

Analyzing Framing through the Casts of Characters in the News

Dallas Card¹ Justin H. Gross² Amber E. Boydston³ Noah A. Smith⁴

¹School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA

²Department of Political Science, University of Massachusetts, Amherst, MA 01003, USA

³Department of Political Science, University of California, Davis, CA 95616, USA

⁴Computer Science & Engineering, University of Washington, WA 98195, USA

dcard@cmu.edu aboydstun@ucdavis.edu jhgross@polsci.umass.edu
nasmith@cs.washington.edu

Abstract

We present an unsupervised model for the discovery and clustering of latent “personas” (characterizations of entities). Our model simultaneously clusters documents featuring similar collections of personas. We evaluate this model on a collection of news articles about immigration, showing that personas help predict the coarse-grained framing annotations in the Media Frames Corpus. We also introduce automated model selection as a fair and robust form of feature evaluation.

1 Introduction

Social science tells us that communication almost inescapably involves **framing**—choosing “a few elements of perceived reality and assembling a narrative that highlights connections among them to promote a particular interpretation” (Entman, 2007). Memorable examples include loaded phrases (*death tax*, *war on terror*), but the literature attests a much wider range of linguistic means toward this end (Pan and Kosicki, 1993; Greene and Resnik, 2009; Choi et al., 2012; Baumer et al., 2015).

Framing is associated with several phenomena to which NLP has been applied, including *ideology* (Lin et al., 2006; Hardisty et al., 2010; Iyyer et al., 2014), *sentiment* (Pang and Lee, 2008; Feldman, 2013), and *stance* (Walker et al., 2012; Hasan and Ng, 2013). Although such author attributes are interesting, framing scholarship is concerned with persistent patterns of representation of particular issues—without necessarily tying these to the states or intentions of authors—and the effects that such patterns

may have on public opinion and policy. We also note that NLP has often been used in large-scale studies of news and its relation to other social phenomena (Leskovec et al., 2009; Gentzkow and Shapiro, 2010; Smith et al., 2013; Niculae et al., 2015).

Can framing be automatically recognized? If so, social-scientific studies of framing will be enabled by new *measurements*, and new applications might bring framing effects to the consciousness of everyday readers. Several recent studies have begun to explore unsupervised framing analysis of political text using autoregressive and hierarchical topic models (Nguyen et al., 2013; Nguyen et al., 2015; Tsur et al., 2015), but most of these conceptualize framing along a single dimension. Rather than trying to place individual articles on a continuum from liberal to conservative or positive to negative, we are interested in discovering broad-based patterns in the ways in which the media communicate about issues.

Here, our focus is on the narratives found in news stories, specifically the participants in those stories. Insofar as journalists make use of archetypal narratives (e.g., the struggle of an individual against a more powerful adversary), we expect to see recurring representations of characters in these narratives (Schneider and Ingram, 1993; Van Gorp, 2010). A classic example is the contrast between “worthy” and “unworthy” victims (Herman and Chomsky, 1988). More recently, Glenn Greenwald has pointed out how he was repeatedly characterized as an *activist* or *blogger*, rather than a *journalist* during his reporting on the NSA (Greenwald, 2014).

Our model builds on the “Dirichlet persona model” (DPM) introduced by Bamman et al. (2013)

for the unsupervised discovery of what they called “personas” in short film summaries (e.g., the “dark hero”). As in the DPM, we operationalize personas as mixture of textually-expressed characteristics: what they do, what is done to them, and their descriptive attributes. We begin by providing a description of our full model, after which we highlight the differences from the DPM.

This paper’s main contributions are:

- We strengthen the DPM’s assumptions about the *combinations* of personas found in documents, applying a Dirichlet process prior to infer patterns of cocurrence (§3). The result is a clustering of documents based on the collections of personas they use, discovered simultaneously with those personas.
- Going beyond named characters, we allow Bamman-style personas to account for entities like institutions, objects, and concepts (§5).
- We find that our model produces interpretable clusters that provide insight into our corpus of immigration news articles (§6).
- We propose a new kind of evaluation based on Bayesian optimization. Given a supervised learning problem, we treat the inclusion of a candidate feature set (here, personas) as a hyperparameter to be optimized alongside other hyperparameters (§7).
- In the case of U.S. news stories about immigration, we find that personas are, in many cases, helpful for automatically inferring the coarse-grained framing and tone employed in a piece of text, as defined in the Media Frames Corpus (Card et al., 2015) (§7).

2 Model Description

The plate diagram for the new model is shown in Figure 1 (right), with the original DPM (Bamman et al., 2013) shown on the left.

As evidence, the model considers tuples $\langle w, r, e, i \rangle$, where w is a word token and r is the category of syntactic relation¹ it bears to an entity with index e mentioned in document with index i . The model’s generative story explains this evidence

¹We adopt the terminology from Bamman et al. (2013) of “agent”, “patient”, and “attribute”, even though these categories of relations are defined in terms of syntactic dependences.




Figure 1: Plate diagrams for the DPM (left), and for the new model (right).

as follows:

1. Let there be K topics as in LDA (Blei et al., 2003). Each topic $\phi_k \sim \text{Dir}(\gamma)$ is a multinomial over the V words in the vocabulary, drawn from a Dirichlet parameterized by γ .
2. For each of P personas p , and for each syntactic relation type r , define a multinomial $\psi_{p,r}$ over the K topics, each drawn from a Dirichlet parameterized by β .
3. Assume an infinite set of distributions over personas drawn from a base distribution H . Each of these $\theta_j \sim \text{Dir}(\alpha)$ is a multinomial over the P personas, with an associated probability of being selected π_j , drawn from the stick-breaking process with hyperparameter λ .
4. For each document i :
 - (a) Draw a cluster assignment $s_i \sim \pi$, with corresponding multinomial distribution over personas θ_{s_i} .
 - (b) For each entity e participating in i :
 - i. Draw e ’s persona $p_e \sim \theta_{s_i}$.
 - ii. For every $\langle r, w \rangle$ tuple associated with e in i , draw $z \sim \psi_{p_e,r}$ then $w \sim \phi_z$.

The DPM (Figure 1, left) has a similar generative story, except that each document has a unique distribution over personas. As such, step 4(a) is replaced with a draw from a symmetric Dirichlet distribution $\theta_i \sim \text{Dir}(\alpha)$.

3 Clustering Stories

The DPM assumes that each document has a unique distribution (θ_i) from which its personas are drawn. However, for entities mentioned in news articles (as well as for the dramatis personae of films), we would expect certain types of personas to occur together frequently, such as articles about lawmakers and laws. Thus we would like to cluster documents based on their “casts” of personas. To do this, we have added a Dirichlet process (DP) prior on the document-specific distribution over personas (step 3), which allows the number of clusters to adapt to the size and complexity of the corpus (Antoniak, 1974; Escobar and West, 1994).

Although the model admits an unbounded number of distributions over personas, the properties of DPs are such that the number used by D documents will tend to be much less than D . As a result, inference under this model provides topics ϕ (distributions over words) interpretable as textual descriptors of entities, personas ψ (distributions over reusable topics), and clusters of articles s with associated distributions over personas θ .

Following Bamman et al. (2013), we perform inference using collapsed Gibbs sampling, collapsing out the distributions over words (ϕ), topics (ψ), and personas (θ), as well as π . On each iteration, we first sample a cluster for each document, followed by a persona for each entity, followed by a topic for each tuple. Because we assume a conjugate base measure, sampling clusters can be done efficiently using the Chinese restaurant process (Aldous, 1985) for story types, personas, and topics, with slice sampling for hyperparameters ($\alpha, \beta, \gamma, \lambda$). Because such algorithms are well known to NLP readers, we have relegated details to the supplementary material.

During sampling, we discard samples from the first 10,000 iterations, and collect one sample from every tenth iteration for following 1,000 iterations. We sample hyperparameters every 20 iterations for the first 500 iterations, and every 100 thereafter.

4 Dataset

The Media Frames Corpus (MFC; Card et al., 2015) consists of annotations for approximately 4,200 articles about immigration taken from 13 U.S. newspapers over the years 1980–2012. The annotations

for these articles are in terms of a set of 15 general-purpose “framing dimensions” (such as Politics and Legality), developed to be broadly applicable to a variety of issues, and to be recognizable in text (by trained annotators). Each article has been annotated with a “primary frame” (the overall dominant aspect of immigration being emphasized), as well as an overall “tone” (pro, neutral, or anti), which is the extent to which a pro-immigration advocate would like to see the article in print, without implying any stance taken by the author.² The MFC contains at least two independent annotations for each article; agreement on the primary frame and tone was established through discussion in cases of initial disagreement. A complete list of these framing dimensions is given in the supplementary material.

In order to train our model on a larger collection of articles, we use the original corpus of articles from which the annotated articles in the MFC were drawn. This produces a corpus of approximately 37,000 articles about immigration; we train the persona model on this larger dataset, only using the smaller set for evaluation on a secondary task. Note that the MFC annotations are not used by our model; rather, we hypothesize that the personas it discovers may serve as features to help predict framing—this serves as one of our evaluations (§7).

5 Identifying Entities

The original focus of the DPM was on *named* characters in movies, which could be identified using named entity recognition and pronominal coreference (Bamman et al., 2013), or name matching for pre-defined characters (Bamman et al., 2014). Here, we are interested in applying our model to entities about which we assume no specific prior knowledge.

In order to include a broader set of entities, we preprocess the corpus and apply a series of filters. First, we obtain lemmas, part-of-speech tags, dependencies, coreference resolution, and named entities from the Stanford CoreNLP pipeline (Manning et al., 2014), as well as supersense tags from the AMALGrAM tagger (Schneider and Smith, 2015). For each document, we consider all tokens with a

²The MFC also contains more fine-grained annotations of spans of text which cue each of the framing dimensions, but we do not make use of those here.

NN* or PRP part of speech as possible entities, partially clustered by coreference. We then merge all clusters (including singletons) within each document that share a non-pronominal mention word.

Next, we exclude all clusters lacking at least one mention classified as a person, organization, location, group, object, artifact, process, or act (by CoreNLP or AMALGrAM). From these, we extract $\langle w, r, e, i \rangle$ tuples using extraction patterns lightly adapted from (Bamman et al., 2013). (The complete set of patterns are given in the supplementary material.) To further restrict the set of entities to those that have sufficient evidence, we construct a vocabulary for each of the three relations, and exclude words that appear less than three times in the corresponding vocabulary.³ We then apply one last filter to exclude entities that have fewer than three qualifying tuples across all mentions. From the dataset described in §4, we extract 128,655 entities, mentioned using 11,262 different mention words, with 575,910 tuples and 11,104 distinct $\langle r, w \rangle$ pairs.

6 Exploratory Analysis

Here we discuss our model, as estimated on the corpus of 37,000 articles discussed in §4 with 50 personas and 100 topics; these values were not tuned. A cursory examination of topics shows that each tends to be a group of either verbs or attributes. Personas, on the other hand, blend topics to include all three relation types. The estimated Dirichlet hyperparameters are all $\ll 1$, giving sparse (and hence easily scanned) distributions over personas, topics, and words.

Table 1 shows all 50 personas. For each p , we show (i) the mention words most strongly associated with p , and (ii) $\langle r, w \rangle$ pairs associated with the persona. (To save space, “I” denotes *immigrant*.) Recall that, like the Dirichlet persona model, our model says nothing about the mention words; they are *not* included as evidence during inference.⁴ Nonetheless, each persona is strongly associated with a

³We also exclude the lemma “say” as a stopword, as it is the most common verb in the corpus by an order of magnitude

⁴We did explore adding mention words as evidence, but they tended to dominate the relation tuples. Because our interest is in a richer set of framing devices than simply the words used to refer to people (and other entities), we consider here only the model based on the surrounding context.

sparse handful of mention words, and we find that labeling each persona by its most strongly associated mention word (excluding *immigrant*) is often sensible (these are capitalized in Table 1, though in some cases the relation words differentiate strongly (e.g., the *group* personas, IDs 17 and 18 in Table 1).

The model finds expected participants (such as *workers*, political *candidates*, and *refugees*), but also more conceptual entities, such as *laws*, *bills* (IDs 3, 37), and the U.S.-Mexican *border* (ID 5), which looms large in the immigration debate. Some interesting distinctions are discovered, such as two of the *worker* personas, one high-skilled and residing legally (ID 48), the other illegal (ID 49).

Using the original publication dates of the articles, we can estimate the frequency of appearance of each persona within immigration coverage by summing the posterior distribution over personas for each entity mention, and plotting these frequencies across time. (Note that time metadata is not given to the model as evidence.) We find immediately that personas can signal events. Figure 2 shows these temporal trajectories for a small, selected set of personas. Although *bills* and *laws* are conceptually similar, and have similar trajectories from 1980 to 2005, they are strongly divergent in 2006 and 2010. These are particularly notable years for immigration policy, corresponding to the failed Comprehensive Immigration Reform Act of 2006 (Senate bill S.2611) and Arizona’s controversial anti-immigration laws from 2010.⁵ Refugees, by contrast, show a marked spike around the year 2000. Inspection showed this persona to be strongly tied to the case of Elián González, which received a great deal of media attention in that year.

The main advantage of the extended model over the DPM is being able to cluster articles by “casts.” During sampling, thousands of clusters are created (and mostly destroyed). Ultimately, our inference procedure settled on approximately 110 clusters, and we consider two examples. Figure 3 shows the temporal trajectories of the two clusters with the greatest representation of the *refugee* persona. Both show the characteristic spike around the year 2000. The top personas for these two clusters are given in Table

⁵Other notable events which appear to be represented include the Illegal Immigration Reform and Immigrant Responsibility Act of 1996, and the Secure Fence Act of 2006.

ID	Mention words	Relations
1	AGENT police official authority	federal _m tell _p find _a arrest _a local _m tell _a
2	ASYLUM crime refugee asylum_seeker	political _m seek _p grant _p commit _p serious _m deny _p
3	BILL law immigration_reform measure	comprehensive _m pass _a pass _p make _a have _a support _p
4	BOAT van crime document	criminal _m other _m have _p use _a use _p be _a
5	BORDER border_patrol border_agent	mexican _m cross _p secure _p southern _m u.s.-mexico _m close _p
6	BUSH official mcnyary people I	have _a tell _a want _a tell _p former _m call _a
7	CANDIDATE bush romney leader	republican _m presidential _m democratic _m have _a call _a support _a
8	CARD document visa status	green _m new _m get _p temporary _m fake _m permanent _m
9	CARD visa state document	consular _m federal _m have _a mexican _m receive _p get _p
10	COMPANY country I state nation	have _a regional _m global _m rural _m take _a require _p
11	COUNTRY people I citizen united_states	american _m other _m enter _p have _a leave _p central _m
12	COUPLE marriage people I class	gay _m bilingual _m same-sex _m have _a prime _m seasonal _m
13	COURT lawsuit suit ruling	federal _m file _p rule _a civil _m file _a have _a
14	EMPLOYER company people business	hire _a have _a many _m require _p employ _a local _m
15	FENCE amendment law wall	real _m 14th _m virtual _m build _p be _a have _a
16	GOVERNMENT court judge official	federal _m local _m have _a rule _a ask _p other _m
17	GROUP deportation attack country	terrorist _m civil _m face _p armed _m islamic _m muslim _m
18	GROUP I voter people bush	hispanic _m immigrant _m local _m many _m want _a have _a
19	I ALIEN immigration people worker	illegal _m allow _p have _a legal _m undocumented _m live _a
20	I ALIEN people criminal inmate	illegal _m criminal _m deport _p immigrant _m detain _p release _p
21	I ALIEN worker immigration employer	illegal _m hire _p undocumented _m employ _p legal _m hire _a
22	I ALIEN worker people immigration	illegal _m arrest _p undocumented _m arrest _a charge _p transport _p
23	I CHILD worker people student	immigrant _m foreign-born _m have _a many _m come _a new _m
24	I GROUP people population business	new _m immigrant _m other _m many _m asian _m have _a
25	I GROUP program center city	new _m have _a first _m be _a other _m make _a
26	I IMMIGRATION alien worker	illegal _m legal _m hire _p have _a allow _p undocumented _m
27	I IMMIGRATION alien worker people	illegal _m legal _m have _a be _a come _a immigrant _m
28	I JEWS refugee israel child	soviet _m jewish _m russian _m have _a vietnamese _m israeli _m
29	I MAN alien refugee people	illegal _m chinese _m cuban _m arrest _p haitian _m find _p
30	I PEOPLE child student worker	many _m young _m have _a illegal _m come _a be _a
31	I PEOPLE country woman man	black _m muslim _m african _m have _a come _a korean _m
32	I WORKER people citizen job	american _m new _m have _a mexican _m illegal _m many _m
33	I WORKER resident student people	legal _m foreign _m permanent _m have _a allow _p skilled _m
34	I WORKER student people child	undocumented _m illegal _m immigrant _m have _a allow _p live _a
35	JOB I people immigration law	have _p have _a be _a take _p good _m make _a
36	JOB study survey I labor	find _a new _m find _p show _a fill _p take _p
37	LAW immigration_law bill measure	new _m federal _m enforce _p require _a pass _p allow _a
38	MAN I woman people haitians	deport _p have _a arrest _p hold _p release _p face _a
39	MAN people agent official I	arrest _p charge _p other _m former _m have _a face _a
40	MAN woman I people girl	tell _a kill _p have _a other _m young _m take _p
41	PEOPLE I child man woman	have _a come _a live _a go _a tell _p work _a
42	PROFILING violence abuse discrimination	racial _m domestic _m safe _m physical _m be _a affordable _m
43	PROGRAM system law agency	new _m national _m federal _m create _p use _p special _m
44	REFUGEE I boy people elian	cuban _m haitian _m chinese _m have _a allow _p return _p
45	SCHOOL people I family english	have _a high _m see _a come _a go _a be _a
46	SERVICE school care college	public _m medical _m provide _p deny _p receive _p attend _p
47	TRAFFICKING rights group flight	human _m international _m commercial _m be _a have _a take _a
48	WORKER I immigration student company	foreign _m legal _m skilled _m hire _p american _m have _a
49	WORKER I people woman man	mexican _m immigrant _m undocumented _m migrant _m illegal _m
50	YEAR program month income	fiscal _m last _m end _a next _m previous _m begin _a

Table 1: Personas with their associated mention words and relation tuples (*a* = agent, *p* patient, *m* = modifier/attribute); I denotes “immigrant.”




Figure 2: Temporal patterns of the mentions of selected personas.




Figure 3: Temporal patterns of two clusters with the greatest overall representation of the *refugee* persona.

2. Type A, which includes a story with the headline “Protesters vow to keep Elián in U.S.,” emphasizes political aspects, while type B (e.g., “Court says no to rights for refugees”) emphasizes legal aspects. Note that *Political* and *Legality* are two of the framing dimensions used in the MFC.

Do these persona-cast clusters relate to frames? For the five most common story clusters, (which have no overlap with the two refugee story types), Figure 4 shows the number of annotated articles with each of the primary frames if we assign each article to its most likely cluster. The second and fifth clusters correlate particularly well with primary frames (*Political* and *Crime*, respectively). This is further reinforced by looking at the most frequent persona for each of these story clusters which are *candidate* (ID 7) for the second and *immigrant* (ID 22), characterized by $illegal_m$ and $arrest_p$, for the fifth.

Refugee story cluster A		
Frequency	Persona	ID
0.49	REFUGEE immigrant boy	44
0.10	BUSH official mcinary	6
0.06	IMMIGRANT man alien	29
0.05	ASYLUM crime refugee	2
Refugee story cluster B		
Frequency	Persona	ID
0.29	MAN immigrant woman	38
0.23	REFUGEE immigrant boy	44
0.12	COURT lawsuit suit	13
0.10	GOVERNMENT court judge	16

Table 2: Truncated distribution over personas for the two clusters depicted in Figure 3. IDs index into Table 1.




Figure 4: Number of annotated articles in each of the five most frequent clusters, with colors showing the proportion of articles annotated with each primary frame.

7 Experiments: Personas and Framing

We evaluate personas as features for automatic analysis of framing and tone, as defined in the MFC (§4). Specifically, we build multi-class text classifiers (separately) for the primary frame and the tone of a news article, for which there are 15 and 3 classes, respectively. Because there are only a few thousand annotated articles, we applied 10-fold cross-validation to estimate performance.

Features are derived from our model by considering each persona and each story cluster as a potential feature. A document’s feature values for story types are the proportion of samples in which it was assigned to each cluster. Persona feature values are similarly derived by the proportion of samples in which each entity was assigned to each persona, with the persona values for each entity in each document summed into a single set of persona values per

Primary frame					
Features:	MF	(W)	(W,P ₁)	(W,P ₂)	(W,P ₂ ,S)
Accuracy:	0.174	0.529	0.537	*0.540	0.537
# Features:	0	3.9k	3.5k	3.5k	2.8k
Tone					
Accuracy:	0.497	0.628	0.631	0.628	0.630
# Features:	0	5.0k	5.0k	5.0k	4.0k

Table 3: Evaluation using a direct comparison to a simple baseline. Each model uses the union of listed features. (W = unigrams and bigrams, P₁ = personas from DPM, P₂ = personas from our model, S = story clusters; MF = always predict most frequent class.) * indicates a statistically significant difference compared to the (W) baseline ($p < 0.05$).

document. We did not use the topics (z) discovered by our model as features.

7.1 Experiment 1: Direct Comparison

For the first experiment, we train independent multi-class logistic regression classifiers for predicting primary frame and tone. We consider adding persona and/or story cluster features to baseline classifiers based only on unigrams and bigrams with binarized counts, a simple but robust baseline (Wang and Manning, 2012).⁶ In all cases, we use L_1 regularization and use 5-fold cross validation within each split’s training set to determine the strength of regularization. We then repeat this for each of the 10 folds, thereby producing one prediction (of primary frame and tone) for every annotated article. The results of this experiment are given in Table 3; for predicting the primary frame, classifiers that used persona and/or story cluster features achieve higher accuracy than the bag-of-words baseline (W); the classifier using personas from our model but not story clusters is significantly better than the baseline.⁷ The enhanced models are also more compact, on average, using fewer effective features. A benefit to predicting tone is also observed, but it did not reach statistical significance.

7.2 Experiment 2: Automatic Evaluation

Although bag-of- n -grams models are known to be a strong baseline for text classification, researchers familiar with the extensive catalogue of features of-

⁶We also binarized the persona feature values.

⁷Two-tailed McNemar’s test ($p < 0.05$).

ferred by NLP will potentially see them as a straw man. We propose a new and more rigorous method of comparison, in which a wide range of features are offered to an automatic model selection algorithm for each of the prediction tasks, with the features to be evaluated withheld from the baseline.

Because no single combination of features and regularization strength is best for all situations, it is an empirical question which features are best for each task. We therefore make use of Bayesian optimization (Bayesopt) to make as many modeling decisions as possible (Pelikan, 2005; Snoek et al., 2012; Bergstra et al., 2015; Yogatama et al., 2015).

In particular, let F be the set of features that might be used as input to any text classification algorithm. Let f be a new feature that is being proposed. Allow the inclusion or exclusion of each feature in the feature set to be a hyperparameter to be optimized, along with any additional decisions such as input transformations (e.g., lowercasing), and feature transformations (e.g., normalization). Using an automatic model selection algorithm such as Bayesian optimization, allow the performance on the validation set to guide choices about all of these hyperparameters on each iteration, and set up two independent experiments.

For the first condition, A_1 , allow the algorithm access to all features in F . For the second, A_2 , allow the algorithm access to all features in $F \cup f$. After R iterations of each, choose the best model or the best set of models from each of A_1 and A_2 (M_1 and M_2 , respectively), based on performance on the validation set. Finally, compare the selected models in terms of performance on the test set (using an appropriate metric such as F_1), and examine the features included in each of the best models. If f is a helpful feature, we should expect to see that, a) $F_1(M_2) > F_1(M_1)$, and b), f is included in the best model(s) found by A_2 .

If $F_1(M_2) > F_1(M_1)$ but f is not included in the best models from A_2 , this suggests that the performance improvement may simply be a matter of chance, and there is no evidence that f is helpful. By contrast, if f is included in the best models, but $F_1(M_2)$ is not significantly better than $F_1(M_1)$, this suggests that f is offering some value, perhaps in a more compressed form of the useful signal from other features, but does not actually offer better per-

Features:	(B)	(B,P ₁)	(B,P ₂)	(B,P ₂ ,S)
Primary frame	0.566	0.568	0.568	0.567
Tone	0.667	0.671	0.667	0.671

Table 4: Mean accuracy of the best three iterations from Bayesian optimization (chosen based on validation accuracy). (B = features from many NLP tools, P₁=personas from the DPM, P₂ = personas from our model, S=story clusters.)

formance.

For this experiment, we use the tree-structured Parzen estimator for Bayesian optimization (Bergstra et al., 2015), with L_1 -regularized logistic regression as the underlying classifier, and set $R = 40$. In addition to the entities and story clusters identified by these models, we allow these classifiers access to a large set of features, including unigrams, bigrams, parts of speech, named entities, dependency tuples, ordinal sentiment values (Manning et al., 2014), multi-word expressions (Justeson and Katz, 1995), supersense tags (Schneider and Smith, 2015), Brown clusters (Brown et al., 1992), frame semantic features (Das et al., 2010), and topics produced by standard LDA (Blei et al., 2003). The inclusion or exclusion of each feature is determined automatically on each iteration, along with feature transformations (removal of rare words, lowercasing, and binary or normalized counts).

The baseline, denoted “B,” offers all features except personas and story clusters to Bayesopt; we consider adding DPM personas, our model’s personas, and our model’s personas and story clusters. Table 4 shows test-set accuracy for each setup, averaged across the three best models returned by Bayesopt.

Using this more rigorous form of evaluation, approximately the same accuracy is obtained in all experimental conditions. However, we can still gain insight into which features are useful by examining those selected by the best models in each condition. For primary frame prediction, both personas and story clusters are included by the best models in every case where they have been offered as possible features, as are unigrams, dependency tuples, and semantic frames. Other commonly-selected features include bigrams and part of speech tags. For predicting tone, personas are only included by half of the best models, with the most common features be-

ing unigrams, bigrams, semantic frames, and Brown clusters. As expected, the best models in each condition obtain better performance than the models from experiment 1, thanks to the inclusion of additional features and transformations.

This secondary evaluation suggests that for this task, persona features are useful in predicting the primary frame, but are unable to offer improved performance over existing features, such as semantic frames. However, the fact that both personas and story clusters are included by all the best models for predicting the primary frame suggests that they are competitive with other features, and perhaps offer useful information in a more compact form.

8 Qualitative Evaluation

Prior to exposure to any output of our model, one of the co-authors on this paper (Gross, who has expertise in both framing and the immigration issue) prepared a list of personas he expected to frequently occur in American news coverage of immigration. Given the example of the “skilled immigrant,” he listed 22 additional named personas, along with a few examples of things they do, things done to them, and attributes.

The list he prepared includes several different characterizations of immigrants (low-skilled, unauthorized, legal, citizen children, undocumented children, refugees, naturalized citizens), non-immigrant personas (U.S. workers, smugglers, politicians, officials, border patrol, vigilantes), related pairs (pro / anti advocacy groups, employers / guest workers, criminals / victims), and a few more conceptual entities (the border, bills, executive actions). Of these, almost all are arguably represented in the personas we have discovered. However, there is rarely a perfect one-to-one mapping: predefined personas are sometimes merged (e.g., “the border” and “border patrols”) or split (e.g., legislation, employers, and various categories of immigrants). Personas which don’t emerge from our model include smugglers, guest workers, vigilantes, and victims of immigrant criminals. On the other hand, our model proposes far more non-person entities, such as ID cards, courts, companies, jobs, and programs.

These partial matchings between predefined personas and the results of our model are generally

identifiable by comparing the names given to the predefined personas to the the most commonly occurring mention words and attributes of our discovered personas. The attributes and action words given to the predefined personas are harder to evaluate, as many of them are rare (e.g. politicians “vacillate”) or compound phrases (e.g. low-skilled immigrants “do jobs Americans won’t do”) that tend to miss the more obvious properties captured by our model. For example, the *employer* persona captured by our model engages in actions like *hire*, *employ*, and *pay*. By contrast, the terms given for the predefined “business owners” persona are “lobby” and “rely on immigrant labor.” Our unsupervised discovery of this persona can clearly be matched to the predefined persona in this case, but doesn’t provide such fine-grained insight into how they might be characterized.

The best match between predefined and discovered personas is the U.S.-Mexican border. Of the words given for the predefined persona, almost all are more frequently associated with *border* than with any other discovered persona (“Mexican-U.S.,” “lawless,” “porous,” “unprotected,” “guarded,” and “militarized”). The most commonly associated words discovered by our model that are missing from the predefined description include *crossed*, *secured*, *southern*, and *closed*.

While this qualitative evaluation helps to demonstrate the face validity of our model, it would be better to have a more comprehensive set of predefined personas, based on input from additional experts. Moreover, it also illustrates the challenge of trying to match the output of an unsupervised model to expected results. Not only is some merging and splitting of categories inevitable, there was a mismatch in this case in the types of entities to be described (people as opposed to more abstract entities), and the ways of describing them (rare but specific words as opposed to more generic but potentially obvious terms).

9 Related Work

Much NLP has focused on identifying entities or events (Ratinov and Roth, 2009; Ritter et al., 2012), analyzing schemes or narrative events in terms of characters (Chambers and Jurafsky, 2009), inferring

the relationships between entities (O’Connor et al., 2013; Iyyer et al., 2016), and predicting personality types from text (Flekova and Gurevych, 2015). Bamman also applied variants of the DPM to characters in novels (Bamman et al., 2014).

Previous work on sentiment, stance, and opinion mining has focused on recognizing stance or political sentiment in online ideological debates (Somasundaran and Wiebe, 2010; Hasan and Ng, 2014; Sridhar et al., 2015), and other forms of social media (O’Connor et al., 2010; Agarwal et al., 2011), and recently through the lens of connotation frames (Rashkin et al., 2016). Opinion mining and sentiment analysis are the subject of ongoing research in NLP and have long served as test platforms for new methodologies (Socher et al., 2013; İrsoy and Cardie, 2014; Tai et al., 2015)

Framing is arguably one of the most important concepts in the social sciences, with roots in to sociology, psychology, and mass communication (Gitlin, 1980; Benford and Snow, 2000; D’Angelo and Kuypers, 2010); the scope and relevance of framing is widely debated (Rees et al., 2001), with many authors applying the concept of framing to analyzing documents on particular issues (Baumgartner et al., 2008; Berinsky and Kinder, 2006).

10 Conclusion

We have extended models for discovering latent personas to simultaneously cluster documents by their “casts” of personas. Our exploration of the model’s inferences and their incorporation into a challenging text analysis task—characterizing coarse-grained framing in news articles—demonstrate that personas are a useful abstraction when applying NLP to social-scientific inquiry. Finally, we introduced a Bayesian optimization approach to rigorously assess the usefulness of new features in machine learning tasks.

Acknowledgments

The authors thank members of the ARK group and anonymous reviewers for helpful feedback on this work. This research was made possible by a Natural Sciences and Engineering Research Council of Canada Postgraduate Scholarship (to D.C.), a Bloomberg Data Science Research Grant (to J.H.G., A.E.B., and N.A.S.), and a University of Washington Innovation Award (to N.A.S.).

References

- Apoorv Agarwal, Boyi Xie, Ilia Vovsha, Owen Rambow, and Rebecca J. Passonneau. 2011. Sentiment analysis of twitter data. In *Proc. of Frame Semantics in NLP: A Workshop in Honor of Chuck Fillmore (1929-2014)*.
- D. Aldous. 1985. Exchangeability and related topics. In *École d'Été St Flour 1983*, pages 1–198. Springer-Verlag.
- Charles E. Antoniak. 1974. Mixtures of dirichlet processes with applications to bayesian nonparametric problems. *Annals of Statistics*, 2(6), November.
- David Bamman, Brendan O'Connor, and Noah A. Smith. 2013. Learning latent personas of film characters. In *Proc. of ACL*.
- David Bamman, Ted Underwood, and Noah A. Smith. 2014. A bayesian mixed effects model of literary character. In *Proc. of ACL*.
- 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 *Proc. of NAACL*.
- 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.
- Robert D. Benford and David A. Snow. 2000. Framing processes and social movements: An overview and assessment. *Annual Review of Sociology*, 26:611–639.
- James Bergstra, Brent Komer, Chris Eliasmith, Dan Yamins, and David D Cox. 2015. Hyperopt: a python library for model selection and hyperparameter optimization. *Computational Science and Discovery*, 8(1).
- Adam J. Berinsky and Donald R. Kinder. 2006. Making sense of issues through media frames: Understanding the Kosovo crisis. *Journal of Politics*, 68(3):640–656.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet allocation. *J. Mach. Learn. Res.*, 3:993–1022.
- Peter F. Brown, Peter V. deSouza, Robert L. Mercer, Vincent J. Della Pietra, and Jenifer C. Lai. 1992. Class-based N-gram models of natural language. *Computational Linguistics*, 18(4):467–479.
- 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 *Proc. of ACL*.
- Nathanael Chambers and Dan Jurafsky. 2009. Unsupervised learning of narrative schemas and their participants. In *Proc. of ACL*.
- Eunsol Choi, Chenhao Tan, Lillian Lee, Cristian Danescu-Niculescu-Mizil, and Jennifer Spindel. 2012. Hedge detection as a lens on framing in the GMO debates: A position paper. In *Proc. of Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics*, pages 70–79.
- Paul D'Angelo and Jim A. Kuypers. 2010. *Doing News Framing Analysis*. Routledge.
- Dipanjana Das, Nathan Schneider, Desai Chen, and Noah A. Smith. 2010. Probabilistic frame-semantic parsing. In *Proc. of NAACL*.
- Robert M. Entman. 2007. Framing bias: Media in the distribution of power. *Journal of Communication*, 57(1):163–173.
- Michael D. Escobar and Mike West. 1994. Bayesian density estimation and inference using mixtures. *J. Amer. Statist. Assoc.*, 90:577–588.
- Ronen Feldman. 2013. Techniques and applications for sentiment analysis. *Commun. ACM*, 56(4):82–89.
- Lucie Flekova and Iryna Gurevych. 2015. Personality profiling of fictional characters using sense-level links between lexical resources. In *Proc. of EMNLP*.
- Matthew Gentzkow and Jesse M. Shapiro. 2010. What drives media slant? Evidence from U.S. daily newspapers. *Econometrica*, 78(1):35–71.
- Todd Gitlin. 1980. *The Whole World is Watching*. Berkeley: University of California Press.
- Stephan Greene and Philip Resnik. 2009. More than words: Syntactic packaging and implicit sentiment. In *Proc. of ACL*.
- Glenn Greenwald. 2014. *No Place to Hide*. Picador.
- Eric Hardisty, Jordan L. Boyd-Graber, and Philip Resnik. 2010. Modeling perspective using adaptor grammars. In *Proc. of EMNLP*.
- Kazi Saidul Hasan and Vincent Ng. 2013. Stance classification of ideological debates: Data, models, features, and constraints. In *Proc. of IJCNLP*.
- Kazi Saidul Hasan and Vincent Ng. 2014. Why are you taking this stance? identifying and classifying reasons in ideological debates. In *Proc. of EMNLP*.
- Edward S. Herman and Noam Chomsky. 1988. *Manufacturing Consent*. Vintage.
- Ozan İrsoy and Claire Cardie. 2014. Opinion mining with deep recurrent neural networks. In *Proc. of EMNLP*.
- Mohit Iyyer, Peter Enns, Jordan L. Boyd-Graber, and Philip Resnik. 2014. Political ideology detection using recursive neural networks. In *Proc. of ACL*.
- 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 *Proc. of NAACL*.
- J. Justeson and S. Katz. 1995. Technical terminology: some linguistic properties and an algorithm for identification in text. *Natural Language Engineering*.
- Jure Leskovec, Lars Backstrom, and Jon Kleinberg. 2009. Meme-tracking and the dynamics of the news cycle. In *Proc. of KDD*.

- Wei-Hao Lin, Theresa Wilson, Janyce Wiebe, and Alexander Hauptmann. 2006. Which side are you on? Identifying perspectives at the document and sentence levels. In *Proc. of CoNLL*.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proc. of ACL*.
- Viet-An Nguyen, Jordan Boyd-Graber, and Philip Resnik. 2013. Lexical and hierarchical topic regression. In *Proc. of NIPS*.
- Viet-An Nguyen, Jordan Boyd-Graber, Philip Resnik, and Kristina Miler. 2015. Tea party in the house: A hierarchical ideal point topic model and its application to Republican legislators in the 112th congress. In *Proc. of ACL*.
- Vlad Niculae, Caroline Suen, Justine Zhang, Cristian Danescu-Niculescu-Mizil, , and Jure Leskovec. 2015. QUOTUS: The structure of political media coverage as revealed by quoting patterns. In *Proceedings of WWW 2015*.
- Brendan T. O'Connor, Ramnath Balasubramanian, Bryan R. Routledge, and Noah A. Smith. 2010. From tweets to polls: Linking text sentiment to public opinion time series. In *ICWSM*.
- Brendan O'Connor, Brandon M. Stewart, and Noah A. Smith. 2013. Learning to extract international relations from political context. In *Proc. of ACL*.
- Zhongdang Pan and Gerald M. Kosicki. 1993. Framing analysis: An approach to news discourse. *Political communication*, 10(1):55–75.
- Bo Pang and Lillian Lee. 2008. Opinion mining and sentiment analysis. *Found. Trends Inf. Retr.*, 2(1-2).
- M Pelikan. 2005. Bayesian optimization algorithm. In *Hierarchical Bayesian optimization algorithm*, pages 31–48. Springer.
- Hannah Rashkin, Sameer Singh, and Yejin Choi. 2016. Connotation frames: A data-driven investigation. In *Proc. of ACL*.
- Lev Ratinov and Dan Roth. 2009. Design challenges and misconceptions in named entity recognition. In *Proc. of CoNLL*.
- Stephen D. Rees, Oscar H. Gandy Jr., , and August E. Grant, editors. 2001. *Framing Public Life*. Routledge.
- Alan Ritter, Mausam, Oren Etzioni, and Sam Clark. 2012. Open domain event extraction from twitter. In *KDD*.
- Anne Schneider and Helen Ingram. 1993. Social construction of target populations: Implications for politics and policy. *The American Political Science Review*, 87(2):334–347.
- Nathan Schneider and Noah A. Smith. 2015. A corpus and model integrating multiword expressions and supersenses. In *Proc. of ACL*.
- Sameer Singh, Amarnag Subramanya, Fernando Pereira, and Andrew McCallum. 2012. Wikilinks: A large-scale cross-document coreference corpus labeled via links to Wikipedia. Technical Report UM-CS-2012-015, University of Massachusetts, Amherst.
- David A. Smith, Ryan Cordell, and Elizabeth Maddock Dillon. 2013. Infectious texts: modeling text reuse in nineteenth-century newspapers. In *Proc. of IEEE International Conference on Big Data*.
- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. 2012. Practical bayesian optimization of machine learning algorithms. In *Proc. of NIPS*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proc. of EMNLP*.
- Swapna Somasundaran and Janyce Wiebe. 2010. Recognizing stances in ideological on-line debates. In *Proceedings of the Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*.
- Dhanya Sridhar, James Foulds, Bert Huang, Lise Getoor, and Marilyn Walker. 2015. Joint models of disagreement and stance in online debate. In *Proc. of ACL*.
- Kai Sheng Tai, Richard Socher, and Christopher D. Manning. 2015. Improved semantic representations from tree-structured long short-term memory networks. In *Proc. of ACL*.
- Oren Tsur, Dan Calacci, and David Lazer. 2015. Frame of mind: Using statistical models for detection of framing and agenda setting campaigns. In *Proc. of ACL*.
- Baldwin Van Gorp. 2010. Strategies to take subjectivity out of framing analysis. In Paul D'Angelo and Jim A. Kuypers, editors, *Doing News Framing Analysis*, chapter 4, pages 84–109. Routledge.
- Marilyn A. Walker, Pranav Anand, Robert Abbott, and Ricky Grant. 2012. Stance classification using dialogic properties of persuasion. In *Proc. of NAACL*.
- Sida Wang and Christopher D Manning. 2012. Baselines and bigrams: Simple, good sentiment and topic classification. In *Proc. of ACL*.
- Dani Yogatama, Lingpeng Kong, and Noah A. Smith. 2015. Bayesian optimization of text representations. In *EMNLP*.