
Predicting the Politics of an Image Using Webly Supervised Data (Supplementary Material)

Christopher Thomas **Adriana Kovashka**
Department of Computer Science
University of Pittsburgh
Pittsburgh, PA 15213
{chris,kovashka}@cs.pitt.edu

1 Introduction

This document presents supplementary results to our main text. We first present additional details of our new political bias dataset, in Section 2. Next, in Section 3, we provide two additional quantitative results using our test set which shows the differences between our best performing method and the baselines on the various topics within our dataset. We also provide results for an application of predicting the bias of different media sources. In Section 4, we present additional qualitative results to complement our result in Fig. 5 from the main text, i.e. images that most strongly predicted several words from articles. In Section 5, we illustrate what our trained document embedding model learns by showing nearby words for a number of query words. In Section 6, we compare human vs. machine performance by showing images that either our best algorithm or humans failed to classify (or both). In Section 7, we include additional examples of images agreed upon by human annotators, as well as the free-form text reasons our participants gave for their Left / Right guesses. We also include our MTurk data collection interface. Finally, in Section 8, we show example images and articles from our dataset.

2 Dataset Details

In this section, we present additional details of our new political bias dataset to complement our main text. Our dataset contains 1,861,336 images total and 1,559,004 articles total. However, after our deduplication procedure (described in our main text), we are left with 1,079,588 unique images upon which we conduct all experiments. In this section, we break down this *unique* count by politics, topic, and media source. We wish to re-emphasize that even though we exclude duplicates here, the articles associated with duplicate images are not necessarily duplicates (the overwhelming majority are unique). Thus, a large body of potentially useful image-text pairs are excluded from this description because the image associated with the text is not unique.

Figure 1 shows the breakdown of unique images in our dataset by politics. There are more images on the left than on the right, resulting in a slight class imbalance. We correct for this class imbalance during training for all of our experiments by ensuring equal class weight in the loss terms. Figure 2 further breaks down the distribution images by topic. For example, we see our dataset contains 83,145 unique images on the subject of religion (from both L/R), our most frequent category, while we collected 17,073 on the subject of vaccines, our least frequent category.

We also present the frequency distribution of our deduplicated dataset broken down by media source in the attached Microsoft Excel file `media_source_stats.xlsx` as there are too many to include or visualize in this document. Note that we also include the political leaning of the media source, as assigned by Media Bias Fact Check (see our main text for details).

Dataset Counts by Politics (after deduplication)

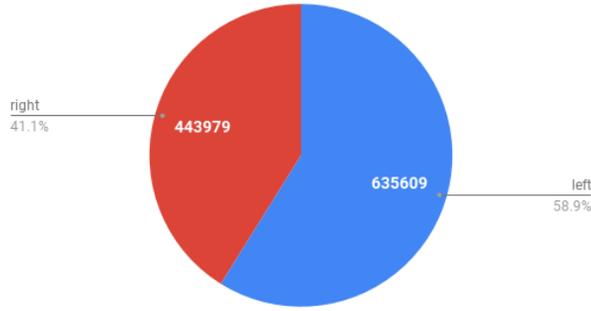


Figure 1: We illustrate the distribution of Left/Right unique images in our deduplicated dataset.

Dataset Counts by Issue (after deduplication)

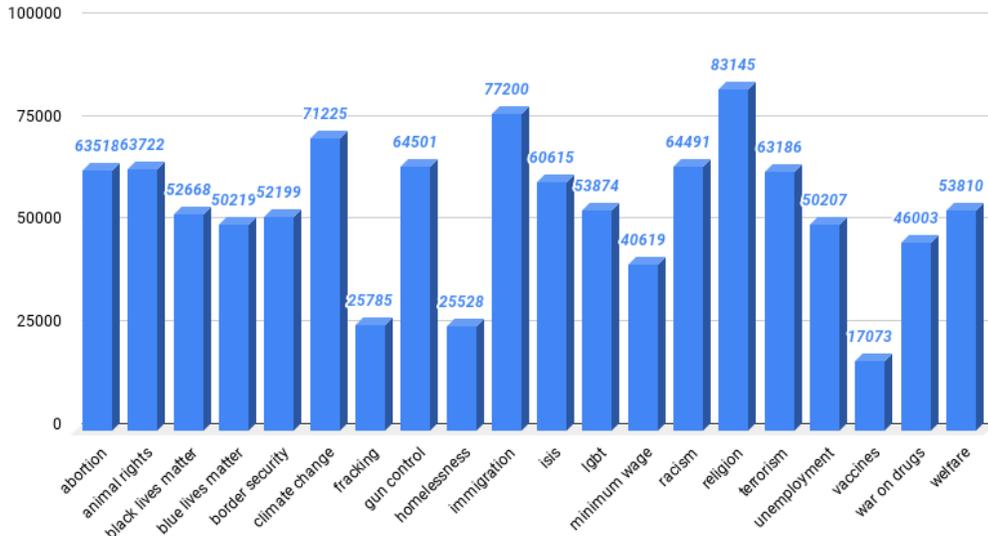


Figure 2: We show the distribution of unique images in our dataset by topic, across both Left/Right.

3 Quantitative Results

We present two quantitative results to supplement our main text. We first wanted to understand on what types of images our best performing method, Ours outperformed the RESNET baseline. In Table 1, we show a result which shows the top-3 topics that our method performed the best (and worst) over the baseline. We notice that for no topic does the baseline outperform our method. Even for those topics on which the baseline performs most competitively with our method, our method still outperforms it by 1-2%. We include complete results including additional baselines, for all topics in the included file, `topic_results.xlsx`.

Method	Vaccines	Fracking	War on Drugs	Border Security	Black Lives Matter	Climate Change
RESNET	0.6768	0.6737	0.6684	0.6922	0.7026	0.6934
Ours	0.7422	0.7209	0.7128	0.7161	0.7269	0.7179

Table 1: Average performance for the three topics where our method achieves the largest vs smallest improvement over the baseline.

Method	Top-20	Top-100	Sun	Change	Breitbart	NewSt	NewYorker	NatRev	Slate	CNN	RevCom
RESNET	0.697	0.690	0.627	0.653	0.527	0.821	0.873	0.718	0.798	0.795	0.875
OURS	0.739	0.724	0.707	0.690	0.607	0.808	0.934	0.758	0.793	0.873	0.781

Table 2: Average performance for the top-20, and top-100 news sources, and individual results for some popular news sources.

In Table 2, we analyze the results as a function of the media source to which the image belongs. We compute the performance of our method on images exclusively from a particular media source, for each media source. We then rank the sources by number of samples in the test set, and check how performance changes as the number of samples decreases. We see that for media sources with more samples, OURS achieves a stronger result than the RESNET baseline (0.739 vs 0.697). We also show results for individual well-known media sources that have many samples in our dataset. The Sun, Breitbart, and National Review are well-known right-leaning sources, while the rest are left-leaning. Our method works well for both right- and left-leaning sources. For a few left-leaning sources, the baseline achieves stronger results. Among common sources, the baseline’s gain is largest on RevCom, a very far-left, “revolutionary communism” website. It is surprising to see how accurate we can infer leaning from images alone; close to or over 80% for many sources shown.

We also provide supplementary results to complement this result in `media_source_results.xlsx`, including for other baselines. We break down the performance for each of our methods by media source. We observe that our method, OURS consistently outperforms the baselines, often substantially.

4 Image to Word Prediction Results

In our main text, we described a model trained to predict words from images. We trained this model to predict which words, from a fixed dictionary of the 1000 most visual words (see main text for details), would be in the article paired with the image. For this result only, we also conditioned the model on the document embedding of the article paired with the image. After training, we ran our entire large weakly-supervised test set through this model and predicted words for all images. For each word, we then sorted all test set images by the score the model assigned for the prediction of that word and show the 100 images for each that have the highest overall probability. We include results for several words in the `image_to_word` folder. We include results for several words, including “immigrant”, “lgbt”, “antifa”, and “nationalist”. We see that the images which strongly predicted the word “immigrant” often feature Hispanic people, children, or law enforcement symbols / personnel. For “lgbt”, we notice that many images feature rainbow flags. “Antifa” often features street scenes with protestors wearing black. We also observe fascist symbols, such as swastikas or Nazi salutes in these photos. “Nationalist” features numerous examples of white supremacist imagery, including Ku Klux Klan garbs, swastikas, and Celtic crosses: symbolism associated with white supremacist and neo-nazi movements. Collectively these results indicate that, although the articles paired with the text are lengthy and much more weakly aligned than traditional image-text embedding tasks (i.e. captions, descriptions, etc.), a consistent visual signal exists that our model is able to grasp.

5 Textual Embedding Word Retrieval Results

We trained a text embedding [1] model on articles from our dataset. In Table 3 we show an example of what our model learned for a number of query words. We compute the embedding of the query words using our model, then find the nearest words in embedding space from the learned dictionary and rank them. We observe that for “Donald Trump”, several of the top words are in Spanish, which are likely coming from articles related to immigration concerning Trump. The translation of these words is fitting in this context, i.e. *intensa* means “intense”, while “*ultraderecha*” means far-right. “*Horripilantes*” means “horrifying / terrifying.” We also notice a “#” sign associated with Trump, likely coming from his use of Twitter. Importantly, we noticed for *events*, like Charlottesville (a protest event in which a protestor was run over by a car in a hate crime), relevant concepts that our *image* classifiers could potentially pick up on appear. For example, “riots”, “antifa” (a protest group), “rally”, etc. are all visualizable concepts associated with the event. We observe for another event, “Parkland” (a mass school shooting event involving 17 deaths), nearby concepts are “Newtown” (another school shooting), “Hogg” (a survivor of the Parkland shooting), “NRA” (the National Rifle Association, which opposed gun measures following the event and was the subject of significant

Query phrase:	donald trump	charlottesville	liberal	fascist	parkland
Results:	auxiliar: 0.4155	charleston: 0.7303	leftist: 0.2721	fascism: 0.7861	newtown: 0.7640
	intensa: 0.4132	parkland: 0.7189	progressive: 0.2650	fascists: 0.7494	hogg: 0.7635
	macron: 0.4102	antifa: 0.7135	conservative: 0.2583	nazi: 0.7169	stoneman: 0.7501
	kkk: 0.4042	putin: 0.7117	liberals: 0.2541	racists: 0.7128	nra: 0.7455
	ultraderecha: 0.4010	ferguson: 0.7038	much: 0.2516	racist: 0.7068	charlottesville: 0.7189
	horripilantes: 0.4005	dallas: 0.6998	wing: 0.2516	totalitarian: 0.6903	shooting: 0.7161
	billionaire: 0.3991	confederate: 0.6995	mainstream: 0.2514	repressive: 0.6866	walkout: 0.7135
	pence: 0.3980	richmond: 0.6956	centrist: 0.2420	terrorist: 0.6862	walkouts: 0.7029
	shooting: 0.3937	shooting: 0.6879	moderate: 0.2323	filmado: 0.6791	charleston: 0.7002
	cruz: 0.3928	horrific: 0.6844	emerged: 0.2312	imperialist: 0.6771	tragedy: 0.6991
	duterte: 0.3924	portland: 0.6828	dismal: 0.2309	communist: 0.6729	orlando: 0.6986
	erdogan: 0.3919	riots: 0.6826	steadily: 0.2269	nazis: 0.6666	emma4change: 0.6931
	continuado: 0.3898	cleveland: 0.6817	radical: 0.2263	globalist: 0.6659	msd: 0.6844
	mueller: 0.3876	heyer: 0.6806	portrayed: 0.2256	nationalist: 0.6655	sandyhook: 0.6841
	tonight: 0.3874	protest: 0.6782	conservatives: 0.2253	genocidal: 0.6630	shootings: 0.6795
	inauguration: 0.3869	rally: 0.6779	shifted: 0.2248	rogue: 0.6627	gun: 0.6752
	gop: 0.3852	nfl: 0.6760	defeating: 0.2245	authoritarian: 0.6620	marjory: 0.6739
america: 0.3848	tragedy: 0.6757	plummeted: 0.2244	extremist: 0.6603	senseless: 0.6701	
potus: 0.3835	islamophobia: 0.6727	outflanked: 0.2219	vanguard: 0.6599	kasky: 0.6688	
brexit: 0.3834	anticom: 0.6721	progressives: 0.2218	antifascist: 0.6583	neveragain: 0.6665	
presidency: 0.3819	spike: 0.6719	leftwing: 0.2217	avakian: 0.6579	traynick: 0.6654	
alabama: 0.3817	berkeley: 0.6718	touted: 0.2209	ahoies: 0.6571	7to: 0.6644	
marcharse: 0.3814	counterprotesters: 0.6702	democrat: 0.2209	waok: 0.6566	sarasota: 0.6613	
cabinet: 0.3812	barcelona: 0.6692	12,030: 0.2204	troutdale: 0.6565	columbine: 0.6610	
netanyahu: 0.3779	memphis: 0.6679	long: 0.2196	clown: 0.6564	horrific: 0.6597	
milo: 0.3770	heaphy: 0.6669	corporatists: 0.2194	supremacist: 0.6556	gun: 0.6596	
republicans: 0.3766	alt: 0.6665	served: 0.2186	democrat: 0.6548	manjares: 0.6596	
opioid: 0.3757	weekend: 0.6662	framed: 0.2186	supremacy: 0.6548	florida: 0.6583	
comey: 0.3737	mcauliffe: 0.6657	hardline: 0.2182	lunatic: 0.6545	loesch: 0.6576	
#: 0.3736	spencer: 0.6654	leftward: 0.2176	misogynist: 0.6533	nationalwalkoutday: 0.6574	

Table 3: We show examples of our learned text embedding. At the top, we show several “query phrases” which we embed using our method. We then compute the distance from each query phrase to all other learned words in our dataset’s vocabulary and rank the words in order of increasing distance. Thus, retrieved words near the top are more closely related to the query phrase in the learned space than words below.

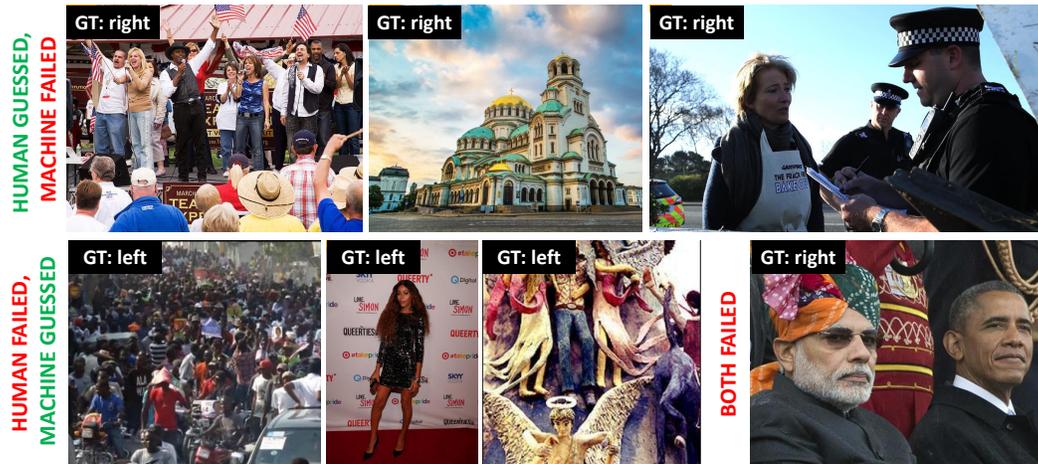


Figure 3: Images that either our best algorithm failed to classify (top), humans failed (bottom left) or both human and machine failed (bottom right). Please see text for our explanation.

press), and a variety of other hashtags and concepts associated with the event. We believe that these results illustrate *how* leveraging text helps our method ultimately perform better by forcing our classifiers to learn to pay attention to certain visual concepts, after being conditioned on the latent document embedding at training time. The representation our image classifiers learn guided by this latent, weak supervision ultimately proves to be superior to the other approaches we tested. We include many additional word query results in `learned_word_embeddings.xlsx`.

6 Human vs. Machine Abilities

In Fig. 3, we show images that humans and/or our best-performing algorithm (OURS) were able/unable to classify. At the top, we show the gap between human reasoning abilities and our classifier. The first image at the top has a subtle country vibe (associated with the right), which was imperceptible

for our algorithm. Next is an image of a non-western church, which was likely too different from churches in the training set. The third image shows British actress Emma Thompson campaigning for Greenpeace and getting cited; our algorithm is unable to analyze such complex scenes. At the bottom left of the figure, we show images that humans were unable to classify, but bias in the data helped our algorithm classify. Images of protests, Hollywood, and art, are common in left-leaning images. Finally, we show an image that neither human nor algorithm were able to classify, as it depends on context from the article, which is unavailable at test time.

7 MTurk Responses

In Figures 4-6, we show example images which at least a majority (2/3) of humans were able to guess the politics of correctly. We note that many times, when a politician of a particular party is shown, human annotators assume the image is the same leaning as the politician’s party (e.g. image of Trump is right). Annotators often rely on racial stereotypes as well (“black women are more liberal,” “most rappers are left,” “left muslims”, “older white man” is right-leaning). Relying on these stereotypical concepts in our HUMAN CONCEPTS model explains why that model performs best on those images containing humans (see main text), though it doesn’t perform best in the dataset at large. We also observe that humans tend to associate the right with guns, patriotic symbols, and religion, whereas they tend to associate peace, compassion, diversity, protests, and minorities with the left. Humans also recognized some of the people appearing in the images and relied on their external knowledge of that person’s political leanings to guess the image’s label. We also include a complete listing of the concepts that were extracted from the MTurk free-form text explanations in `human_concepts.xlsx`.

We also include our MTurk data collection interface in HTML file `MTurk_Interface.html`. Note that as you answer the questions, additional questions will appear. We first asked annotators to determine if the image met certain exclusionary criteria, i.e. text, blurry, etc. We then asked annotators to classify the image as left / right / ambiguous. We then asked what features of the image were relied on by the annotator to make their decision. We then showed annotators the article text going with the document and asked whether the text met certain exclusionary criteria, mainly originating from HTML scraping errors. We also asked annotators if the image and text were related to one another and to paste the text from the article that most aligned with the image. We then asked the workers to predict the politics of the image-text *pair* (as opposed to the image alone) and finally asked workers to state political topic(s) of the image-text pair.

8 Example Images and Documents from Dataset

In Figures 7-9, we show example images and some text from their associated articles from our dataset. Note that the text we include for each image is truncated, as many of the articles are quite lengthy. We also annotate each image with a “L” or “R” depending on whether the image comes from the left or right respectively, as well as the original source for the image and article text.

We believe these images highlight how extreme some of our media sources are. For example, in Fig. 7, we see an image of apparent Hispanic gang members with Obama’s head cropped as one of them. The article discusses illegal immigration and alleges Obama has facilitated allowing “illegals” to stay. In Fig. 8, we see several images of protests, one of which is associated with the left and one of which is associated with the right. We note, however, that the protest image associated with the right (bottom left) actually shows protesters carrying signed *supportive* of Planned Parenthood, a topic associated with the left. A similar situation is found in Fig. 9 where we see an image of a transgender man (bottom row, middle) labeled right. These images’ labels only makes sense in conjunction with their paired article text, which are describing circumstances related to the image. These examples underscore one primary challenge of learning visual classifiers on our dataset: *images’ labels often depend upon the context on which they appear as much as they depend on what is in the image itself*. Thus, learning to predict the politics from an image is highly challenging due to the inherent high-level semantic nature of the problem as well as the presence of noisy data. We believe our method, guided by privileged information from the text domain helps guide the training and alleviates some of these problems.

			
Republican president Guess: R	A heroic memeified photo of Obama comes from liberals Guess: L	Liberal stance: Anti-discrimination for Hispanics Guess: L	This picture is showing trump supporters at a rally. Guess: R
positive picture of Trump Guess: R	A positive picture of Obama Guess: L	Supporting a liberal policy Guess: L	Gun rights supporter are generally right leaning. Guess: R
trump smiling Guess: R	PIC OF OBAMA, LIBERAL PRESIDENT Guess: L	Pro immigration Guess: L	Second Amendment shirt would lean right. Guess: R
			
THE LEFT LOVES TO PROTEST. Guess: L	Looks like a man cross dressing so that would only be supported by a left winger Guess: L	many black women are more liberal than conservative Guess: L	the image involves voters and the Republicans are very concerned about the threat of voter fraud Guess: R
they like protesting a lot Guess: L	Weirdness embraced Guess: L	Most african american women lean left Guess: L	i chose right because it looks like a voting booth Guess: R
Looks like a leftist political rally Guess: L	Looks like a gay person Guess: L	Guessed incorrectly	Guessed incorrectly
			
the Democratic candidates Guess: L	this attorney is very liberal Guess: L	I believe they are trying to illustrate "police brutality" Guess: L	Obama Guess: L
Sanders and Clinton Guess: L	gloria allred is a leftist and represents womens causes Guess: L	seems to be talking about the struggle of african americans Guess: L	Obama is there Guess: L
This is from the democratic primary debate Guess: L	Guessed incorrectly	Guessed incorrectly	For sure left Obama Guess: L

Figure 4: Examples of images whose politics were correctly guessed by at least a majority (2/3) of MTurkers. We also include the reasons given for their prediction by the MTurkers below each image. MTurkers who guessed the image incorrectly are indicated by “Guessed incorrectly.”

References

- [1] Q. Le and T. Mikolov. Distributed representations of sentences and documents. In *International Conference on Machine Learning*, pages 1188–1196, 2014.

			
There's a gun Guess: R	promotes violence Guess: R	shows care and compassion Guess: L	Military leans right. Guess: R
guntoting romanticized Guess: R	Flags and knives and who knows what. Seems right leaning. Guess: R	The left mourning, probably another school shooting. Guess: L	military source is often more right Guess: R
Guessed incorrectly	its an anti-drug image Guess: R	trying to show rallies against current policies Guess: L	Guessed incorrectly
			
Right loves the flag Guess: R	older white man Guess: R	This person looks liberal. Guess: L	Lots and lots of guns. Guess: R
People on the right tend to cling to national symbols. Guess: R	Military Guess: R	I think this is one of the kids from one of the school shootings. Guess: L	Gun and flag Guess: R
Guessed incorrectly	Gender and race and flag background. Guess: R	Guessed incorrectly	The right would like a more militant George Washington. Guess: R
			
confederate flag Guess: R	Orange popsicle of death Guess: R	African Americans are usually democrats Guess: L	Looks like a working class white man who is angry with liberals Guess: R
Only a person on the right would use a confederate flag. Guess: R	it's not an awful picture of him. Guess: R	African Americans tend to be more to the left. Guess: L	Looks like a right not sure Guess: R
confederate flag Guess: R	This is Donald Trump, a republican, though he is making an odd face. Guess: R	Music industry Guess: L	Looks like a white supremacist Guess: R

Figure 5: Examples of images whose politics were correctly guessed by at least a majority (2/3) of MTurkers. We also include the reasons given for their prediction by the MTurkers below each image. MTurkers who guessed the image incorrectly are indicated by “Guessed incorrectly.”

			
Blue surroundings Guess: L	left muslims Guess: L	People on the right love church. Guess: R	smiling trump Guess: R
I have seen his speeches Guess: L	Minorities tend to fall to the left. Guess: L	Religious windows Guess: R	Trump is shown as happy and giving thumbs up, so this is pro Republican AKA right leaning. Guess: R
I think that's Cory Booker, a Democrat hero. Guess: L	Guesed incorrectly	Guesed incorrectly	President is a Republican Guess: R
			
Looks like a peace rally, a common theme on the left. Guess: L	Screaming woman, must be on the left. Guess: L	Due to the LGBT flag in the street. Guess: L	These are trump supporters Guess: R
looks like a candlelight protest. Guess: L	looks like some sort of protest. Guess: L	Definitely left because this is the colors for the gay flag and i am pretty sure that the gay community lean towards left more than the right. Guess: L	because trump is a republican and he would never support a democratic candidate Guess: R
fight terrorist Guess: L	the speakers voice. Guess: L	Colors on the road Guess: L	positive image of trump rally Guess: R
			
most rappers are left Guess: L	The images contain solar panels and windmills that are green political Guess: L	I used my instincts and my knowledge on certain people / things. Guess: R	He wants everyone to know he supports LGBT Guess: R
Appears to be a liberal which is less conservative and leans to the left. Guess: L	Clean energy Guess: L	photoshopping a giant american flag on a location where it doesn't seem very applicable is just patriotism whoring Guess: R	It is attempting to curry favor for Trump, by showing that he has LGBTQ supporters. Guess: R
Freedom of expression Guess: L	Guesed incorrectly	It's of the American flag. Guess: R	supportive image of trump Guess: R

Figure 6: Examples of images whose politics were correctly guessed by at least a majority (2/3) of MTurkers. We also include the reasons given for their prediction by the MTurkers below each image. MTurkers who guessed the image incorrectly are indicated by “Guesed incorrectly.”



Figure 7: Example images and articles (truncated) from our dataset. We annotate each image with the media source from which it and the article came, as well as the politics of that media source, as determined by Media Bias Fact Check (see our main text for details).

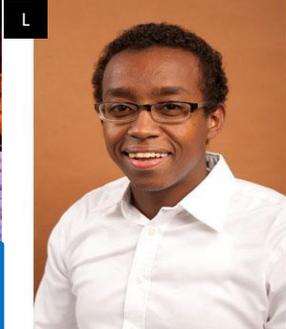
		
<p>“When Jedi Jimenez approached the podium at the People’s State of the City Thursday, he faced hundreds of Long Beach residents, crammed shoulder to shoulder in the pews of the First Congregational Church located downtown. Attendees shared one unifying goal: to take their city’s issues head on. When Jimenez finally spoke, he didn’t just ask for the crowd’s attention – he commanded it. “Over the past year, our country has faced some of the biggest threats to our values of democracy, inclusion and justice.” Source: Daily49er.com</p>	<p>“Like a scout for scholars, the Point Foundation searches out the best and brightest LGBT students. The Los Angeles-based organization grants tens of thousands of dollars in tuition to dozens of collegians each year, some cut off financially from their families because of their sexual orientation. Following a six-month search, the Point Foundation recently announced its 25 scholars of 2010. The diverse group includes a former janitor, a young man who underwent an exorcism at his mother’s hands, and a woman, previously fired for being gay, now entering her third year of law school.” Source: Advocate.com</p>	<p>“Two years after disengagement Israel has put a blockade on the Gaza Strip not allowing goods and other necessities into the region making Gazans almost completely aid dependent. Two years ago, Israel completed its unilateral withdrawal from the Gaza Strip. We all remember the intense media campaign shamelessly portraying the settlers as dispossessed victims of a bold move for peace. Among others, Harvard economist Sara Roy argued that Israel’s version of disengagement would bring disaster to an already desperate Gaza. Today, we are witnessing emergence of an unparalleled economic catastrophe in the Gaza Strip and with it, the evaporation of the last remaining hopes for a Palestinian state.” Source: Electronicintifada.net</p>
		
<p>“On January 23, 2017, the Senate confirmed Rep. Mike Pompeo , a Republican congressman from Kansas, as director of the Central Intelligence Agency. Pompeo, 53, has served in the House of Representatives since 2011. He succeeds a 25-year veteran of the CIA, John Brennan, who’s served as the agency’s chief since 2013. Advertisement - Continue Reading Below Here’s what you need to know about Pompeo: 1. He served in the Army. Mike Pompeo during a TV appearance while he was a member of the Army. Pompeo graduated first in his class from West Point in 1986, according to his congressional biography.” Source: Cosmopolitan.com</p>	<p>“On Sunday, Senator Susan Collins (R-ME) said she would not vote for President Trump’s nominee to the Supreme Court if the nominee was “hostile” to Roe v. Wade . This shouldn’t come as a surprise; Collins showed how callous she was to the rights of the unborn child in 2003. On October 21, 2003, voting with ...” Source: DailyWire.com</p>	<p>“A Presbyterian chaplain in Maine penned an op-ed this month in a local newspaper in which he described Planned Parenthood as “blessed” and lauded the nation’s largest abortion provider for promoting “life-affirming values.” The Rev. Marvin Ellison, who ministers to patients at a Planned Parenthood facility in Portland, recently joined with other chaplains to host ...” Source: IJR.com</p>

Figure 8: Example images and articles (truncated) from our dataset. We annotate each image with the media source from which it and the article came, as well as the politics of that media source, as determined by Media Bias Fact Check (see our main text for details).



Figure 9: Example images and articles (truncated) from our dataset. We annotate each image with the media source from which it and the article came, as well as the politics of that media source, as determined by Media Bias Fact Check (see our main text for details).