

Representation Learning for Narratives in Social Media

Serge Belongie

DIKU and Pioneer Centre for AI



The Beginnings of a Moonshot

A grayscale photograph of an astronaut in a full spacesuit standing on the lunar surface. The astronaut is positioned on the left side of the frame, facing away from the camera. In the background, a lunar lander is visible, with its descent stage and ascent stage. The lunar surface is covered in dust and small rocks, and the sky is a uniform dark gray.

- Overarching goal: stop the spread of **misinformation** online
- Approach
 - Identify the **prevailing narratives** in multimedia content
 - Learn a narrative-informed **metric**
 - **Human-in-the-loop** model training
- Challenges
 - Narratives are **unfalsifiable**
 - Complex relationship with fact checking
- Opportunities
 - Detect bot & sock puppet activity in **misinformation campaigns**
 - Narrative **prediction in novel settings**

What is a Narrative?

- In literature: story w/protagonist, villain, transformation, etc.
- In social media:



Jay Kapoor ✓
@JayKapoorNYC



NARRATIVE VIOLATION



wilson k @wilsonkriegel · Dec 1

Silicon Valley, New York City, and Boston received 66% of all VC 💰 in the US in Q3, 2021. The exact same percentage as in 2017.

Austin & Miami each had less than 2% of VC \$ & Philadelphia beats them both

10:50 PM · Dec 1, 2021 · Twitter for iPhone

Facts vs. Narratives



Paul Graham  @paulg · 4h

The US, Canada, and all of Europe combined account for .87% of plastic emissions. The Philippines accounts for 36%.

Possible Narratives Advanced by Tweet

- Banning plastic straws...
 - achieves nothing good
 - is performative
 - misleads people about the true sources of this kind of waste



Possible Narratives in Opposition

- Most of this plastic is exported from the US to the Philippines for recycling
- Even if this is “sustainability theater,” it keeps an important issue on people’s minds

Sometimes people distill narratives on their own:



(((yoav' (())(())

@yoavgo



the colonialists monster entity which was established for the sole purpose of enslaving and killing poor peace-loving palestinians for funzies

vs

the arch-evil palestinian terrorist children who deserve to die before they were even born

10:44 AM · May 13, 2021 · Twitter for Android

Motivating example: headline w/thumbnail

CNN

Bill and Melinda Gates are ending their marriage

Bill and Melinda Gates are ending their marriage after 27 years, the pair in a statement on their verified Twitter accounts.



BBC

Bill and Melinda Gates agreed 'separation contract' before announcing divorce

The billionaire couple reportedly agreed how to divide their assets before announcing their divorce.
11 hours ago



NY Daily News

Bill and Melinda Gates divorcing

Philanthropists Bill and Melinda Gates announced Monday they are getting divorced.
1 day ago



The New York Times

Bill and Melinda Gates Are Divorcing

Bill and Melinda Gates, two of the richest people in the world, who reshaped philanthropy and public health with the fortune Mr. Gates made as...
3 hours ago



Bill and Melinda Gates are divorcing: What they've said about their marriage

Bill and Melinda Gates are divorcing: What they've said about their marriage · On how they first met · On the decision to get married · On having ' ...
8 hours ago



What narratives does the choice of photo suggest?

Motivating Example: Out-of-Context Photo

 **Kimberly Klacik** ✓
@kimKBaltimore

...

A look at [#BuildBackBetter](#)



Narratives are often the quiet part (i.e., not said out loud)

**I Support my Black Friends,
but not BLM.**

**I Support my White Friends,
but not the KKK.**

**I Don't Support
Hate Groups.**



**Skin color means NOTHING.
You're either a good person or
you're not.**



PeerAagaard 🇩🇰
@AagaardPeer



Hvis dette havde været muslimer, der fejrede Eid eller LGBT'ere til Pride eller en hvilken som helst anden befolkningsgruppe var politikere og medier gået amok i fordømmelser. Godt at det var bare var fordrukne fodboldfans, for så er det nemlig helt i orden.

[Translate Tweet](#)





Looking back on how Visipedia got its start

CUB-200-2011






Illustration: Leon Edler

The long read

Hand dryers v paper towels: the surprisingly dirty fight for the right to dry your hands

For a century, the humble paper towel has dominated public toilets. But a new generation of hand dryers has sparked a war for loo supremacy

by [Samanth Subramanian](#)


Thu 25 Apr 2019 06.01 BST



In the summer of 2005, a Chicago marketing consultant named George Campbell received a tantalising call from a headhunter. Was he open to an interview at Dyson? The company was secretly preparing to launch a new appliance, and it needed a sales strategy for the US: that was all the headhunter would divulge. Campbell was excited; he saw Dyson as “a company with the iconic quality of Apple, and an ability to take a basic product like a vacuum cleaner and make an 80% margin on it”.

Hand Drying: Prevailing Narratives

- Air dryers are more hygienic than paper towels
- Air dryers are better for the environment than paper towels
- Paper towels are a waste of paper
- Air dryers cause skin chapping
- Air dryers are unhygienic since they blow germs around the room
- Air dryers take too long to dry your hands
- Air dryers are loud
- Paper towels can also be used to protect your hand when opening the door
- Air dryers waste energy
- People depend on public bathrooms to have paper towels for uses beyond drying hands
- Paper towels can help to wipe hands clean
- Air dryers circulate fecal matter throughout the bathroom





Touch the image to zoom in



[Visit the JITTY Store](#)

JITTY Disposable Bath Towels Portable Soft Cotton Towel Set for Hotel Bathroom Spa Travel Highly Absorbent 27x55 Inch 4 Individual Pack

★★★★★ [3 ratings](#)

\$15⁹⁹

prime

& **FREE Returns**



Related Concepts

Useful idiots

Meme epidemiology

Quantitative Debating vs. Holistic/Gut/Gestalt Debating

Blaise Pascal: “To tell someone they're wrong, first tell them how they're right”

Hedonic preference vs. objective similarity

Document summarization

Question Answering systems

Push-poll detection

AI-mediated communication

Federated Learning (for enterprise case)

“No one is the villain in their own story” vs. AITA

Marketplace of ideas

Dark money tracking: troll farm discovery

$k \ll N$ (number of narratives vs. number of social media users)

Data Journalism

Dawkins meme

Framing

Street Epistemology

Power relationships/shifts

Cooperative AI

Topic modeling

Bayesian Truth Serum & The List Experiment

SOCIAL CHEMISTRY 101: Learning to Reason about Social and Moral Norms

Maxwell Forbes^{†‡} Jena D. Hwang[‡] Vered Schwartz^{†‡} Maarten Sap[†] Yejin Choi^{†‡}

[†]Paul G. Allen School of Computer Science & Engineering, University of Washington

[‡]Allen Institute for AI

{mbforbes, msap, yejin}@cs.washington.edu, {jenah, vereds}@allenai.org

maxwellforbes.com/social-chemistry

Abstract

Social norms—the unspoken commonsense rules about acceptable social behavior—are crucial in understanding the underlying causes and intents of people’s actions in narratives. For example, underlying an action such as “wanting to call cops on my neighbor” are social norms that inform our conduct, such as “It is expected that you report crimes.”

We present SOCIAL CHEMISTRY, a new conceptual formalism to study people’s everyday social norms and moral judgments over a rich spectrum of real life situations described in net



Rules-of-Thumb

Our study centers around cultural *rules-of-thumb*. Each rule-of-thumb is inspired by a *situation*:

SITUATION

My roommate ran the blender at 5am

RULE-OF-THUMB

It's rude to make loud noises early in the morning

A rule-of-thumb has a simple structure: it is the *judgment* of an *action*.

<u>It's rude</u>	<u>to make loud noises early in the morning</u>
JUDGMENT	ACTION

We use rules-of-thumb to capture *cultural norms*. These include moral, ethical, and social norms. We treat rules-of-thumb as explanations of everyday social expectations.

ESPgame, Taboo, Family Feud, Newlywed Game, Bingo, etc.

The ESP Game

2:05 Time Left

The ESP Game

0090 score



Taboo Words

- USO
- BANNER
- PEOPLE

Your Guesses

- CROWD
- STAR
- STARS
- BLUE
- WHITE
- BLACK

Type your next guess:

Pass

Flag

Your partner has entered a guess



© 2002-2003 Carnegie Mellon University, all rights reserved. Patent Pending.

THE *Casual* CLASSROOM

TABOO CARDS

WORD GUESSING GAME

TUESDAY	WATCH	SUMMER
CANT SAY	CANT SAY	CANT SAY
WEEK DAY MONDAY WEDNESDAY	LOOK TIME CLOCK WRIST	SEASON AUTUMN WINTER SPRING
PLANET	MOVIE	LUNCH
CANT SAY	CANT SAY	CANT SAY
SUN SPACE EARTH ORBIT	FILM ACTOR HOLLYWOOD POPCORN	EAT CAFE DINNER BREAKFAST

Intentionomy: a Dataset and Study towards Human Intent Understanding

Menglin Jia^{1,2} Zuxuan Wu^{2,3} Austin Reiter² Claire Cardie¹ Serge Belongie¹ Ser-Nam Lim²
¹Cornell University ²Facebook AI ³Fudan University

Abstract

An image is worth a thousand words, conveying information that goes beyond the mere visual content therein. In this paper, we study the intent behind social media images with an aim to analyze how visual information can facilitate recognition of human intent. Towards this goal, we introduce an intent dataset, Intentionomy, comprising 14K images covering a wide range of everyday scenes. These images are manually annotated with 28 intent categories derived from a social psychology taxonomy. We then systematically study whether, and to what extent, commonly used visual information, i.e., object and context, contribute to human motive understanding. Based on our findings, we conduct further study to quantify the effect of attending to object and context classes as well as textual information in the form of hashtags when training an intent classifier. Our results quantitatively and qualitatively shed light on how visual and textual information can produce observable effects when predicting intent.¹

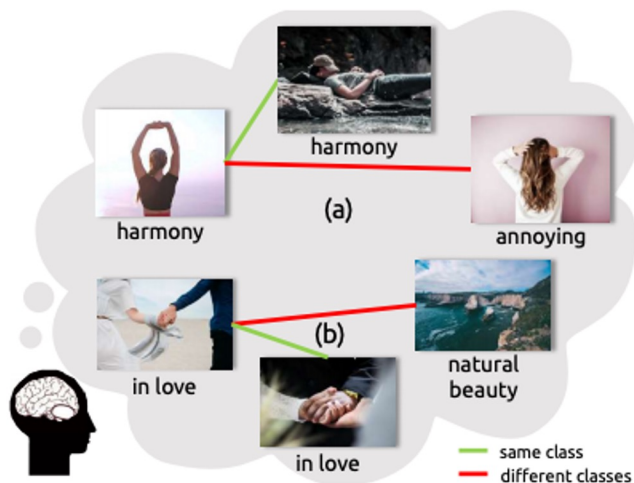


Figure 1. Intent behind images: while (b) shows that the visual motif of holding hands aligns with the common intent of “in love”, (a) illustrates that similarity based on visual appearance alone of texten would lead to an incorrect match with respect to intent.

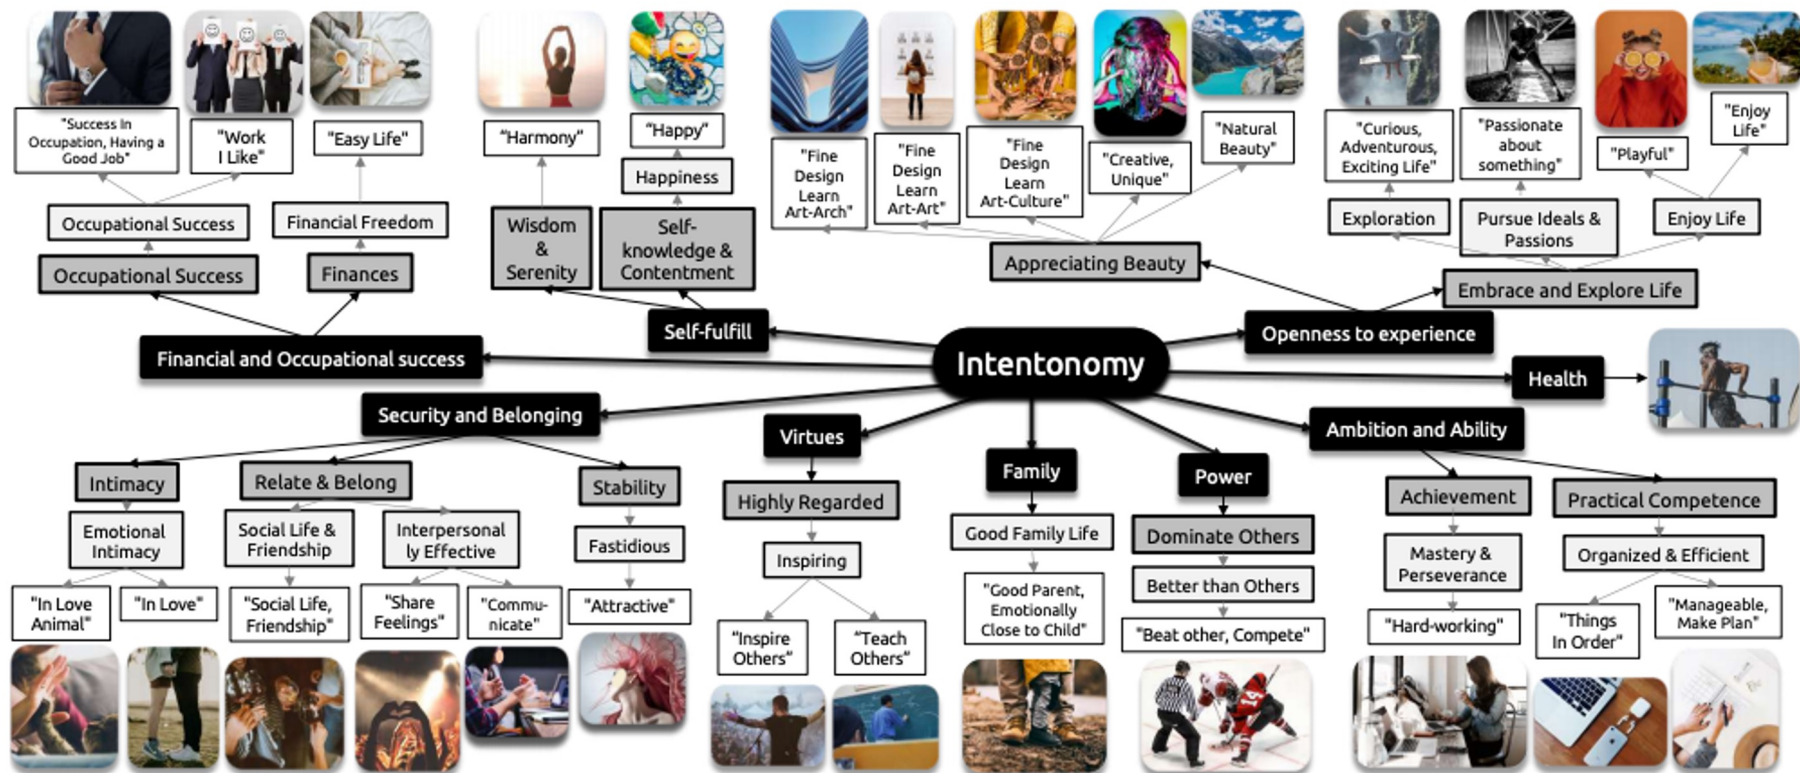


Figure 2. Ontology visualization. We select 28 labels from a general human motive taxonomy used in psychology research [72]. There are 9 super-categories in total (*in black box*), namely “virtues”, “self-fulfill”, “openness to experience”, “security and belonging”, “power”, “health”, “family”, “ambition and ability”, “financial and occupational success”. See the Appendix D for dataset statistics.

Learning Concept Embeddings with Combined Human-Machine Expertise

Michael J. Wilber^{1,2}

Iljung S. Kwak³

David Kriegman^{3,4}

Serge Belongie^{1,2}

¹ Department of Computer Science, Cornell University ² Cornell Tech

³ Department of Computer Science and Engineering, UC San Diego ⁴ Dropbox

Abstract

This paper presents our work on “SNaCK,” a low-dimensional concept embedding algorithm that combines human expertise with automatic machine similarity kernels. Both parts are complimentary: human insight can capture relationships that are not apparent from the object’s visual similarity and the machine can help relieve the human from having to exhaustively specify many constraints. We show that our SNaCK embeddings are useful in several tasks: distinguishing prime and nonprime numbers on MNIST, discovering labeling mistakes in the Caltech UCSD Birds (CUB) dataset with the help of deep-learned features, creating training datasets for bird classifiers, capturing subjective human taste on a new dataset of 10,000 foods, and qualitatively exploring an unstructured set of pictographic characters. Comparisons with the state-of-the-art in these tasks show that SNaCK produces better concept embeddings that require less human supervision than the leading methods.

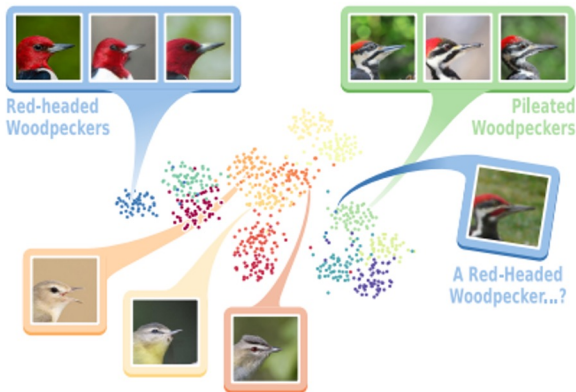
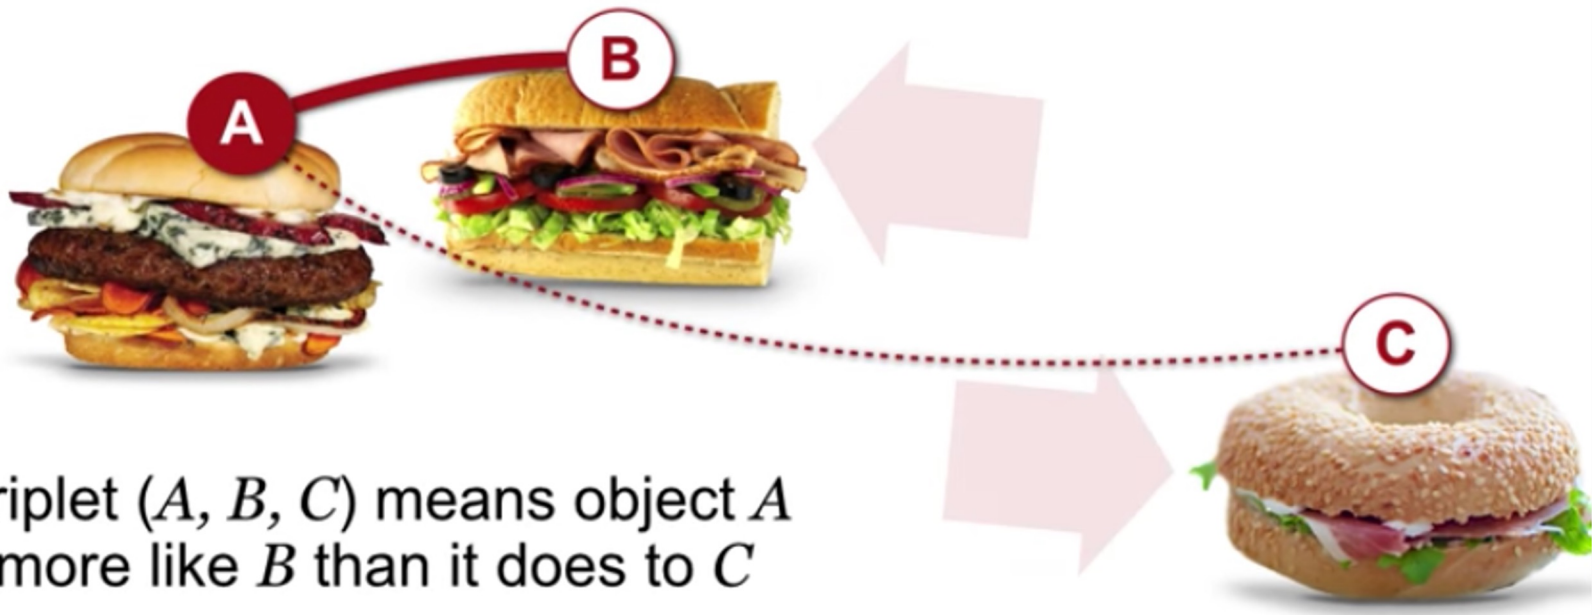


Figure 1. Our SNaCK embeddings capture human expertise with the help of machine similarity kernels. For example, an expert can use this concept embedding of a subset of CUB-200 to quickly find labeling mistakes. Red-headed Woodpeckers are visually dissimilar to Pileated Woodpeckers, but SNaCK moved a Red-headed

Let's ask experts directly!

Workers on Mechanical Turk give us thousands of **triplet constraints** that directly tell us about taste.




Each triplet (A, B, C) means object A tastes more like B than it does to C

Our design: **Grid questions**

To collect constraints efficiently, we ask:



“Please select the **four foods** that taste similar to the one on the left.”



Our design: **Grid questions**

To collect constraints efficiently, we ask:

“Please select the **four foods** that taste similar to the one on the left.”



Our design: **Grid questions**

Selected points are more similar to the reference food than unselected points.



Near



Far

Mixing humans and machines

Human Expertise



Our SNaCK embedding algorithm combines the expert's **triplet constraints** with a **visual similarity kernel**.


Deep Learning



Our Algorithm



Concept Embedding



Mixing humans and machines

Human Expertise



Our SNaCK embedding algorithm combines the expert's **triplet constraints** with a **visual similarity kernel**.

$$C_{SNaCK} = \lambda \cdot C_{tSTE} + (1 - \lambda) \cdot C_{tSNE}$$


Deep Learning



Our Algorithm



Concept Embedding



t-STE for Music Genre Discovery

[van der Maaten & Weinberger 2012]

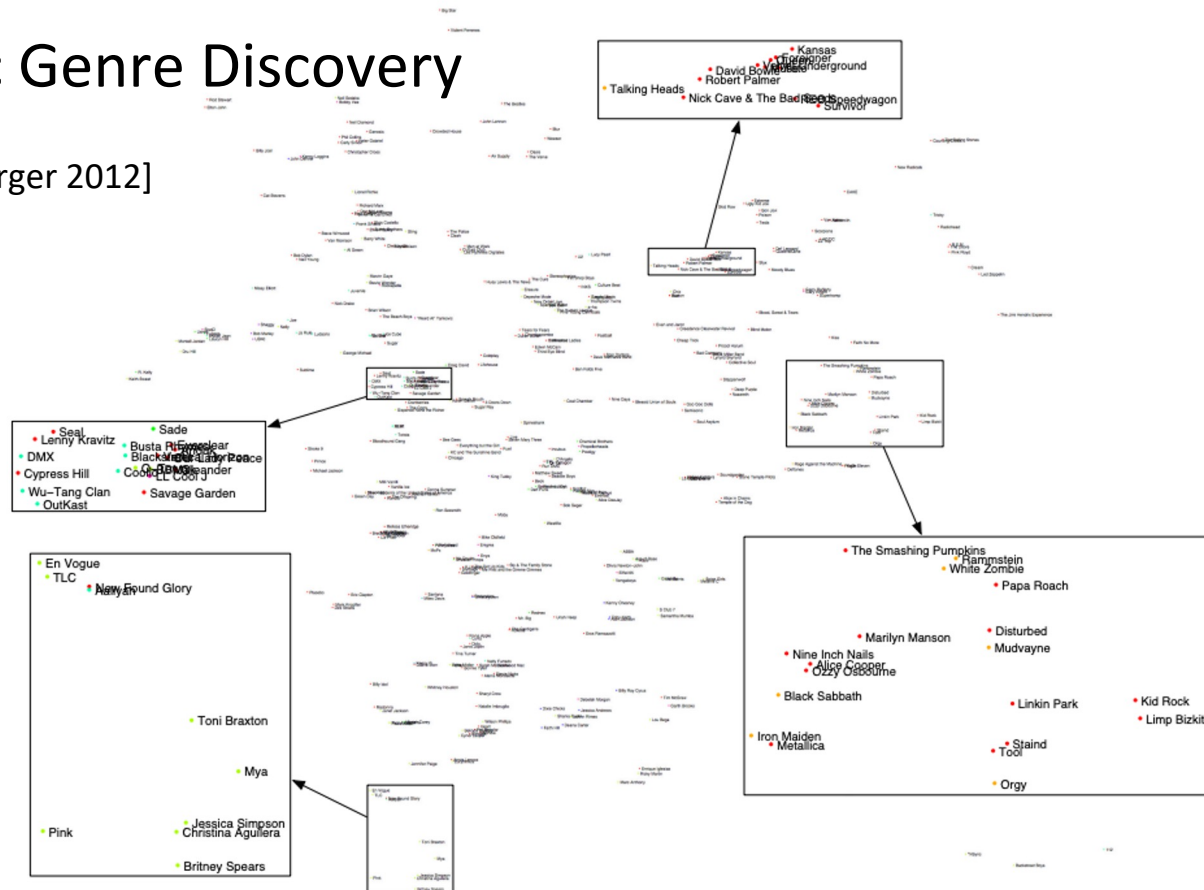
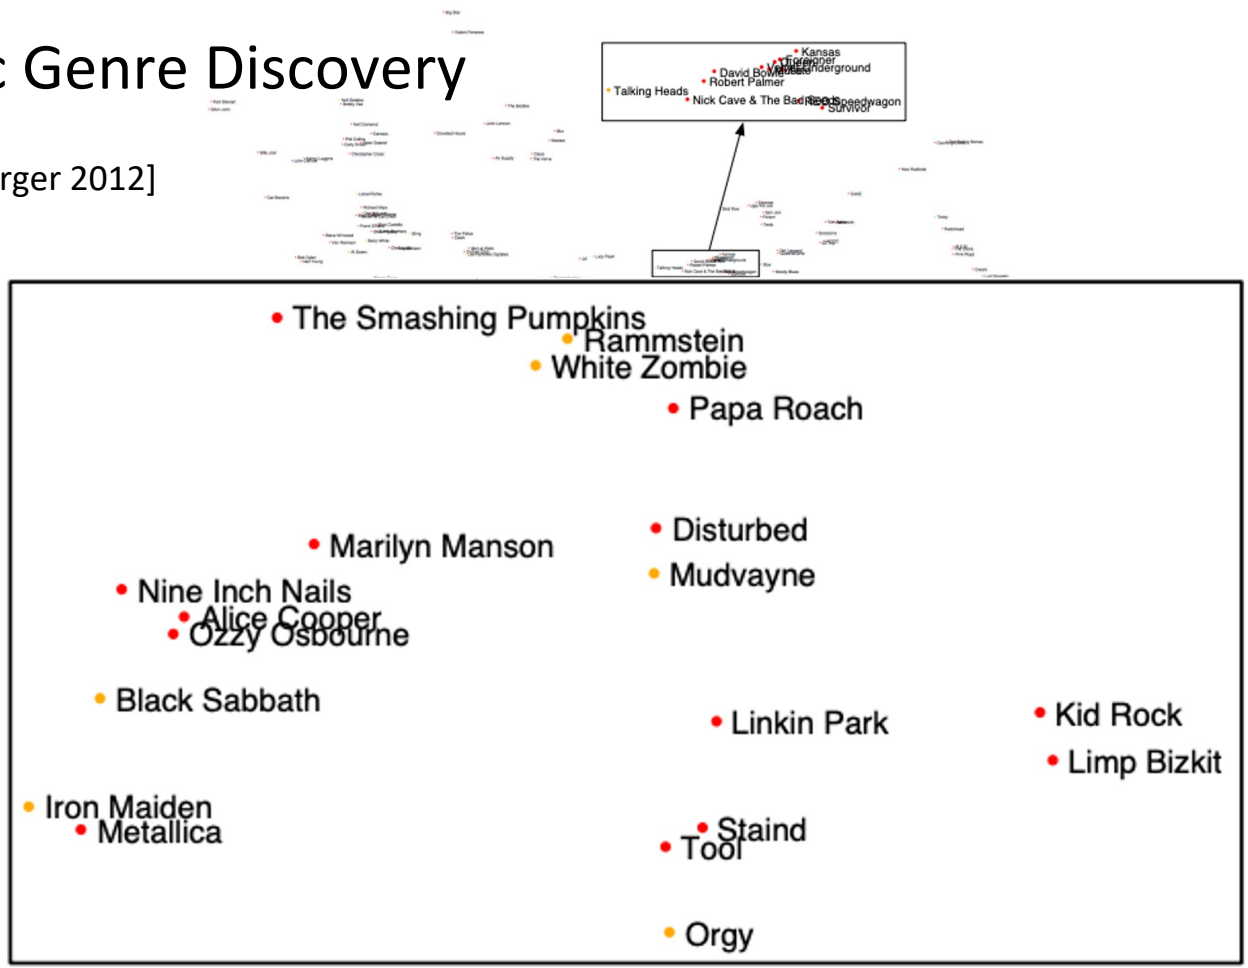
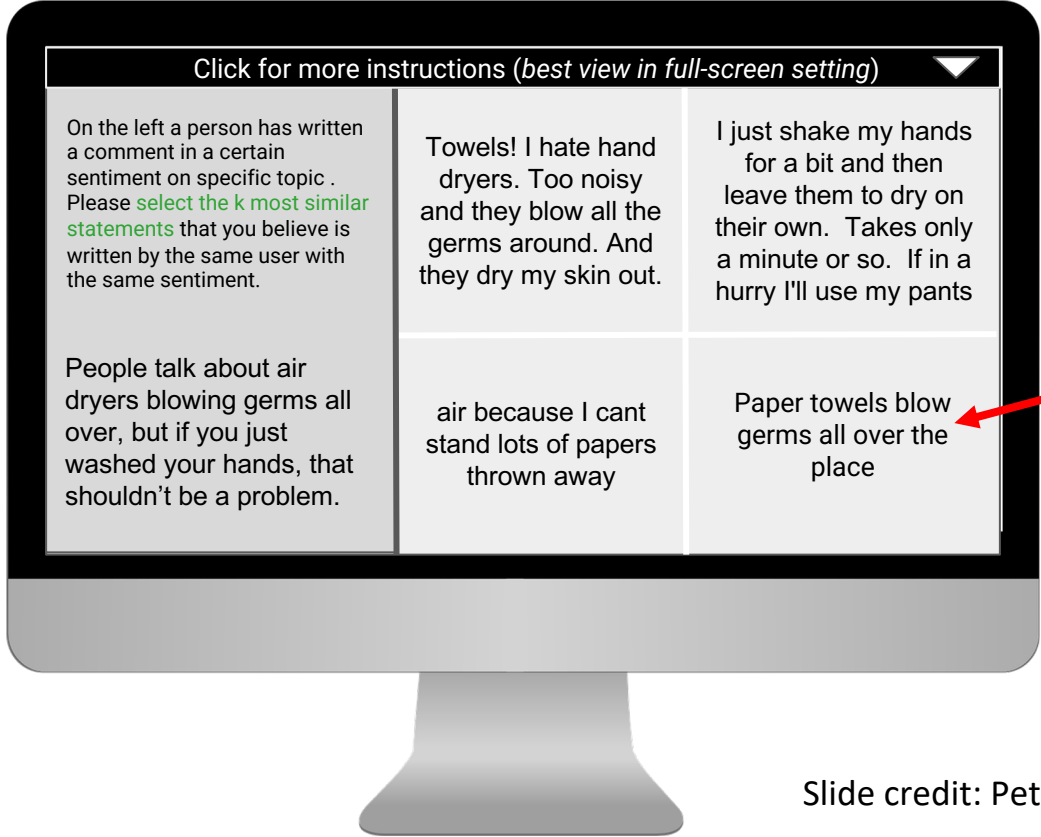


Fig. 4. Two-dimensional music artist map constructed by t-STE based on all triplets in the music artists data set. A larger version of the map is available on <http://homepage.tudelft.nl/19j49/ste>.

t-STE for Music Genre Discovery

[van der Maaten & Weinberger 2012]






\s alert!


Slide credit: Peter Ebert Christensen

Mediated machine vision


Datasets registered,
openly licensed, issued
DOI




Training datasets
prepared, openly
licensed, issued DOI



Models used in
tools, cited using
DOI



Versioned models
published, openly
licensed, issued
DOI



Key Questions

- Narrative sourcing
- Efficient crowd annotator mobilization
- Interplay with mainstream fact-checking research
- Accounting for sarcasm, humor, nuance
- Can a “Narrative2Vec” approach suffice?
- The role of image meme templates
- Methods of surfacing/presenting “x-ray goggles” information
- The long tail of misinformation
- Adversarial attacks
- Sustainable workflows w/journalists, social media, users
- Kitchen Sink vs. Cohesive Intelligence

The New York Times

What Makes a ‘Narrative Violation’?

Venture capitalists have a favorite new term.



Photo Illustration by The New York Times; Getty Images (Derrida, Foucault), Shutterstock (face)



By Erin Griffith

Nov. 26, 2019