

“Big data pragmatics!”
or
“Putting the ACL in computational social science”
or

*If you think these title alternatives could turn people on,
turn people off,
or otherwise have an effect,
this talk might be for you.*

Lillian Lee, Cornell University
<http://www.cs.cornell.edu/home/llee>

The one equation in this talk

Lots of on-line conversations

(Facebook, Twitter, ...; YouTube comments...; Yelp reviews...;...)

=

Many systems with humans and language as key components

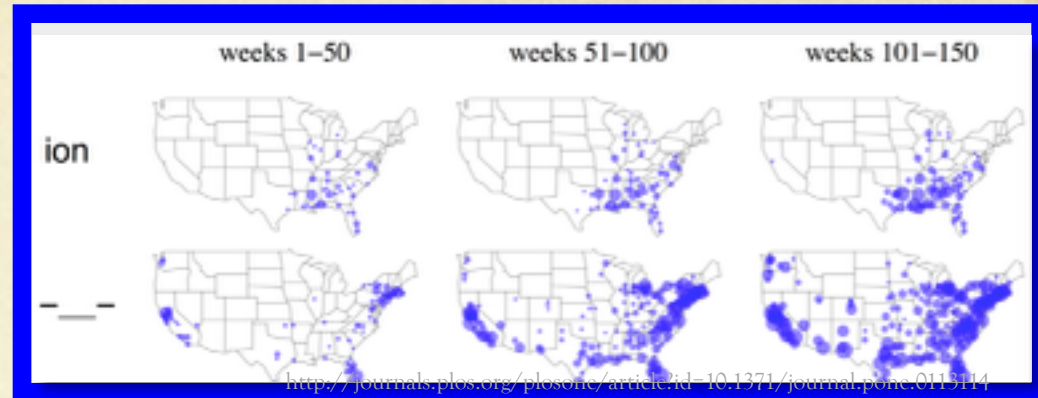
=

Fantastic opportunities for NLP + the social sciences
to build better systems and learn more about people

A sampling

Lexical diffusion:

Jacob Eisenstein, Brendan O'Connor,
Noah Smith, Eric P. Xing, 2014. ↗



Power relationships from language:

Vinod Prabhakaran, Owen Rambow.

Best short paper hon. mention, ACL 2014 ↗

Language matching and community engagement:

Cristian Danescu-Niculescu-Mizil,
Bob West, Dan Jurafsky,
Jure Leskovec, Chris Potts.
Best paper, WWW 2013. ↗






"Lillian"

What about the effect of language choice?

One aspect of phrasing: framing

The *framing* of an arguments emphasizes certain principles or perspectives.




“One of the most important concepts in the study of public opinion” [James Druckman, 2001]

Hedging and framing in GMO debates: Eunsol Choi, Chenhao Tan, Lillian Lee, Cristian Danescu-Niculescu-Mizil, Jennifer Spindel 2012 




"green revolution"

"Frankenfood"

Other *ACL framing work includes: Viet-An Nguyen, Jordan Boyd-Graber, Philip Resnik 2013 ; Eric Baumer, Elisha Elovic, Ying Qin, Francesca Polletta, Geri Gay, 2015 ; Oren Tsur, Dan Calacci, David Lazer 2015. 

Past research: phrasing may not matter

Daniel Hopkins, SSRN 2013: “...there is no evidence that groups targeted by specific frames [such as ‘death panels’ in the health care debates] respond accordingly.” 


Justin Grimmer, Solomon Messing, Sean Westwood, *The Impression of Influence*, 2014: *total number of messages* mattered more than amount of money the messages described. 

Either Sasa Petrovic, Miles Osborne, or Victor Lavrenko, slashdot 2014: “...a famous person can write anything and it will be retweeted. An unknown person can write the same tweet and it will be ignored.”

Purely hypothetical NAACL Board meeting

I knew I should have said “arf”.

sourced from the Cartoon Bank



Still, can wording alone be influential?

Non-options: Have better ideas.

(Instantaneously) become alpha dog.

Be a dog at all.

“Parallel universe” experimental paradigm

Exploit situations with *many* instances of:

...the same speaker




...in the same situation, or

conveying the same info...

...**varying their wording** (beyond a fixed set of lexical choices)

and see the effects.



Relates to work on *style* (e.g., Annie Louis and Ani Nenkova, 2013 ) and *paraphrasing* (e.g., Regina Barzilay and Kathy McKeown 2001 , Wei Xu, Alan Ritter, Chris Callison-Burch, Bill Dolan, Yangfeng Ji, 2014 )

Outline

Memorability and cultural penetration:

Cristian Danescu-Niculescu-Mizil, Justin Cheng, Jon Kleinberg, Lillian Lee, ACL 2012 ↗



<https://www.flickr.com/photos/hyku/3614261299/in/photostream/>



Information sharing and spread:

Chenhao Tan, Lillian Lee, Bo Pang, ACL 2014 ↗

Claim strength and its effects:

- Chenhao Tan, Lillian Lee, ACL (short) 2014 ↗
- Chenhao Tan, Lillian Lee, work in progress



<http://www.jamaicaohserver.com/assets/11552774/views.jpg>
<http://www.bekesayfasi.com/wp-content/uploads/2013/11/empathie-326x235.png>

Other *ACL work includes: Marco Guerini, Gödke Ötzbal, Carlo Strapparava, 2015 ↗
Tim Althoff, Cristian Danescu-Niculescu-Mizil, Dan Jurafsky 2014. ↗

<http://pixabay.com/en/twitter-tweet-twitter-bird-312464/>

Aside on presentation style

Paraphrasing Stuart Shieber: ↗

Your goal is *not* to convince people that you are brilliant, but that your solution is trivial.

It takes a certain strength of character to take that as one's goal.

Purely hypothetical reviewing situation



Aside on presentation style

Paraphrasing Stuart Shieber: ↗

Your goal is *not* to convince people that you are brilliant, but that your solution is trivial.

It takes a certain strength of character to take that as one's goal.


But if people think your findings are obvious,
they must also think that you are correct.



Part I: Does phrasing affect memorability?

[Much related work in many fields; see paper for refs. Our direct inspiration: Jure Leskovec, Lars Backstrom, Jon Kleinberg 2009 [!\[\]\(3dfb8d66e81160ad61421a3452093d1b_img.jpg\)](#),
Meme modification: Matthew Simmons, Lada Adamic & Eytan Adar '11 [!\[\]\(21ece2018b00c7267b3324c50bbed633_img.jpg\)](#)]

http://www.afi.com/100years/quotes.aspx



"FRANKLY, MY DEAR, I DON'T GIVE A DAMN" TOPS AFI'S LIST OF 100 GREATEST MOVIE QUOTES OF ALL TIME

OTHER WINNERS INCLUDE:

THE GODFATHER, **"I'M GOING TO MAKE HIM AN OFFER HE CAN'T REFUSE"**

THE WIZARD OF OZ, **"TOTO, I'VE GOT A FEELING WE'RE NOT IN KANSAS ANYMORE"**

AND CASABLANCA, **"HERE'S LOOKING AT YOU, KID"**


AFI's 100 YEARS 100 "Movie Quotes"

AFI'S 100 Years...100 Movie Quotes: America's Greatest Quips, Comebacks and Catchphrases

LOS ANGELES, June 22, 2005 — The American Film Institute revealed the top movie quotes of all time in **AFI's 100 Years...100 Movie Quotes**, **a three-hour special television event on CBS** hosted by actor and action star Pierce Brosnan with commentary from many of Hollywood's most celebrated actors and filmmakers. A jury of 1,500 film artists, critics and historians selected "Frankly, my dear, I don't give a damn," spoken by Clark Gable in the celebrated Civil War epic, GONE WITH THE WIND as the most memorable movie quote of all time.

Movie quotes:
massively,
permanently viral

Motivations

Broad motivation: what achieves massive cultural uptake?
Does it only depend on contextual factors? (cf. Salganik, Dodds, Watts, “MusicLab” experiment, *Science* 2006 )

Practical motivation: which material to promote?

- Ad slogans, political slogans

The (Jedi mind-) trick



Obi-Wan: You don't need to see his identification.

Stormtrooper: [ditto]

Obi-Wan: These aren't the droids you're looking for.

Stormtrooper: [ditto]

Obi-Wan: He can go about his business.

Stormtrooper: [ditto]

Bye, daddy. I hope
you find the droids
you're looking for.



**“WE DO OTHER THINGS
BESIDES LOOK FOR
DROIDS.
BUT THAT'S ALL
ANYONE EVER
REMEMBERS.”**



These aren't
the droids
you're looking for



Data


From ~ 1000 movie scripts (many lines long),
pair IMDB “memorable quotes” with
~ adjacent,
same-length,
same-speaker
“non-memorable quotes”.
Filter with Google/Bing counts: 2200 pairs.

Pilot study

Subjects were shown 12 pairs from movies they *hadn't* seen.

<http://www.cs.cornell.edu/~cristian/memorability.html>

First quote	Second quote
<i>Half a million dollars will always be missed</i>	<i>I know the type, trust me on this.</i>
<i>I think it's time to try some unsafe velocities.</i>	<i>No cold feet, or any other parts of our anatomy.</i>
<i>A little advice about feelings kiddo; don't expect it always to tickle.</i>	<i>I mean there's someone besides your mother you've got to forgive.</i>



(context/actor effects
explain all, bad labels, etc.)

Thirteen minutes of fame

Cornell University -- they're always doing research at Cornell! Thank goodness for Cornell University...

It's a complex study; we've got a link to it in the description--

--Don't read it, though
-- it's boring.



On average, memorable quotes (significantly)...

... contain more **surprising** combinations of words

according to 1-,2-,3-gram lexical language models trained on the Brown corpus

“...aren’t the droids...”

... are built on a more **common** syntactic scaffolding

according to 1-,2-,3-gram part-of-speech language models trained on Brown


“You’re gonna need a bigger boat”

[vs. *“You’re gonna need a boat that is bigger”*]

Our classifier, with these + other features (10-fold xval): 64.27%

Applications to social-media UI

[Lars Backstrom, Jon Kleinberg, Lillian Lee, Cristian Danescu-Niculescu-Mizil, 2013]



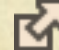
More-unusual Facebook posts get more comments
(under certain circumstances), but not so with Wikipedia.



Part II: Information diffusion

Other *ACL work includes: Yoav Artzi, Patrick Pantel, Michael Gamon 2012 

Marco Guerini, Carlo Strapparava, Gödke Ötzbal, 2011 

Sasa Petrovic, Miles Osborne, Victor Lavrenko 2011 

Oren Tsur, Ari Rappoport 2012  ; Aobo Wang, Tao Chen, Min-Yen Kan 2012 

The parallel universe

Many Twitter users re-post about the *same* URL w/in 12 hours, varying their text, with significantly different retweet results



cactus_music @cactus_music · 15 Sep 2012

I know at some point you've have been saved from hunger by our rolling food trucks friends. Let's help support them! bit.ly/P6GYCq



cactus_music @cactus_music · 15 Sep 2012

Food trucks are the epitome of small independently owned LOCAL businesses! Help keep them going! Sign the petition bit.ly/P6GYCq



GeorgeMonbiot @GeorgeMonbiot · 4 Oct 2012

read @ameliagentleman's report today, then tell me Tories are no longer the nasty party: guardian.co.uk/society/2012/o...



GeorgeMonbiot @GeorgeMonbiot · 4 Oct 2012

Work capability tests: designed by bastards, performed by idiots. guardian.co.uk/society/2012/o...

Example classification results

- Estimate of human accuracy:
(sample of 100 pairs; 106 judges; 39 judgments/pair)

61.3% per-human average
- Our classifier on 11K pairs of truly* held-out data:
65.6%

*We ran only one experiment on it, and that was at submission time

Example feature results

On average ...

Don't be too different from the community, as defined by scoring against a general Twitter bigram LM.

But also be true to yourself, as defined by scoring against a user-specific unigram LM.



Part III: claim strength

Much related work on hedging: see the CoNLL 2010 shared task [!\[\]\(c507f772dba2b921f86777f01218e570_img.jpg\)](#)

Example: perils of underclaiming

The US embassy initially referred to the attacks at Kunming as:
“the terrible and senseless act of violence”.



Weibo user Cao Fan: “If you say that the Kunming attack is a ‘terrible and senseless act of violence’, then the 9/11 attack can be called a ‘regrettable traffic incident’”

After Prodding, U.S. State Department Labels Kunming Attack ‘Terrorism’

By DIDI KIRSTEN TATLOW MARCH 4, 2014 1:05 AM 19 Comments

Example: perils of overclaiming

The authors claim that they are addressing a document classification problem without using any prior linguistic knowledge - to which I am tempted to ask, what is this paper doing being submitted to the A C *L* conference?

Nonetheless I think the paper is a reasonable fit, especially since the technique actually does make use of several facts about language that are different from other sorts of data (photographs, etc.):

Written English can be tokenized into meaningful 'words' at whitespace ...

[much more follows]

Example: perils of overclaiming

The authors claim that they are addressing a document classification problem without using any prior linguistic knowledge - to which I am tempted to ask, what is this paper doing being submitted to the A C *I * conference?

...

Lines and lines and lines of sarcasm

...

(photographs, etc.):

...

(But reviewers are always right.)

...

Case study: strength revisions

On the arxiv e-print archive, authors post LaTeX source for *different versions* of the same paper.

In order to overcome this inconsistency, an additional constraint due to the requirement of extensivity is needed in the maximization procedure.




Therefore, an additional constraint due to the requirement of extensivity is needed in the maximization procedure, leading to a novel generalized maximization procedure.

Circadian pattern and burstiness in human communication activity



Circadian pattern and burstiness in mobile phone communication

Strength-labeled corpus

- 500 pairs received 9 labels each.
398 had an absolute-majority label: 93 weaker, 194 stronger, 99 a change not affecting strength
- One interesting finding: participants are swayed by details, even if their addition makes the statement less general.
(cf. Bell, Loftus '89, courts )

we also proved that if [math] is sufficiently homogeneous then ...



we also proved that if [math] is **not** **totally disconnected and** sufficiently homogeneous then ...

Hedging and influence at the Fed



AG: I assume iron ore is in [the CRB]?

K: I don't know if iron ore is in there
but copper is: copper scrap is in there, I think.

AG: That couldn't have done that much. Steel, for
example, is actually down.

K: I don't think steel is in the CRB.

Hedges: expressions of tentativeness


Preliminary results with words occurring in hedged contexts vs.
words in the *same utterance* that are not hedged.

The hedged words have less “impact” on the immediately
subsequent utterances, but greater impact later on in the discussion.

Summary: Putting the ACL in computational social science



Almost all our datasets can be found from my homepage. 
If you beat our results, everybody wins!



Summary: Putting the ACL in computational social science



Almost all our data
you beat our results



Looking forward

Deeper interplay between
NLP
and

how people use and are affected by language
is a huge opportunity for all concerned.

Image: Warner Bros.



I think this is the
beginning of a
beautiful friendship.

Thanks!