

# To Protect and To Serve? Analyzing Entity-Centric Framing of Police Violence

Caleb Ziems

School of Interactive Computing  
Georgia Institute of Technology  
cziems@gatech.edu

Diyi Yang

School of Interactive Computing  
Georgia Institute of Technology  
dyang888@gatech.edu

## Abstract

Framing has significant but subtle effects on public opinion and policy. We propose an NLP framework to measure entity-centric frames. We use it to understand media coverage on police violence in the United States in a new POLICE VIOLENCE FRAME CORPUS of 82k news articles spanning 7k police killings. Our work uncovers more than a dozen framing devices and reveals significant differences in the way liberal and conservative news sources frame both the *issue* of police violence and the *entities* involved. Conservative sources emphasize when the victim is armed or attacking an officer and are more likely to mention the victim’s criminal record. Liberal sources focus more on the underlying systemic injustice, highlighting the victim’s race and that they were unarmed. We discover temporary spikes in these injustice frames near high-profile shooting events, and finally, we show protest volume correlates with and precedes media framing decisions. <sup>1</sup>

## 1 Introduction

The normative standard in American journalism is for the news to be neutral and objective, especially regarding politically charged events (Schudson, 2001). Despite this expectation, journalists are unable to report on all of an event’s details simultaneously. By choosing to include or exclude details, or by highlighting salient details in a particular order, journalists unavoidably induce a preferred interpretation among readers (Iyengar, 1990). This selective presentation is called **framing** (Entman, 2007). Framing influences the way people think by “telling them what to think about” (Entman, 2010). In this way, frames impact both public opinion (Chong and Druckman, 2007; Iyengar, 1990; McCombs, 2002; Price et al., 2005; Rugg, 1941;

<sup>1</sup>Data and code available at: <https://github.com/GT-SALT/framing-police-violence>

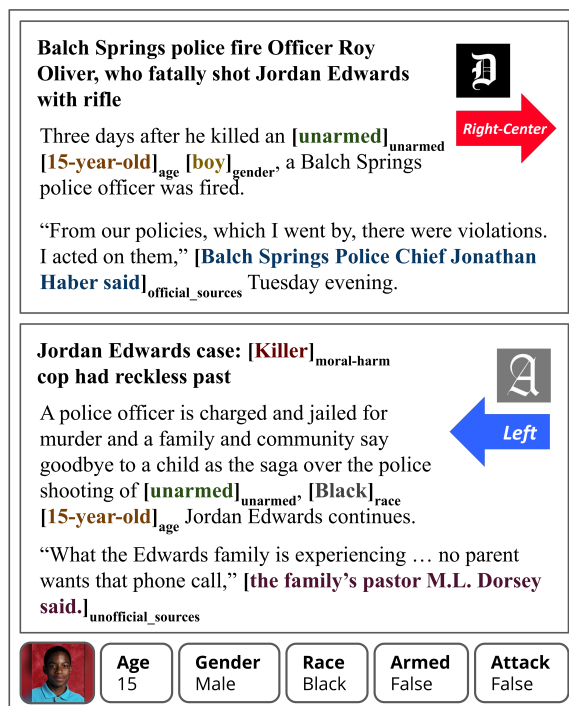


Figure 1: **Framing the murder of Jordan Edwards.** Our system automatically identifies key details or *frames* that shape a reader’s understanding of the shooting. Importantly, we can distinguish the victim’s attributes from the descriptions of the officer, like *killer* in “killer cop.” Only the left-leaning article uses this morally-weighted term, *killer*, and also takes care to mention the victim’s *race*. While the left-leaning article highlights a quote from the Edwards’ pastor, an *unofficial source*, the right-center article cites only *official sources* (namely the Chief of Police). Both mention the victim’s *age* and *unarmed* status.

Schuldt et al., 2011) and policy decisions (Baumgartner et al., 2008; Dardis et al., 2008).

Prior work has revealed an abundance of politically effective framing devices (Bryan et al., 2011; Gentzkow and Shapiro, 2010; Price et al., 2005; Rugg, 1941; Schuldt et al., 2011), some of which have been operationalized and measured at scale using methods from NLP (Card et al., 2015; Demszky et al., 2019; Field et al., 2018; Greene and Resnik, 2009; Recasens et al., 2013; Tsur et al., 2015). While these works extensively cover *issue*

frames in broad topics of political debate (e.g. *immigration*), they overlook a wide array of *entity* frames (how an individual is represented; e.g. a particular undocumented worker described as lazy), and these can have huge policy implications for target populations (Schneider and Ingram, 1993).

In this paper, we introduce an NLP framework to understand *entity framing* and its relation to *issue framing* in political news. As a case study, we consider news coverage of police violence. Though we choose this domain for the stark contrast between two readily-discernible entities (police and victim), our framing measures can also be applied to other major social issues (Luo et al., 2020b; Mendelsohn et al., 2021), and salient entities involved in these events, like protesters, politicians, migrants, etc.

We make several novel contributions. First, we introduce the POLICE VIOLENCE FRAME CORPUS that contains 82k news articles on over 7k police shooting incidents. See Figure 1 for example articles with annotated frames. Next, we build a set of syntax-aware methods for extracting 15 issue and entity frames, implemented using entity co-reference and the syntactic dependency parse. Unlike bag-of-words methods (e.g. topic modeling) our **entity-centric methods** can distinguish between a white *man* and a white *car*. In this example, we identify race frames by scanning the attributive and predicative adjectives of any VICTIM tokens. Such distinctions can be crucial, especially in a domain where *officer* aggression will have different ramifications than aggression from a *suspected criminal*. By exact string-matching, we can also extract, for the first time, differences in the **order that frames appear** within each document.

We find that liberal sources discuss *race* and *systemic racism* much earlier, which can prime readers to interpret all other frames through the lens of injustice. Furthermore, we quantify and statistically confirm what smaller-scale content analyses in the social sciences have previously shown (Drakulich et al., 2020; Fridkin et al., 2017; Lawrence, 2000), that conservative sources highlight law-and-order and focus on the victim’s criminal record or their harm or resistance towards the officer, which could justify police conduct. Finally, we rigorously examine the broader interactions between media framing and offline events. We find that high-profile shootings are correlated with an increase in *systemic* and *racial* framing, and that increased protest activity *Granger-causes* or precedes media attention

towards the victim’s race and unarmed status.

## 2 Related Work

A large body of related work in NLP focuses on detecting stance, ideology, or political leaning (Baly et al., 2020; Bamman and Smith, 2015; Iyyer et al., 2014; Johnson et al., 2017; PreoŃiuc-Pietro et al., 2017; Luo et al., 2020a; Stefanov et al., 2020). While we show a relationship between framing and political leaning, we argue that frames are often more subtle than overt expressions of stance, and cognitively more salient than other stylistic differences in the language of political actors, thus more challenging to be measured.

Specifically, we distinguish between *issue* frames (Iyengar, 1990) and *entity* frames (van den Berg et al., 2020). Entity frames are descriptions of individuals that can shape a reader’s ideas about a broader issue. The entity frames of interest here are the victim’s age, gender, race, criminality, mental illness, and attacking/fleeing/unarmed status. One related work found a *shooter identity* cluster in their topic model that contained broad descriptors like “crazy” (Demszky et al., 2019). However, their bag-of-words method would not differentiate a crazy *shooter* from a crazy *situation*. To follow up, there is need for a *syntax-aware* analysis.

In a systematic study of issue framing, Card et al. (2015) applied the Policy Frames Codebook of Boydston et al. (2013) to build the Media Frames Corpus (MFC). They annotated spans of text from discussions on tobacco, same-sex marriage, and immigration policy with broad meta-topic framing labels like *health and safety*. Field et al. (2018) built lexicons from the MFC annotations to classify issue frames in Russian news, and Roy and Goldwasser (2020) extended this work with sub-frame lexicons to refine the broad categories of the MFC. Some have considered the way Moral Foundations (Haidt and Graham, 2007) can serve as issue frames (Kwak et al., 2020; Mokhberian et al., 2020; Priniski et al., 2021), and others have built issue-specific typologies (Mendelsohn et al., 2021). While issue framing has been well-studied (Ajjour et al., 2019; Baumer et al., 2015), entity framing remains *under-examined* in NLP with a few exceptions. One line of work used an unsupervised approach to identify *personas* or clusters of co-occurring verbs and adjective modifiers (Card et al., 2016; Bamman et al., 2013; Iyyer et al., 2016; Joseph et al., 2017). Another line of work



Figure 2: Construction process of POLICE VIOLENCE FRAME CORPUS.

combined psychology lexicons with distributional methods to measure implicit differences in power, sentiment, and agency attributed to male and female entities in news and film (Sap et al., 2017; Field and Tsvetkov, 2019; Field et al., 2019).

Social scientists have experimentally manipulated framing devices related to police violence, including *law and order*, *police brutality*, or *racial stereotypes*, revealing dramatic effects on participants’ perceptions of police shootings (Fridell, 2017; Dukes and Gaither, 2017; Porter et al., 2018). Criminologists have trained coders to manually annotate on the order of 100 news articles for framing devices relevant to police use of force (Hirschfield and Simon, 2010; Ash et al., 2019). While these studies provide great theoretical insight, their manual coding schemes and small corpora are not suited for large scale real-time analysis of news reports nationwide. While many recent works in NLP have started to detect police shootings (Nguyen and Nguyen, 2018; Keith et al., 2017) and other gun violence (Pavlick et al., 2016), we are the first to model entity-centric framing around police violence.

### 3 POLICE VIOLENCE FRAME CORPUS

To study media framing of police violence, we introduce POLICE VIOLENCE FRAME CORPUS (PVFC) which contains over 82,000 news reports of police shooting events. We now describe the corpus construction as it is shown in Figure 2.

#### 3.1 Identifying Shooting Events

We first use Mapping Police Violence (Sinyangwe et al., 2021) to identify shooting events. It is a representative, reliable, and detailed record, as it cross-references the three most complete databases available: Fatal Encounters (Burghart, 2020), the U.S. Police Shootings Database (Tate et al., 2021), and Killed by Police, all of which have been validated by the Bureau of Justice Statistics (Banks et al., 2016). Prior works in sociology and criminology (Gray and Parker, 2020) use it as an alterna-

tive to official police reports because local police departments significantly underreport shootings (Williams et al., 2019). At the time of our retrieval, the Mapping Police Violence dataset identified 8,169 named victims of police shootings between January 1, 2013 and September 4, 2020 and provided the victim’s age, gender, and race, whether they fled or attacked the officer, and whether the victim had a known mental illness or was armed (and with what weapon), as well as the location and date of the shooting, the agency responsible, and whether the incident was recorded on video.

#### 3.2 Collecting News Reports

For each named police shooting or violent encounter in Mapping Police Violence, we query the Google search API for up to 30 news articles relevant to that event. We found this sample size is large enough to represent both sides without introducing too much noise. Our query string includes officer keywords, the victim’s name, and a time window restricted to within one month of the event (see Appendix A for details and design choices). Next, we extracted article publication dates using the Webhose extractor (Geva, 2018), and as a preprocessing step, we used the Dagnet library (Peters and Lecocq, 2013) to automatically filter and remove ads, navigation items, or other irrelevant content from the raw HTML. In the end, the POLICE VIOLENCE FRAME CORPUS contained 82,100 articles across 7,679 events. The per-ideology statistics of reported events are given in Table 1. The racial and ethnic distribution is: White (43.0%), Black (29.7%), Hispanic (15.3%), Asian (1.5%), Native American (1.3%), and Pacific Islander (0.5%), while the other 8.7% of articles report on a victim of unknown race/ethnicity.

#### 3.3 Assigning Media Slant Labels

We associated each news source with a political leaning by matching its URL domain name with the Media Bias Fact Check (2020) record. With more than 1,500 records, the MBFC contains the

Leaning	Articles (#)	Events (#)	Sources (#)	Armed (%)	Attack (%)	Fleeing (%)	Mental Illness (%)	Video (%)
Left	1,090	6730	90	58.0	35.1	24.6	20.6	13.6
Left Center	9,761	3,794	233	66.1	45.9	23.8	18.0	11.1
Least Biased	5,214	3,428	164	73.3	53.7	26.2	19.3	8.8
Right Center	3,993	2,631	105	71.1	50.9	24.9	18.0	9.9
Right	1,009	782	40	64.8	48.7	23.3	19.3	15.5
None	59,739	7,300	12,931	71.0	52.3	24.9	18.4	8.9
Total	80,806	7,647	13,563	70.3	51.3	24.8	18.4	9.4

Table 1: POLICE VIOLENCE FRAME CORPUS statistics. The number of articles and the breakdown of events by whether the victim was *Armed*, *Attacking*, *Fleeing*, had *Mental Illness* or was filmed on *Video* according to Mapping Police Violence data. *Leaning* is decided via Media Bias Fact Check in Section 3.3.

largest available collection of crowdsourced media slant labels, and it has been used as ground truth in other recent work on news bias (Dinkov et al., 2019; Baly et al., 2018, 2019; Nadeem et al., 2019; Stefanov et al., 2020). The MBFC labels are *extreme left*, *left*, *left-center*, *least biased*, *right-center*, *right*, and *extreme right*. For our political framing analysis (Section 6), we consider a source liberal if its MBFC slant label is *left* or *extreme left*, and we consider the source conservative if its label is *right* or *extreme right*. We manually filter this polarized subset to ensure that all articles are on-topic. This led to 1,090 liberal articles and 1,002 conservative articles.

## 4 Media Frames Extraction

We are interested in both the issue and entity frames that structure public narratives on police violence. We will now present our computational framework for extracting both from news text. Throughout this section, we cite numerous prior works from criminology and sociology to motivate our taxonomy, but we are the first to measure these frames computationally. As a preview of the system, Figure 1 shows the key frames extracted from two articles on the murder of 15-year-old Jordan Edwards. Notably, only the left-leaning article mentions the victim’s *race*. Most importantly, our system distinguishes the victim’s attributes from descriptions of the officer. Here, it is the *officer* who is described as a “killer,” and not the victim.

### 4.1 Entity-Centric Frames

Our entity-centric analysis and lexicons are a key contribution in this work. We distinguish the victim’s attributes like *race* and *armed* status from that of the officer or some other entity, and so we move beyond generic and global *issue* frames to understand how the target population is portrayed. These methods require a partitioning of entity tokens into

VICTIM and OFFICER sets. To do so, we first append to each set any tokens matching a *victim* or *officer* regex. The officer regex is general, but the victim regex matches the known name, race, and gender of the victim in PVFC, like *Ronette Morales*, *Hispanic*, *woman* (See Appendix B). Second, we use the huggingface `neuralcoref` for coreference resolution based on Clark and Manning (2016), and append all tokens from spans that corefer to the VICTIM or OFFICER set respectively.

**Age, Gender, and Race.** Following Ash et al. (2019), we consider the age, gender and race of the victim, which are central to an intersectional understanding of unjust police conduct (Dottolo and Stewart, 2008). We extract age and gender frames by string matching on the gender modifier or the numeric age. We extract race frames by searching the attributive or predicative adjectives and predicate nouns of VICTIM tokens and matching these with the victim’s known race.

**Armed or Unarmed.** Knowing whether the victim was armed or unarmed is a crucial variable for measuring structural racism in policing (Mesic et al., 2018). We identify mentions of an unarmed victim with the regex `unarm(?:ed|ing|s)?`, and mentions of an armed victim with `arm(ed|ing|s)?`, excluding tokens with noun part-of-speech.

**Attacking or Fleeing.** Since *Tennessee v. Garner* (1985), the lower courts have ruled that police use of deadly force is justified against felons in flight only when the felon is dangerous (Harmon, 2008). Since *Plumhoff v. Rickard* (2014), deadly force is justified by the risk of the fleeing suspect. Thus whether the victim fled or attacked the officer can inform the officer’s judgment on the appropriateness of deadly force. We propose an entity-specific string-matching method to extract attack frames, where a VICTIM token must be the head of a verb like *injure*, and we use expressions like



\bfllee(:?ing) to extract fleeing mentions.

**Criminality.** Whether the article frames the victim as someone who has engaged in criminal activity may serve to justify police conduct (Hirschfield and Simon, 2010). To capture this frame, we used Empath (Fast et al., 2016) to build a novel lexicon of unambiguously criminal behaviors (e.g. *cocaine*, *robbed*), and searched for these terms.

**Mental Illness.** Police are often the first responders in mental health emergencies (Patch and Arigo, 1999), but there is growing concern that the police are not sufficiently trained to de-escalate crisis situations involving persons with mental illness (Kerr et al., 2010). Mentioning a victim’s mental illness may also highlight evidence of this structural shortcoming. We again used Empath to build a custom lexicon for known mental illnesses and their correlates (e.g. *alcoholic*, *bipolar*, *schizophrenia*). As for *Criminality*, this is not an exhaustive list; we balance precision and recall by ensuring that terms are unambiguous in the context of police violence. Still, we may not capture other signs of mental illness, like descriptions of erratic behaviors.

## 4.2 Issue Frames

**Legal Language.** Similar to Ash et al. (2019), we investigate frames which emphasize legal outcomes for police conduct. To capture this frame, we used a public lexicon of legal terms from the Administrative Office of the U.S. Courts 2021.

**Official and Unofficial Sources.** Many news accounts favor official reports which frame police violence as the state-authorized response to dangerous criminal activity (Hirschfield and Simon, 2010; Lawrence, 2000). Others may include unofficial sources like interviews with first-hand witnesses. We identify official and unofficial sources with the following Hearst-like patterns: <source> <verb> <clause> or *according to* <source>, <clause>. Our unique entity-centric approach lets us exclude the victim’s quotes and focus on witness testimony.

**Systemic.** While the news has historically favored episodic (Iyengar, 1990) fragmented (Bennett, 2016), or decontextualized narratives (Lawrence, 2000), there has been an increase in systemic framing since the 1999 shooting of Amadou Diallo (Hirschfield and Simon, 2010). Such articles identify police shootings as the product of structural or institutional racism. To extract this frame, we look for sentences that (1) mention *other* police shooting incidents or (2) use keywords related to

the national or global scope of the problem.

**Video.** Video evidence was a catalyst for the Rodney King protests (Lawrence, 2000). Psychology studies have found that subjects who witnessed a police shooting on video were significantly more likely to consider the shooting unjustified compared with those who observed through news text or audio (McCamman and Culhane, 2017). We identify reports of body or dash camera footage using the simple regex `(body(?: )?cam|dash(?: )?cam)`

## 4.3 Moral Foundations

Moral Foundations Theory (Haidt and Graham, 2007) (MFT) is a framework for understanding universal values that underlie human judgments of right and wrong. These values form five dichotomous pairs: care/harm, fairness/cheating, loyalty/betrayal, authority/subversion, and purity/degradation. While MFT is rooted in psychology, it has since been applied in political science to differentiate liberal and conservative thought Graham et al. (2009). We quantify the moral foundations that media invoke to frame the virtues or vices of the officer and the victim in a given report using the extended MFT dictionary of Rezapour et al. (2019).

## 4.4 Linguistic Style

To supplement our understanding of overtly topical *entity* and *issue* frames, we investigate two relevant linguistic structures: passive verbs and modals.

**Passive Constructions.** Prior works identify framing effects that arise from passive phrases in narratives of police violence (Hirschfield and Simon, 2010; Ash et al., 2019). In this work, we distinguish *agentive* passives (e.g. “He was killed by police.”) from *agentless* passives (e.g. “He was killed.”). While both deprive actors of agency (Richardson, 2006), only the latter obscures the actor entirely, effectively removing any blame from them (Greene and Resnik, 2009). We specifically contrast liberal and conservative use of VICTIM-headed agentless passives (passive verbs whose patient belongs to the VICTIM set).

**Modal Verbs.** Modals are used *deontically* to express necessity and possibility, and in this way, they are often used to make moral arguments, suggest solutions, or assign blame (Portner, 2009). Following Demszky et al. (2019), we count the document-level frequency of tokens belonging to four modal categories: MUST, SHOULD (*should / shouldn’t / should’ve*), NEED and HAVE TO.

## 5 Validating Frame Extraction Methods

One coder labeled 50 randomly-sampled news articles with indices to mark the order of frames present. Against this ground truth, our binary frame extraction system achieves high precision and recall scores above 70%, with only *race* and *unofficial sources* at 66% and 65% precision. Accuracy is no less than 70% for any frame (see Table 5 in Appendix C). One advantage of our system is that it is not a black box – it gives us the precise location of each frame in the document. When we sort the indices of the predicted frame locations, we find that our system achieves a 0.752 Spearman correlation with the ground truth frame ordering. Finally, the annotator gave us a rank order of the officer and victim moral foundations most exemplified in the document. When we sort, for each document, the foundations by score, our system achieves 0.66 mAP for the officer and 0.40 mAP for the victim.

## 6 Political Framing of Police Violence

### 6.1 Frame Inclusion Aligns with Slant

As shown in Table 2 (Left), we find that liberal sources frame the issue of police violence as more of a systemic issue, using *race*, *unarmed*, and *mental illness* entity frames, while conservative sources frame police conduct as justified with regard to an uncooperative victim. Specifically, conservative sources more often mention a victim is *armed* (.588 vs. .439, +34%), *attacking* (+46%), and *fleeing* (+47%). These strategies serve to justify police conduct since *Tennessee v. Garner* (1985) affirmed the use of deadly force on dangerous suspects in flight (Harmon, 2008), and this narrative is furthered by *official sources* (+38%), *legal language* (+5%), and the victim’s *criminal record* (+7%). Liberal news instead emphasizes the victim’s *race* (+100%), *mental illness* (+25%), and that the victim was *unarmed* (+25%). Cumulatively, these details reinforce the prominent *systemic racism* narrative that appears 47% more often in liberal media. The victim’s mental illness may signal police failure to handle mental health emergencies (Kerr et al., 2010), and the police killing of an unarmed Black victim provides evidence of institutional racism in law enforcement (Aymer, 2016; Tolliver et al., 2016). Together with gender and age, the victim’s race informs an intersectional account of police discrimination (Dottolo and Stewart, 2008). Surprisingly, we find that liberal sources mention age

Framing Device	Inclusion			Ordering		
	Lib.	Cons.	<i>p</i>	Lib.	Cons.	<i>p</i>
Age	0.472	<b>0.764</b>	***	<b>0.480</b>	0.467	
Armed	0.439	<b>0.588</b>	***	0.313	<b>0.358</b>	*
Attack	0.369	<b>0.539</b>	***	0.267	<b>0.306</b>	
Criminal record	0.613	<b>0.655</b>	*	<b>0.294</b>	0.278	**
Fleeing	0.228	<b>0.336</b>	***	<b>0.246</b>	0.217	
Gender	0.610	<b>0.620</b>	***	0.611	<b>0.622</b>	
Legal language	0.875	<b>0.919</b>	***	<b>0.523</b>	0.419	***
Mental illness	<b>0.181</b>	0.145	*	0.301	0.296	
Official sources	0.586	<b>0.808</b>	***	<b>0.194</b>	0.163	***
Race	<b>0.428</b>	0.214	***	<b>0.296</b>	0.233	***
Systemic	<b>0.428</b>	0.291	***	<b>0.410</b>	0.283	***
Unarmed	<b>0.195</b>	0.110	***	0.408	<b>0.470</b>	
Unofficial sources	0.708	<b>0.780</b>	**	<b>0.184</b>	0.163	***
Video	0.164	<b>0.191</b>		0.283	<b>0.436</b>	***

Table 2: (Left) **Frame inclusion aligns with political slant.** The proportion of liberal and conservative news articles that include the given framing device. (Right) **Frame ordering aligns with media slant.** The average inverse document frame order in liberal and conservative news articles where the *frame is present*. Significance given by Mann-Whitney rank test: \* ( $p < 0.05$ ), \*\* ( $p < 0.01$ ), \*\*\* ( $p < 0.001$ )

and gender significantly *less* often than do conservative sources. We find no significant differences in the mention of video evidence, possibly because this detail is broadly newsworthy.

**Controlling for confounds amplifies ideological differences.** One potential confound is *agenda setting*, or ideological differences in the amount of coverage that is devoted to different events (McCombs, 2002). In fact, conservative sources were significantly more likely to cover cases in which the victim was *armed* and *attacking* overall. However, our findings in this section are actually **magnified** when we level these differences and consider only news sources where the ground truth metadata reflects the framing category (see Appendix D).

**Framing decisions are a function of slant.** Finally, we expect that news sources will differ not only diametrically at the political poles, but also linearly in the degree of their polarization. To examine this, we collected an integer score ranging from -35 (extreme left) to +35 (extreme right), which we scraped from the MBFC using an open source tool (*car*). We aggregated articles by their political leaning scores and found the proportion of articles in each bin that express the frame. Linear regressions reveal a statistically significant negative correlation between conservatism and the *criminal record* ( $r=-0.319$ ), *unarmed* ( $r=-0.303$ ), *race* ( $r=-0.667$ ) and *systemic* frames ( $r=-0.283$ ).




Figure 3: **Framing as a function of political leaning.** The MBFC political leaning score vs. the document frame proportion.

## 6.2 Frame Ordering Aligns with Slant

The Inverted Pyramid, one of the most popular styles of journalism, dictates that the most important information in an article should come first (Pötter, 2003; Upadhyay et al., 2016). We hypothesize that the ordering of frames will reflect the author’s judgment on which details are most important, so we should observe ideological differences in frame ordering. In Table 2 (Right) we compare, for each frame, its average inverse document rank in liberal and conservative news articles in which the frame was *already present*. We find that conservative sources highlight that the victim is armed and attacking by placing these details earlier in the report when they are mentioned (avg. inverse rank .358 vs. .313 for *armed*, and .306 vs. .267 for *attacking*). By prioritizing these details early in the article, conservative sources further highlight the need for law and order (Drakulich et al., 2020; Fridkin et al., 2017).

Although in Section 6.1 we found conservative sources favored police reports, we now observe a liberal bias favoring early quotations from these *official sources* (.194 vs. .163 avg. inverse rank). At the same time, liberal sources highlight *unofficial sources* like eyewitnesses who may identify police brutality as a “pervasive and endemic problem” (Lawrence, 2000). Most notably, liberal sources prioritize *legal language* (0.523 vs. 0.419) and *systemic framing* (0.410 vs. 0.283). Liberal sources place these frames, on average, *second* in the total frame ordering (inverse rank  $\approx 1/2$ ), which primes readers to interpret almost all other remaining frames through the lens of injustice and structural racism. This confirms prior work (Graham et al., 2009; Hirschfield and Simon, 2010).

## 6.3 Moral Framing Differences

Prior work (Graham et al., 2013; Haidt and Graham, 2007) suggests that liberals emphasize the care/harm and fairness/cheating dimensions, especially as vice in the officer (Lawrence, 2000), while

conservatives might defend the foundations more equally, especially as virtue in the officer or vice in the victim (Drakulich et al., 2020). We test this by computing, for each moral foundation and for each entity (victim, officer), the proportion of liberal and conservative articles in which either a modifier or agentive verb from the Rezapour et al. (2019) moral foundation dictionary is used to describe that entity. Figure 4 shows that conservative sources unsurprisingly place more emphasis on the victim’s *harmful* behaviors (+48%). Only liberal sources mention the officer’s unfairness or *cheating*. Liberal articles also include more mentions of officer *subversion* (+130%) and, surprisingly, *fairness* (+135%). These results are all significant with  $p < 0.05$ ; no other ideological differences are significant.

## 6.4 The Politics of Linguistic Style

Liberal politicians largely support Black Lives Matter and its calls for police reform (Hill and Marion, 2018), while conservative politicians have historically opposed the Black Lives Matter movement or any anti-police sentiment (Drakulich et al., 2020). We hypothesize that there will be significant differences in the use of agentless passive constructions and modal verbs of necessity (Greene and Resnik, 2009; Portner, 2009) between conservative and liberal sources. When we compare the average document-level frequency for each framing device, normalized by the length of the document in words, we find these hypotheses supported in Table 3. Liberal sources use modal verbs of necessity like SHOULD and HAVE TO more frequently. Conservative sources use agentless passive constructions 61% more than liberal sources ( $2.55 \times 10^{-3}$  vs.  $1.58 \times 10^{-3}$ ), and *violent passives*<sup>2</sup> 31% more. However, we also find that conservative sources discuss the victim more overall (+34%). To remove this confound, we re-normalize the *Pas-*

<sup>2</sup>To indicate violence, we check that the lemma is in {‘attack’, ‘confront’, ‘fire’, ‘harm’, ‘injure’, ‘kill’, ‘lunge’, ‘murder’, ‘shoot’, ‘stab’, ‘strike’}




Figure 4: **Differences in moral foundation frames.** The average moral framing proportions in liberal and conservative articles.

Framing Device	Lib.	Cons.	Cohen's $d$
MUST **	<b>2.09e-4</b>	1.03e-5	0.132
SHOULD ***	<b>4.44e-4</b>	2.04e-4	0.262
NEED ***	<b>2.64e-4</b>	1.15e-4	0.190
HAVE TO ***	<b>4.09e-4</b>	1.75e-4	0.196
Passive ***	1.58e-3	<b>2.55e-3</b>	0.367
Passive Violence ***	6.21e-4	<b>8.13e-4</b>	0.121

Table 3: **Politics and linguistic styles.** The average document frequency of linguistic structures, normalized by the length of the document. \*\*\*  $p < 0.001$

sive and *Violent Passive* counts by the number of victim tokens instead, and the results still hold: +39% and +22% respectively.

## 7 The Broader Scope of Media Framing

Prior works have algorithmically tracked collective attention in the news cycle (Leskovec et al., 2009), and measured the correlation between media attention and offline political action (De Choudhury et al., 2016; Holt et al., 2013; Mooijman et al., 2018). This section examines the broader scope of media framing via two studies.

### 7.1 Peaks Near High-Profile Killings

Do news framing strategies co-evolve across time? We expect to see coordinated peaks in the prevalence of *race*, *unarmed*, and *systemic* frames across U.S. news media, especially near high-profile killings of unarmed Black American citizens. To investigate this hypothesis, we took, for each salient frame, the proportion of articles that mention that frame out of the 82,000 news articles in our dataset, *excluding* any articles that report one of the 15 high-profile police killings listed in Figure 5. Then we found the Pearson correlation between each of the time series in a pairwise manner: 0.49 systemic/unarmed, 0.56 race/unarmed, and 0.70 race/systemic; all statistically significant with  $p < 2.0 \times 10^{-167}$ .

The time series in Figure 5, smoothed over a 15-day rolling window, reveal local spikes near each

of the high-profile killings, with the largest surge in racial and systemic framing near the shootings of ALTON STERLING and PHILANDO CASTILE. Two of the earliest surges appear near the killing of ERIC GARNER and MICHAEL BROWN, which largely ignited the Black Lives Matter movement (Carney, 2016). Recent spikes also appear near the killing of BREONNA TAYLOR and GEORGE FLOYD, which sparked record-setting protests in 2020 (Buchanan et al., 2020). We quantify this with an intervention test (Toda and Yamamoto, 1995) on each time series  $X = (X_1, X_2, \dots, X_t, \dots)$  by fitting an AR(1) model defined by

$$X_t = \beta_0 X_{t-1} + \beta_1 P(t) + c + \epsilon_t$$

with parameters  $\beta$ , constant  $c$ , error  $\epsilon_t$ , and a pulse function  $P(t)$  to indicate the intervention

$$P(t) = \begin{cases} 1, & \text{there was a high-profile shooting at } t \\ 0, & \text{else} \end{cases}$$

The AR(1) is an auto-regressive model where only the previous term  $X_{t-1}$  influences the prediction for  $X_t$ , and the intervention  $P(t)$  allows us to test the null hypothesis that a high-profile killing does not impact framing proportions ( $\beta_1 = 0$ ). We find the coefficient on the intervention  $\beta_1$  is positive for each frame, and significant only in the *unarmed* regression ( $\beta_1 = 2.03$ ,  $p < 0.01$ ). Given this and the high correlation between the three framing categories, we conclude that high-profile killings influence media decisions to frame other killings.

### 7.2 Political Action Precedes Media Framing

We predict that protest volume will positively correlate with media attention on the race and unarmed status of recent victims and the underlying systemic injustice of police killings. Using the CountLove (Leung and Perkins, 2021) protest volume estimates from January 15, 2017 through December 1, 2020, we aligned the per-day national volume with the *race*, *unarmed*, and *systemic* time series (Figure 5) and found low but positive Pearson correlations of 0.098, 0.073, and 0.088 respectively, each






Figure 5: **Peaks near high-profile shootings.** Per-day proportion of articles with *race*, *unarmed*, and *systemic* frames included across time, excluding articles for the high-profile police shootings listed. This reveals local spikes near 15 high-profile incidents. For example, on the *Left* we see *race* framing spikes after the death of Michael Brown.

Frame	Pearson $r$	Granger 1-lag $p$	Granger 2-lag $p$
Race	0.098	0.0185	0.0381
Unarmed	0.073	0.1646	0.0024
Systemic	0.088	0.0801	0.0600

Table 4: **Media Framing and Political Action.** Correlation and  $p$  values for political protests *Granger-causing* media attention towards the race, unarmed, and systemic frames.

statistically significant. These correlations were directed, with protest volume *Granger-causing* an increase in these framing strategies. For two aligned time series  $X$  and  $Y$ , we say  $X$  *Granger-causes*  $Y$  if past values  $X_{t-\ell} \in X$  lead to better predictions of the current  $Y_t \in Y$  than do the past values  $Y_{t-\ell} \in Y$  alone (Granger, 1980). Here,  $\ell$  is called the lag. We considered a lag of 1 and a lag of 2 days. The rightmost column of Figure 5 shows that, with  $\ell = 2$ , protest volume *Granger-causes* *race* and *unarmed* framing with statistical significance ( $p < 0.05$ ) by the SSR F-test. The reverse direction is *not* statistically significant. This reveals that offline protest behaviors *precede* these important media framing decisions, not the other way around. This echoes similar findings on media shifts after the Ferguson protests (Arora et al., 2019) and social media engagement after protests like Arab Spring (Wolfsfeld et al., 2013).

## 8 Discussion and Conclusion

In this work, we present new tools for measuring entity-centric media framing, introduce the POLICE VIOLENCE FRAME CORPUS, and use them to understand media coverage on police violence in the United States. Our work uncovers 15 domain-relevant framing devices and reveals significant differences in the way liberal and conservative news

sources frame both the *issue* of police violence and the *entities* involved. We also show that framing strategies co-evolve, and that protest activity precedes or anticipates crucial media framing decisions.

We should carefully consider the limitations of this work and the potential for bias. Since we matched age, gender and race directly with the MPV, we expect minimal bias, but acknowledge that our exact string-matching methods will miss context clues (e.g. drinking age), imprecise referents (e.g. “teenager”), and circumlocution. We also rely on lexicons derived from expert sources (e.g. U.S. Courts 2021) or from data (Empath), both of which are inherently incomplete. Even the most straightforward keywords (e.g. armed, fleeing) are prone to error. However, biases could also appear in discriminative text classifiers. The advantage of our approach is that it is interpretable and extractive, allowing us to identify matched spans of text and quantify differences in frame ordering. Furthermore, it is grounded heavy in the social science literature. Similar methods could be applied to other major issues such as climate change (Luo et al., 2020a) and immigration (Mendelsohn et al., 2021), where entities include politicians, protesters, and minorities, and where race, mental illness, and unarmed status may all be salient framing devices (e.g. describing the perpetrator or victim of anti-Asian abuse or violence; Gover et al. 2020; Chiang 2020; Ziems et al. 2020; Vidgen et al. 2020).

## Acknowledgments

The authors would like to thank the members of SALT lab and the anonymous reviewers for their thoughtful feedback.

## Ethical Considerations

To respect copyright law and the intellectual property, we withhold full news text from the public data repository. Outside of the victim metadata, POLICE VIOLENCE FRAME CORPUS does not contain private or sensitive information. We do not anticipate any significant risks of deployment. However, we caution that our extraction methods are fallible (see Section 5).

## References

- A fully scrubbed csv of all of media bias fact check's primary categories.
- Yamen Ajjour, Milad Alshomary, Henning Wachsmuth, and Benno Stein. 2019. [Modeling frames in argumentation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2922–2932, Hong Kong, China. Association for Computational Linguistics.
- Maneesh Arora, Davin L Phoenix, and Archie Delshad. 2019. Framing police and protesters: assessing volume and framing of news coverage post-ferguson, and corresponding impacts on legislative activity. *Politics, Groups, and Identities*, 7(1):151–164.
- Erin Ash, Yiwei Xu, Alexandria Jenkins, and Chenjerai Kumanyika. 2019. Framing use of force: An analysis of news organizations' social media posts about police shootings. *Electronic News*, 13(2):93–107.
- Samuel R Aymer. 2016. “i can't breathe”: A case study—helping black men cope with race-related trauma stemming from police killing and brutality. *Journal of Human Behavior in the Social Environment*, 26(3-4):367–376.
- Ramy Baly, Giovanni Da San Martino, James Glass, and Preslav Nakov. 2020. [We can detect your bias: Predicting the political ideology of news articles](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4982–4991, Online. Association for Computational Linguistics.
- Ramy Baly, Georgi Karadzhov, Dimitar Alexandrov, James Glass, and Preslav Nakov. 2018. [Predicting factuality of reporting and bias of news media sources](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3528–3539, Brussels, Belgium. Association for Computational Linguistics.
- Ramy Baly, Georgi Karadzhov, Abdelrhman Saleh, James Glass, and Preslav Nakov. 2019. [Multi-task ordinal regression for jointly predicting the trustworthiness and the leading political ideology of news media](#). 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 2109–2116, Minneapolis, Minnesota. Association for Computational Linguistics.
- David Bamman, Brendan O'Connor, and Noah A. Smith. 2013. [Learning latent personas of film characters](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 352–361, Sofia, Bulgaria. Association for Computational Linguistics.
- David Bamman and Noah A. Smith. 2015. [Open extraction of fine-grained political statements](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 76–85, Lisbon, Portugal. Association for Computational Linguistics.
- Duren Banks, Paul Ruddle, Erin Kennedy, and Michael G Planty. 2016. *Arrest-related deaths program redesign study, 2015-16: Preliminary findings*. Department of Justice, Office of Justice Programs, Bureau of Justice . . . .
- Eric Baumer, Elisha Elovic, Ying Qin, Francesca Polletta, and Geri Gay. 2015. [Testing and comparing computational approaches for identifying the language of framing in political news](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1472–1482, Denver, Colorado. Association for Computational Linguistics.
- Frank R Baumgartner, Suzanna L De Boef, and Amber E Boydston. 2008. *The decline of the death penalty and the discovery of innocence*. Cambridge University Press.
- W Lance Bennett. 2016. *News: The politics of illusion*. University of Chicago Press.
- Esther van den Berg, Katharina Korfhage, Josef Ruppenhofer, Michael Wiegand, and Katja Markert. 2020. [Doctor who? framing through names and titles in German](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4924–4932, Marseille, France. European Language Resources Association.
- Amber E Boydston, Justin H Gross, Philip Resnik, and Noah A Smith. 2013. Identifying media frames and frame dynamics within and across policy issues. In *New Directions in Analyzing Text as Data Workshop, London*.
- Christopher J Bryan, Gregory M Walton, Todd Rogers, and Carol S Dweck. 2011. Motivating voter turnout by invoking the self. *Proceedings of the National Academy of Sciences*, 108(31):12653–12656.
- Larry Buchanan, Quoc Trung Bui, and Jugal K Patel. 2020. Black lives matter may be the largest movement in us history. *The New York Times*, 3.

- Brian Burghart. 2020. [Fatal encounters](#).
- California Legislature. 2021. [Sec. 2800.1](#).
- Dallas Card, Amber E. Boydston, Justin H. Gross, Philip Resnik, and Noah A. Smith. 2015. [The media frames corpus: Annotations of frames across issues](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 438–444, Beijing, China. Association for Computational Linguistics.
- Dallas Card, Justin Gross, Amber Boydston, and Noah A. Smith. 2016. [Analyzing framing through the casts of characters in the news](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1410–1420, Austin, Texas. Association for Computational Linguistics.
- Nikita Carney. 2016. All lives matter, but so does race: Black lives matter and the evolving role of social media. *Humanity & Society*, 40(2):180–199.
- Pamela P Chiang. 2020. Anti-asian racism, responses, and the impact on asian americans’ lives: A social-ecological perspective. In *COVID-19*, pages 215–229. Routledge.
- Dennis Chong and James N Druckman. 2007. Framing theory. *Annu. Rev. Polit. Sci.*, 10:103–126.
- Kevin Clark and Christopher D. Manning. 2016. [Deep reinforcement learning for mention-ranking coreference models](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2256–2262, Austin, Texas. Association for Computational Linguistics.
- Frank E Dardis, Frank R Baumgartner, Amber E Boydston, Suzanna De Boef, and Fuyuan Shen. 2008. Media framing of capital punishment and its impact on individuals’ cognitive responses. *Mass Communication & Society*, 11(2):115–140.
- Munmun De Choudhury, Shagun Jhaver, Benjamin Sugar, and Ingmar Weber. 2016. Social media participation in an activist movement for racial equality. In *Proceedings of the... International AAAI Conference on Weblogs and Social Media. International AAAI Conference on Weblogs and Social Media*, volume 2016, page 92. NIH Public Access.
- Dorotya Demszky, Nikhil Garg, Rob Voigt, James Zou, Jesse Shapiro, Matthew Gentzkow, and Dan Jurafsky. 2019. [Analyzing polarization in social media: Method and application to tweets on 21 mass shootings](#). 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 2970–3005, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yoan Dinkov, Ahmed Ali, Ivan Koychev, and Preslav Nakov. 2019. Predicting the leading political ideology of youtube channels using acoustic, textual, and metadata information. *Proc. Interspeech 2019*, pages 501–505.
- Andrea L Dottolo and Abigail J Stewart. 2008. “don’t ever forget now, you’re a black man in america”: Intersections of race, class and gender in encounters with the police. *Sex Roles*, 59(5-6):350–364.
- Kevin Drakulich, Kevin H Wozniak, John Hagan, and Devon Johnson. 2020. Race and policing in the 2016 presidential election: Black lives matter, the police, and dog whistle politics. *Criminology*, 58(2):370–402.
- Kristin Nicole Dukes and Sarah E Gaither. 2017. Black racial stereotypes and victim blaming: Implications for media coverage and criminal proceedings in cases of police violence against racial and ethnic minorities. *Journal of Social Issues*, 73(4):789–807.
- Robert M Entman. 2007. Framing bias: Media in the distribution of power. *Journal of Communication*, 57(1):163–173.
- Robert M Entman. 2010. Media framing biases and political power: Explaining slant in news of campaign 2008. *Journalism*, 11(4):389–408.
- Ethan Fast, Binbin Chen, and Michael S. Bernstein. 2016. [Empath: Understanding topic signals in large-scale text](#). In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, San Jose, CA, USA, May 7-12, 2016*, pages 4647–4657. ACM.
- Anjalie Field, Gayatri Bhat, and Yulia Tsvetkov. 2019. Contextual affective analysis: A case study of people portrayals in online# metoo stories. In *Proceedings of the international AAAI conference on web and social media*, volume 13, pages 158–169.
- Anjalie Field, Doron Kliger, Shuly Wintner, Jennifer Pan, Dan Jurafsky, and Yulia Tsvetkov. 2018. [Framing and agenda-setting in Russian news: a computational analysis of intricate political strategies](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3570–3580, Brussels, Belgium. Association for Computational Linguistics.
- Anjalie Field and Yulia Tsvetkov. 2019. [Entity-centric contextual affective analysis](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2550–2560, Florence, Italy. Association for Computational Linguistics.
- Lorie A Fridell. 2017. Explaining the disparity in results across studies assessing racial disparity in police use of force: A research note. *American journal of criminal justice*, 42(3):502–513.



- Kim Fridkin, Amanda Wintersieck, Jillian Courey, and Joshua Thompson. 2017. Race and police brutality: The importance of media framing. *International Journal of Communication*, 11:21.
- Paula Froke, Anna Jo Bratton, Oskar Garcia, Jeff McMillan, David Minthorn, and Jerry Schwartz. 2019. *The Associated Press stylebook 2019 and briefing on media law*. Associated Press.
- Matthew Gentzkow and Jesse M Shapiro. 2010. What drives media slant? evidence from us daily newspapers. *Econometrica*, 78(1):35–71.
- Ran Geva. 2018. [Webhose article date extractor](#).
- Angela R Gover, Shannon B Harper, and Lynn Langton. 2020. Anti-asian hate crime during the covid-19 pandemic: Exploring the reproduction of inequality. *American Journal of Criminal Justice*, 45(4):647–667.
- Jesse Graham, Jonathan Haidt, Sena Koleva, Matt Motyl, Ravi Iyer, Sean P Wojcik, and Peter H Ditto. 2013. Moral foundations theory: The pragmatic validity of moral pluralism. In *Advances in experimental social psychology*, volume 47, pages 55–130. Elsevier.
- Jesse Graham, Jonathan Haidt, and Brian A Nosek. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5):1029.
- Clive WJ Granger. 1980. Testing for causality: a personal viewpoint. *Journal of Economic Dynamics and control*, 2:329–352.
- Andrew C Gray and Karen F Parker. 2020. Race and police killings: examining the links between racial threat and police shootings of black americans. *Journal of Ethnicity in Criminal Justice*, pages 1–26.
- Stephan Greene and Philip Resnik. 2009. [More than words: Syntactic packaging and implicit sentiment](#). In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 503–511, Boulder, Colorado. Association for Computational Linguistics.
- Jonathan Haidt and Jesse Graham. 2007. When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*, 20(1):98–116.
- Rachel A Harmon. 2008. When is police violence justified. *Nw. UL Rev.*, 102:1119.
- Joshua B Hill and Nancy E Marion. 2018. Crime in the 2016 presidential election: a new era? *American Journal of Criminal Justice*, 43(2):222–246.
- Paul J Hirschfield and Daniella Simon. 2010. Legitimizing police violence: Newspaper narratives of deadly force. *Theoretical Criminology*, 14(2):155–182.
- Kristoffer Holt, Adam Shehata, Jesper Strömbäck, and Elisabet Ljungberg. 2013. Age and the effects of news media attention and social media use on political interest and participation: Do social media function as leveller? *European Journal of Communication*, 28(1):19–34.
- Shanto Iyengar. 1990. Framing responsibility for political issues: The case of poverty. *Political behavior*, 12(1):19–40.
- Mohit Iyyer, Peter Enns, Jordan Boyd-Graber, and Philip Resnik. 2014. [Political ideology detection using recursive neural networks](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1113–1122, Baltimore, Maryland. Association for Computational Linguistics.
- Mohit Iyyer, Anupam Guha, Snigdha Chaturvedi, Jordan Boyd-Graber, and Hal Daumé III. 2016. [Feuding families and former Friends: Unsupervised learning for dynamic fictional relationships](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1534–1544, San Diego, California. Association for Computational Linguistics.
- Kristen Johnson, I-Ta Lee, and Dan Goldwasser. 2017. [Ideological phrase indicators for classification of political discourse framing on Twitter](#). In *Proceedings of the Second Workshop on NLP and Computational Social Science*, pages 90–99, Vancouver, Canada. Association for Computational Linguistics.
- Kenneth Joseph, Wei Wei, and Kathleen M Carley. 2017. Girls rule, boys drool: Extracting semantic and affective stereotypes from twitter. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*, pages 1362–1374.
- Katherine Keith, Abram Handler, Michael Pinkham, Cara Magliozzi, Joshua McDuffie, and Brendan O’Connor. 2017. [Identifying civilians killed by police with distantly supervised entity-event extraction](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1547–1557, Copenhagen, Denmark. Association for Computational Linguistics.
- Amy N Kerr, Melissa Morabito, and Amy C Watson. 2010. Police encounters, mental illness, and injury: An exploratory investigation. *Journal of police crisis negotiations*, 10(1-2):116–132.
- Haewoon Kwak, Jisun An, and Yong-Yeol Ahn. 2020. [Frameaxis: Characterizing framing bias and intensity with word embedding](#). *ArXiv preprint*, abs/2002.08608.
- Regina G Lawrence. 2000. *The politics of force: Media and the construction of police brutality*. Univ of California Press.



- Jure Leskovec, Lars Backstrom, and Jon Kleinberg. 2009. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506.
- Tommy Leung and Nathan Perkins. 2021. [\[link\]](#).
- Dara Lind. 2014. Cops do 20,000 no-knock raids a year. civilians often pay the price when they go wrong. *Vox*.
- Li Lucy, Dorottya Demszky, Patricia Bromley, and Dan Jurafsky. 2020. Content analysis of textbooks via natural language processing: Findings on gender, race, and ethnicity in texas us history textbooks. *AERA Open*, 6(3):2332858420940312.
- Yiwei Luo, Dallas Card, and Dan Jurafsky. 2020a. Desmog: Detecting stance in media on global warming. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 3296–3315.
- Yiwei Luo, Dallas Card, and Dan Jurafsky. 2020b. [Detecting stance in media on global warming](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3296–3315, Online. Association for Computational Linguistics.
- Michael McCamman and Scott Culhane. 2017. Police body cameras and us: Public perceptions of the justification of the police use of force in the body camera era. *Translational Issues in Psychological Science*, 3(2):167.
- Maxwell McCombs. 2002. The agenda-setting role of the mass media in the shaping of public opinion. In *Mass Media Economics 2002 Conference, London School of Economics: <http://sticerd.lse.ac.uk/dp-extra/McCombs.pdf>*.
- Media Bias Fact Check. 2020. [Media bias / fact check](#).
- Julia Mendelsohn, Ceren Budak, and David Jurgens. 2021. [Modeling framing in immigration discourse on social media](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2219–2263, Online. Association for Computational Linguistics.
- Aldina Mesic, Lydia Franklin, Alev Cansever, Fiona Potter, Anika Sharma, Anita Knopov, and Michael Siegel. 2018. The relationship between structural racism and black-white disparities in fatal police shootings at the state level. *Journal of the National Medical Association*, 110(2):106–116.
- Minnesota Legislature. 2021. [Sec. 609.487 mn statutes](#).
- Negar Mokherian, Andrés Abeliuk, Patrick Cummings, and Kristina Lerman. 2020. Moral framing and ideological bias of news. In *International Conference on Social Informatics*, pages 206–219. Springer.
- Marlon Mooijman, Joe Hoover, Ying Lin, Heng Ji, and Morteza Dehghani. 2018. Moralization in social networks and the emergence of violence during protests. *Nature human behaviour*, 2(6):389–396.
- Khalil Gibran Muhammad. 2019. *The condemnation of Blackness: Race, crime, and the making of modern urban America, with a new preface*. Harvard University Press.
- Moin Nadeem, Wei Fang, Brian Xu, Mitra Mohtarami, and James Glass. 2019. [FAKTA: An automatic end-to-end fact checking system](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 78–83, Minneapolis, Minnesota. Association for Computational Linguistics.
- Minh Nguyen and Thien Huu Nguyen. 2018. [Who is killed by police: Introducing supervised attention for hierarchical LSTMs](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2277–2287, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Peter C Patch and Bruce A Arrigo. 1999. Police officer attitudes and use of discretion in situations involving the mentally ill: The need to narrow the focus. *International Journal of Law and Psychiatry*.
- Ellie Pavlick, Heng Ji, Xiaoman Pan, and Chris Callison-Burch. 2016. [The gun violence database: A new task and data set for NLP](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1018–1024, Austin, Texas. Association for Computational Linguistics.
- Matthew E Peters and Dan Lecocq. 2013. Content extraction using diverse feature sets. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 89–90.
- Ethan V Porter, Thomas Wood, and Cathy Cohen. 2018. The public’s dilemma: race and political evaluations of police killings. *Politics, Groups, and Identities*.
- Paul Portner. 2009. *Modality*, volume 1. Oxford University Press.
- Horst Pöttker. 2003. News and its communicative quality: The inverted pyramid—when and why did it appear? *Journalism Studies*, 4(4):501–511.
- Daniel Preotjiuc-Pietro, Ye Liu, Daniel Hopkins, and Lyle Ungar. 2017. [Beyond binary labels: Political ideology prediction of Twitter users](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 729–740, Vancouver, Canada. Association for Computational Linguistics.
- Vincent Price, Lilach Nir, and Joseph N Cappella. 2005. Framing public discussion of gay civil unions. *Public Opinion Quarterly*, 69(2):179–212.

- J Hunter Priniski, Negar Mokhberian, Bahareh Harandizadeh, Fred Morstatter, Kristina Lerman, Hongjing Lu, and P Jeffrey Brantingham. 2021. [Mapping moral valence of tweets following the killing of george floyd](#). *ArXiv preprint*, abs/2104.09578.
- Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013. [Linguistic models for analyzing and detecting biased language](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1650–1659, Sofia, Bulgaria. Association for Computational Linguistics.
- Rezvaneh Rezapour, Saamil H. Shah, and Jana Diesner. 2019. [Enhancing the measurement of social effects by capturing morality](#). In *Proceedings of the Tenth Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 35–45, Minneapolis, USA. Association for Computational Linguistics.
- John Richardson. 2006. *Analysing newspapers: An approach from critical discourse analysis*. Palgrave.
- Shamik Roy and Dan Goldwasser. 2020. [Weakly supervised learning of nuanced frames for analyzing polarization in news media](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7698–7716, Online. Association for Computational Linguistics.
- Donald Rugg. 1941. Experiments in wording questions: II. *Public opinion quarterly*.
- Maarten Sap, Marcella Cindy Prasetio, Ari Holtzman, Hannah Rashkin, and Yejin Choi. 2017. [Connotation frames of power and agency in modern films](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2329–2334, Copenhagen, Denmark. Association for Computational Linguistics.
- Anne Schneider and Helen Ingram. 1993. Social construction of target populations: Implications for politics and policy. *American political science review*, 87(2):334–347.
- Michael Schudson. 2001. The objectivity norm in american journalism. *Journalism*, 2(2):149–170.
- Jonathon P Schuldt, Sara H Konrath, and Norbert Schwarz. 2011. “global warming” or “climate change”? whether the planet is warming depends on question wording. *Public opinion quarterly*, 75(1):115–124.
- Samuel Sinyangwe, DeRay McKesson, and Johnetta Elzie. 2021. [Mapping police violence](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 527–537, Online. Association for Computational Linguistics.
- Julie Tate, Jennifer Jenkins, and Steven Rich. 2021. [The u.s. police shootings database](#).
- Hiro Y Toda and Taku Yamamoto. 1995. Statistical inference in vector autoregressions with possibly integrated processes. *Journal of econometrics*, 66(1-2):225–250.
- Willie F Tolliver, Bernadette R Hadden, Fabienne Snowden, and Robyn Brown-Manning. 2016. Police killings of unarmed black people: Centering race and racism in human behavior and the social environment content. *Journal of Human Behavior in the Social Environment*, 26(3-4):279–286.
- Oren Tsur, Dan Calacci, and David Lazer. 2015. [A frame of mind: Using statistical models for detection of framing and agenda setting campaigns](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1629–1638, Beijing, China. Association for Computational Linguistics.
- United States Courts. 2021. [Glossary of legal terms](#).
- Shyam Upadhyay, Christos Christodoulopoulos, and Dan Roth. 2016. [“making the news”: Identifying noteworthy events in news articles](#). In *Proceedings of the Fourth Workshop on Events*, pages 1–7, San Diego, California. Association for Computational Linguistics.
- Bertie Vidgen, Scott Hale, Ella Guest, Helen Margetts, David Broniatowski, Zeerak Waseem, Austin Botelho, Matthew Hall, and Rebekah Tromble. 2020. [Detecting east asian prejudice on social media](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 162–172.
- Howard E Williams, Scott W Bowman, and Jordan Taylor Jung. 2019. The limitations of government databases for analyzing fatal officer-involved shootings in the united states. *Criminal Justice Policy Review*, 30(2):201–222.
- Gadi Wolfsfeld, Elad Segev, and Tamir Sheafer. 2013. Social media and the arab spring: Politics comes first. *The International Journal of Press/Politics*, 18(2):115–137.
- Caleb Ziems, Bing He, Sandeep Soni, and Srijan Kumar. 2020. [Racism is a virus: Anti-asian hate and counterhate in social media during the covid-19 crisis](#). *ArXiv preprint*, abs/2005.12423.
- Peter Stefanov, Kareem Darwish, Atanas Atanasov, and Preslav Nakov. 2020. [Predicting the topical stance and political leaning of media using tweets](#).

## A News Data Collection Methods

Each news article search was made as a request to the Google search API using the following form

```
https://www.google.com/search?q=q
&num=30&hl=en
```

The query string `q` was structured in the following way. We included high-precision *officer* and *shooting* keywords, as well as the victim’s full name string (which may contain a middle name or initial), with `first_name`, `last_name`, the first and last space separated word in the full name field respectively. We restricted the search to recent articles within one month following `date`, or the the day of the shooting event. We also included articles published on `day-days(1)` to account for possible time zone misalignment or imprecision.

```
q = (full_name OR first_name
OR last_name) AND (shooting OR
shot OR killed OR died OR fight
OR gun) AND (police OR officer OR
officers OR law OR enforcement OR
cop OR cops OR sheriff OR patrol)
after: date-days(1) before:date
+days(30)
```

The query returns up to 30 articles, which is equivalent to the first page of Google search results in a browser. We found this sample size of 30 to be large enough to contain a sufficient degree of diversity representing both liberal and conservative articles. A larger sample size could introduce additional noise or false positives in this data collection process.

**Potential Confounds** We are aware of some potential confounds in our data collection that could impact results. Firstly, some sources may not mention victim’s name, and these articles will not be represented in our dataset. Articles that omit the victim’s name may be particularly pro-police. Second, liberal and conservative sources could differ in their rate of publishing editorials, opinion pieces, or other content that is not strictly news-related. To investigate, one annotator labeled 100 randomly selected articles, 50 from the left and 50 from the right, indicating whether the article was news, opinion, or other. With simple binomial test, however, we just fail to reject the null hypothesis that the proportion of opinion pieces is statistically different between liberal and conservative sources (0.18 lib. vs. 0.06 cons.,  $p=0.0648$ ).

## B Frame Extraction

Here we detail our frame extraction methods which come in two varieties. The first variety includes document-level regular expressions, and the second variety involves conditional string matching algorithms that rely on a partitioning of the all entity-related tokens into VICTIM and OFFICER sets. These extractive methods were “debugged” in minor ways after investigating their behavior on a development set, correcting for unexpected false positives and false negatives, but we did not iteratively refine regexes or extraction procedures to maximize precision and recall. Because our methods all extract spans of text, we were also able to verify that these rules were capturing different underlying segments of text. When we compute, for each pair of frames, the proportion of articles in which the difference between respective frame indices was within 25 tokens, we find the highest overlap between *legal language* and *criminal record* (25.5%). However, only 10/91 pairs have >10% overlap.

### B.1 Victim and Officer Partitioning

First, we append to each set any tokens matching a *victim* or *officer* regex respectively. The victim regex matches the known name, race, and gender of the victim in the PVFCdataset. For example, for the hispanic female victim named Ronette Morales, we would match tokens in the set

```
{'daughter', 'female', 'girl',
'hispanic', 'immigrant',
'latina', 'latino', 'mexican',
'mexican-american', 'morales',
'mother', 'ronette', 'sister',
'woman' }
```

The officer regex, on the other hand, is given by

```
police|officer|\blaw\b|\
benforcement\b|\bcop(?:s)?\
b|sheriff|\bpatrol(?:s)?\b
|\bforce(?:s)?\b|\btrooper
(?:s)?\b|\bmarshal(?:s)?\b
|\bcaptain(?:s)?\b|\
blieutenant(?:s)?\b|\
bsergeant(?:s)?\b|\bPD\b|\
bgestapo\b|\bdeput(?:y|ies)
\b|\bmount(?:s)?\b|\
btraffic\b|\bconstabular(?:
y|ies)\b|\bauthorit(?:y|ies)
\b|\bpower(?:s)?\b|\
```

```

buniform(?:s)?\b|\bunit(?:s)
)\b|\bdepartment(?:s)?\b|
agenc(?:y|ies)\b|\bbadge(?:s)
)\b|\bchazzler(?:s)?\b|\
bcobbler(?:s)?\b|\bfuzz\b|\
bpig\b|\bk-9\b|\bnarc\b|\
bSWAT\b|\bFBI\b|\bcoppa\b|\
bfive-o\b|\b5-0\b|\b12\b|\
btwelve\b

```

Second, we run the huggingface neuralcoref 4.0 pipeline for coreference resolution, and append all tokens from spans with coreference to the VICTIM or OFFICER set respectively. As an additional plausibility check, we ensure that at least one token in the span is recognized as being *human*. By human, we mean either a proper noun, pronoun, a token with spaCy entity type PERSON, or a token belonging to the set of “People-Related” nouns extracted in Lucy et al. (2020) using WordNet hyponym relations.

## B.2 Document-Level Regular Expressions

For the following categories, we used regular expression methods, returning the index of the first regex match, which we later sort for our final frame ranking. For categories with a dedicated lexicon, we used an exact string matching regex over these words ‘\bword1\b|\bword2\b|...’ to match word1, word2, and all words in that lexicon. If no match was found, that framing category was said to be absent, and the frame rank was set to *inf*.

### B.2.1 Legal language

We compiled a lexicon of legal terms from the Administrative Office of the U.S. Courts 2021,<sup>3</sup> supplemented with the Law & Order terms listed in an online word list source.<sup>4</sup> We then hand-filtered any polysemous or otherwise ambiguous words like *answer*, *assume*, and *bench*, which could lead to false positives in a general setting. Finally, we employed an exact string matching regex over the words in the lexicon.

### B.2.2 Mental illness.

To create a lexicon of terms related to mental illness, we used the Empath tool (Fast et al., 2016) to generate the words most similar to the token *mental\_illness* in an embedding space derived from contemporary New York Times data.

<sup>3</sup><https://www.uscourts.gov/glossary>

<sup>4</sup>[http://www.efnet.com/vocab/wordlists/law\\_and\\_order](http://www.efnet.com/vocab/wordlists/law_and_order)

We hand-filtered this set to remove generic illnesses and any words not related to mental health. We then employed an exact string matching regex over the words in the lexicon.

### B.2.3 Criminal record.

We again used Empath to create a lexicon of terms related to known crimes. We seeded the NYT similarity search with the terms *abuse*, *arson*, *crime*, *steal*, *trafficking*, and *warrant*. We then expanded this set using unambiguous crime names from the Wikipedia Category:Crimes page,<sup>5</sup> and finally hand-filtered so that the set included only crimes (e.g. *theft*) or criminal substances (e.g. *cocaine*). We then employed an exact string matching regex over the words in the lexicon.

### B.2.4 Fleeing.

To capture reports of a fleeing suspect, we use the following regular expression (`\bflee(?:ing)?\b|\bfled\b|\bspe(?:e)?d(?:ing)?(?:off|away|toward|towards)|(took|take(?:n)?)off|desert(?:get|getting|got|run|running|ran)away|pursu(?:it|ed)`). In this way, we identify fleeing both on foot (e.g. Minnesota 609.487, Subd. 6, 2021) and via motor vehicle (e.g. California 2800.1 VC, 2021). These are the forms of evasion that are explicitly enumerated by law. We include `pursu(?:it|ed)` to account for an evasion that is framed from the officer’s perspective, which is a pursuit.

### B.2.5 Video.

We identify reports of body or dash camera footage using the simple regex (`body(?:)?cam|dash(?:)?cam`). We do not use any other related lemmas like *video*, *film*, *record* because we found these to be highly associated with false positives, especially in web text where embedded videos are common. Similarly, we did not match on the word *camera* alone because of false-positives (e.g. “family members declined on-camera interviews”).

### B.2.6 Age.

According to the Associated Press Style Guide (Froke et al., 2019), journalists should always report ages numerically. To avoid false positives, we do not match their spelled-out forms. We identify

<sup>5</sup><https://en.wikipedia.org/wiki/Category:Crimes>



mention of age with an exact string match on the known numerical age of the victim, separated by \b word boundaries.

### B.2.7 Gender.

Unlike Sap et al. (2017)), we are not interested in simply identifying the gender of the victim, but rather, whether there was specific mention of the victim’s gender where a non-gendered alternative was available. For example, to avoid gendering a female victim, one could replace titles like *mother* with *parent*, *daughter* with *child*, *sister* with *sibling*, and *female*, *woman* or *girl* with *person* or simply with the name of the victim. Thus if the victim is female, we match `\b(woman|girl|daughter|mother|sister|female)\b` and if the victim is male we match `\b(man|boy|son|father|brother|male)\b`. We do not match non-binary genders because we do not have ground truth labels for any non-binary targets.

### B.2.8 Unarmed

We identify mentions of an unarmed victim with the regex `unarm(?:ed|ing|s)?`. Manual inspection of news articles reveals that this simple modifier is the standard adjective to describe unarmed victims, so it is sufficient in most cases. Unfortunately, it cannot capture other more subtle context clues (e.g. the victim was *sleeping*, the victim’s hands were in the air) or forms of circumlocution.

### B.2.9 Armed

We match individual tokens to the `^arm(ed|ing|s)?` regex and only return the matching span for tokens that do not have a NOUN Part of Speech tag. This is necessary to disambiguate the verb *arm* from the noun *arm*. We can be confident that when an article mentions *armed*, it is referring to the victim since an armed officer is not newsworthy. On the other hand, we do not match specific weapons because we cannot immediately infer that discussion about a weapon implies the *victim* was armed (it could be an officer’s weapon). We resolve this ambiguity when we extract ATTACK frames, ensuring that the VICTIM is the agent who is wielding a weapon object dependency.

## B.3 Matching Partitioned Tokens

After partitioning the entity-related tokens into VICTIM and OFFICER sets, we extract the following

frames for each document  $\mathcal{D}$ . In all of the following, we define the set `OBJECT = {dobj, iobj, obj, obl, advcl, pobj}` to indicate object dependencies.

### B.3.1 Race

We are determined to prune false positives from our *race* frame detection. We only match *race* where the *race* term is given as an attributive or predicative modifier of the known victim. To do so, we scan, for each token  $t_k \in \text{VICTIM}$ , all children of the head of  $t_k$  in the dependency parse. This set of children would include predicate adjectives of a copular head verb. If the child matched with any member of the lexicon corresponding to the victim’s race, we return the initial index of  $t_k$ . We also expect to capture adjective modifiers in this way because the VICTIM tokens derive from entity spans that include modifiers.

### B.3.2 Attack.

Intuitively, we infer an article has mentioned an attack from the victim if we find the victim has acted in violence or has wielded an object that matches their known weapon or if the officer has been acted upon by a violent vehicular attack. More specifically, for a given document, if we find an VICTIM *nsubj* token in that document having a verbal head in the ATTACK set `{attack, confront, fire, harm, injure, lunge, shoot, stab, strike}` or having a child with OBJECT dependency that matches the victim’s known weapon type (e.g. *gun*, *knife*, etc.) then we return the token’s index as an **attack** mention. To capture vehicular attacks, we also match tokens whose verbal head is in `{accelerate, advance, drive}` and whose object is in the OFFICER set. This process is detailed in Algorithm 1, with a helper function in Algorithm 2.

### B.3.3 Official Source / Unofficial Source.

We use the same high-level method both to identify interviews from Official Sources (e.g. police), and to determine if the article includes quotations or summarizes the perspective of an Unofficial Source (a bystander or civilian other than the victim). To do so, we consider two basic and representative phrasal forms: (1) `<SOURCE> <VERB> <CLAUSE>`, and (2) `according to <SOURCE>, <CLAUSE>`. To extract Phrase Type 1, we identify tokens of entity type PERSON or part of speech PRON such that the token is an *nsubj* or *nsubjpass* whose head lemma belongs to the verb set `{answer, claim, confirm, declare, explain, reply, report, say,`

---

**Algorithm 1:** `attack( $\mathcal{D}, \mathcal{W}$ )`

---

**Input:** Dependency parsed document  $\mathcal{D}$ , and tokens  $\mathcal{W}$  used to describe the victim’s weapon (may be empty)  
**Output:** Document string index  $i$  of the token used to identify an attack from the victim

`ATTACK`  $\leftarrow$  {*attack, confront, fire, harm, injure, lunge, shoot, stab, strike*};  
`ADVANCE`  $\leftarrow$  {*accelerate, advance, {drive}*};  
`OFFICER, VICTIM`  $\leftarrow$  `partition( $\mathcal{D}$ )`;  
**for**  $t_j \in \mathcal{D}$  **do**  
  **if** `dep( $t_j$ ) = nsubj` **then**  
    **for**  $(v, o) \in \text{verbs\_with\_objs}(t_j, [])$  **do**  
      **if** (`lemma( $v$ )`  $\in$  `ATTACK`) **and**  
         $[(t_j \in \text{VICTIM}) \text{ or } (o \in \text{OFFICER} \cup \mathcal{W})]$  **then**  
          **return** `index( $v, \mathcal{D}$ )`;  
        **end**  
      **if**  $v \in \text{ADVANCE}$  **and**  
         $o \in \text{OFFICER}$  **then**  
          **return** `index( $v, \mathcal{D}$ )`;  
        **end**  
      **end**  
    **end**  
  **end**  
**end**  
**return** `inf`;

---

---

**Algorithm 2:** `verbs_with_objs( $v, \mathcal{L}$ )`

---

**Input:** Verb  $v$  from dependency parsed document, recursively generated list  $\mathcal{L}$  of (verb, object) tuples (initially empty)  
**Output:**  $\mathcal{L}$

**for**  $c \in \text{children}(v)$  **do**  
  **if**  $c \in \text{OBJECT}$  **then**  
    `append(( $v, c$ ),  $\mathcal{L}$ )`;  
  **end**  
  **else if** `dep( $c$ ) = prep` **then**  
    `append(( $v, \text{get\_pobj}(\mathcal{L}, c)$ ),  $\mathcal{L}$ )`;  
  **end**  
  **else if** `dep( $c$ )`  $\in$  {*conj, xcomp*} **then**  
     $\mathcal{L} \leftarrow \text{verbs\_with\_objs}(c, \mathcal{L})$ ;  
  **end**  
**end**  
**return**  $\mathcal{L}$ ;

---

*state, tell*). To extract Phrase Type 2, we identify tokens in a dependency relation<sup>6</sup> with the word *according*. If such a token is found in either case and it is *outside* the VICTIM token set, then we return that token’s index as an **Unofficial Source** match. If the token is found in the OFFICER set or has a lemma in {*authority, investigator, official, source*}, then we return the token’s index as an **Official Source** match.

### B.3.4 Systemic claims.

This category is arguably the most variable, and as a result, possibly the most difficult to identify reliably. Systemic claims are used to frame police shootings as a product of structural or institutional racism. To identify this frame, we look for sentences that (1) mention *other* police shooting incidents or (2) use certain keywords related to the national or global scope of the problem. We decide (2) using `(nation(?:[ -])?wide|wide(?:[ -])?spread|police violence|police shootings|police killings|racism|racial|systemic|reform|no(?:[ -])?knock)` as our regular expression. Here, *nation-wide* and *widespread* indicate scope, *police violence*, *police shootings*, and *police killings* describe the persistent issue, while *racism*, *racial*, *systemic* indicate the root of the issue, and *reform* the solution. We also include *no-knock* since there have been over 20k no-knock raids per year since the start of our data collection, and the failures of this policy have been used heavily as evidence in support of police reform (Lind, 2014). To identify (1), we match tokens  $t_k$  of entity type PERSON with thematic relation PATIENT (a dependency relation in {*nsubjpass, dobj, iobj, obj*}) such that  $t_k \notin \text{VICTIM}$  and  $t_k \notin \text{OFFICER}$  and  $t_k$  is not the object of a VICTIM *nsubj*. If `lemma( $t_k$ )` belongs to the set {*kill, murder, shoot*}, we return the index of  $t_k$  as a match for systemic framing. This process is detailed in Algorithm 3, with a helper function in Algorithm 4.

## C Validating Frame Extraction Methods

We report the accuracy, precision, and recall of our system in Table 5. Ground truth is the binary presence of the frame in the 50 annotated articles above. We observe high precision and recall scores

<sup>6</sup>In spaCy 2.1, we need to consider two-hop relations: `According (prep)  $\rightarrow$  to (pobj)  $\rightarrow$  < SOURCE >`

---

**Algorithm 3:** systemic( $\mathcal{D}$ )

---

**Input:** Dependency parsed document  $\mathcal{D}$ **Output:** Document string index  $i$  of the token used to identify locus of systemic framingSHOOTING  $\leftarrow$  {shoot, kill, murder};OFFICER, VICTIM  $\leftarrow$  partition( $\mathcal{D}$ );**for**  $t_j \in \mathcal{D}$  **do**    **if** ( $dep(t_j) \in \{nsubjpass, dobj, iobj, obj\}$  **and** ( $lemma(head(t_j)) \in \{shoot, kill, murder\}$ ) **and** ( $t_j \notin$  VICTIM) **and** ( $t_j \notin$  OFFICER) **and** ( $ent\_type(t_j) =$  PERSON) **and not**     $has\_victim\_subject(t_j)$  **then**        **return**  $index(head(t_j), \mathcal{D})$ ;    **end****end****return**  $inf$ ;

---

---

**Algorithm 4:** has\_victim\_subject( $o$ )

---

**Input:** Object token  $o$  from dependency parsed document**Output:** boolean**for**  $c \in children(head(o))$  **do**    **if**  $c \in$  VICTIM **and**  $dep(c) = nsubj$  **then**        **return**  $true$ ;    **end****end****return**  $false$ ;

---

Frame	Acc	Prec	Recall
Age	86%	100%	84%
Armed	76%	76%	76%
Attack	72%	73%	73%
Criminal record	84%	77%	96%
Fleeing	96%	89%	100%
Gender	88%	91%	91%
Legal language	88%	86%	100%
Mental illness	100%	100%	100%
Official sources	92%	95%	95%
Race	92%	66%	100%
Systemic	88%	86%	100%
Unarmed	96%	88%	88%
Unofficial sources	70%	65%	88%
Video	90%	93%	78%

Table 5: Frame extraction performance on 50 hand-labeled news articles

Victim Variable	Lib.	Cons.	Cohen’s $d$
Mental illness	0.206	0.194	-0.030
Fleeing	0.246	0.235	-0.027
Video	0.136	0.155	0.054
Armed **	0.580	0.648	0.140
Attack ***	0.351	0.486	0.276

Table 6: **Agenda setting.** Proportion of liberal and conservative articles that report on killings where *Victim Variable* is true (e.g. the victim really was Fleeing). We see that conservative sources report more on cases where the victim is armed and attacking

generally above 70%, with only *race* and *unofficial sources* at 66% and 65% precision.

## D Framing vs. Agenda Setting

One potential confound is *agenda setting*, or ideological differences in the amount of coverage that is devoted to different events (McCombs, 2002). In Table 6, we see that conservative sources were significantly more likely to cover cases in which the victim was *armed* (.648 vs. .580, +12%) and *attacking* (.486 vs. .351, +38%) overall. However, we find that the differences in partisan frame alignment are **magnified** when we consider only news sources where the ground truth metadata reflects the framing category. That is, we observed larger effect sizes (Cohen’s  $d$ ) in Table 7 than we did for the observed differences in Table 2. Furthermore, when conditioning on ground truth *race*, these frames are universally more prevalent when the victim is Black as opposed to when the victim

<b>Framing Device</b>	<b>Lib.</b>	<b>Cons.</b>	<b>Cohen's <i>d</i></b>
Armed (T) ***	0.590	<b>0.701</b>	0.233
Armed (T, black) **	0.639	<b>0.762</b>	0.270
Armed (T, white) **	0.552	<b>0.693</b>	0.293
Attack (T) ***	0.381	<b>0.575</b>	0.395
Attack (T, black) ***	0.407	<b>0.585</b>	0.359
Attack (T, white) ***	0.378	<b>0.573</b>	0.396
Fleeing (T) ***	0.381	<b>0.604</b>	0.458
Fleeing (T, black) *	0.424	<b>0.589</b>	0.334
Fleeing (T, white) ***	0.250	<b>0.542</b>	0.618
Mental illness (T) *	<b>0.433</b>	0.320	0.235
Mental illness (T, black) *	<b>0.480</b>	0.291	0.387
Mental illness (T, white)	<b>0.430</b>	0.347	0.171
Race (black) ***	<b>0.612</b>	0.373	0.492
Race (white)	<b>0.197</b>	0.146	0.139
Unarmed (T) ***	<b>0.365</b>	0.218	0.324
Unarmed (T, black)	<b>0.441</b>	0.337	0.212
Unarmed (T, white) **	<b>0.261</b>	0.118	0.380
Video (T) *	0.486	<b>0.626</b>	0.282
Video (T, black)	0.529	<b>0.639</b>	0.223
Video (T,white)	0.394	<b>0.577</b>	0.368

Table 7: **Frame alignment is magnified when conditioned on ground truth.** The proportion of liberal and conservative news articles that include framing device conditioned on articles where ground truth reflects the framing category (T) and the victim's race is given (black, white).

is white. News reports on white victims thus appear more episodic (Lawrence, 2000), while reports on Black victims appear to be more polarizing in terms of the given framing devices. Policing continues to be a highly racialized issue (Muhammad, 2019).