at above-chance accuracy; (2) LLMs conditioned on race produce structurally different narratives for White and Black profiles from the same seed; and (3) images generated from race-redacted narratives are classified by CLIP at significantly different rates depending on the racial origin of the source text. In all three stages, the signal is spread across vocabulary, scenario templates, and the learned associations of generative models.

A core finding for AI safety is that removing sensitive attributes from inputs does not prevent them from influencing outputs. Racial signal persists through correlated language and can be amplified when systems are chained across modalities, leading to new failure modes. Our results do not suggest intentional bias by individual officers or models; they show how administrative language and generative systems can reproduce structural biases without explicit cues.

6 Limitations

The dataset covers 2015–2019 consistently, with a separately obtained 2022 batch contributing 176 records. The 2022 batch is retained to extend temporal coverage but is not fully comparable to the main corpus in terms of departmental policy context; analyses stratified by year should treat it separately. The broader scope limitation is single-city coverage: all records come from one mid-sized northeastern department, and whether the lexical and structural patterns documented here generalize to other jurisdictions is an open question.

Another limitation is selection bias. These are only incidents that officers deemed serious enough to file formal use-of-force reports. Lower-level physical contact that officers do not officially document will not appear in our dataset. We also cannot verify narrative accuracy since we only have the officer’s written account, not video or civilian testimony.

On the modeling side, classifier accuracy is a conservative lower bound on the racial signal present. A model that fails to exceed chance does not prove the absence of signal; it may simply lack the capacity or training data to recover it. Separately, standardizing subject pronouns to they/them before classification removes gender signal as intended, but may also remove correlated narrative features, potentially deflating classifier performance.

7 Ethical Considerations

New York State FOIL data is public record; the documents we received had sensitive subject information redacted prior to release. We applied an additional redaction pass removing officer names, badge numbers, and exact addresses to protect officer privacy. We see no privacy concern with the released dataset and believe it addresses a real resource gap in computational work on policing.

Our separability results depend on fine-tuned classifiers, and we acknowledge that language models carry biases from pretraining that may influence downstream tasks (Bender et al., 2021). We treat our classifier results as evidence of distributional signal in the corpus, not as a deployment-ready race inference tool, and we complement them with log-odds analysis that does not rely on model internals.

The multimodal experiments require a direct

statement of intent. We are not reviving physiognomy (Lombroso, 1911). The point of the image generation experiment is the opposite: to show that a system with no access to race labels can nonetheless produce racially differentiated outputs because the *text itself* carries the signal. This is a failure mode, not a capability to be used. As generative systems become more capable, these failure modes become harder to anticipate (Dutta et al., 2024; Betley et al., 2025). Recent incidents – such as a commercial chatbot complying with requests to generate non-consensual intimate imagery (Hayes et al., 2026) – illustrate that the sociotechnical risk surface expands faster than governance frameworks. Criminal justice applications are particularly exposed to this dynamic, given the stakes attached to decisions made from institutional records (Basu et al., 2025; Gebru et al., 2021)

Our annotators were graduate research assistants compensated through stipend or course credit and worked under a written codebook with researcher supervision. We used Claude and ChatGPT for code development, experimental infrastructure, and polishing prose during the writing process. All scientific decisions, interpretations, and core arguments were made by the authors.

Acknowledgements

We acknowledge the Upstate New York Police Research Consortium (UNY-PRC) for ongoing institutional support. We extend our sincere thanks to UNY-PRC research interns Lauren Arena, Delaney Baker, Carissa Bayack, Frankie Carlin, Josh Fishkin, Olivia Gerhardt, and Sophia McSmith, alongside community collaborators Kim Rachunok and Ayonna Titus, for their painstaking labor in manually extracting, coding, standardizing, and cleaning the underlying use-of-force dataset. We thank Soumyajit Datta for his input. Dutta and KhudaBukhsh were partly supported by a gift from Lenovo.

References

- Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2022. Machine bias. In *Ethics of data and analytics*, pages 254–264. Auerbach Publications.
- Sagnik Basu, Shubham Prakash, Ashish Maruti Barge, Siddharth D Jaiswal, Abhisek Dash, Saptarshi Ghosh, and Animesh Mukherjee. 2025. *Judging by appearances? auditing and intervening vision-*

- language models for bail prediction. *Preprint*, arXiv:2510.00088.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Jan Betley, Jorio Cocola, Dylan Feng, James Chua, Andy Ardit, Anna Szyber-Betley, and Owain Evans. 2025. Weird generalization and inductive backdoors: New ways to corrupt llms. *arXiv preprint arXiv:2512.09742*.
- Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. 2023. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. In *Proceedings of the 2023 ACM conference on fairness, accountability, and transparency*, pages 1493–1504.
- Su Lin Blodgett, Solon Barocas, Hal Daumé Iii, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of “bias” in nlp. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 5454–5476.
- Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91. PMLR.
- Michele Caianiello. 2021. Dangerous liaisons, potentialities and risks deriving from the interaction between artificial intelligence and preventive justice. *European Journal of Crime, Criminal Law and Criminal Justice*, 29(1):1–23.
- Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- Malcolm Coulthard, Alison Johnson, and David Wright. 2007. *An introduction to forensic linguistics: Language in evidence*, volume 6. Routledge London.
- Graham A. Cutting and Anne-Francoise Cutting-Decelle. 2021. [Intelligent document processing – methods and tools in the real world](#). *Preprint*, arXiv:2112.14070.
- S Del Bino, J Sok, E Bessac, and F Bernerd. 2006. Relationship between skin response to ultraviolet exposure and skin color type. *Pigment cell research*, 19(6):606–614.
- Carl Dement and Melissa Inglis. 2024. Artificial intelligence-assisted criminal justice reporting: An exploratory study of benefits, concerns, and future directions. *Criminology & Criminal Justice*, page 17488958241274296.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. Qlora: Efficient finetuning of quantized llms. *Advances in neural information processing systems*, 36:10088–10115.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pages 4171–4186.
- Arka Dutta, Adel Khorramrouz, Sujana Dutta, and Ashiqur R KhudaBukhsh. 2024. Down the toxicity rabbit hole: A framework to bias audit large language models with key emphasis on racism, antisemitism, and misogyny. In *IJCAI*, pages 7242–7250.
- Sujana Dutta, Beibei Li, Daniel S Nagin, and Ashiqur R KhudaBukhsh. 2022. A murder and protests, the capitol riot, and the chauvin trial: Estimating disparate news media stance. In *IJCAI*, pages 5059–5065.
- Jennifer L Eberhardt, Phillip Atiba Goff, Valerie J Purdie, and Paul G Davies. 2004. Seeing black: race, crime, and visual processing. *Journal of personality and social psychology*, 87(6):876.
- Frank Edwards, Hedwig Lee, and Michael Esposito. 2019. Risk of being killed by police use of force in the united states by age, race–ethnicity, and sex. *Proceedings of the national academy of sciences*, 116(34):16793–16798.
- Biaoyan Fang, Trevor Cohn, Timothy Baldwin, and Lea Frermann. 2023. [More than votes? voting and language based partisanship in the US Supreme Court](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4604–4614, Singapore. Association for Computational Linguistics.
- Xiao Fang, Shangkun Che, Minjia Mao, Hongzhe Zhang, Ming Zhao, and Xiaohang Zhao. 2024. Bias of ai-generated content: an examination of news produced by large language models. *Scientific reports*, 14(1):5224.
- Andrew Guthrie Ferguson. 2025. Generative suspicion and the risks of ai-assisted police reports. *Nw. UL Rev.*, 120:299.
- Anjalie Field, Prateek Verma, Nay San, Jennifer L Eberhardt, and Dan Jurafsky. 2023. Developing speech processing pipelines for police accountability. *arXiv preprint arXiv:2306.06086*.
- Ted Forsyth. 2026. Draft description of foil acquisition and subject resistance reports (srrs). Unpublished manuscript / internal draft (RIT), dated 2026-01-18.

- Roland G Fryer Jr. 2019. An empirical analysis of racial differences in police use of force. *Journal of Political Economy*, 127(3):1210–1261.
- Caroline Gans-Combe. 2022. Automated justice: Issues, benefits and risks in the use of artificial intelligence and its algorithms in access to justice and law enforcement. *Ethics, integrity and policymaking: the value of the case study*, pages 175–194.
- Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92.
- Andrew Gelman, Jeffrey Fagan, and Alex Kiss. 2007. An analysis of the new york city police department’s “stop-and-frisk” policy in the context of claims of racial bias. *Journal of the American statistical association*, 102(479):813–823.
- John S Goldkamp. 1976. Minorities as victims of police shootings: Interpretations of racial disproportionality and police use of deadly force. *The Justice System Journal*, pages 169–183.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Christal Hayes, Osmond Chia, and Liv McMahon. 2026. [Elon musk’s x to block grok from undressing images of real people](#). Accessed: 2026-03-16.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*.
- Valentin Hofmann, Pratyusha Ria Kalluri, Dan Jurafsky, and Sharese King. 2024. Ai generates covertly racist decisions about people based on their dialect. *Nature*, 633(8028):147–154.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Liang Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3.
- Clayton Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media*, volume 8, pages 216–225.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. *Mistral 7b*. *Preprint*, arXiv:2310.06825.
- Kimmo Karkkainen and Jungseock Joo. 2021. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1548–1558.
- Dean Knox, Will Lowe, and Jonathan Mummolo. 2020. Administrative records mask racially biased policing. *American Political Science Review*, 114(3):619–637.
- VI Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet Physics Doklady*, volume 10.
- Gina Lombroso. 1911. *Criminal man, according to the classification of Cesare Lombroso*, volume 27. Putnam.
- Burt L Monroe, Michael P Colaresi, and Kevin M Quinn. 2008. Fightin’ words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4):372–403.
- Khalil Gibran Muhammad. 2019. The condemnation of blackness: Race, crime, and the making of modern urban america, with a new preface. In *The Condemnation of Blackness*. Harvard University Press.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. Stereoset: Measuring stereotypical bias in pretrained language models. In *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (volume 1: long papers)*, pages 5356–5371.
- Pranav Narayanan Venkit, Christopher Graziul, Miranda Ardith Goodman, Samantha Nicole Kenny, and Shomir Wilson. 2024. Race and privacy in broadcast police communications. *Proceedings of the ACM on Human-computer Interaction*, 8(CSCW2):1–26.
- New York State Committee on Open Government. a. [Explanation of time limits for response](#). Open Government (NY). Accessed 2026-01-19.
- New York State Committee on Open Government. b. [Freedom of information law \(foil\)](#). Open Government (NY). Accessed 2026-01-19.
- New York State Department of State. [Freedom of information law \(public officers law, article 6, sections 84–90\)](#). Committee on Open Government, NY Department of State. Accessed 2026-01-19.
- Linke Ouyang, Yuan Qu, Hongbin Zhou, Jiawei Zhu, Rui Zhang, Qunshu Lin, Bin Wang, Zhiyuan Zhao, Man Jiang, Xiaomeng Zhao, and 1 others. 2025. Omnidocbench: Benchmarking diverse pdf document parsing with comprehensive annotations. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24838–24848.
- Emma Pierson, Camelia Simoiu, Jan Overgoor, Sam Corbett-Davies, Daniel Jenson, Amy Shoemaker, Vignesh Ramachandran, Phoebe Barghouty, Cheryl

- Phillips, Ravi Shroff, and 1 others. 2020. A large-scale analysis of racial disparities in police stops across the united states. *Nature human behaviour*, 4(7):736–745.
- Vinodkumar Prabhakaran, Camilla Griffiths, Hang Su, Prateek Verma, Nelson Morgan, Jennifer L Eberhardt, and Dan Jurafsky. 2018. Detecting institutional dialog acts in police traffic stops. *Transactions of the Association for Computational Linguistics*, 6:467–481.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR.
- Inioluwa Deborah Raji and Joy Buolamwini. 2019. Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial ai products. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 429–435.
- Eugenia H Rho, Maggie Harrington, Yuyang Zhong, Reid Pryzant, Nicholas P Camp, Dan Jurafsky, and Jennifer L Eberhardt. 2023. Escalated police stops of black men are linguistically and psychologically distinct in their earliest moments. *Proceedings of the National Academy of Sciences*, 120(23):e2216162120.
- Rashida Richardson, Jason M Schultz, and Kate Crawford. 2019. Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *NYUL Rev. Online*, 94:15.
- Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 59–68.
- Evan Shieh, Faye-Marie Vassel, Cassidy R Sugimoto, and Thema Monroe-White. 2026. Intersectional biases in narratives produced by open-ended prompting of generative language models. *Nature Communications*, 17(1):1243.
- Robin Staab, Mark Vero, Mislav Balunovic, and Martin Vechev. 2024. Beyond memorization: Violating privacy via inference with large language models. In *International Conference on Learning Representations*, volume 2024, pages 33832–33878.
- Supriti Vijay, Aman Priyanshu, and Ashiqur R KhudaBukhsh. 2025. When neutral summaries are not that neutral: Quantifying political neutrality in llm-generated news summaries (student abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 29514–29516.
- Rob Voigt, Nicholas P Camp, Vinodkumar Prabhakaran, William L Hamilton, Rebecca C Hetey, Camilla M Griffiths, David Jurgens, Dan Jurafsky, and Jennifer L Eberhardt. 2017. Language from police body camera footage shows racial disparities in officer respect. *Proceedings of the national Academy of sciences*, 114(25):6521–6526.
- Loïc Wacquant. 2020. *Punishing the poor: The neoliberal government of social insecurity*. duke university Press.

A Appendix

A.1 RPD-UoF additional details

A.1.1 Demographics and Racial Distribution.

Stage	2015–19	2022
Raw FOIL release (SRR records)	5,297	1,698
Scanned PDF documents	3,453	12 monthly
After dedup. by CR# + Gemini OCR	3,160	—
– unknown race in admin. data	–86	—
– Hispanic ($n=151$), Asian ($n=8$)	–159	—
Subtotal	2,915	176
Final analysis corpus		3,091
Black subjects	2,277 (73.7%)	
White subjects	814 (26.3%)	

Table 7: From raw FOIL release to analysis corpus. The 5,297 raw records include duplicate filings and OCR failures; the 1,698 figure for 2022 is the raw extracted count before race-label matching.

We verified demographics by matching case numbers to official departmental records, as we require actual race labels from administrative data rather than officer perceptions written in the narratives themselves. Table 7 gives the full accounting from the 5,297 raw SRR records to the 3,091-record analysis corpus.

The racial breakdown shows a stark pattern. Among the 3,160 successfully extracted 2015–2019 records, 2,137 subjects are Black (67.6%), 777 are White (24.6%), 151 are Hispanic (4.8%), 8 are Asian (0.3%), and 87 (2.8%) have no race recorded in administrative data. Restricted to the Black/White analysis corpus, 2,277 subjects are Black (73.7%) and 814 are White (26.3%). Black subjects therefore appear in use-of-force reports at nearly three times the rate of White subjects. This 2.8:1 disparity matches national patterns documented in other policing research, but it also creates substantial class imbalance that affects our modeling choices.

For gender, 74.1% of records in the analysis corpus list the subject as male, 25.4% as female, and 0.6% have no gender information. Figure 1 shows the demographic details.

A.1.2 Narrative structure and statistics:

Narratives average **522 words** ($\sigma = 356$, median = 421, range 0–3,668; nine records fall below fifteen words and are retained but flagged in the release). The distribution is right-skewed: most reports are moderate in length, with a minority of multi-addenda records. Taking the primary recorded value of the structured `force_used` field, the most common force types are Verbal Commands ($n=797$), Physical Force ($n=437$), OC Spray ($n=129$), and Ground Stabilization ($n=50$); the field is free text and a substantial share of records carry compound or non-standard entries, so these counts are lower bounds on each category. For subject resistance, officers most often record Avoiding Custody ($n=524$), Verbal Resistance ($n=260$), and Assaultive behavior ($n=59$). These structured fields let us relate narrative framing to formally recorded incident severity, and we use them in Section 4.1.1 to check that the lexical contrast is not an artifact of incident mix.

A.1.3 IDP quality:

Figure 4 shows the efficacy of LLM labelling by comparing it against the adjudicated human annotation on the 750-record verification sample: the pipeline agrees on 87.7% of race labels and 91.1% of gender labels. Disagreements concentrate in non-standard race abbreviations and in records whose scan quality is poor enough that the human annotators themselves disagreed; those records are flagged in the release.




Figure 4: LLM vs. human annotation agreement for race and gender labels.

Figure 5 shows the completeness of intelligent document processing pipeline using Gemini model based on different fields cross-checked with human annotation.




Figure 5: Field completeness of RPD-UoF IDP using Gemini vision model

A.2 Human Annotation details

Annotation was conducted by two graduate students working from scanned PDFs of the original case reports. Both annotators were supported by research assistantships and received no additional per-record compensation. Annotation was supervised by an expert social scientist with prior experience in police administrative data research.

Variables coded for each record include subject demographics (sex, race, ethnicity), incident meta-data (date, time, location, arrest status, charges), subject resistance type (verbal, passive, active, armed), officer force type (verbal, hands-on, OC, Taser, bean bag, firearm), and medical outcomes (injury, hospitalization). Records where the human-override label differs from the pipeline extraction are flagged in the released dataset; we recommend downstream users exclude these from race-stratified analyses.

A.3 Hyperparameter details

All models are trained on a balanced corpus of 1,408 narratives (704 per class, downsampled from the original 2,656-record Black/White subset), split 80/20 into train and test sets using stratified sampling. Encoder models (DeBERTa-v3-small, BERT-base-uncased) are fine-tuned for 5 epochs with a learning rate of 3×10^{-5} , batch size 16, and early stopping (patience = 3) against validation balanced accuracy. Training is conducted in FP32 throughout; FP16 mixed-precision is disabled for DeBERTa-v3-small due to numerical instability in its disentangled attention mechanism. Large decoder models (Mistral-7B-v0.1, LLaMA-3-8B) are fine-tuned for 3 epochs using QLoRA (Dettmers et al., 2023) with 4-bit NF4 quantization, targeting the query and value projection layers ($r=16$, $\alpha=32$, $\approx 0.096\%$ of parameters trainable). An effective batch size of 16 is achieved via gradient accumulation (per-device batch size of 1, 16 accumulation steps). All models are trained with three

random seeds (42, 123, 999), and the best check-point per seed is selected by validation balanced accuracy. Experiments are run on four NVIDIA A6000 (48 GB) GPUs.

A.4 Additional prompts

P1. ZERO-SHOT NARRATIVE GENERATION PROMPT

You are tasked with generating realistic police use-of-force incident narratives. Below is an example of how such narratives are typically written:

Example Narrative:

example_narrative

Now generate a realistic police use-of-force incident narrative for a subject with the following profile:

- Race: target_tuple['race']
- Sex: target_tuple['sex']
- Age: target_tuple['age']

Write a detailed and realistic narrative describing the incident, including the situation context, officer response, subject behavior, and any force used. The narrative should be between 150 and 300 words. Do not explicitly mention the subject's race, sex, or age in the narrative.

Narrative:

P2. COUNTERFACT. GEN. PROMPT (ZERO-SHOT)

You are a police officer writing an official use-of-force report. The subject is a {race} {gender} in their {age_range}.

Continue this incident narrative in standard police report style (3–5 sentences):

{seed_text}

Write the continuation as the responding officer. Be specific about actions taken.

P2B. COUNTERFACT. GEN. PROMPT (FEW-SHOT)

You are a police officer writing an official use-of-force report. The subject is a {race} {gender} in their {age_range}.

Here are example narrative continuations:

Seed: {example_seed_i}
Continuation:
{example_continuation_i}
(three fixed exemplars, identical for both race conditions)

Now continue this narrative in the same style (3–5 sentences):

Seed: {seed_text}
Continuation:

P3. MUGSHOT IMAGE GENERATION PROMPT

Police composite sketch style portrait of a {gender} person.

Style: Realistic pencil-sketch, frontal view, head and shoulders only.

Character Description:

{sketch_text[:300]}

Constraints:

- Neutral background
- No text or watermarks

P4. SUBJECT CHARACTER INFERENCE PROMPT

Based on this police incident narrative, create a brief character description of the subject.

The subject's race and personal identifiers have been redacted — base your description **only** on behaviors, context, and incident details.

Narrative:

{narrative}

Return JSON in the following format:

```
{
  "personality": 2–3 sentences about
    likely personality traits inferred from be-
    havior in the narrative,
  "likely_job": most likely occupation
    based on context (time of day, location, de-
    meanor),
  "likely_hobby": one or two likely hob-
    bies,
  "general_blurb": 2–3 sentence overall
    characterization of this person
}
```

A.5 Additional Results

A.5.1 RQ1 additional results

Figure 6 shows the log-odds lexical difference between Black-subject and White-subject narratives in RPD-UoF. Explicit race terms, officer and subject names, and geographic proper nouns are removed before estimation, so every term shown is ordinary incident vocabulary.

A.5.2 RQ2 additional results

Table 8 shows a summary of the classifier results on LLM-generated narratives based on template.

Classifier LLM	BA	F1	W-R	B-R	$\Delta P(B)$	p	q
<i>DeBERTa-v3-small</i>							
Claude Sonnet	.495	.494	.452	.538	+0.001	.860	.948
DeepSeek-V3	.500	.493	.614	.386	+0.002	.768	.948
Gemini Flash	.512	.494	.703	.322	+0.008	.168	.335
Gemma3-27b	.493	.474	.300	.686	−0.008	.163	.335
GPT-5.2	.513	.511	.452	.574	+0.011	.075	.299
LLaMA-3.1-70b	.482	.476	.593	.371	+0.000	.948	.948
Mistral-Large	.496	.495	.447	.545	+0.010	.040	.299
Mistral-Small-3.2	.495	.483	.648	.341	−0.008	.279	.446
<i>BERT-base-uncased</i>							
Claude Sonnet	.466	.456	.332	.600	− 0.042	7.1×10^{-6}	5.7×10^{-5} ***
DeepSeek-V3	.524	.523	.536	.511	+0.019	.070	.139
Gemini Flash	.497	.490	.615	.380	+0.004	.675	.675
Gemma3-27b	.516	.511	.408	.624	+0.007	.473	.631
GPT-5.2	.461	.456	.364	.559	− 0.040	4.0×10^{-5}	1.6×10^{-4} ***
LLaMA-3.1-70b	.437	.382	.737	.138	− 0.023	1.2×10^{-2}	3.2×10^{-2} *
Mistral-Large	.520	.520	.498	.542	+0.013	.137	.220
Mistral-Small-3.2	.504	.467	.240	.767	+0.006	.563	.643

Table 8: Classifier-based bias audit of LLM-generated narratives. Each row reports balanced accuracy (BA), macro-F1, per-class recall (W-R = White, B-R = Black), and the mean within-pair $\Delta P(B)$ gap ($P(\text{Black}|\text{Black-conditioned}) - P(\text{Black}|\text{White-conditioned})$). p -values are from paired t -tests; q -values are FDR-corrected (Benjamini-Hochberg). Significance: * $q < 0.05$, ** $q < 0.01$, *** $q < 0.001$.




Figure 6: Log-odds lexical difference between Black-subject and White-subject narratives in RPD-UoF, after removing race terms, person names, and geographic proper nouns.

A.5.3 RQ3 Additional Results

Figure 7 shows the distribution of Individual Typology Angle across the 875 generated images, split by the racial origin of the source narrative. Images generated from Black-origin narratives are rendered darker (lower ITA). The distributions overlap heavily, which is expected for stylised pencil-sketch renders, but the shift is consistent and agrees in sign and magnitude with the CLIP result in Section 4.3.2.




Figure 7: Individual Typology Angle of GPT-5-generated images by racial origin of the source narrative (lower = darker). Dashed lines mark group means.

On the FairFace validation set, CLIP skin-tone scores separate cleanly by race group: mean 2.24 for Black subjects, 2.81 for Latino, and 3.76 for White ($p < 0.001$). This confirms that the CLIP-based scoring pipeline is discriminative on real pho-

tographs, establishing the baseline against which generated-image results are interpreted.




Figure 8: Mean CLIP skin-tone scores for Black and White subjects in the FairFace validation set (1 = very dark, 5 = very light). The 1.52-point gap confirms that the scoring approach is discriminative on real photographs.




Figure 9: Distribution of CLIP skin-tone scores across Black, Latino, and White groups in FairFace. Distributions are well-separated and monotonically ordered, validating the prompt-based scoring scale prior to application on generated images.