



蘇州大學

SOOCHOW UNIVERSITY

A Stylometric Inquiry into Hyperpartisan and Fake News

Martin Potthast, Johannes Kiesel, Kevin Reinartz,
Janek Bevendorff, Benno Stein
Leipzig University, Bauhaus-Universität Weimar

Longxiang Shen
NLP Lab, Soochow University





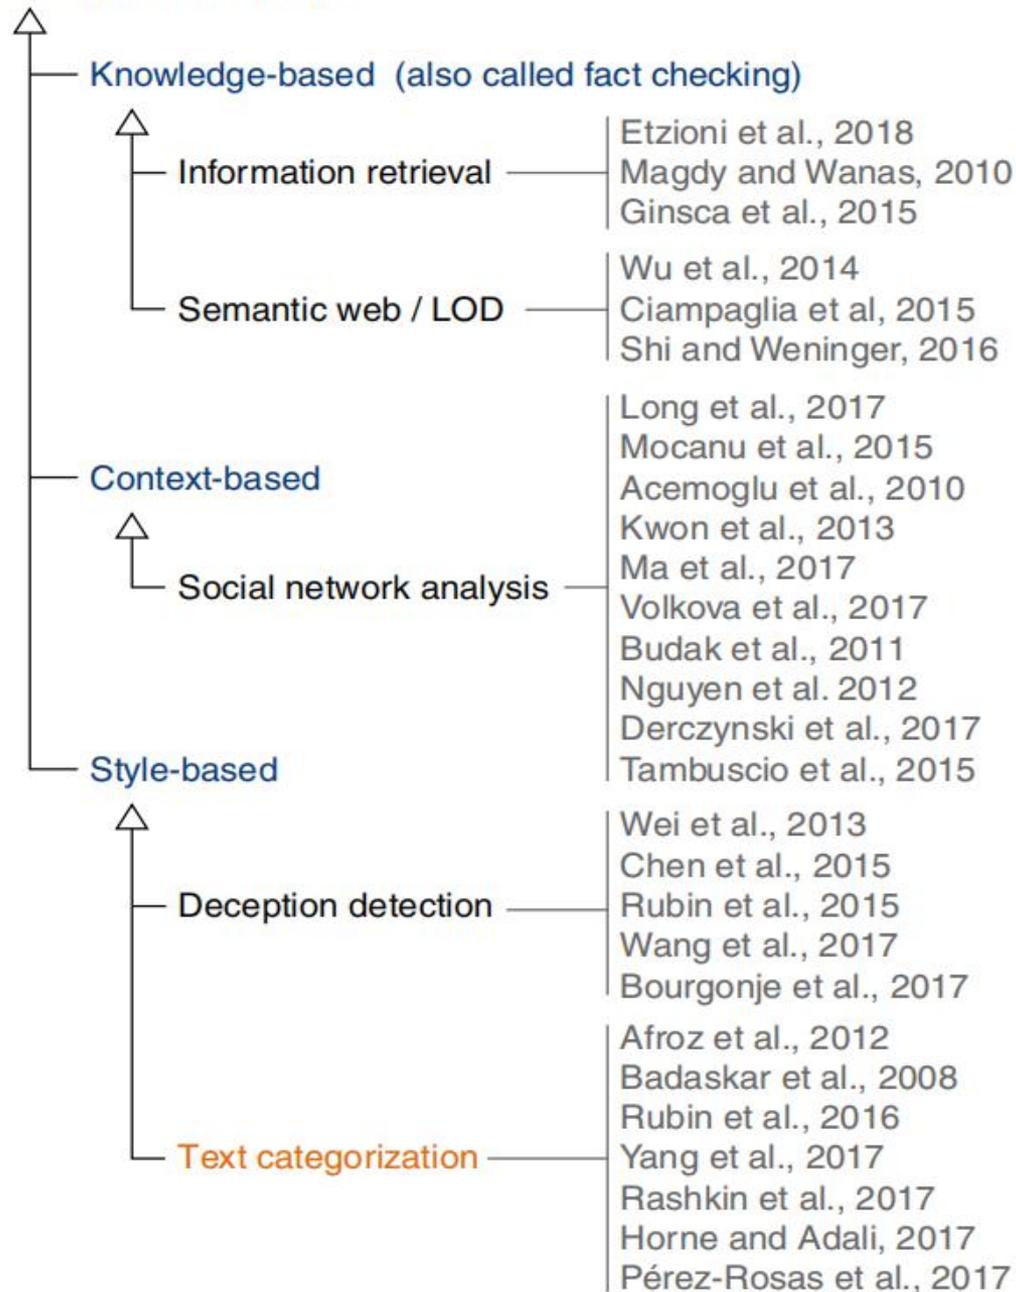
Task

- Fake news detection
 - Knowledge-based (also called fact checking)
 - Context-based
 - Style-based



- Fake news detection
 - Knowledge-based
 - Context-based
 - Style-based

Fake news detection





Motivation

- Construct a large new corpus with respect to veracity and hyperpartisanship
- Extensive experiments on discriminating fake news, hyperpartisan news and satire
- Validation experiments to verify the finding that the writing style of the left and the right have in common



Corpus Construction

- Manual fact-checking
 - Mostly true
 - Mixture of true and false (mix, for short)
 - Mostly false
 - No factual content (n/a, for short)
- fake news
-
- A diagram consisting of two arrows. One arrow originates from the text 'Mostly false' and points towards the text 'fake news'. A second arrow originates from the text 'No factual content (n/a, for short)' and also points towards the text 'fake news'.



Corpus Statistics

<i>Orientation</i> Publisher	Fact-checking results					Key statistics per article				
	true	mix	false	n/a	Σ	Paras.	Links		Words	
							extern	all	quoted	all
<i>Mainstream</i>	806	8	0	12	826	20.1	2.2	3.7	18.1	692.0
ABC News	90	2	0	3	95	21.1	1.0	4.8	21.0	551.9
CNN	295	4	0	8	307	19.3	2.4	2.5	15.3	588.3
Politico	421	2	0	1	424	20.5	2.3	4.3	19.9	798.5
<i>Left-wing</i>	182	51	15	8	256	14.6	4.5	4.9	28.6	423.2
Addicting Info	95	25	8	7	135	15.9	4.4	4.5	30.5	430.5
Occupy Democrats	55	23	6	0	91	10.9	4.1	4.7	29.0	421.7
The Other 98%	32	3	1	1	30	20.2	6.4	7.2	21.2	394.5
<i>Right-wing</i>	276	153	72	44	545	14.1	2.5	3.1	24.6	397.4
Eagle Rising	107	47	25	36	214	12.9	2.6	2.8	17.3	388.3
Freedom Daily	48	24	22	4	99	14.6	2.2	2.3	23.5	419.3
Right Wing News	121	82	25	4	232	15.0	2.5	3.6	33.6	396.6
Σ	1264	212	87	64	1627	17.2	2.7	3.7	20.6	551.0



Methodology

- Style Feature selection (common features)

- N-grams, n in [1, 3]

- Stop words

- Parts-of-speech

- 10 readability scores →

Automated Readability Index, Coleman Liau Index, Flesh Kincaid Grade Level and Reading Ease, Gunning Fog Index, LIX, McAlpine EFLAW Score, RIX, SMOG Grade, Strain Index

- Dictionary features →

Frequency of words from a tailor-made dictionary in a document



Methodology

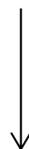
- Style Feature selection (domain-specific features)
 - Ratios of quoted words
 - External links
 - Number of paragraphs
 - Average length



Methodology

- Unmasking

- Takes two documents as input and outputs its confidence whether they have been written by the same author



authorship verification

Moshe Koppel, Jonathan Schler, and Elisheva Bonchek-Dokow. 2007. Measuring differentiability: Unmasking pseudonymous authors.



Methodology

- Baselines
 - A topic-based bag of words model
 - A model using only the domain-specific news style features
 - Naive baselines that classify all items into one of the categories in question



Outline

- Can (left/right) hyperpartisanship be distinguished from the mainstream ?
- Is style-based fake news detection feasible ?
- Can fake news be distinguished from satire ?



Outline

- Can (left/right) hyperpartisanship be distinguished from the mainstream ?
- Is style-based fake news detection feasible ?
- Can fake news be distinguished from satire ?



Experiments and Results

- Hyperpartisanship vs. Mainstream
 - Orientation cannot be discriminated based on the basic corpus characteristics observed with respect to paragraphs, quotations, and hyperlinks.

Features	Accuracy	Precision			Recall			F ₁		
	all	left	right	main.	left	right	main.	left	right	main.
Style	0.60	0.21	0.56	0.75	0.20	0.59	0.74	0.20	0.57	0.75
Topic	0.64	0.24	0.62	0.72	0.15	0.54	0.86	0.19	0.58	0.79
News style	0.39	0.09	0.35	0.59	0.14	0.36	0.49	0.11	0.36	0.53
All-left	0.16	0.16	-	-	1.00	0.0	0.0	0.27	-	-
All-right	0.33	-	0.33	-	0.0	1.00	0.0	-	0.50	-
All-main.	0.51	-	-	0.51	0.0	0.0	1.00	-	-	0.68

Performance of predicting orientation.



Experiments and Results

- Hyperpartisanship vs. Mainstream
 - Style-based classifier, outperforming the topic baseline

Features	Accuracy	Precision		Recall		F ₁	
	all	hyp.	main.	hyp.	main.	hyp.	main.
Style	0.75	0.69	0.86	0.89	0.62	0.78	0.72
Topic	0.71	0.66	0.79	0.83	0.60	0.74	0.68
News style	0.56	0.54	0.58	0.65	0.47	0.59	0.52
All-hyp.	0.49	0.49	-	1.00	0.0	0.66	-
All-main.	0.51	-	0.51	0.0	1.00	-	0.68

Performance of predicting hyperpartisanship.



Experiments and Results

- Hyperpartisanship vs. Mainstream

- Riddle

How could it be that hyperpartisanship in general can be much better discriminated from the mainstream than individual orientation?

- Hypothesis

Perhaps, the writing style of hyperpartisan left and right are not altogether different, despite their opposing agendas.



Experiments and Results

- Hyperpartisanship vs. Mainstream
 - Full style-based classifiers have a tendency of classifying left as right and right as left

Features	Left		Right	
	Trained on: right+main.	all	left+main.	all
Style	0.74	0.90	0.66	0.89
Topic	0.68	0.79	0.48	0.85
News style	0.52	0.61	0.47	0.66

Ratio of left articles misclassified right when omitting left articles from training, and vice versa.



Experiments and Results

- Hyperpartisanship vs. Mainstream

- Explanation

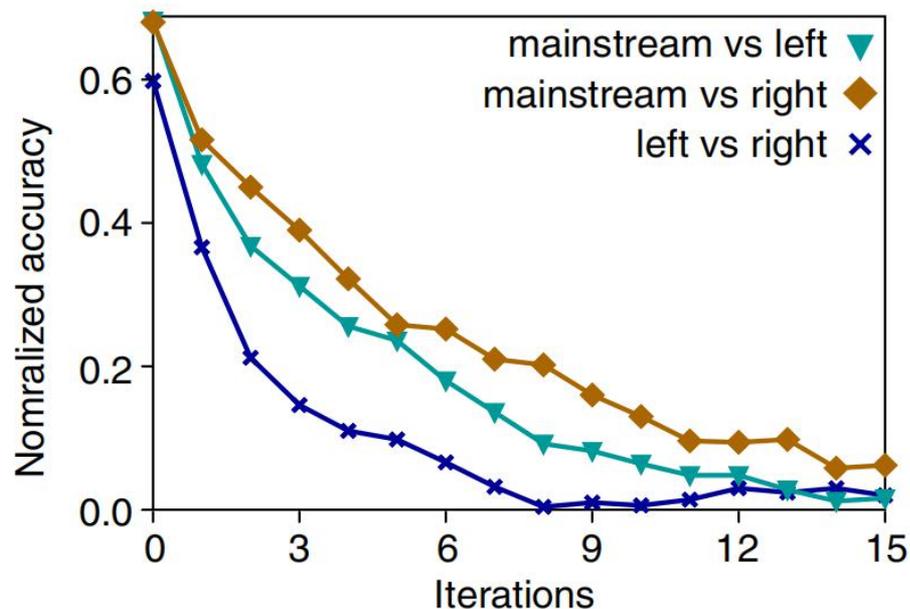
Leading up to the elections, the hyperpartisan left was often merely reacting to topics prompted by the hyperpartisan right, instead of bringing up their own.

Features	Left		Right	
	Trained on: right+main.	all	left+main.	all
Style	0.74	0.90	0.66	0.89
Topic	0.68	0.79	0.48	0.85
News style	0.52	0.61	0.47	0.66



Experiments and Results

- Hyperpartisanship vs. Mainstream
 - the curve comparing left vs. right has a distinctly steeper slope than either of the others



Unmasking applied to pairs of political orientations.
The steeper a curve, the more similar the respective styles.



Experiments and Results

- Discussion

- Experiments support the hypothesis that the hyperpartisan left and the hyperpartisan right have more in common in terms of writing style than any of the two have with the mainstream.



- Hyperpartisan news can apparently be discriminated well from the mainstream.



Outline

- Can (left/right) hyperpartisanship be distinguished from the mainstream ?
- Is style-based fake news detection feasible ?
- Can fake news be distinguished from satire ?



Experiments and Results

- Fake vs. Real

- Style-based fake news classification simply does not work in general

Features	Accuracy	Precision		Recall		F ₁	
	all	fake	real	fake	real	fake	real
<i>Generic classifier</i>							
Style	0.55	0.42	0.62	0.41	0.64	0.41	0.63
Topic	0.52	0.41	0.62	0.48	0.55	0.44	0.58
<i>Orientation-specific classifier</i>							
Style	0.55	0.43	0.64	0.49	0.59	0.46	0.61
Topic	0.58	0.46	0.65	0.45	0.66	0.46	0.66
All-fake	0.39	0.39	-	1.00	0.0	0.56	-
All-real	0.61	-	0.61	0.0	1.00	-	0.76

Performance of predicting veracity.



Experiments and Results

- Fake vs. Real

- Style-based fake news classification simply does not work in general.

Features	Accuracy	Precision		Recall		F ₁	
	all	fake	real	fake	real	fake	real
<i>Generic classifier</i>							
Style	0.55	0.42	0.62	0.41	0.64	0.41	0.63
Topic	0.52	0.41	0.62	0.48	0.55	0.44	0.58
<i>Orientation-specific classifier</i>							
Style	0.55	0.43	0.64	0.49	0.59	0.46	0.61
Topic	0.58	0.46	0.65	0.45	0.66	0.46	0.66
All-fake	0.39	0.39	-	1.00	0.0	0.56	-
All-real	0.61	-	0.61	0.0	1.00	-	0.76

Individually trained on articles from either orientation



Outline

- Can (left/right) hyperpartisanship be distinguished from the mainstream ?
- Is style-based fake news detection feasible ?
- Can fake news be distinguished from satire ?



Experiments and Results

- Fake vs. Real vs. Satire

- Satire ← Not all fake news are the same

Takes the form of news but lies more or less obviously to amuse its readers



Should be discriminated from other fakes



Experiments and Results

- Fake vs. Real vs. Satire
 - Style-based model significantly outperforms all baselines across the board

Features	Accuracy	Precision		Recall		F ₁	
	all	sat.	real	sat.	real	sat.	real
Style	0.82	0.84	0.80	0.78	0.85	0.81	0.82
Topic	0.77	0.78	0.75	0.74	0.79	0.76	0.77
All-sat.	0.50	0.50	-	1.00	0.0	0.67	-
All-real	0.50	-	0.50	0.00	1.00	-	0.67
Rubin et al.	n/a	0.90	n/a	0.84	n/a	0.87	n/a

Performance of predicting satire (sat.)



Experiments and Results

- Fake vs. Real vs. Satire
 - Style-based model significantly outperforms all baselines across the board.

Features	Accuracy	Precision		Recall		F ₁	
	all	sat.	real	sat.	real	sat.	real
Style	0.82	0.84	0.80	0.78	0.85	0.81	0.82
Topic	0.77	0.78	0.75	0.74	0.79	0.76	0.77
All-sat.	0.50	0.50	-	1.00	0.0	0.67	-
All-real	0.50	-	0.50	0.00	1.00	-	0.67
Rubin et al.	n/a	0.90	n/a	0.84	n/a	0.87	n/a



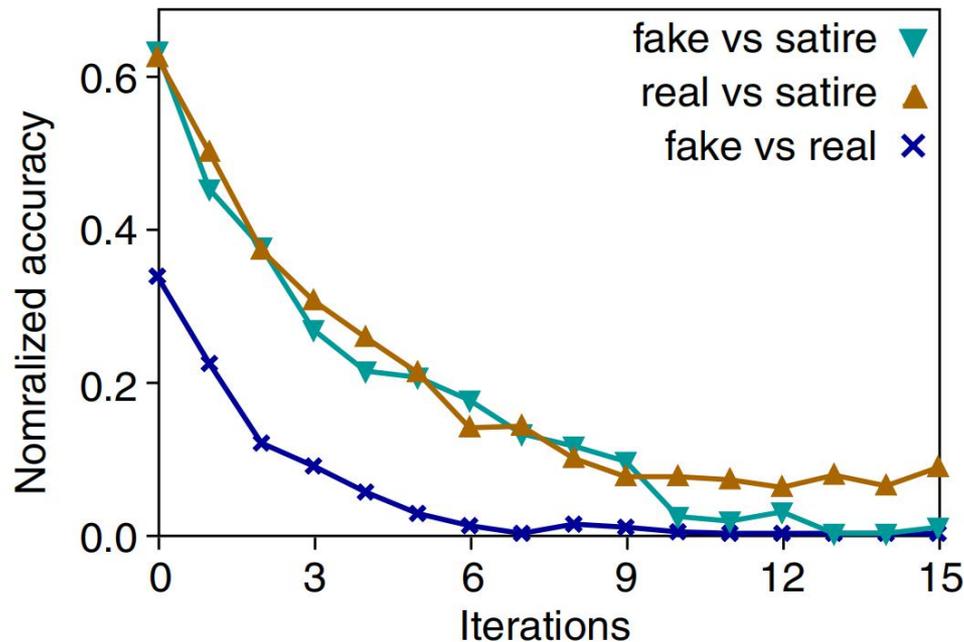
Includes features based on topic, absurdity, grammar, and punctuation



Experiments and Results

- Fake vs. Real vs. Satire

- The style of fake news has more in common with that of real news than either of the two have with satire.



Unmasking applied to pairs of sets of news that are fake, real, and satire.



Experiments and Results

- Discussion

- Satire is distinct enough from fake and real news and it can be discriminated with reasonable accuracy.



Just like with hyperpartisan news compared to mainstream news



Conclusion

- Notwithstanding the many attacks on fake news by developing one way or another of fact-checking, we believe it worthwhile to mount our attack from another angle: writing style.
- News articles conveying a hyperpartisan world view can be distinguished from more balanced news by writing style alone.
- The writing styles of news of the two opposing orientations are in fact very similar.



Conclusion

- Notwithstanding the many attacks on fake news by developing one way or another of fact-checking, we believe it worthwhile to mount our attack from another angle: writing style.
- News articles conveying a hyperpartisan world view can be distinguished from more balanced news by writing style alone.
- The writing styles of news of the two opposing orientations are in fact very similar.



There appears to be a common writing style of left and right extremism.



Conclusion

- Notwithstanding the many attacks on fake news by developing one way or another of fact-checking, we believe it worthwhile to mount our attack from another angle: writing style.
- News articles conveying a hyperpartisan world view can be distinguished from more balanced news by writing style alone.
- The writing styles of news of the two opposing orientations are in fact very similar.
- Satire can be distinguished well from other news.



Conclusion

- Notwithstanding the many attacks on fake news by developing one way or another of fact-checking, we believe it worthwhile to mount our attack from another angle: writing style.
- News articles conveying a hyperpartisan world view can be distinguished from more balanced news by writing style alone.
- The writing styles of news of the two opposing orientations are in fact very similar.
- Satire can be distinguished well from other news.

↓
Ensuring that humor will not be outcast by fake news detection technology.



Thank you!

