

A log-linear model for unsupervised text normalization

Yi Yang and Jacob Eisenstein

Georgia Institute of Technology

Lexical variation

- ▶ Social media language is different from standard English.
 - ▶ **Tweet:** ya ur website suxx Brah
 - ▶ **Standard:** yeah, your website sucks bro

Lexical variation

- ▶ Social media language is different from standard English.
 - ▶ **Tweet:** ya ur website suxx Brah
 - ▶ **Standard:** yeah, your website sucks bro
- ▶ **Word variants:** phonetic substitutions, abbreviations, unconventional spellings, etc.

Lexical variation

- ▶ Social media language is different from standard English.
 - ▶ **Tweet:** ya ur website suxx Brah
 - ▶ **Standard:** yeah, your website sucks bro
- ▶ **Word variants:** phonetic substitutions, abbreviations, unconventional spellings, etc.
- ▶ **Normalization:** map from orthographic variants to standard spellings.

Lexical variation

- ▶ Social media language is different from standard English.
 - ▶ **Tweet:** ya ur website suxx Brah
 - ▶ **Standard:** yeah, your website sucks bro
- ▶ **Word variants:** phonetic substitutions, abbreviations, unconventional spellings, etc.
- ▶ **Normalization:** map from orthographic variants to standard spellings.
 - ▶ Large label space.
 - ▶ No labeled data.

Why normalization?

- ▶ Improve downstream NLP applications.
 - ▶ Part-of-speech tagging [6]
 - ▶ Dependency parsing [11]
 - ▶ Machine translation [7, 8]

Why normalization?

- ▶ Improve downstream NLP applications.
 - ▶ Part-of-speech tagging [6]
 - ▶ Dependency parsing [11]
 - ▶ Machine translation [7, 8]
- ▶ Is there more systematic variation in social media?
 - ▶ Mine patterns in how words are spelled.
 - ▶ Discover coherent orthographic styles.

Related work

Method	Surface	Context	Final System
Han & Baldwin 2011	edit distance LCS	language model syntax	linear combination
Liu et al. 2012	character-based translation	distributional similarity	decoding with language model
Hassan et al. 2013	edit distance LCS	random walk	decoding with language model
Our approach	edit distance LCS, word pairs	language model	unified model with features



Our approach

- ▶ **Unsupervised**
- ▶ **Featurized**
- ▶ **Context-driven**
- ▶ **Joint**

Characteristics

Unsupervised

- ▶ Labeled data for Twitter normalization is limited.
- ▶ Language changes, annotations may become stale and ill-suited to new spellings and words.



Characteristics

Featurized: permitting overlapping features

- ▶ String similarity features.
- ▶ Lexical features that memorize conventionalized word pairs, e.g. you/*u*, to/*2*.

Characteristics

Context-driven

Give me suttin to Believe in

- ▶ Unsupervised normalization needs strong cue of local context.
- ▶ String similarity needs to be overcome by contextual preference.
 - ▶ $\text{Edit-Dist}(\text{button/suiting}, \text{suttin}) = 2$
 - ▶ $\text{Edit-Dist}(\text{something}, \text{suttin}) = 5$

Characteristics

Joint

Gimme suttin 2 beleive innnn

- ▶ No words are in the standard vocabulary.
- ▶ Joint inference is needed to take advantage of context information.

Normalization in a probabilistic model

Maximize likelihood of observed (Twitter) text,
marginalizing over normalizations.

$$P(s) \quad P(\text{ya ur website suxx Brah})$$

Normalization in a probabilistic model

Maximize likelihood of observed (Twitter) text,
marginalizing over normalizations.

$$\begin{aligned} P(s) & P(\text{ya ur website suxx Brah}) \\ = \sum_t P(s, t) & P(\text{ya ur website } \dots, \text{yam urn website } \dots) \\ & + P(\text{ya ur website } \dots, \text{yak your website } \dots) \\ & + P(\text{ya ur website } \dots, \text{yea cure website } \dots) \\ & + \dots \end{aligned}$$

A noisy channel model

$$P(\mathbf{s}, \mathbf{t}) = P_t(\mathbf{t})P_e(\mathbf{s}|\mathbf{t})$$

$$P_e(\mathbf{y}\mathbf{a} \ \mathbf{u}\mathbf{r} \ \dots | \mathbf{y}\mathbf{a}\mathbf{m} \ \mathbf{u}\mathbf{r}\mathbf{n} \ \dots) \\ \times P_t(\mathbf{y}\mathbf{a}\mathbf{m} \ \mathbf{u}\mathbf{r}\mathbf{n} \ \mathbf{w}\mathbf{e}\mathbf{b}\mathbf{s}\mathbf{i}\mathbf{t}\mathbf{e} \ \dots)$$

A noisy channel model

$$P(\mathbf{s}, \mathbf{t}) = P_t(\mathbf{t})P_e(\mathbf{s}|\mathbf{t})$$

$$P_e(\mathbf{y}\mathbf{a} \ \mathbf{u}\mathbf{r} \ \dots | \mathbf{y}\mathbf{a}\mathbf{m} \ \mathbf{u}\mathbf{r}\mathbf{n} \ \dots) \\ \times P_t(\mathbf{y}\mathbf{a}\mathbf{m} \ \mathbf{u}\mathbf{r}\mathbf{n} \ \mathbf{w}\mathbf{e}\mathbf{b}\mathbf{s}\mathbf{i}\mathbf{t}\mathbf{e} \ \dots)$$

- ▶ The language model can be estimated offline.
- ▶ The emission model is locally normalized and log-linear.

$$P_e(\mathbf{s}|\mathbf{t}) = \prod_n P_e(s_n|t_n)$$

$$P_e(\mathbf{y}\mathbf{a}|\mathbf{y}\mathbf{a}\mathbf{m})P_e(\mathbf{u}\mathbf{r}|\mathbf{u}\mathbf{r}\mathbf{n}) \dots$$

$$P_e(s_n|t_n) = \frac{\exp(\boldsymbol{\theta}' \mathbf{f}(s_n, t_n))}{Z(t_n)}$$

$$\frac{\exp(\boldsymbol{\theta}' \mathbf{f}(\mathbf{u}\mathbf{r}, \mathbf{y}\mathbf{o}\mathbf{u}\mathbf{r}))}{Z(\mathbf{y}\mathbf{o}\mathbf{u}\mathbf{r})}$$

Features

String similarity

- ▶ Combine edit distance and longest-common subsequence (LCS) [3]:
- ▶ Bin to create binary features, e.g.,

$$\text{top5}(s, t) = \begin{cases} 1, & t \in \text{Top5}(s) \\ 0, & \text{otherwise} \end{cases}$$

Features

String similarity

- ▶ Combine edit distance and longest-common subsequence (LCS) [3]:
- ▶ Bin to create binary features, e.g.,

$$top5(s, t) = \begin{cases} 1, & t \in Top5(s) \\ 0, & \text{otherwise} \end{cases}$$

Word pairs

- ▶ Assign weights to commonly occurring word pairs, e.g. you/**u**, sucks/**suxx**.
- ▶ This allows the model to “memorize” frequent substitutions.

Learning

Our goal is to learn to weight these features.
We compute the gradient of the likelihood...

Learning

Our goal is to learn to weight these features.
We compute the gradient of the likelihood...

CRF

$$\ell(\boldsymbol{\theta}) = \log P(\mathbf{y}|\mathbf{x})$$

$$\frac{\partial \ell}{\partial \boldsymbol{\theta}} = f(\mathbf{x}, \mathbf{y}) - E_{\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}}[f(\mathbf{x}, \mathbf{y})]$$

Learning

Our goal is to learn to weight these features.
We compute the gradient of the likelihood...

CRF

$$\ell(\boldsymbol{\theta}) = \log P(\mathbf{y}|\mathbf{x})$$

$$\frac{\partial \ell}{\partial \boldsymbol{\theta}} = f(\mathbf{x}, \mathbf{y}) - E_{\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}}[f(\mathbf{x}, \mathbf{y})]$$


Our model (MRF)

$$\ell(\boldsymbol{\theta}) = \log P(\mathbf{s})$$

$$\frac{\partial \ell}{\partial \boldsymbol{\theta}} = E_{\mathbf{t}|\mathbf{s}}[f(\mathbf{s}, \mathbf{t}) - E_{\mathbf{s}'|\mathbf{t}}[f(\mathbf{s}, \mathbf{t})]]$$


Dynamic programming

In a locally-normalized model, these expectations can be computed from the marginals $P(t_n | s_{1:N})$ [1].



Dynamic programming

In a locally-normalized model, these expectations can be computed from the marginals $P(t_n | s_{1:N})$ [1].



- ▶ **Outer expectation** $E_{t|s}$: forward-backward
- ▶ **Inner expectation** $E_{s'|t}$: sum over all possible s_n at each n .
- ▶ **Time complexity**: $\mathcal{O}(NV_T^2 + NV_T^2 V_S)$...
But $V_T, V_S > 10^4$

Sequential Monte Carlo

A randomized algorithm to approximate $P(\mathbf{t}|\mathbf{s})$ as a weighted sum,

$$P(\mathbf{t}|\mathbf{s}) \approx \sum_k \omega^{(k)} \delta_{\mathbf{t}^{(k)}}(\mathbf{t})$$

$$P(\mathbf{t}|\text{ya ur website} \dots) \approx \begin{cases} 0.2 \times \delta(\text{you your website} \dots) \\ 0.6 \times \delta(\text{yeah your website} \dots) \\ \dots \\ 0.1 \times \delta(\text{you you're website} \dots) \end{cases}$$

- ▶ efficient and simple
- ▶ easy to parallelize
- ▶ number of samples K provides intuitive tuning between accuracy and speed

Sequential Importance Sampling (SIS)

At each n , for each sample k

- ▶ Draw $t_n^{(k)}$ from the **proposal distribution** $Q(t)$.
- ▶ Update $\omega_n^{(k)}$ so that

$$\begin{aligned}\omega_n^{(k)} &\propto \frac{P(t_{1:n}^{(k)} | s_{1:n})}{Q(t_{1:n}^{(k)})} \\ &= \dots \\ &= \frac{\overbrace{P_e(s_n | t_n^{(k)})}^{\text{Emission model}} \overbrace{P_t(t_n^{(k)} | t_{n-1}^{(k)})}^{\text{Language model}}}{\underbrace{Q(t_n^{(k)} | t_{n-1}^{(k)}, s_n)}_{\text{Proposal distribution}}} \omega_{n-1}^{(k)}\end{aligned}$$

Computing the gradient from SIS

- ▶ **Outer expectation:**

$$E_{t|s}[f(\mathbf{s}, \mathbf{t})] = \sum_k^K \omega_N^{(k)} \sum_n^N \mathbf{f}(s_n, t_n^{(k)})$$

Computing the gradient from SIS

- ▶ **Outer expectation:**

$$E_{t|s}[f(\mathbf{s}, \mathbf{t})] = \sum_k^K \omega_N^{(k)} \sum_n^N \mathbf{f}(s_n, t_n^{(k)})$$

- ▶ **Inner expectation:** draw L samples of $s'_n | t_n$:

$$E_{t|s}[E_{s'|t}f(\mathbf{s}', \mathbf{t})] = \sum_k^K \omega_N^{(k)} \sum_n^N \frac{1}{L} \sum_\ell^L \mathbf{f}(s_n^{(\ell,k)}, t_n^{(k)})$$

Computing the gradient from SIS

- ▶ **Outer expectation:**

$$E_{t|s}[f(\mathbf{s}, \mathbf{t})] = \sum_k^K \omega_N^{(k)} \sum_n^N \mathbf{f}(s_n, t_n^{(k)})$$

- ▶ **Inner expectation:** draw L samples of $s'_n | t_n$:

$$E_{t|s}[E_{s'|t}f(\mathbf{s}', \mathbf{t})] = \sum_k^K \omega_N^{(k)} \sum_n^N \frac{1}{L} \sum_\ell^L \mathbf{f}(s_n^{(\ell,k)}, t_n^{(k)})$$

(In practice, we set $K = 10$ and $L = 1$)

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your		...
	yeah	you're		...

$\omega_n^{(k)}$				
$s_n^{(\ell)}$				

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your		...
	yeah	you're		...

$\omega_n^{(k)}$				
$s_n^{(\ell)}$				

Sample $t_n^{(k)}$ according to $Q(t|\circ, ya)$.

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your		...
	yeah	you're		...

$\omega_n^{(k)}$	$\frac{P(\text{yeah}, \text{ya} \circ)}{Q(\text{yeah} \circ, \text{ya})}$			
$s_n^{(\ell)}$				

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your		...
	yeah	you're		...

$\omega_n^{(k)}$	$\frac{P(\text{yeah}, \text{ya} \circ)}{Q(\text{yeah} \circ, \text{ya})}$			
$s_n^{(\ell)}$	yea			

Sample $s_n^{(l)}$ according to $P(s|\text{yeah})$.

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your		...
	yeah	you're		...

$\omega_n^{(k)}$	$\frac{P(\text{yeah}, \text{ya} \circ)}{Q(\text{yeah} \circ, \text{ya})}$	$\omega_1^{(k)} \frac{P(\text{your}, \text{ur} \text{yeah})}{Q(\text{your} \text{yeah}, \text{ur})}$		
$s_n^{(\ell)}$	yea	youu		

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your		...
	yeah	you're		...

$\omega_n^{(k)}$	$\frac{P(\text{yeah}, \text{ya} \circ)}{Q(\text{yeah} \circ, \text{ya})}$	$\omega_1^{(k)} \frac{P(\text{your}, \text{ur} \text{yeah})}{Q(\text{your} \text{yeah}, \text{ur})}$	$\omega_2^{(k)} \frac{P_t(\text{website}, \text{website} \text{your})}{Q(\text{website} \text{your}, \text{website})}$...
$s_n^{(\ell)}$	yea	youu	website	...

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your	website	...
	yeah	you're		...

$\omega_n^{(k)}$	$\frac{P(\text{yeah}, \text{ya} \circ)}{Q(\text{yeah} \circ, \text{ya})}$	$\omega_1^{(k)} \frac{P(\text{your}, \text{ur} \text{yeah})}{Q(\text{your} \text{yeah}, \text{ur})}$	$\omega_2^{(k)} \frac{P_t(\text{website}, \text{website} \text{your})}{Q(\text{website} \text{your}, \text{website})}$...
$s_n^{(\ell)}$	yea	youu	website	...

Update:

$$\omega_N^{(k)} (f(\text{yeah}, \text{ya}) - f(\text{yeah}, \text{yea}) \\ + f(\text{your}, \text{ur}) - f(\text{your}, \text{youu}) + \dots)$$

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your	website	...
	yeah	you're		...

$\omega_n^{(k)}$	$\frac{P(\text{yeah}, \text{ya} \circ)}{Q(\text{yeah} \circ, \text{ya})}$	$\omega_1^{(k)} \frac{P(\text{your}, \text{ur} \text{yeah})}{Q(\text{your} \text{yeah}, \text{ur})}$	$\omega_2^{(k)} \frac{P_t(\text{website}, \text{website} \text{your})}{Q(\text{website} \text{your}, \text{website})}$...
$s_n^{(\ell)}$	yea	youu	website	...

Update:

$$\omega_N^{(k)} (f(\text{yeah}, \text{ya}) - f(\text{yeah}, \text{yea}) \\ + f(\text{your}, \text{ur}) - f(\text{your}, \text{youu}) + \dots)$$

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your	website	...
	yeah	you're		...

$\omega_n^{(k)}$	$\frac{P(\text{yeah}, \text{ya} \circ)}{Q(\text{yeah} \circ, \text{ya})}$	$\omega_1^{(k)} \frac{P(\text{your}, \text{ur} \text{yeah})}{Q(\text{your} \text{yeah}, \text{ur})}$	$\omega_2^{(k)} \frac{P_t(\text{website}, \text{website} \text{your})}{Q(\text{website} \text{your}, \text{website})}$...
$s_n^{(\ell)}$	yea	youu	website	...

Update:

$$\omega_N^{(k)} (f(\text{yeah}, \text{ya}) - f(\text{yeah}, \text{yea}) + f(\text{your}, \text{ur}) - f(\text{your}, \text{youu}) + \dots)$$

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your	website	...
	yeah	you're		...

$\omega_n^{(k)}$	$\frac{P(\text{yeah}, \text{ya} \circ)}{Q(\text{yeah} \circ, \text{ya})}$	$\omega_1^{(k)} \frac{P(\text{your}, \text{ur} \text{yeah})}{Q(\text{your} \text{yeah}, \text{ur})}$	$\omega_2^{(k)} \frac{P_t(\text{website}, \text{website} \text{your})}{Q(\text{website} \text{your}, \text{website})}$...
$s_n^{(\ell)}$	yea	youu	website	...

Update:

$$\omega_N^{(k)} (f(\text{yeah}, \text{ya}) - f(\text{yeah}, \text{yea}) \\ + f(\text{your}, \text{ur}) - f(\text{your}, \text{youu}) + \dots)$$

Example

s_n	ya	ur	website	...
	yam	urn	website	...
$t_n^{(k)}$	you	your	website	...
	yeah	you're		...

$\omega_n^{(k)}$	$\frac{P(\text{yeah}, \text{ya} \circ)}{Q(\text{yeah} \circ, \text{ya})}$	$\omega_1^{(k)} \frac{P(\text{your}, \text{ur} \text{yeah})}{Q(\text{your} \text{yeah}, \text{ur})}$	$\omega_2^{(k)} \frac{P_t(\text{website}, \text{website} \text{your})}{Q(\text{website} \text{your}, \text{website})}$...
$s_n^{(\ell)}$	yea	youu	website	...

Update:

$$\omega_N^{(k)} (f(\text{yeah}, \text{ya}) - f(\text{yeah}, \text{yea}) \\ + f(\text{your}, \text{ur}) - f(\text{your}, \text{youu}) + \dots)$$

The proposal distribution

A proposal distribution can guide sampling focuses on high-probability region.

The proposal distribution

A proposal distribution can guide sampling focuses on high-probability region.

1. $Q(t_n|s_n, t_{n-1}) = P(t_n|s_n, t_{n-1}) \propto P(s_n|t_n)P(t_n|t_{n-1})$
 - ▶ Accurate (optimal!)
 - ▶ Not fast: computing $P(s_n|t_n)$ costs $\mathcal{O}(V_T V_S)$

The proposal distribution

A proposal distribution can guide sampling focuses on high-probability region.

1. $Q(t_n|s_n, t_{n-1}) = P(t_n|s_n, t_{n-1}) \propto P(s_n|t_n)P(t_n|t_{n-1})$
 - ▶ Accurate (optimal!)
 - ▶ Not fast: computing $P(s_n|t_n)$ costs $\mathcal{O}(V_T V_S)$

2. Our proposal

$$\begin{aligned} Q(t_n|s_n, t_{n-1}) &\propto P(s_n|t_n)Z(t_n)P(t_n|t_{n-1}) \\ &= \exp(\boldsymbol{\theta}'\mathbf{f}(s_n, t_n)) P(t_n|t_{n-1}) \end{aligned}$$

- ▶ Fast: $\mathcal{O}(V_S + V_T)$ to compute $\omega_n^{(k)}$
- ▶ Fairly accurate: biased by a factor of $Z(t_n)$

Implementation details

unLOL: **u**nsupervised **n**ormalization in a **LOg-Linear** model

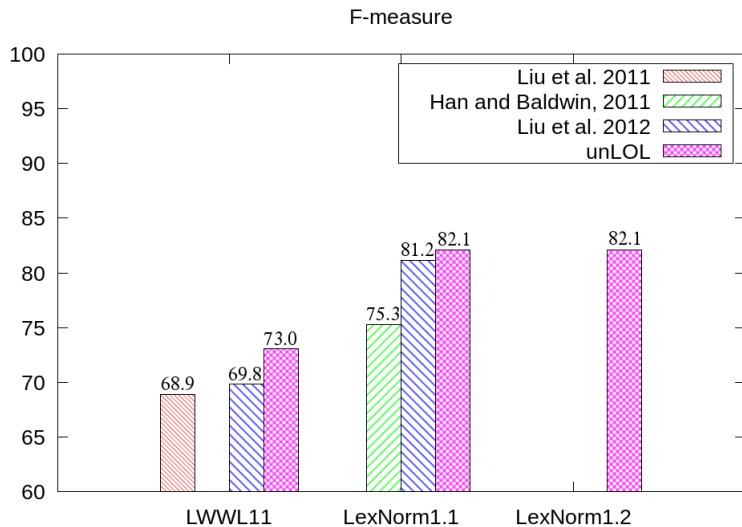
- ▶ **Decoding**: propose K sequences $\mathbf{t}^{(k)}$, perform Viterbi within this limited set.
- ▶ **Normalization targets**:
 - ▶ Training: all alphanumeric strings not in aspell
 - ▶ Test: normalization targets are given (standard in this task)
- ▶ **Language model**: Kneser-Ney smoothed trigram model, from tweets with no OOV words
- ▶ **Training data**: For each OOV word in the test set, draw 50 tweets from Edinburgh corpus [10] and Twitter API.

Evaluation

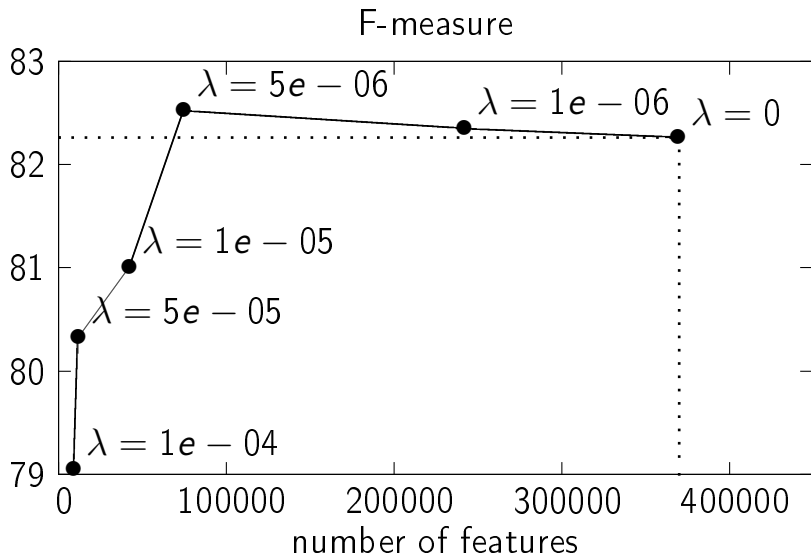
Datasets

- ▶ LWWL11 [9]: 3802 isolated words
- ▶ LexNorm1.1 [5]: 549 tweets with 1184 nonstandard tokens
- ▶ LexNorm1.2: 172 manual corrections to LexNorm1.1 (www.cc.gatech.edu/~jeisenst/lexnorm.v1.2.tgz)

Results



Adding L1 Regularization




Analysis


Can normalization help identify orthographic styles?

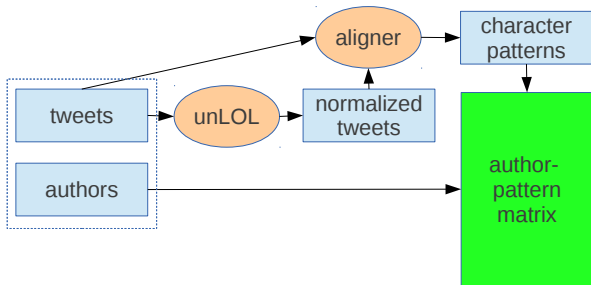
- ▶ Use unLOL to automatically label lots of tweets
- ▶ Use Levenshtein alignment between original and normalized tweets to find substitution patterns, e.g., ng/n_
- ▶ Use matrix factorization to find sets of patterns that are used by the same set of authors.

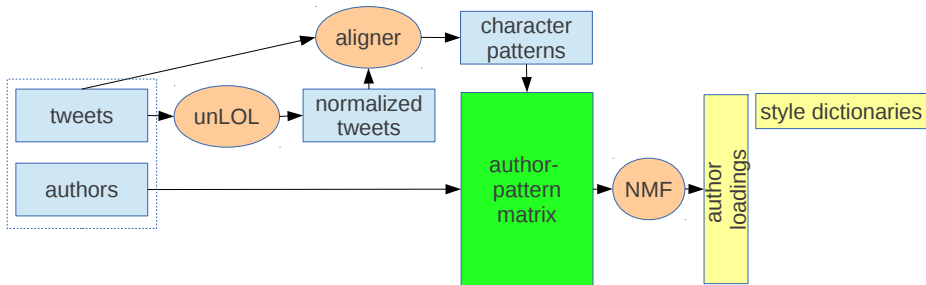
tweets

authors









Observations

Some styles mirror phonological variables:
g-dropping, (TH)-stopping, t-deletion... [4]

style	rules	examples
g-dropping	$g^*/_*$ ng/n_ g/_	goin, talkin, watchin, feelin, makin
t-deletion	$t^*/_*$ st/s_ t/_	jus, bc, shh, wha, gota, wea, mus, firts, jes, subsistutes
th-stopping	h/_ $*t/*d$ th/d_ t/d	dat, de, skool, fone, dese, dha, shid, dhat, dat's

Observations

Others are known

“netspeak” phenomena, like expressive lengthening [2]

style	rules	examples
(kd)-adding	i_/id _/_k _/d _*/k*	idk, fuckk, okk, backk, workk, badd, andd, goodd, bedd, elidgible, pidgeon
o-adding	o_/oo _*/o* _/_o	soo, noo, doo, oohh, loove, thoo, helloo
e-adding	_/_i e_/ee _/_e _*/e*	mee, ive, retweet, bestie, lovee, nicee, heey, likee, iphone, homie, ii, damnit

Observations





We get meaningful styles even from mistaken normalizations! (e.g., i'm/**ima**, out/**outta**)

style	rules	examples
a-adding	<hr/> <ul style="list-style-type: none">_/a__/ma_/m*/a*	<hr/> ima, outta, needa, shoulda, woulda, mm, comming, tomm, Boutt, ppreci-ate

Summary

- ▶ unLOL: joint, unsupervised, and log-linear
 - ▶ **Features** balance between edit distance and conventionalized word pairs.
 - ▶ **Main challenge**: massive label space, in which quadratic algorithms are not practical.
 - ▶ **Solution**: approximate gradient with SMC
 - ▶ **Proposal distribution** focuses the samples on the high-likelihood part of the search space.
- ▶ Normalization enables the study of systematic orthographic variation.

References I

-  T. Berg-Kirkpatrick, A. Bouchard-Côté, J. DeNero, and D. Klein.
Painless unsupervised learning with features.
In *Proceedings of NAACL*, pages 582–590, 2010.
-  S. Brody and N. Diakopoulos.
Cooooooooooooooooo!!!!!!!!!!!!!!!!!!!!!!: using word lengthening to detect sentiment in microblogs.
In *Proceedings of EMNLP*, 2011.
-  D. Contractor, T. A. Faruque, and L. V. Subramaniam.
Unsupervised cleansing of noisy text.
In *Proceedings of COLING*, pages 189–196, 2010.
-  J. Eisenstein.
Phonological factors in social media writing.
In *Proceedings of the NAACL Workshop on Language Analysis in Social Media*, 2013.

References II



B. Han and T. Baldwin.

Lexical normalisation of short text messages: makin sens a #twitter.
In *Proceedings of ACL*, pages 368–378, 2011.



B. Han, P. Cook, and T. Baldwin.

Lexical normalization for social media text.

ACM Transactions on Intelligent Systems and Technology, 4(1):5,
2013.



H. Hassan and A. Menezes.

Social text normalization using contextual graph random walks.
In *Proceedings of ACL*, 2013.

References III



W. Ling, C. Dyer, I. Trancoso, and A. Black.

Paraphrasing 4 microblog normalization.

In Proceedings of the 2013 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, Seattle, USA, October 2013. Association for Computational Linguistics.



F. Liu, F. Weng, B. Wang, and Y. Liu.

Insertion, deletion, or substitution?: normalizing text messages without pre-categorization nor supervision.

In Proceedings of ACL, pages 71–76, 2011.



S. Petrović, M. Osborne, and V. Lavrenko.

The edinburgh twitter corpus.

In Proceedings of the NAACL HLT Workshop on Computational Linguistics in a World of Social Media, pages 25–26, 2010.

References IV



C. Zhang, T. Baldwin, H. Ho, B. Kimelfeld, and Y. Li.

Adaptive parser-centric text normalization.

In *Proceedings of ACL*, pages 1159–1168, 2013.