Digital humanities: modeling semi-structured data from traditional scholarship

Tom Lippincott (tom@cs.jhu.edu)
NLP Fall 2019
Human Language Technology Center of Excellence
Center for Language and Speech Processing

Outline

Intro: A few thoughts on “Digital humanities”

Graphs and Autoencoders

Motivating study: Post-Atlantic Slave Trade

Another application: Authorship attribution of the Hebrew bible

Ongoing work

Intro: A few thoughts on “Digital humanities”

What is “digital humanities”?

Some responses:

- “an idea that will increasingly become invisible” -Stanford
- “a term of tactical convenience” -UMD
- “I don’t: I’m sick of trying to define it” -GMU
- “a convenient label, but fundamentally I dont believe in it” -NYU
- “an unfortunate neologism” -Library of Congress

What is “digital humanities”?

Themes at DH2019

- Visualization
- Geographic information systems
- Social and ethical issues
- Education
- VR, maker spaces
- OCR
- Machine learning

Working definitions

Digital humanities

Traditional scholar

(Traditional) scholarly dataset

Computational researcher

Working definitions

Digital humanities

Traditional inquiries enabled by computational intelligence

Traditional scholar

(Traditional) scholarly dataset

Computational researcher

Working definitions

Digital humanities
Traditional inquiries enabled by computational intelligence

Traditional scholar
Academic from field that doesn’t typically employ quantitative methods (History, Literary Criticism, . . .)

(Traditional) scholarly dataset

Computational researcher

Working definitions

Digital humanities
Traditional inquiries enabled by computational intelligence

Traditional scholar
Academic from field that doesn’t typically employ quantitative methods (History, Literary Criticism, . . .)

(Traditional) scholarly dataset
Data assembled by a traditional researcher in the field

Computational researcher

Working definitions

Digital humanities
Traditional inquiries enabled by computational intelligence

Traditional scholar
Academic from field that doesn't typically employ quantitative methods (History, Literary Criticism, . . .)

(Traditional) scholarly dataset
Data assembled by a traditional researcher in the field

Computational researcher
Design and bring machine learning models to bear on datasets

Why is collaboration rare?

Traditional scholars have insight into the *data*

Computational researchers can pair *data* with appropriate models

Why is collaboration rare?

Traditional scholars have insight into the *data*

- Data is painstakingly gathered and coveted
- Hypotheses are *subtle* but not numerically evaluated
- May publish one or two papers during PhD, but *dissertation* is primary focus

Computational researchers can pair *data* with appropriate models

Why is collaboration rare?

Traditional scholars have insight into the *data*

- Data is painstakingly gathered and coveted
- Hypotheses are *subtle* but not numerically evaluated
- May publish one or two papers during PhD, but *dissertation* is primary focus

Computational researchers can pair *data* with appropriate models

- Data is aggressively shared to encourage rigorous evaluation
- Tasks are often *shallow* and *prespecified*
- Publish multiple papers per year

Topic models: a success story

Topic models: a success story

Widely used

- Low barrier to entry: everyone has “documents”
- Little expertise required to train
- Output easy to visualize and interpret

Topic models: a success story

Widely used

- Low barrier to entry: everyone has “documents”
- Little expertise required to train
- Output easy to visualize and interpret

Widely abused

- Deceptively easy to use: it will always give you *something*
- You can always find “patterns”: confirmation bias abounds
- Older than some undergrads: LDA from early 2000s

A guiding challenge:

Can we leverage sophisticated modeling techniques without losing the advantages that popularize topic models and recreating some of the same bad community practices?

Aside: Traditional Scholars are Knowledge Workers

Financial analysts, investigative reporters . . .

- Concerned with specific domains
- Need to gather and understand datasets
- Construct and reason over knowledge bases
- Wide range of technical abilities
- The DH story is relevant to industry, government, etc

Graphs and Autoencoders

General relational data

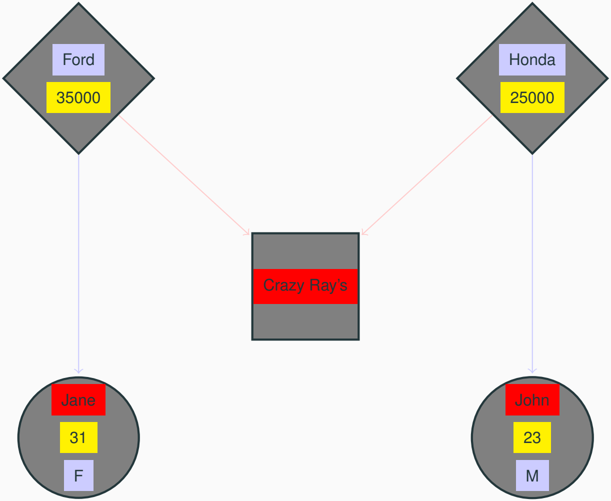

Common format: JSON

```
[
  {”name” : ”John”, ”age” : 23, ”gender” : ”M”},
  {”name” : ”Jane”, ”age” : 31, ”gender” : ”F”},
  {”business” : ”Crazy Ray’s”},
  {”make” : ”Honda”, ”price” : 25000,
       ”owned_by” : 0, ”sold_by” : 2},
  {”make” : ”Ford”, ”price” : 35000,
       ”owned_by” : 1, ”sold_by” : 2}
]
```

Let’s design a model that naturally adapts to the data structure

Building-blocks for modeling relational data

Encoders, decoders, and autoencoders

Capture the *entities* and *fields*

Graph convolutional networks

Capture the *relationships*

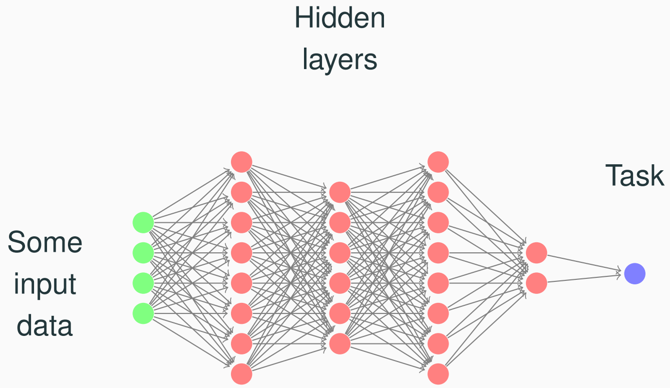
Hidden
layers
Task
Some
input
data

Hidden
layers

Some
input
data

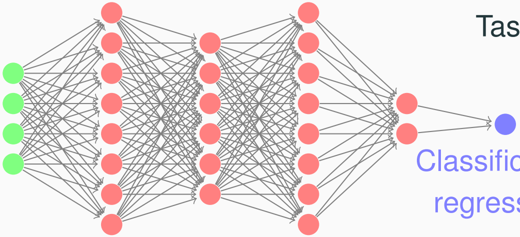

Task

Classification,
regression,
generation,
some combo ...

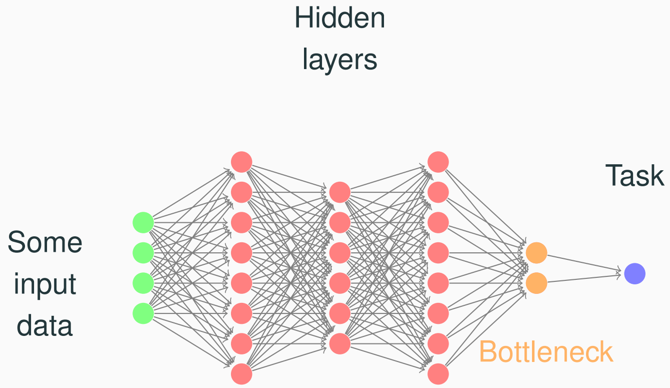
Hidden
layers
Some
input
data
Task
Bottleneck

Encoder

Hidden
layers

Some
input
data

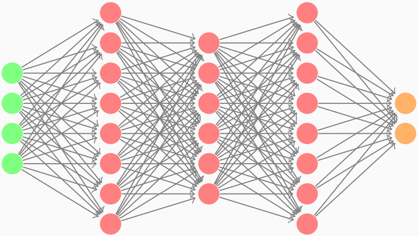

Decoder

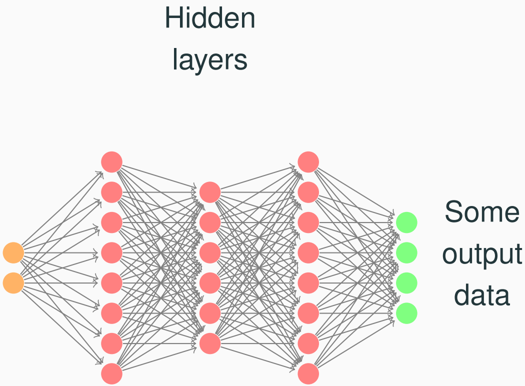

Encoders and decoders are often paired

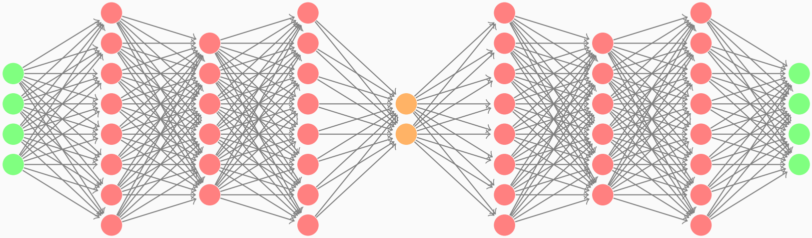

If the goal is to reconstruct the input, it's an autoencoder

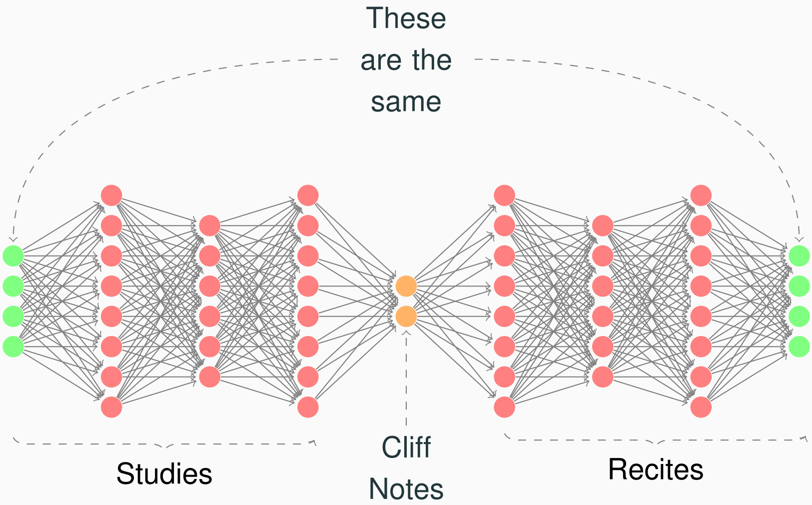

*coder summary

- An **encoder** transforms data into a fixed-length representation
- A **decoder** takes a fixed-length representation and generates data
- An **autoencoder** is an encoder and decoder working together to preserve data through a **bottleneck**

On to *graph convolutions*. . .

Normal 1D CNN

Grid (image, text . . .)

Normal 1D CNN

Grid (image, text . . .)

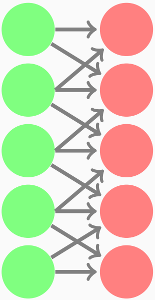

Normal 1D CNN

Grid (image, text . . .)

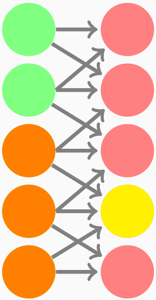

Each position incorporates its "receptive field"

Normal 1D CNN

Grid (image, text . . .)

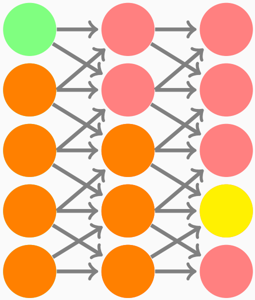

Repeat process, expand field

Graph convolutional network (GCN)

Graph nodes
(e.g. entities)

Graph convolutional network (GCN)

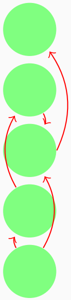

Graph nodes
(e.g. entities)

Adjacent
nodes
(related
entities)

Graph convolutional network (GCN)

Graph nodes
(e.g. entities)

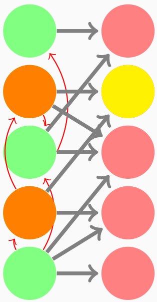

Each node
incorporates
its neighbors

Graph convolutional network (GCN)

Graph nodes
(e.g. entities)

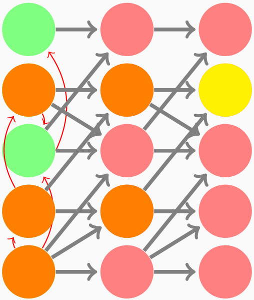

Info spreads
according
to graph

GCN summary

- Extends CNNs from grids to graphs
- Information passes along edges
- Each GCN layer allows nodes to see one further “hop”

Modeling relational data

- Encoders, decoders, autoencoders
- Graph convolutional mechanism
- **Combine these to match the data being modeled**

Graph Entity Autoencoder (GEA)

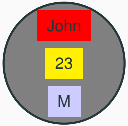

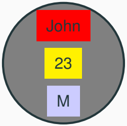

Graph Entity Autoencoder (GEA)

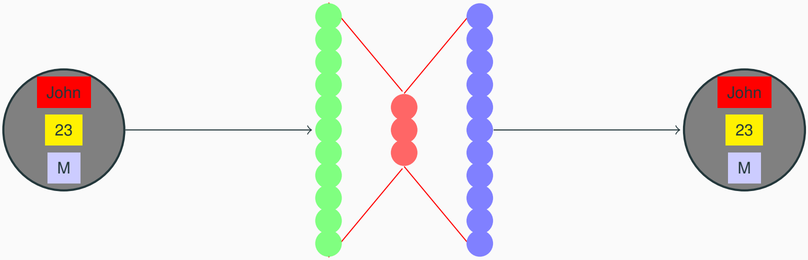

Graph Entity Autoencoder (GEA)

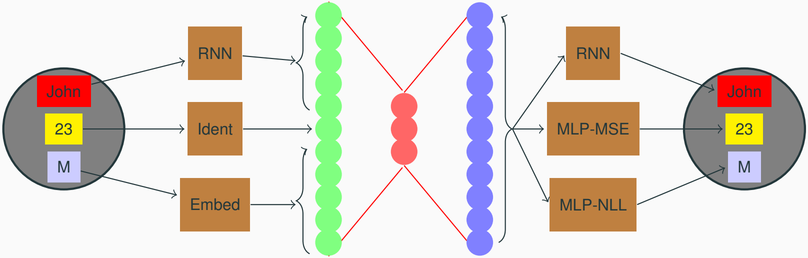

Graph Entity Autoencoder (GEA)

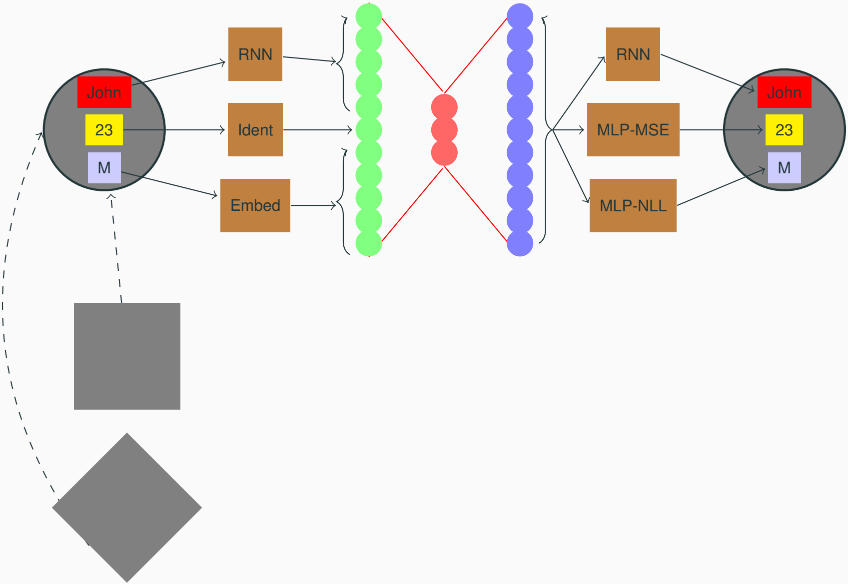

Graph Entity Autoencoder (GEA)

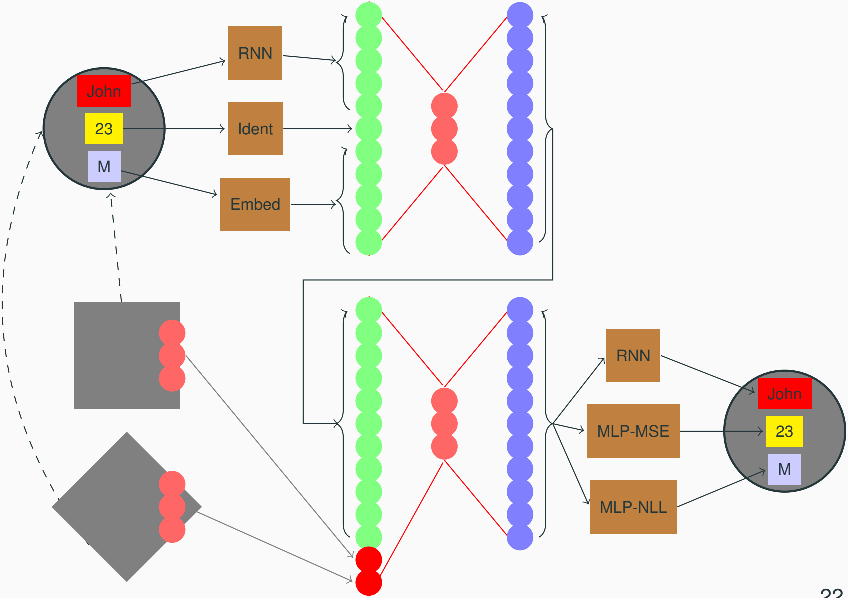

Graph Entity Autoencoder (GEA)

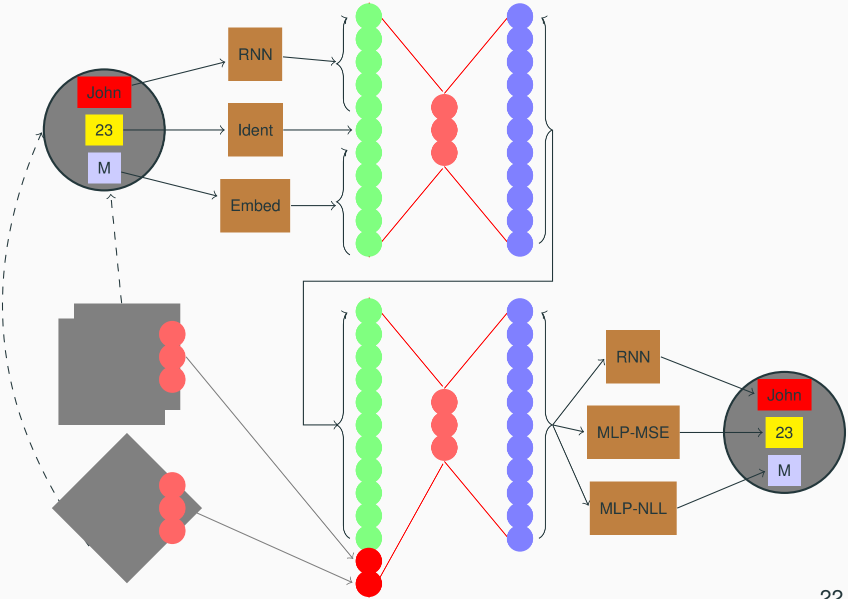

Graph Entity Autoencoder (GEA)

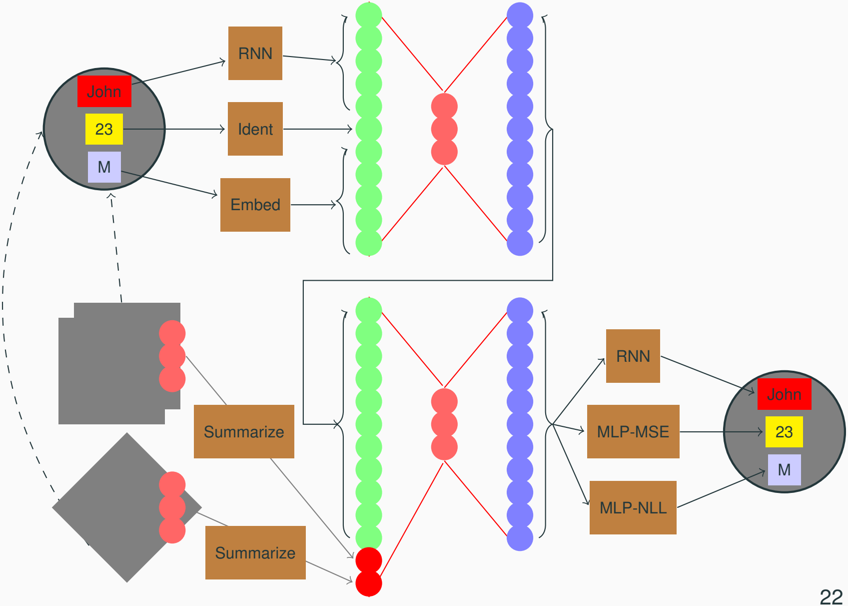

How can we use a trained model?

- Compute distance between two entities
- Find flat or hierarchical clusters of entities
- Generate likely value of missing field
- Detect an improbable value of a present field
- Observe response of one field to another

Motivating study: Post-Atlantic Slave Trade

Shipping manifests

Shipping manifests

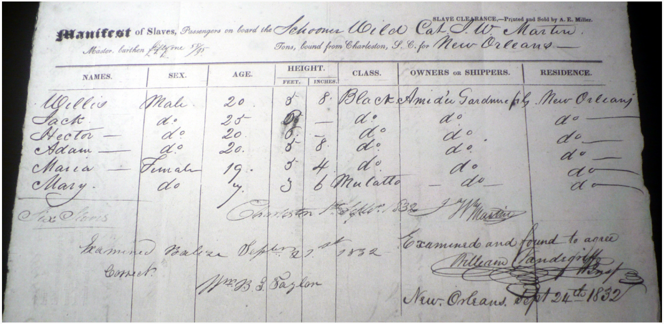

SLAVE CLEARANCE,—Printed and Sold by A. E. Miller.

Manifest of Slaves, Passengers on board the Schooner Wild Cat J. W. Martin, Master, burthen fifty one 5/95 Tons, bound from Charleston, S. C. for New Orleans —

NAMES.	SEX.	AGE.	HEIGHT. FEET.	HEIGHT. INCHES.	CLASS.	OWNERS OR SHIPPERS.	RESIDENCE.
Willis	Male	20	5	8	Black	Amidée Gardner fils	New Orleans
Jack	do	25	5	—	do	do	do
Hector	do	20	5	—	do	do	do
Adam	do	20	5	8	do	do	do
Maria	Female	19	5	4	do	do	do
Mary	do	7	3	6	Mulatto	do	do

Six Slaves

Charleston 1st Septr. 1832 J. W. Martin

Examined Balize Septr 27th 1832 Correct.

Wm. B. G. Taylor

Examined and found to agree William Vandegriff

New Orleans Sept 24th 1832

Shipping manifests

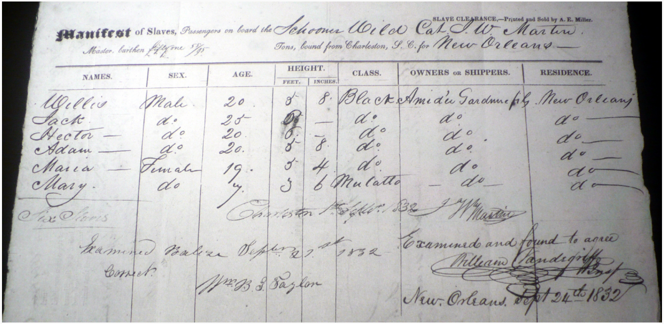

SLAVE CLEARANCE,—Printed and Sold by A. E. Miller.

Manifest of Slaves, Passengers on board the Schooner Wild Cat J. W. Martin, Master, burthen fifty one 5/95 Tons, bound from Charleston, S. C. for New Orleans —

NAMES.	SEX.	AGE.	HEIGHT. FEET.	HEIGHT. INCHES.	CLASS.	OWNERS OR SHIPPERS.	RESIDENCE.
Willis	Male	20	5	8	Black	Amidée Gardner fils	New Orleans
Jack	do	25	[illegible]	—	do	do	do
Hector	do	20	6	—	do	do	do
Adam	do	20	5	8	do	do	do
Maria	Female	19	5	4	do	do	do
Mary	do	7	3	6	Mulatto	do	do

Six Slaves

Charleston 1st Septr 1832 J. W. Martin

Examined Balize Septr 21st 1832 Correct. Wm B. G. Taylor

Examined and found to agree William Vandegriff

New Orleans Sept 24th 1832

slave name	slave sex	slave age	owner name	journey date	vessel type

Shipping manifests

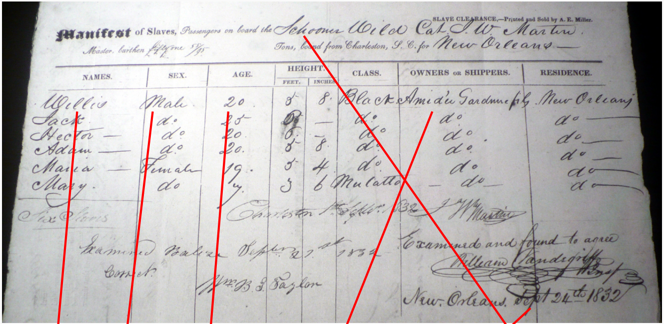

slave name	slave sex	slave age	owner name	journey date	vessel type
Willis	m	20	Amidu	1832/9/24	Schooner

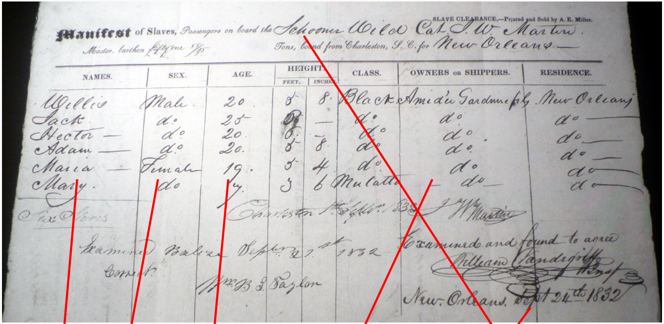
SLAVE CLEARANCE.—Printed and Sold by A. E. Miller.

Manifest of Slaves, Passengers on board the Schooner Wild Cat J. W. Martin Master, burthen fifty one 5/95 Tons, bound from Charleston, S. C. for New Orleans

NAMES.	SEX.	AGE.	HEIGHT. FEET.	HEIGHT. INCHES.	CLASS.	OWNERS OR SHIPPERS.	RESIDENCE.
Willis	Male	20	5	8	Black	Amidu [illegible]	New Orleans
Jack	do	25	[illegible]	—	do	do	do
Hector	do	20	6		do	do	do
Adam	do	20	5	8	do	do	do
Maria	Female	19	5	4	do	do	do
Mary	do	14	3	6	Mulatto	do	do

Six Slaves

Charleston [illegible] 1832 J. W. Martin

Examined Balize Septr 27th 1832 Correct. Wm B. G. Taylor

Examined and found to agree William Vandegriff

New Orleans Sept 24th 1832

slave name	slave sex	slave age	owner name	journey date	vessel type
Willis	m	20	Amidu	1832/9/24	Schooner
Maria	f	19	Amidu	1832/09/24	Schooner

Fugitive notices

Fugitive notices

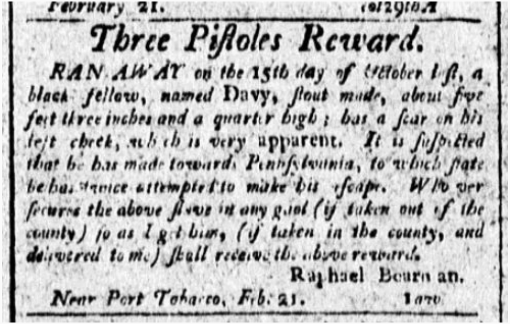

February 21. [illegible]

Three Pistoles Reward.

RAN AWAY on the 15th day of October last, a black fellow, named Davy, stout made, about five feet three inches and a quarter high; has a scar on his left cheek, which is very apparent. It is suspected that he has made towards Pennsylvania, to which state he has twice attempted to make his escape. Whoever secures the above slave in any gaol (if taken out of the county) so as I get him, (if taken in the county, and delivered to me) shall receive the above reward.

Raphael Bourman.

Near Port Tobacco, Feb. 21. [illegible]

Fugitive notices

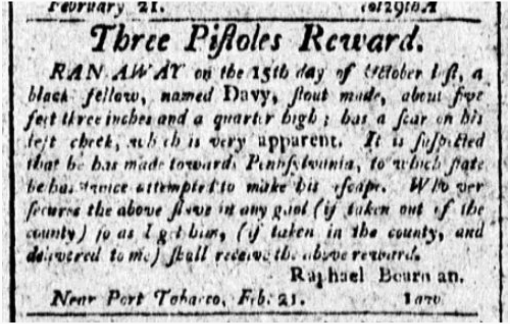

February 21. [illegible]

Three Pistoles Reward.

RAN AWAY on the 15th day of October last, a black fellow, named Davy, stout made, about five feet three inches and a quarter high; has a scar on his left cheek, which is very apparent. It is suspected that he has made towards Pennsylvania, to which state he has twice attempted to make his escape. Whoever secures the above slave in any gaol (if taken out of the county) so as I get him, (if taken in the county, and delivered to me) shall receive the above reward.

Raphael Bourman.

Near Port Tobacco, Feb. 21. [illegible]

slave name	slave sex	escape date	escape location	owner name	notice reward	notice date

Fugitive notices

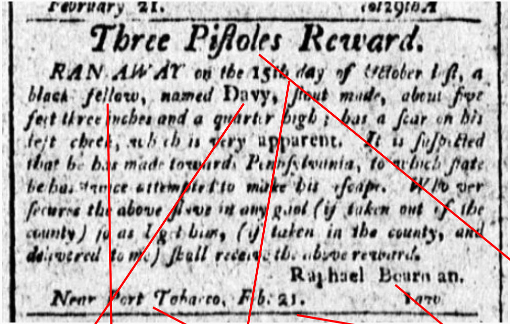

February 21.

Three Pistoles Reward.

RAN AWAY on the 15th day of October last, a black fellow, named Davy, stout made, about five feet three inches and a quarter high; has a scar on his left cheek, which is very apparent. It is suspected that he has made towards Pennsylvania, to which state he has twice attempted to make his escape. Whoever secures the above slave in any gaol (if taken out of the county) so as I get him, (if taken in the county, and delivered to me) shall receive the above reward.

Raphael Bourman.

Near Port Tobacco, Feb. 21.

slave name	slave sex	escape date	escape location	owner name	notice reward	notice date
Davy	m	1795/10/15	Port Tobacco	Bourman	3 Pistoles	1796/02/21

Some numbers

- 45k manifest entries spanning five cities
- 11k fugitive notices from 70 gazettes
- 28k unique slave names
- 7k unique owner names
- Not big data, but thousands of studies like this at a research university!

Difficulties with data in the wild

Difficulties with data in the wild

- Unnormalized
  - People/places/things recorded *many* times
  - “What’s the age/height/sex distribution of escapees?”

Difficulties with data in the wild

- Unnormalized
  - People/places/things recorded *many* times
  - “What’s the age/height/sex distribution of escapees?”
- Noisy
  - Vessel type: Bark, Barke, BArque, Barque, Barques
  - Slave name: “Nelly’?, Nelly’s child”, “not visible”
  - Owner sex: 3k missing

Difficulties with data in the wild

- Unnormalized
  - People/places/things recorded *many* times
  - “What’s the age/height/sex distribution of escapees?”
- Noisy
  - Vessel type: Bark, Barke, BArque, Barque, Barques
  - Slave name: “Nelly’?, Nelly’s child”, “not visible”
  - Owner sex: 3k missing
- Underspecified entities
  - *Majority* of slaves have no last name
  - Can’t tell if two “Johns” are the same person

What might a historian want to do with this data?

- Follow one slave throughout their life
- Group owners according to the nature of their workforce
- Map out trade “ecosystems” of sellers, shippers, owners, etc
- Determine what drove valuation in transactions and rewards
- Reconstruct slave families when there are no last names

Ask the traditional scholar to follow some simple guidelines when gathering data

Traditional scholarly data

slave_name	*Jim*
slave_age	20
owner_name	*Jane*
owner_sex	*F*
vessel_name	*Uncas*
vessel_type	*Brig*
voyage_date	6/2/1823
voyage_dest	29.9, 90.0
...	...

Numbers

slave_name	*Jim*
slave_age	20
owner_name	*Jane*
owner_sex	*F*
vessel_name	*Uncas*
vessel_type	*Brig*
voyage_date	6/2/1823
voyage_dest	29.9, 90.0
...	...

Entities, field types, and relations

Categories

slave_name	*Jim*
slave_age	20
owner_name	*Jane*
owner_sex	*F*
vessel_name	*Uncas*
vessel_type	*Brig*
voyage_date	6/2/1823
voyage_dest	29.9, 90.0
...	...

Entities, field types, and relations

Strings

slave_name	*Jim*
slave_age	20
owner_name	*Jane*
owner_sex	*F*
vessel_name	*Uncas*
vessel_type	*Brig*
voyage_date	6/2/1823
voyage_dest	29.9, 90.0
...	...

More complex fields

slave_name	*Jim*
slave_age	20
owner_name	*Jane*
owner_sex	*F*
vessel_name	*Uncas*
vessel_type	*Brig*
voyage_date	6/2/1823
voyage_dest	29.9, 90.0
...	...

Entities, field types, and relations

Entities

slave_name	*Jim*
slave_age	20
owner_name	*Jane*
owner_sex	*F*
vessel_name	*Uncas*
vessel_type	*Brig*
voyage_date	6/2/1823
voyage_dest	29.9, 90.0
...	...

Entities, field types, and relations

Entities

slave_name	*Jim*
slave_age	20
owner_name	*Jane*
owner_sex	*F*
vessel_name	*Uncas*
vessel_type	*Brig*
voyage_date	6/2/1823
voyage_dest	29.9, 90.0
...	...

Entities, field types, and relations

Slave-to-owner

slave_name	*Jim*
slave_age	20
owner_name	*Jane*
owner_sex	*F*
vessel_name	*Uncas*
vessel_type	*Brig*
voyage_date	6/2/1823
voyage_dest	29.9, 90.0
...	...

Vessel-to-voyage, slave-to-voyage

slave_name	*Jim*
slave_age	20
owner_name	*Jane*
owner_sex	*F*
vessel_name	*Uncas*
vessel_type	*Brig*
voyage_date	6/2/1823
voyage_dest	29.9, 90.0
...	...

Entities, field types, and relations

Fewer assumptions

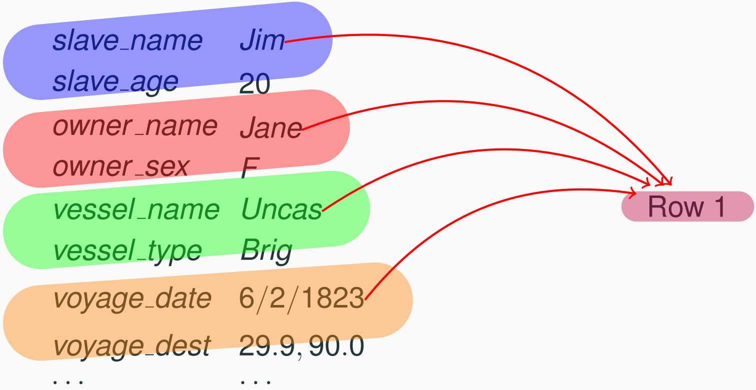

Example data point: one graph component

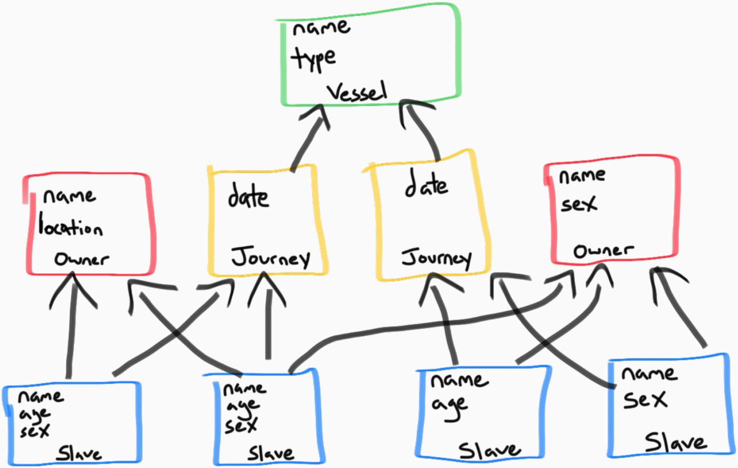

Train a GEA model . . .

Example insights looking at most-similar entities

Mistranscriptions

Baltiomre, Austin Woolfolk	$\Longleftrightarrow$	Baltimore, Austin Woolfolk
New Orleans, William Wiliams	$\Longleftrightarrow$	New Orleans, William Williams

Semantically-equivalent variants

Baltimore, George Y. Kelso	$\Leftrightarrow$	Baltimore, Kelso & Ferguson
New Orleans, Leon Chabert	$\Leftrightarrow$	Louisiana, Leon Chabert

Same slave transported multiple times

Louisa, F, 16yo	$\Leftrightarrow$	Louisa, F, 17yo
Waters, F, 14yo	$\Leftrightarrow$	Waters, F, 15yo
Kesiah, F, 20yo	$\Leftrightarrow$	Kesiah, F, 22yo
Taylor, F, 15yo	$\Leftrightarrow$	Taylor, F, 16yo

Another application: Authorship attribution of the Hebrew bible

Transmission of a text: the "Documentary Hypothesis"

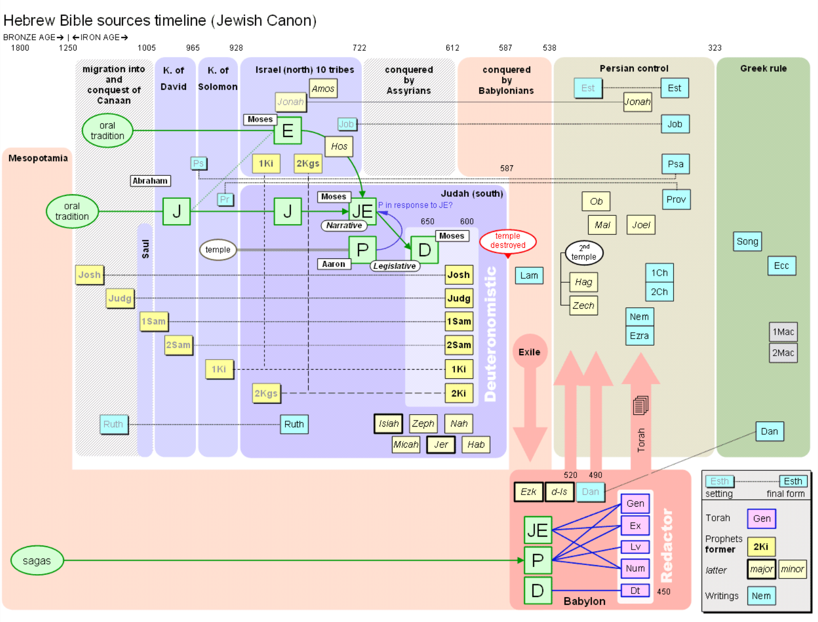

Hypothesis as pointers into document structure

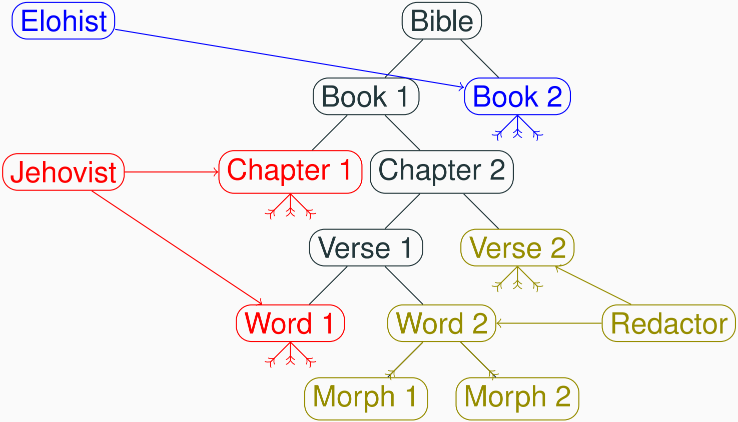

Assume the hypothesis, see how various models and features learn it as a supervised classification problem

Thomas Mendenhall: The Characteristic Curves of Composition

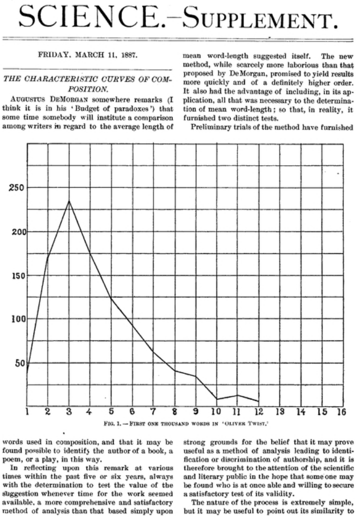

SCIENCE.—SUPPLEMENT.

FRIDAY, MARCH 11, 1887.

THE CHARACTERISTIC CURVES OF COMPOSITION.

AUGUSTUS DEMORGAN somewhere remarks (I think it is in his 'Budget of paradoxes') that some time somebody will institute a comparison among writers in regard to the average length of words used in composition, and that it may be found possible to identify the author of a book, a poem, or a play, in this way.

In reflecting upon this remark at various times within the past five or six years, always with the determination to test the value of the suggestion whenever time for the work seemed available, a more comprehensive and satisfactory method of analysis than that based simply upon mean word-length suggested itself. The new method, while scarcely more laborious than that proposed by DeMorgan, promised to yield results more quickly and of a definitely higher order. It also had the advantage of including, in its application, all that was necessary to the determination of mean word-length; so that, in reality, it furnished two distinct tests.

Preliminary trials of the method have furnished strong grounds for the belief that it may prove useful as a method of analysis leading to identification or discrimination of authorship, and it is therefore brought to the attention of the scientific and literary public in the hope that some one may be found who is at once able and willing to secure a satisfactory test of its validity.

The nature of the process is extremely simple, but it may be useful to point out its similarity to

FIG. 1.—FIRST ONE THOUSAND WORDS IN 'OLIVER TWIST.'

Mosteller and Wallace: Inference in an Authorship Problem

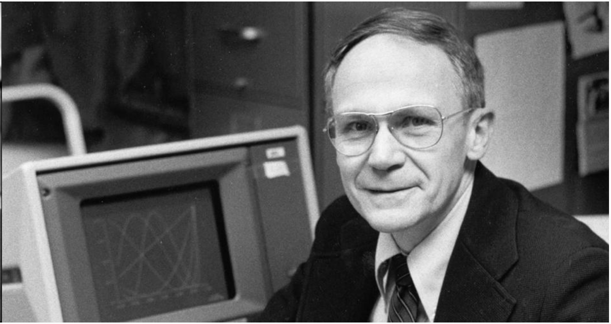

The Federalist papers

- 85 articles written by Hamilton, Madison, and Jay
- 12 are unattributed
- Frequency analysis of *function words* determined Madison as author

Back to the Documentary Hypothesis

Problems

- The “authors” are also editors, redactors, synthesizers . . . they interact in context-dependent ways
- There is no predefined segmentation into “articles”
- We *know* more than function-words are important (e.g. name of God)

Solutions

- Limit vocabulary to words that are used frequently by all authors
- Employ a GCN to exploit the document structure

GEA predicts the author *slightly* better . . .

Model	F-score
LR	41.39
MLP	47.45
GEA	48.60

Gold	Guess							
	J	E	P	1D	2D	nD	R	O
J	100	8	7	0	0	0	3	0
E	22	53	8	0	0	0	0	0
P	13	5	77	0	1	0	4	0
1D	2	0	2	7	1	0	0	0
2D	2	2	1	0	5	0	0	0
nD	0	0	0	1	0	0	0	0
R	3	3	11	0	0	0	33	0
O	2	0	1	0	0	0	1	0

Error analysis

Sentiment and in-context word senses

- “wife” shows up as polygamous in older but monogamous in newer sources
- Redactor’s positive view of Aaron+Moses, violent story of rebellion

Narrative continuity

- Abraham and Isaac story thought to *originally end with sacrifice*, changed by the Redactor
- “it was the season for grapes”
  (travel and geographic locations)
  “They broke off some grapes.”

Ongoing work

Visualizing results

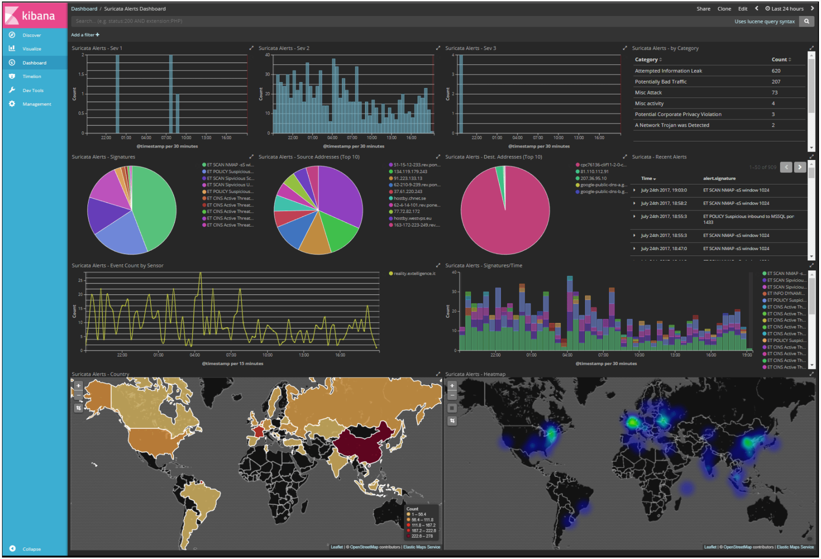

Other applications

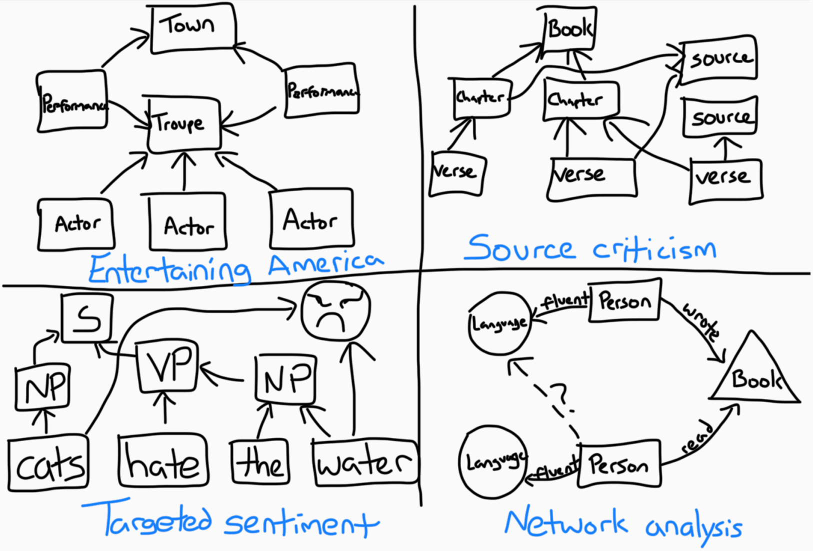

Thanks!

Quick plug: come to David Mimno's talk!

- Nov. 15 at noon (Hackerman B17)
- CS Professor at Cornell
- Rare CS faculty working in DH (topic modeling)

References

- Embedding Multimodal Relational Data for Knowledge Base Completion, Singh et al., 2018
- Inductive Representation Learning on Large Graphs, Hamilton et al., 2017
- Semi-Supervised Classification with Graph Convolutional Networks, Kipf et al., 2016
- Reducing the Dimensionality of Data with Neural Networks, Hinton et al., 2006