Statistical Topic Models for Computational Social Science

Hanna M. Wallach

University of Massachusetts Amherst

wallach@cs.umass.edu

Complex Social Processes

~~TOP SECRET~~

21 FEB 67

2541

OUTLINE

Page 49

I. Military actions against North Vietnam and In Laos

A. Present program 1

B. Options for increased military programs 2

1. Destroy modern industry 3
   - Thermal power (7-plant grid)?
   - Steel and cement
   - Machine tool plant
   - Other

SANITIZED
E.O. 12356, Sec. 3.4
NLJ 90-192
By ~~LRP~~, NARA, Date 4-6-93

(2.) Destroy dikes and levees 6

“Traditional” Social Science

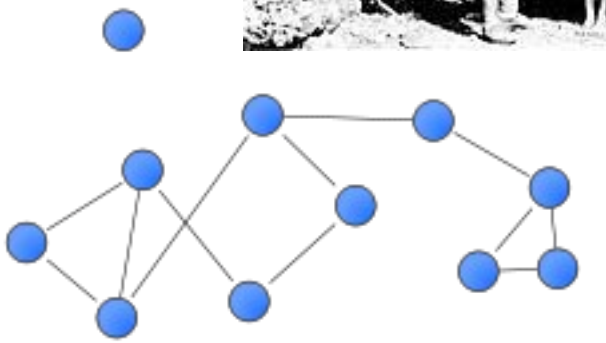

- Case studies
- Interviews
- Participant observation
- Survey research
- Social network analysis

⇒ Self-reports, one-time snapshots, small scale

The Computer "Revolution"

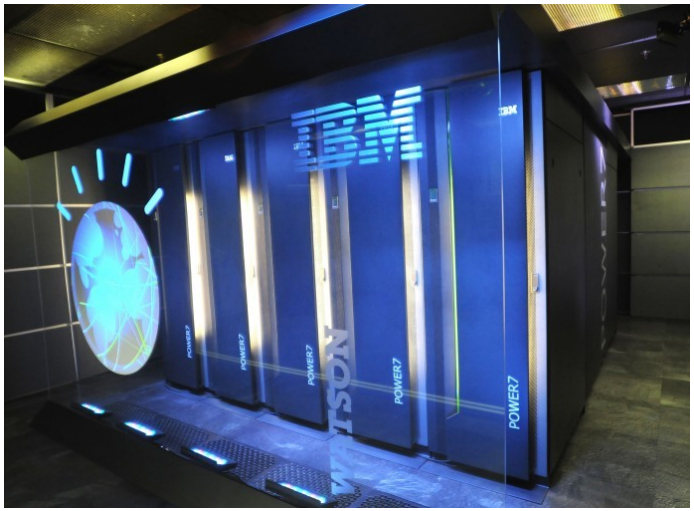

hadoop MapReduce

Google

The Unreasonable Effectiveness of Data

Computational Social Science

Candida Hofer

"A computational social science is emerging that leverages the capacity to collect and analyze data with an unprecedented breadth and depth and scale and may reveal patterns of individual and group behaviors."

— Lazer et al., 2009

Structure vs. Content

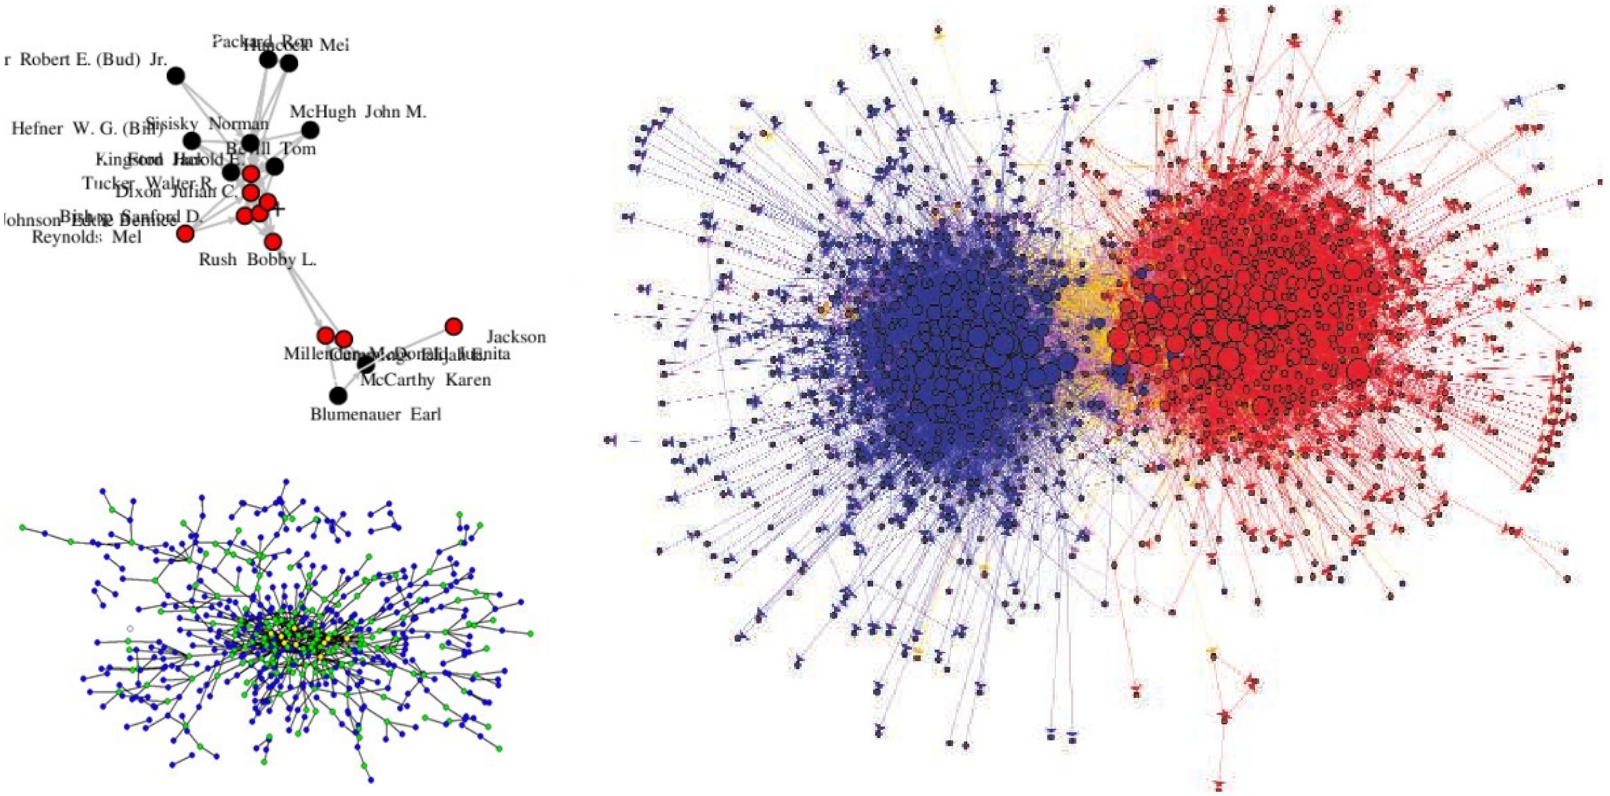

Products of Interactions

“Scientific information is both the basic raw material for, and one of the principal products of, scientific research [...] Scientists find out what other scientists are accomplishing through [...] journals, books, abstracts and indexes, bibliographies, reviews.”

— NSF Brochure, 1962

Text as Data

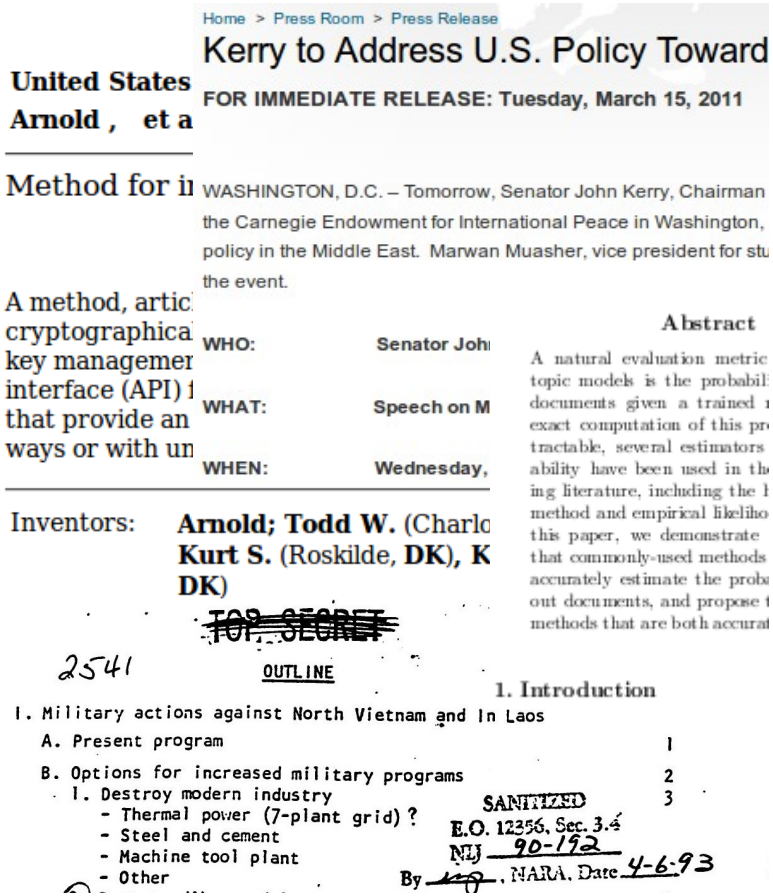

- Structured and formal: e.g., publications, patents, press releases
- Messy and unstructured: e.g., chat logs, OCRed documents, transcripts

⇒ Large scale, robust methods for analyzing text

Collaborate to Study Collaboration

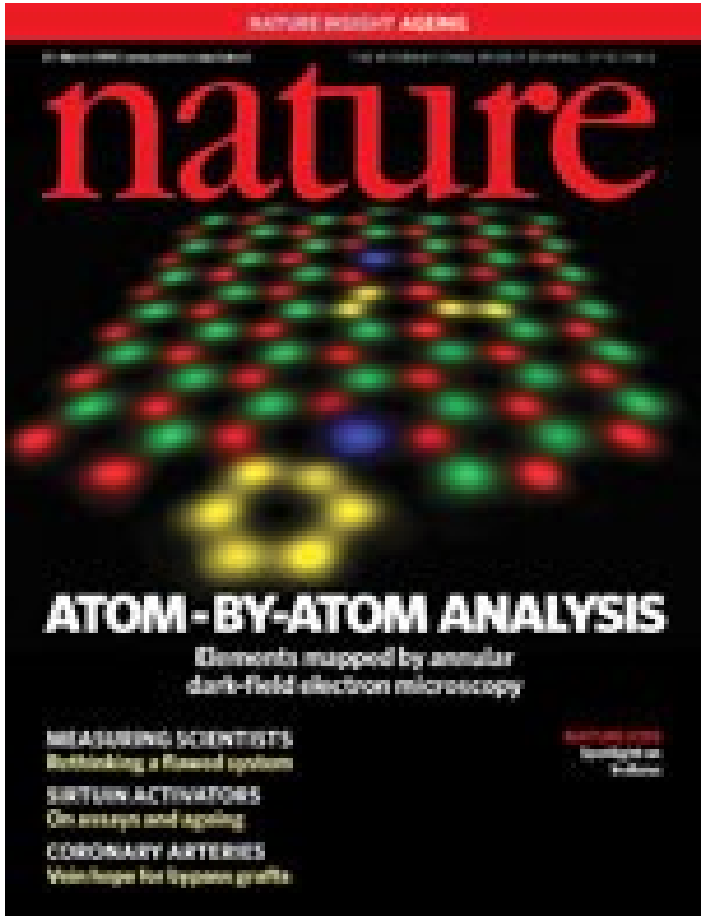

"There needs to be a greater focus on what these [interaction] data mean [...] This requires the input of social scientists, rather than just those more traditionally involved in data capture, such as computer scientists."

— Julia Lane, NSF, 24 March 2010

Different (But Overlapping) Roles

- **Social science:** specific models for specific applications, extensive post-analysis work
- **Computer science:** novel classes of models, mathematical and computational properties of models that extend across applications

This Talk

- Statistical topic models for text analysis
- “Off-the-shelf” topic models: priors, stop words
- Studying formerly-classified government documents

Statistical Modeling

- Modeling challenges:
  - Aggregating and representing large data sets
  - Handling data from sources with disparate emphases
  - Reasoning under uncertain information
  - Performing efficient inference
- Bayesian latent (hidden) variable models:
  - Powerful and flexible [Wallach et al. & Adams et al., AISTATS '10]
  - This talk: statistical topic models

Statistical Topic Modeling

- Three fundamental assumptions:
  - Documents have latent semantic structure (“topics”)
  - We can infer topics from word–document co-occurrences
  - Can simulate this inference algorithmically
- Given a data set, the goal is to
  - Learn the composition of the topics that best represent it
  - Learn which topics are used in each document

Why Topic Models?

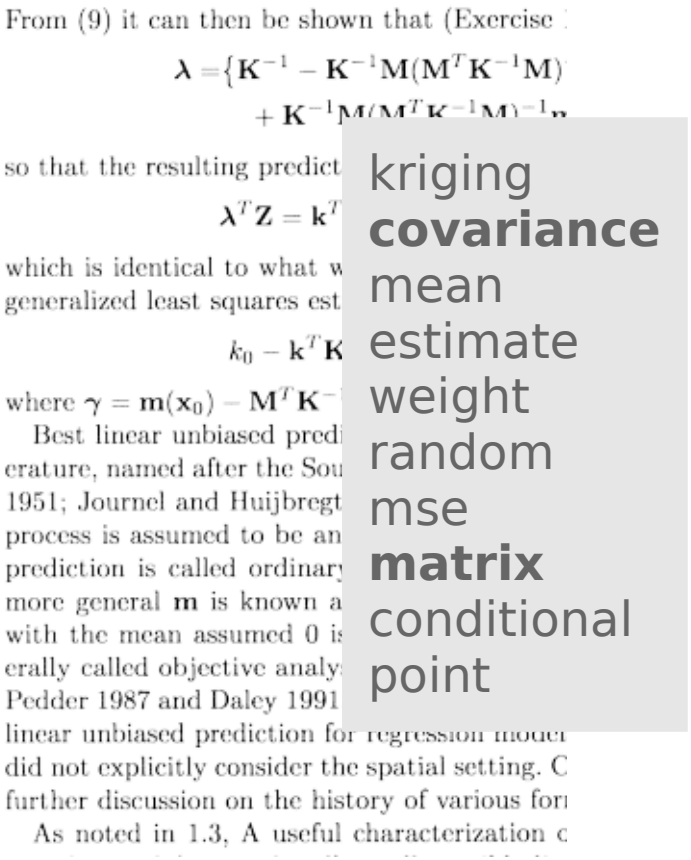

kriging
covariance
mean
estimate
weight
random
mse
matrix
conditional
point

vs.

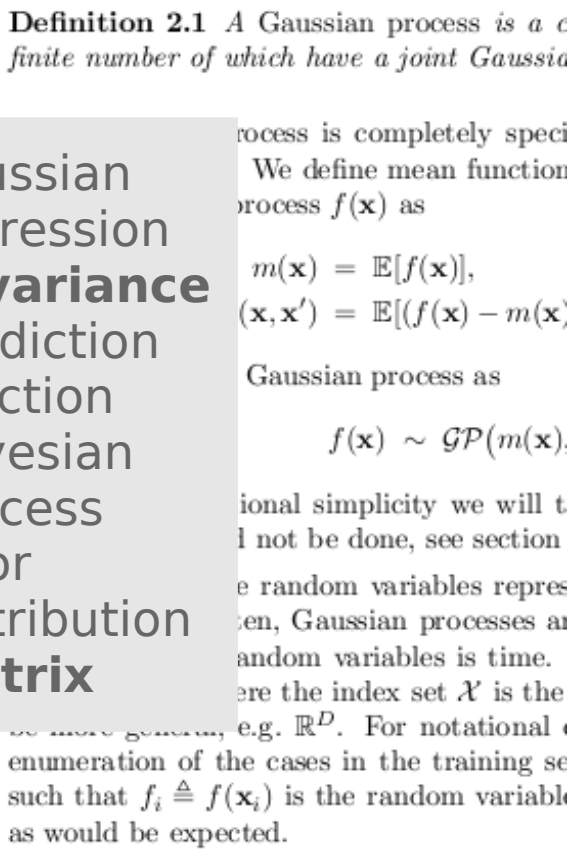

gaussian
regression
covariance
prediction
function
bayesian
process
prior
distribution
matrix

Topics and Words

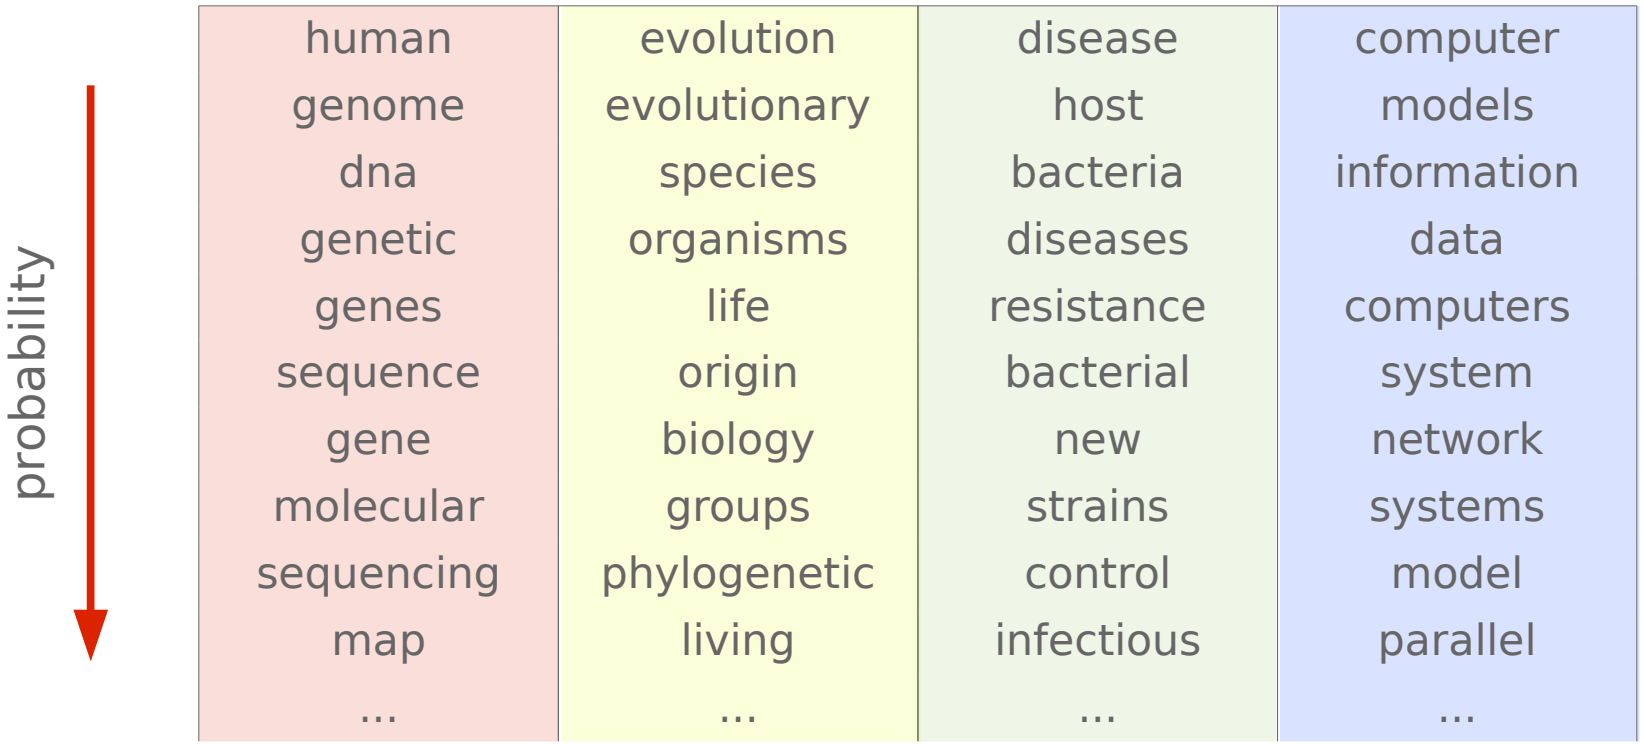

Documents and Topics

Seeking Life's Bare (Genetic) Necessities

COLD SPRING HARBOR, NEW YORK—How many genes does an organism need to survive? Last week at the genome meeting here,* two genome researchers with radically different approaches presented complementary views of the basic genes needed for life. One research team, using computer analyses to compare known genomes, concluded that today's organisms can be sustained with just 250 genes, and that the earliest life forms required a mere 128 genes. The other researcher mapped genes in a simple parasite and estimated that for this organism, 800 genes are plenty to do the job—but that anything short of 100 wouldn't be enough.

Although the numbers don't match precisely, those predictions "are not all that far apart," especially in comparison to the 75,000 genes in the human genome, notes Siv Andersson of Uppsala University in Sweden, who arrived at the 800 number. But coming up with a consensus answer may be more than just a genetic numbers game, particularly as more and more genomes are completely mapped and sequenced. "It may be a way of organizing any newly sequenced genome," explains Arcady Mushegian, a computational molecular biologist at the National Center for Biotechnology Information (NCBI) in Bethesda, Maryland. Comparing an

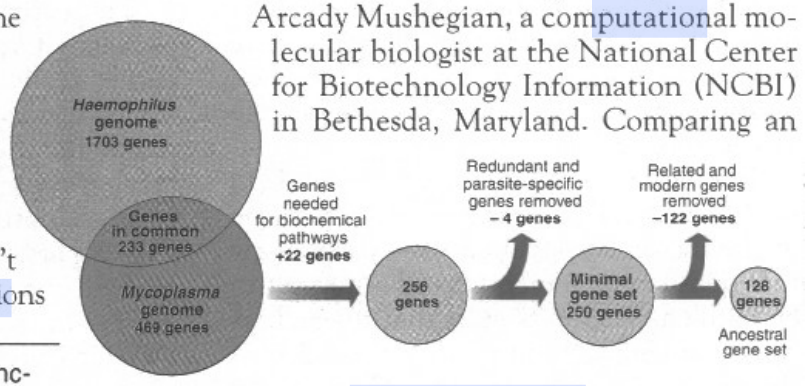

* Genome Mapping and Sequencing, Cold Spring Harbor, New York, May 8 to 12.

Stripping down. Computer analysis yields an estimate of the minimum modern and ancient genomes.

SCIENCE • VOL. 272 • 24 MAY 1996

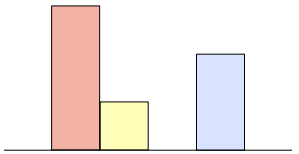

Mixtures vs. Admixtures

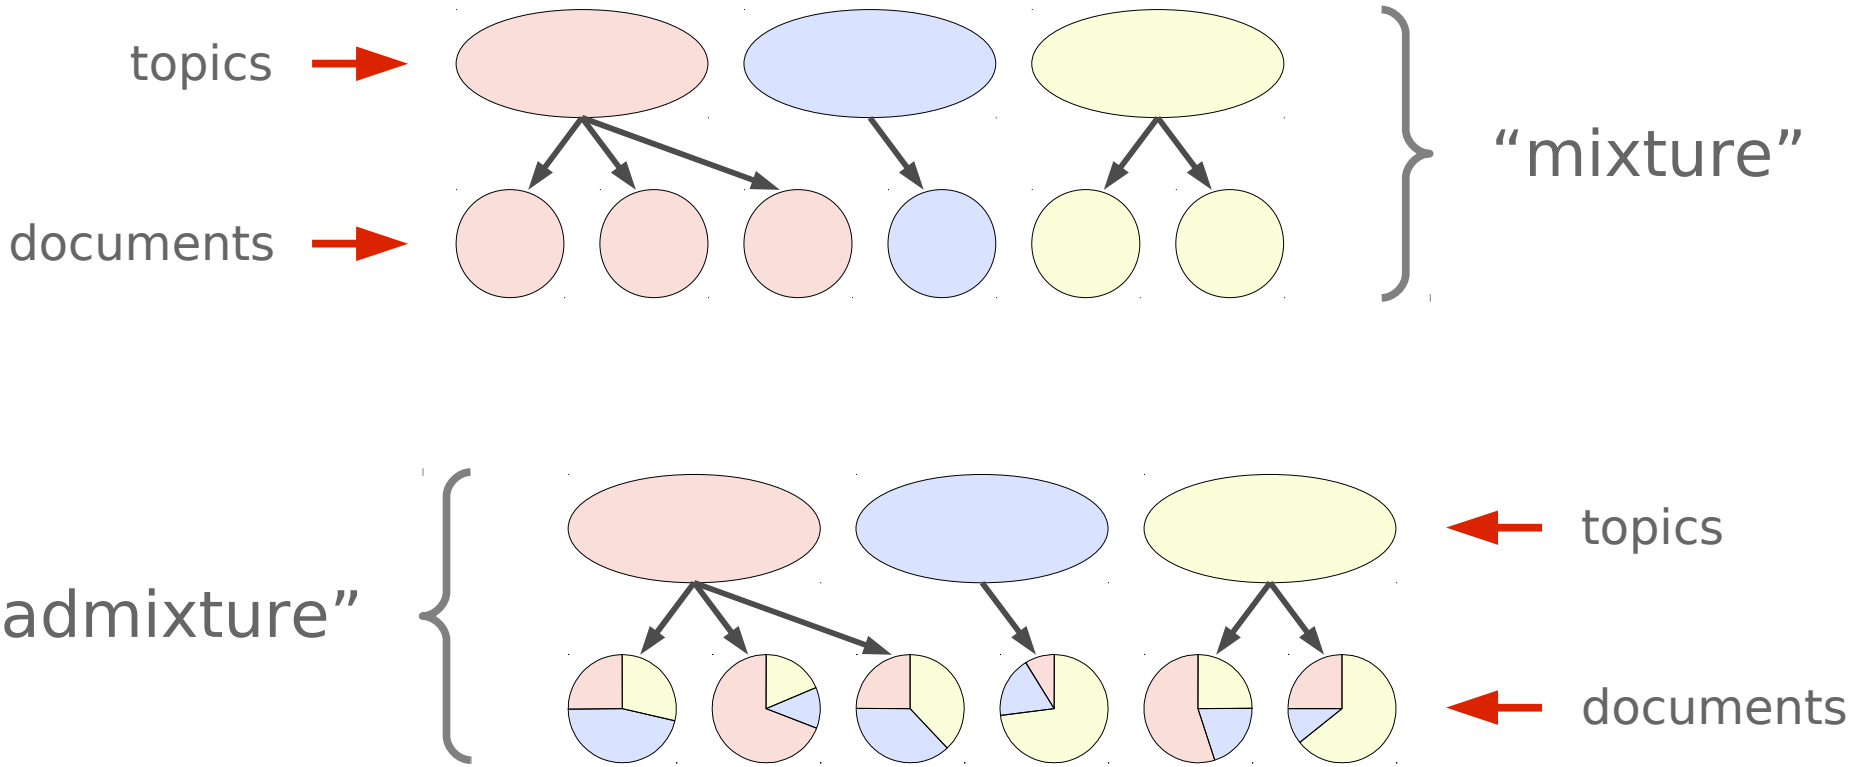

Generative Statistical Modeling

- Assume data was generated by a probabilistic model:
  - Model may have hidden structure (latent variables)
  - Model defines a joint distribution over all variables
  - Model parameters are unknown
- Infer hidden structure and model parameters from data
- Situate new data in estimated model

Generative Process

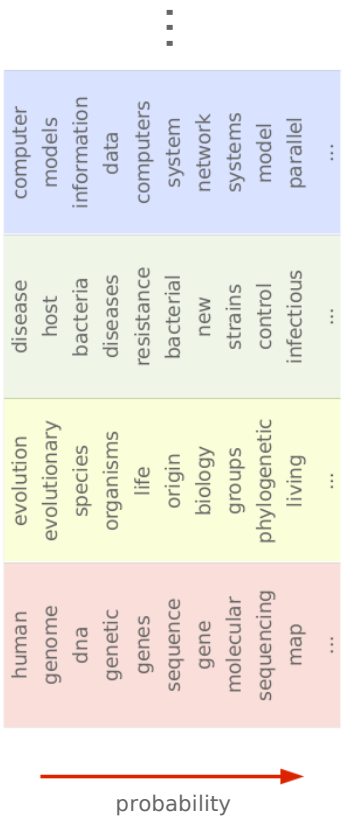

Choose a Distribution Over Topics

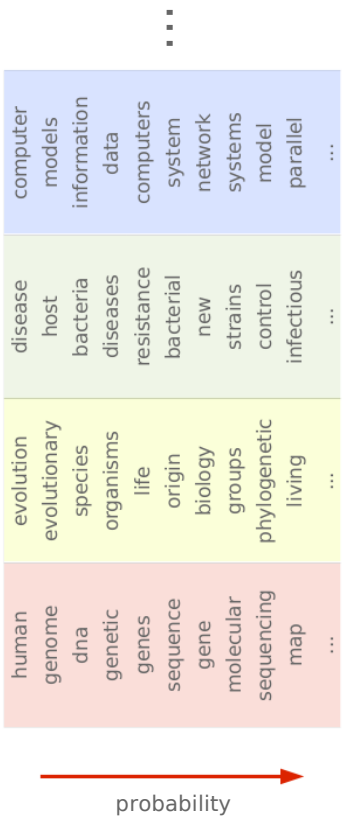

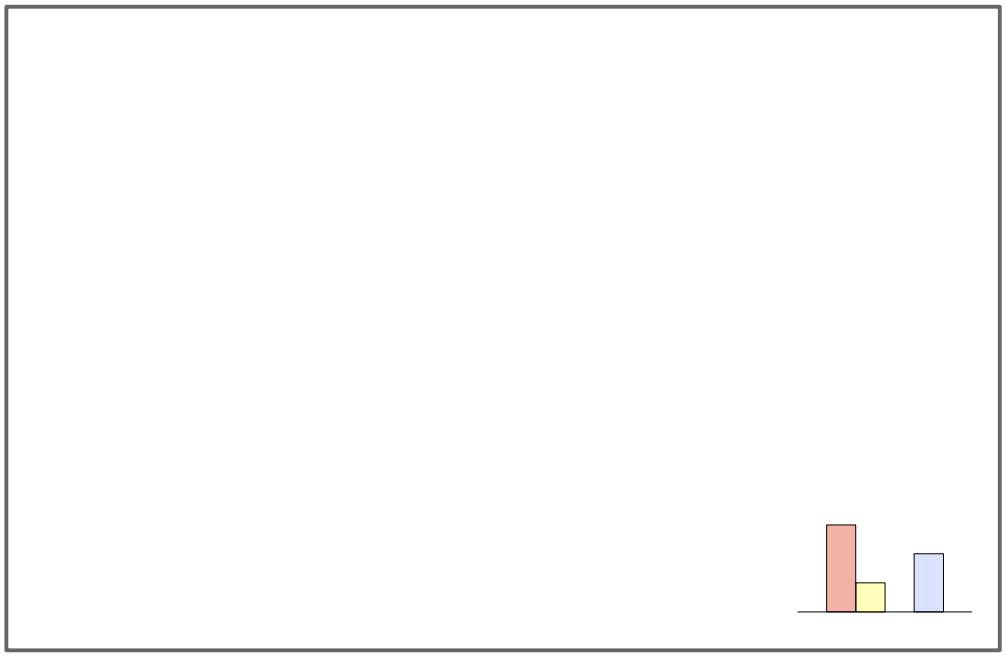

Choose a Topic

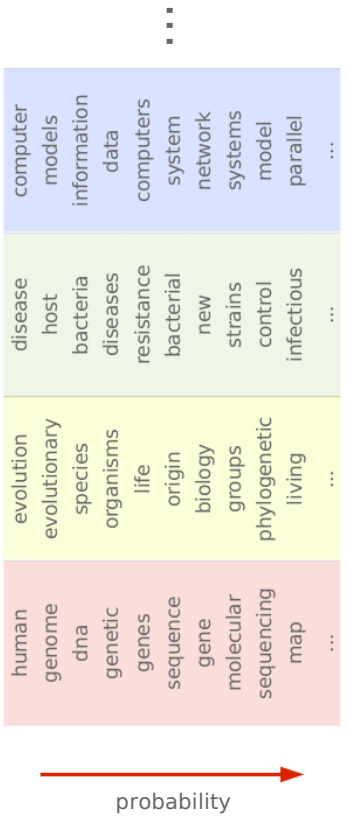

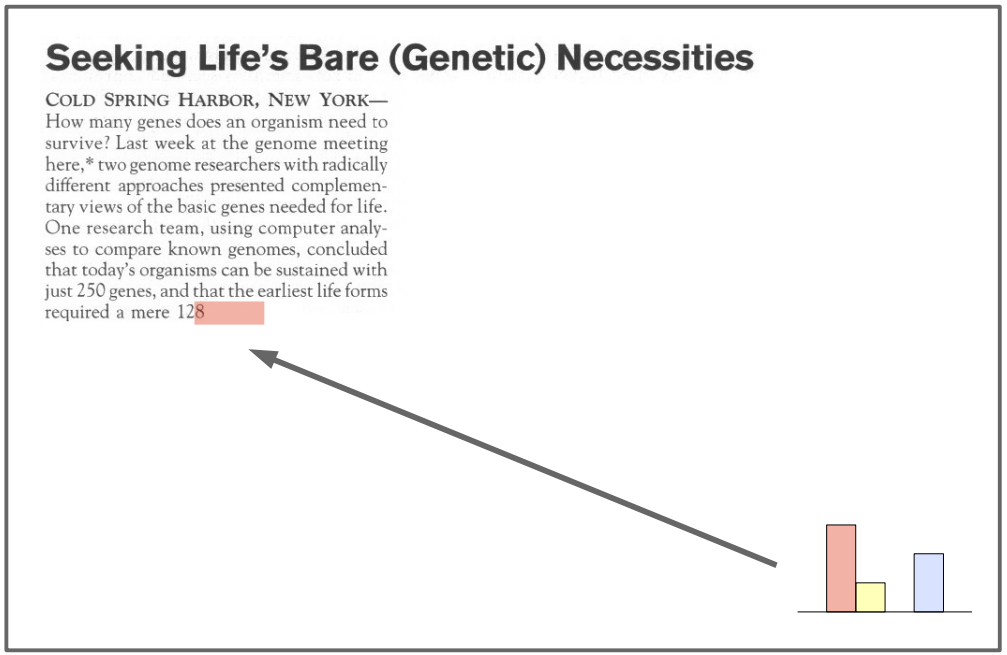

Seeking Life's Bare (Genetic) Necessities

COLD SPRING HARBOR, NEW YORK— How many genes does an organism need to survive? Last week at the genome meeting here,* two genome researchers with radically different approaches presented complementary views of the basic genes needed for life. One research team, using computer analyses to compare known genomes, concluded that today's organisms can be sustained with just 250 genes, and that the earliest life forms required a mere 128

Choose a Word

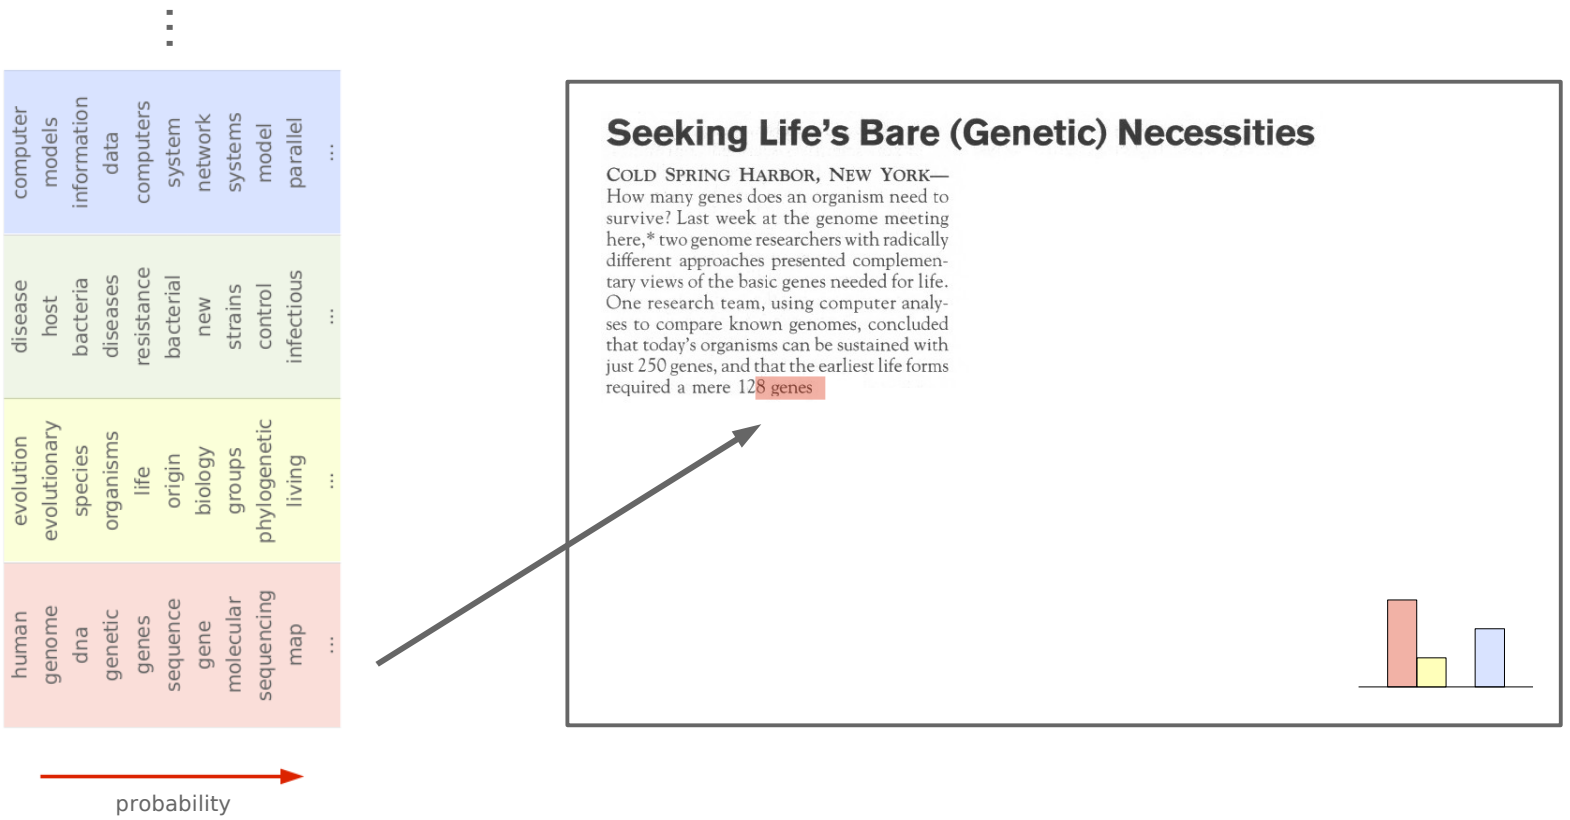

... And So On

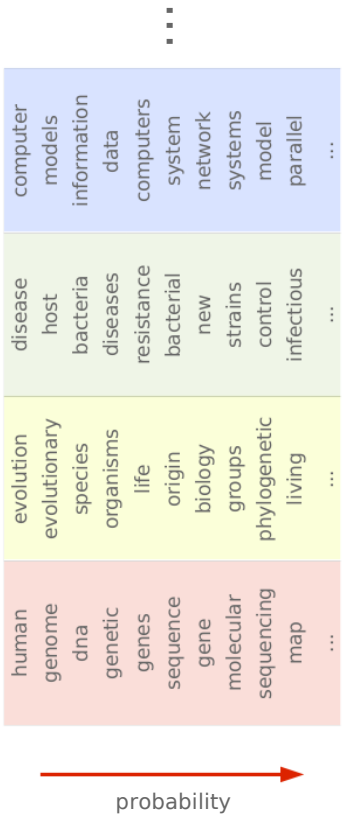

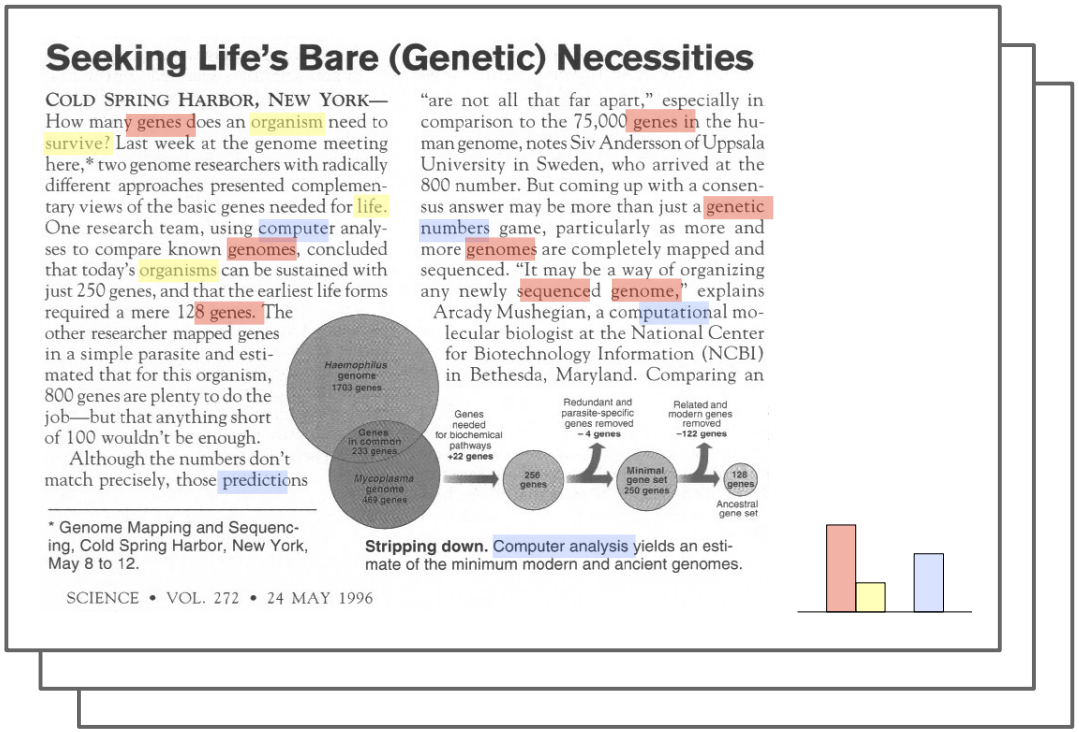

Seeking Life's Bare (Genetic) Necessities

COLD SPRING HARBOR, NEW YORK—How many genes does an organism need to survive? Last week at the genome meeting here,* two genome researchers with radically different approaches presented complementary views of the basic genes needed for life. One research team, using computer analyses to compare known genomes, concluded that today's organisms can be sustained with just 250 genes, and that the earliest life forms required a mere 128 genes. The other researcher mapped genes in a simple parasite and estimated that for this organism, 800 genes are plenty to do the job—but that anything short of 100 wouldn't be enough.

Although the numbers don't match precisely, those predictions "are not all that far apart," especially in comparison to the 75,000 genes in the human genome, notes Siv Andersson of Uppsala University in Sweden, who arrived at the 800 number. But coming up with a consensus answer may be more than just a genetic numbers game, particularly as more and more genomes are completely mapped and sequenced. "It may be a way of organizing any newly sequenced genome," explains Arcady Mushegian, a computational molecular biologist at the National Center for Biotechnology Information (NCBI) in Bethesda, Maryland. Comparing an

* Genome Mapping and Sequencing, Cold Spring Harbor, New York, May 8 to 12.

Stripping down. Computer analysis yields an estimate of the minimum modern and ancient genomes.

SCIENCE • VOL. 272 • 24 MAY 1996

Real Data: Statistical Inference

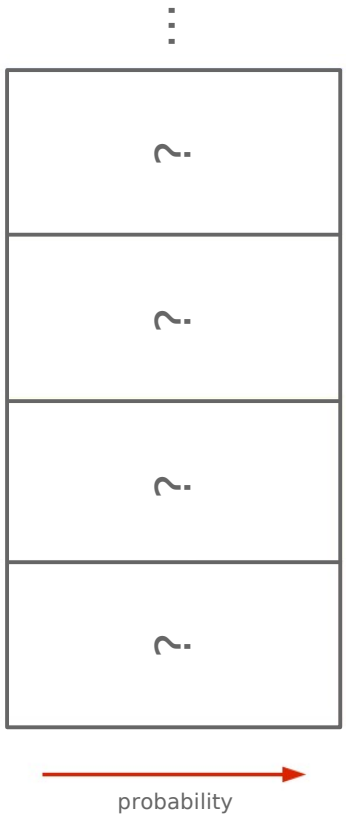

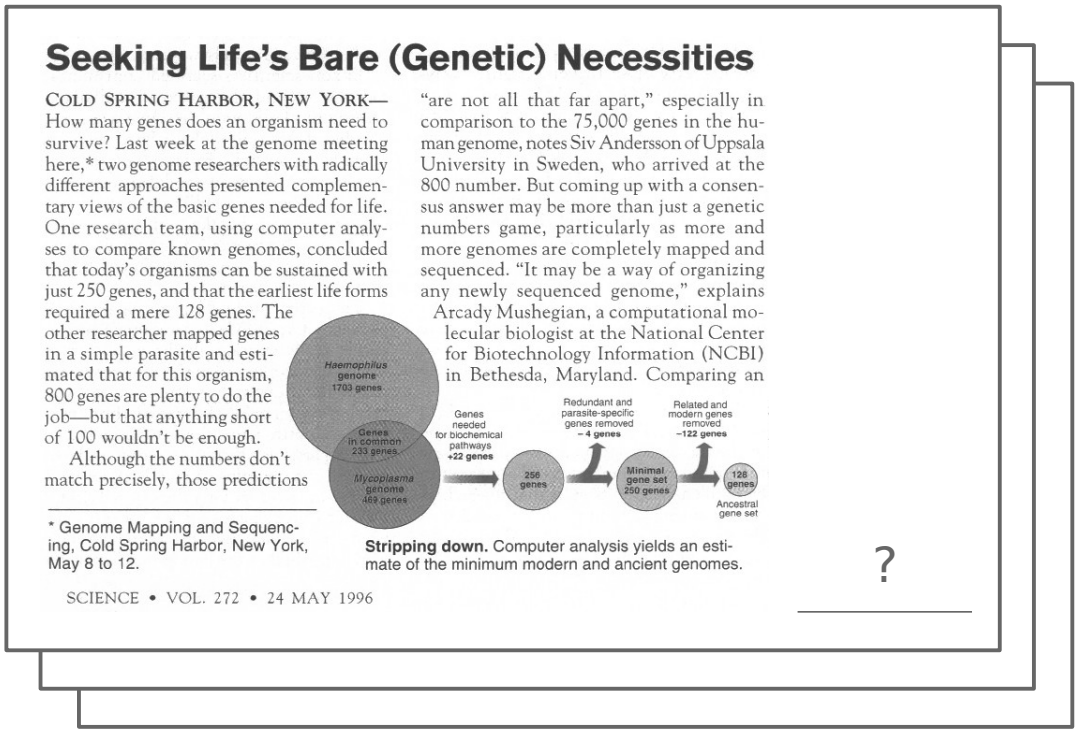

Seeking Life's Bare (Genetic) Necessities

COLD SPRING HARBOR, NEW YORK—How many genes does an organism need to survive? Last week at the genome meeting here,* two genome researchers with radically different approaches presented complementary views of the basic genes needed for life. One research team, using computer analyses to compare known genomes, concluded that today's organisms can be sustained with just 250 genes, and that the earliest life forms required a mere 128 genes. The other researcher mapped genes in a simple parasite and estimated that for this organism, 800 genes are plenty to do the job—but that anything short of 100 wouldn't be enough.

Although the numbers don't match precisely, those predictions "are not all that far apart," especially in comparison to the 75,000 genes in the human genome, notes Siv Andersson of Uppsala University in Sweden, who arrived at the 800 number. But coming up with a consensus answer may be more than just a genetic numbers game, particularly as more and more genomes are completely mapped and sequenced. "It may be a way of organizing any newly sequenced genome," explains Arcady Mushegian, a computational molecular biologist at the National Center for Biotechnology Information (NCBI) in Bethesda, Maryland. Comparing an

Stripping down. Computer analysis yields an estimate of the minimum modern and ancient genomes.

* Genome Mapping and Sequencing, Cold Spring Harbor, New York, May 8 to 12.

SCIENCE • VOL. 272 • 24 MAY 1996

?

The End Result...

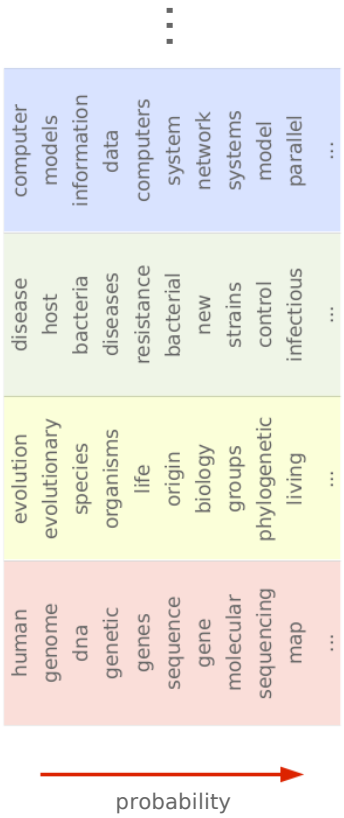

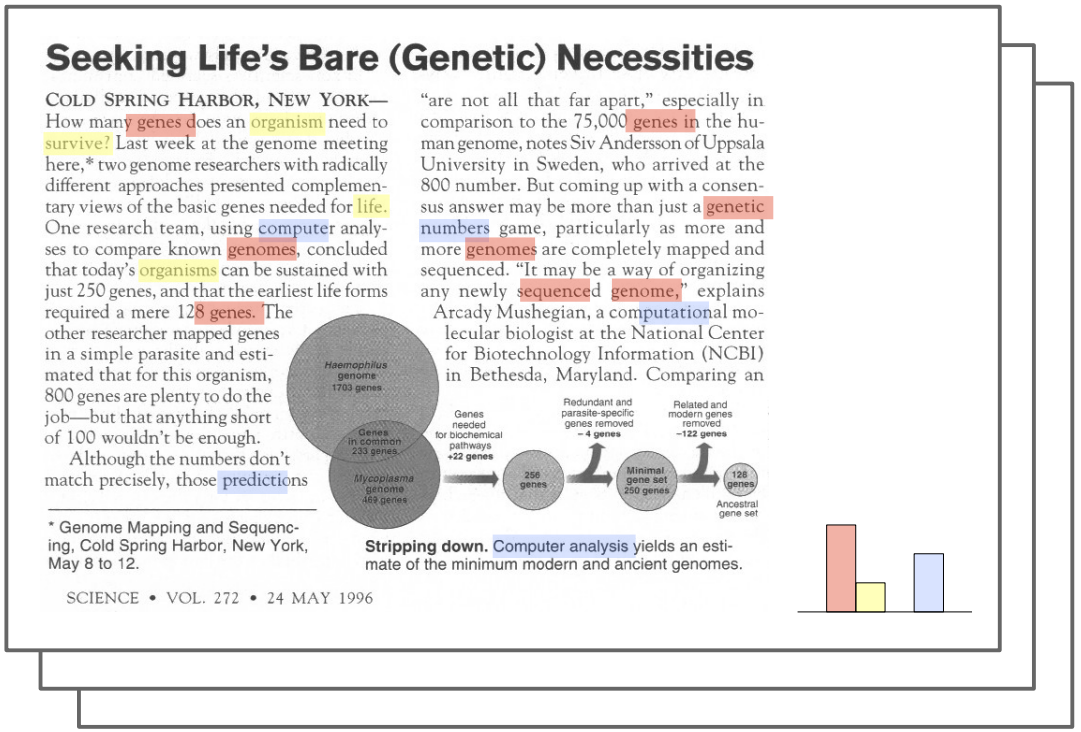

Seeking Life's Bare (Genetic) Necessities

COLD SPRING HARBOR, NEW YORK— How many genes does an organism need to survive? Last week at the genome meeting here,* two genome researchers with radically different approaches presented complementary views of the basic genes needed for life. One research team, using computer analyses to compare known genomes, concluded that today's organisms can be sustained with just 250 genes, and that the earliest life forms required a mere 128 genes. The other researcher mapped genes in a simple parasite and estimated that for this organism, 800 genes are plenty to do the job—but that anything short of 100 wouldn't be enough.

Although the numbers don't match precisely, those predictions "are not all that far apart," especially in comparison to the 75,000 genes in the human genome, notes Siv Andersson of Uppsala University in Sweden, who arrived at the 800 number. But coming up with a consensus answer may be more than just a genetic numbers game, particularly as more and more genomes are completely mapped and sequenced. "It may be a way of organizing any newly sequenced genome," explains Arcady Mushegian, a computational molecular biologist at the National Center for Biotechnology Information (NCBI) in Bethesda, Maryland. Comparing an

Stripping down. Computer analysis yields an estimate of the minimum modern and ancient genomes.

* Genome Mapping and Sequencing, Cold Spring Harbor, New York, May 8 to 12.

SCIENCE • VOL. 272 • 24 MAY 1996

This Talk

- Statistical topic models for text analysis
- “Off-the-shelf” topic models: priors, stop words
- Studying formerly-classified government documents

The State of The Art

- Topic models are extremely appealing
- ... but they're not always usable by non-experts
- Need to bridge this gap between producers and consumers of topic modeling technology:
  - Address problems/challenges faced by practitioners
  - Question unquestioned assumptions
  - Explore the interplay between theory and practice

“Off-the-Shelf” Topic Modeling

I want to model technology emergence by analyzing patent abstracts...

I have a statistical model that you can use...

“Off-the-Shelf” Topic Modeling

I want to model technology emergence by analyzing patent abstracts...

I have a statistical model that you can use...

a	a	the	the
field	the	of	**invention**
emission	**carbon**	a	of
an	and	to	to
electron	**gas**	and	**present**
...	...	...	...

“Off-the-Shelf” Topic Modeling?

Help! All my topics consist of “the, and of, to, a ...”

Now they all consist of “invention, present, thereof ...”

Wait, but how do I choose the right number of topics?

Preprocess your data to remove stop words...

Make a domain-specific list of stop words...

Evaluate the probability of unseen data for different numbers...

Directed Graphical Models

$$P(y, x_1, \ldots, x_N) = P(y) \prod_{n=1}^{N} P(x_n \mid y)$$

- Nodes: random variables (latent or observed)
- Edges: probabilistic dependencies between variables
- Plates: “macros” that allow subgraphs to be replicated

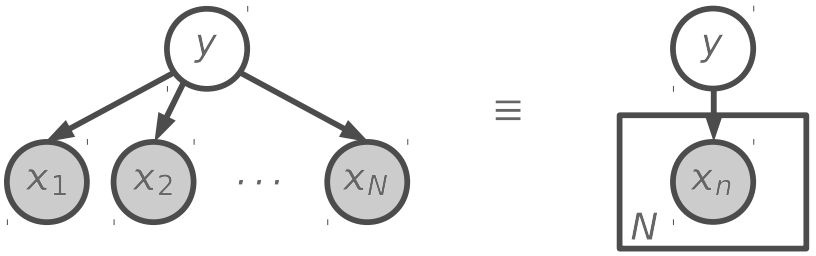

Statistical Topic Modeling

[Hofmann, '99]

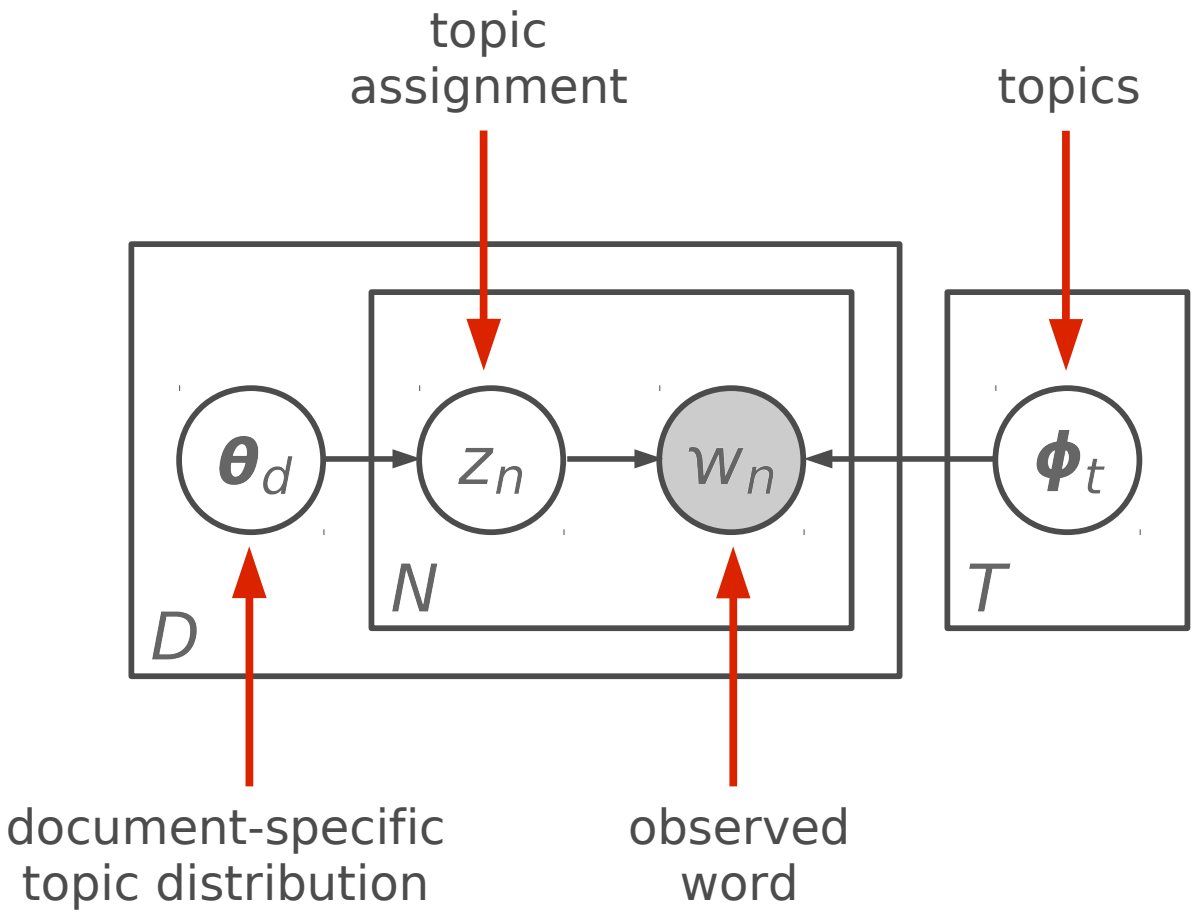

Latent Dirichlet Allocation (LDA)

[Blei, Ng & Jordan, '03]

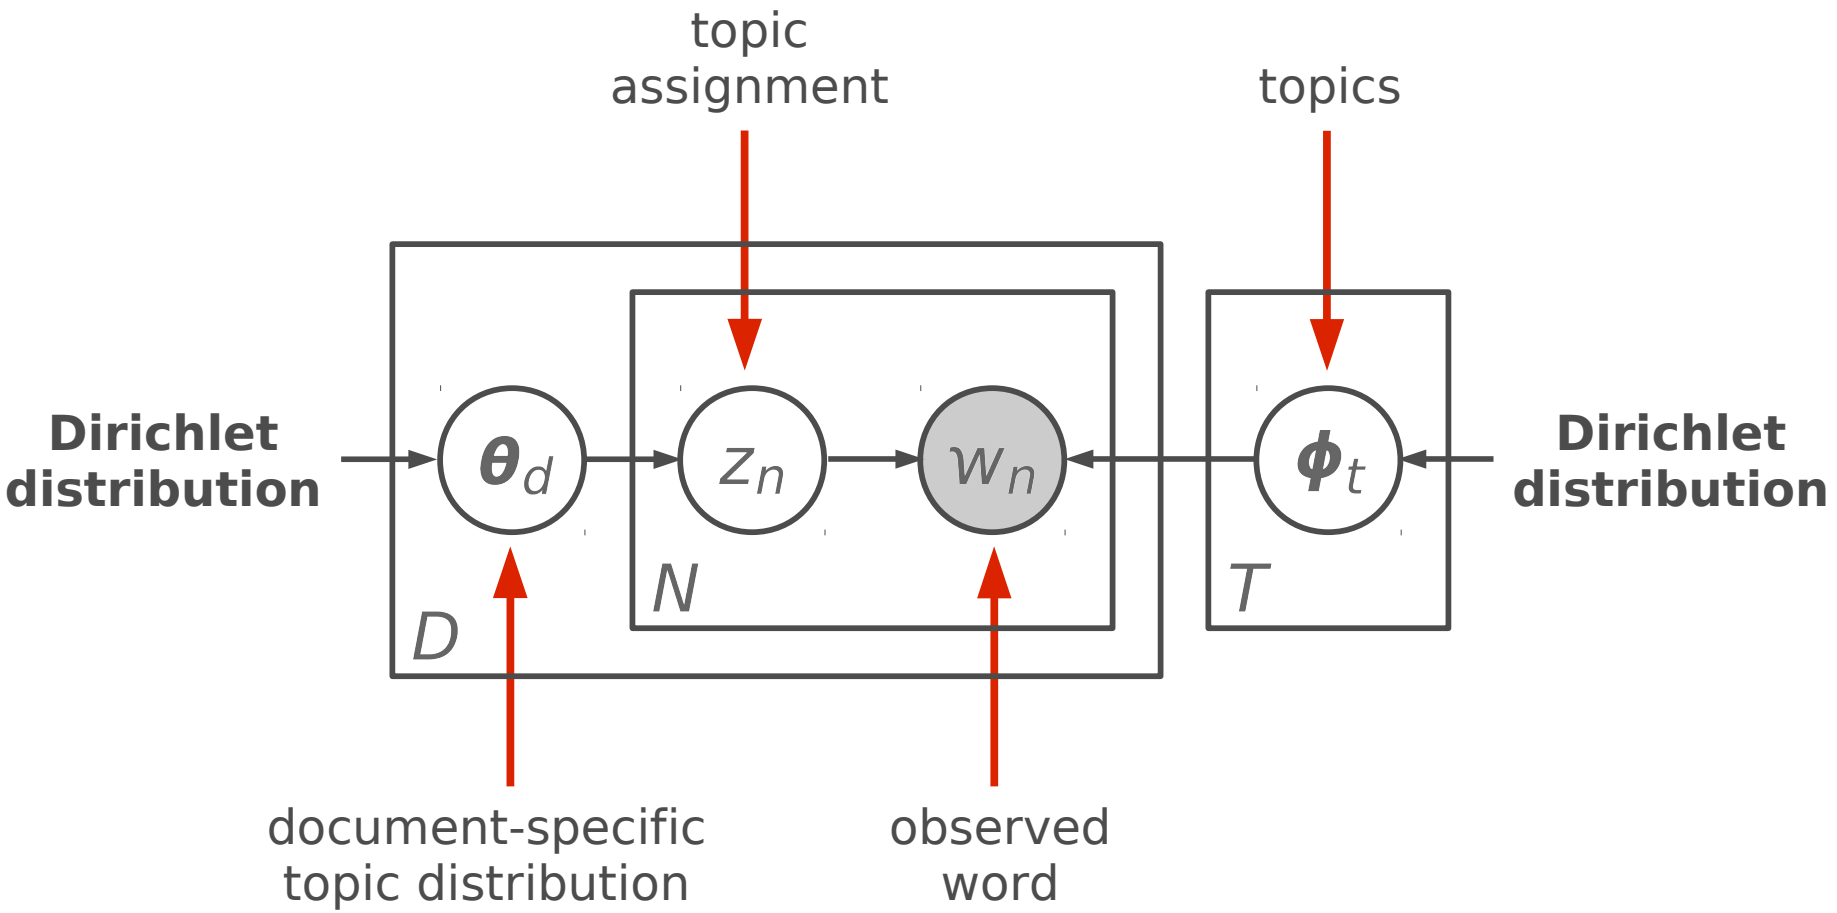

Discrete Probability Distributions

- 3-dimensional discrete probability distributions can be visually represented in 2-dimensional space:

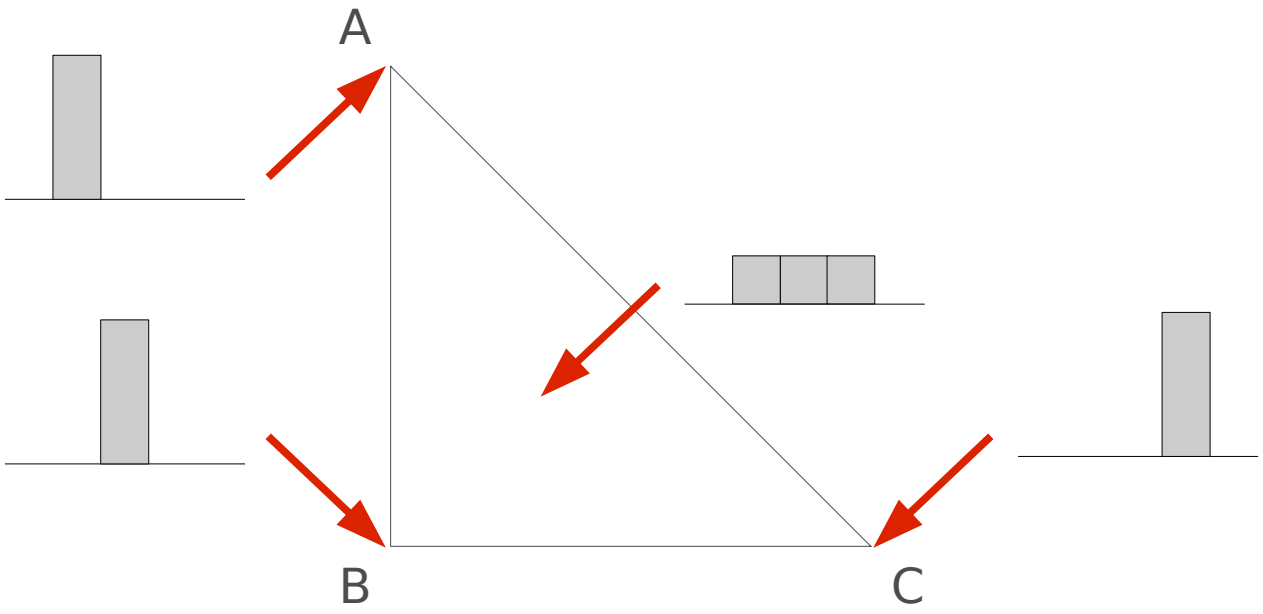

Dirichlet Distribution

- Distribution over discrete probability distributions:

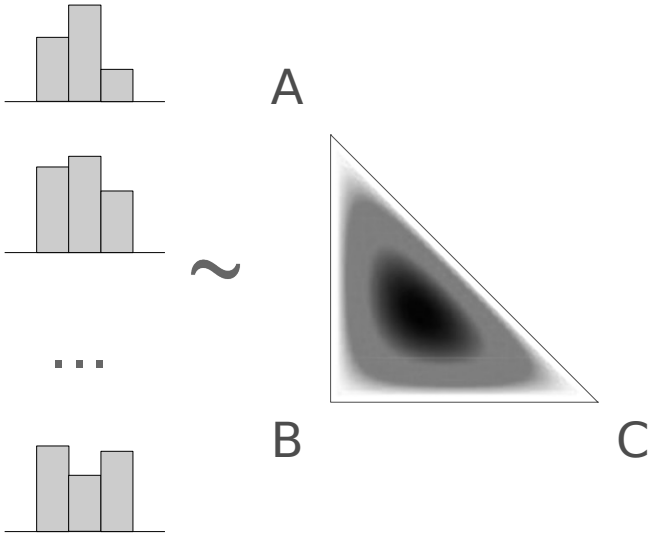

$\sim$ $\equiv$ $\boldsymbol{p} \sim \mathrm{Dir}(\alpha \boldsymbol{m})$

base measure (mean)

concentration parameter

Dirichlet Parameters

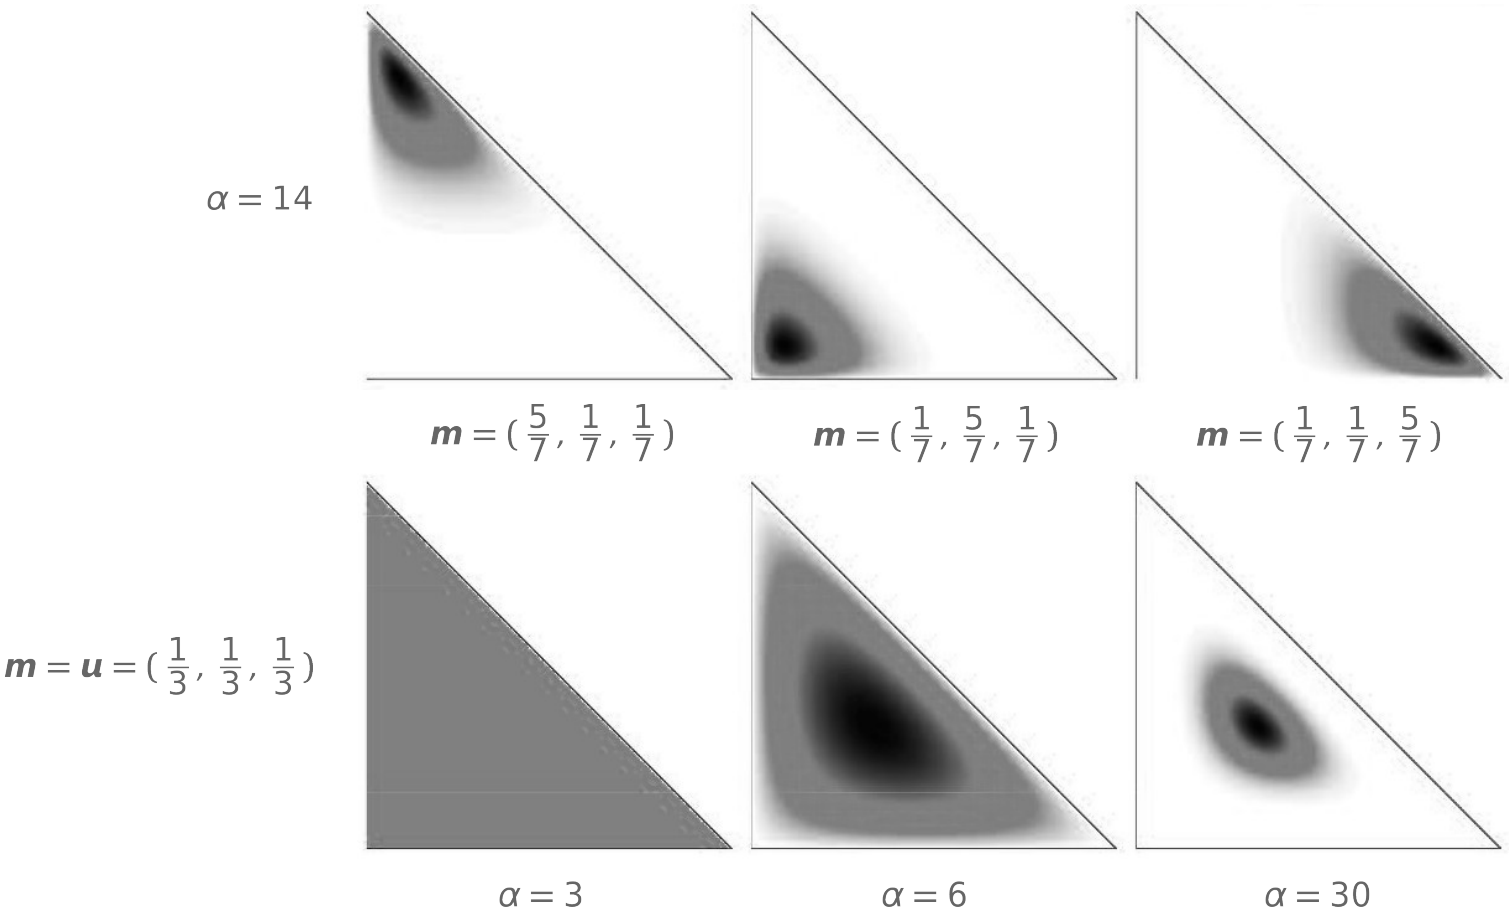

Dirichlet Priors for LDA

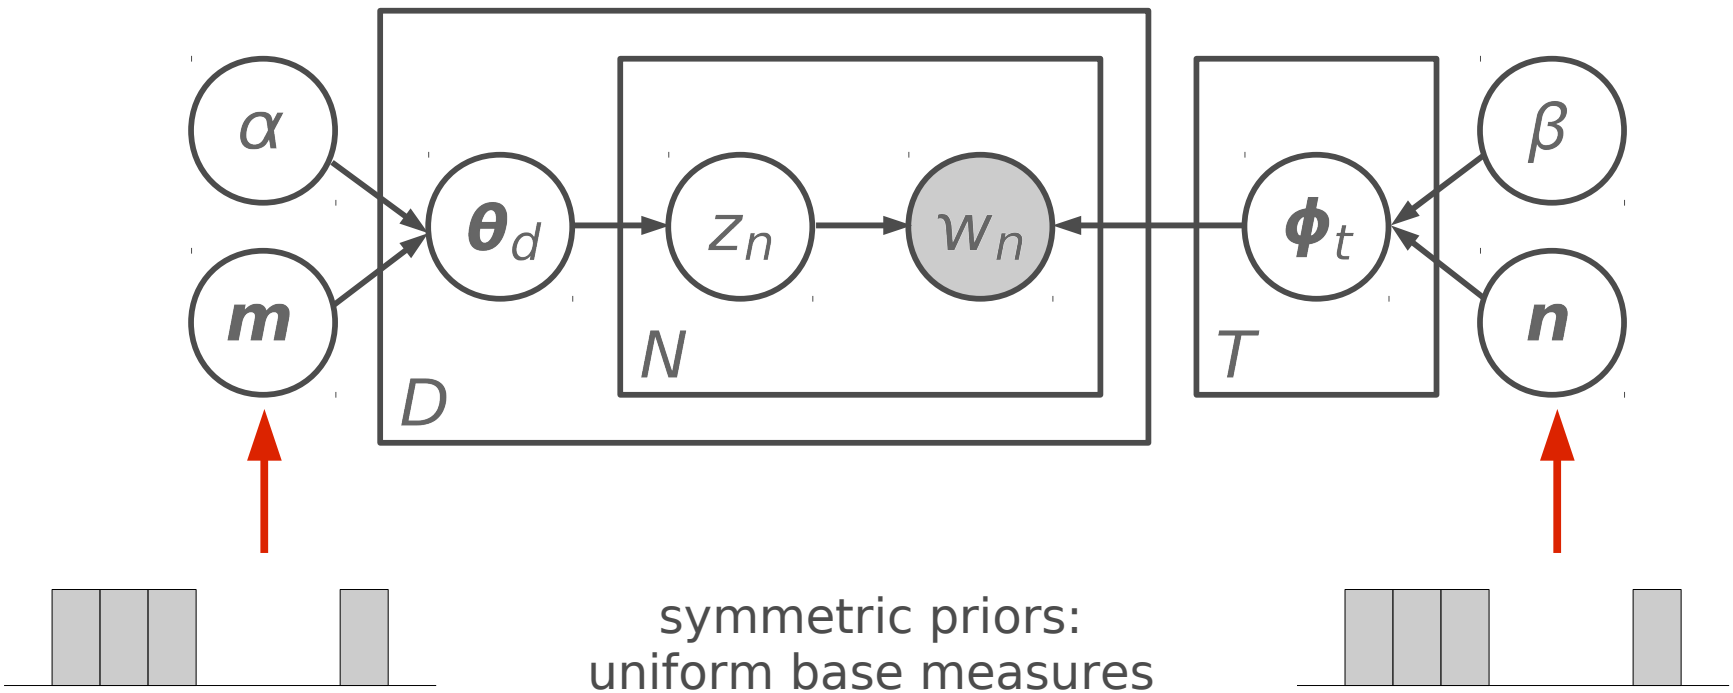

Dirichlet Priors for LDA

- Two scalar concentration parameters: α and β
- Concentration parameters are usually set heuristically
  - e.g., $\alpha = 50$ and $\beta = 0.01W$
- Some recent work on learning optimal values for the concentration parameters from data
- No rigorous study of the Dirichlet priors:
  - e.g., asymmetric vs. symmetric base measures
  - Effects of the base measures on the inferred topics

Symmetric → Asymmetric

[Wallach et al., '09]

- Use prior over $\Theta = \{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_D\}$ as a running example
- Uniform base measure → nonuniform base measure

$\Theta \sim \mathrm{Dir}(\alpha \boldsymbol{m})$ $\Theta \sim \mathrm{Dir}(\alpha \boldsymbol{m})$

- Asymmetric prior: some topics more likely a priori

Hierarchical Asymmetric Dirichlet

- Which topics should be more probable a priori?
  - Draw ***m*** from a Dirichlet distribution:

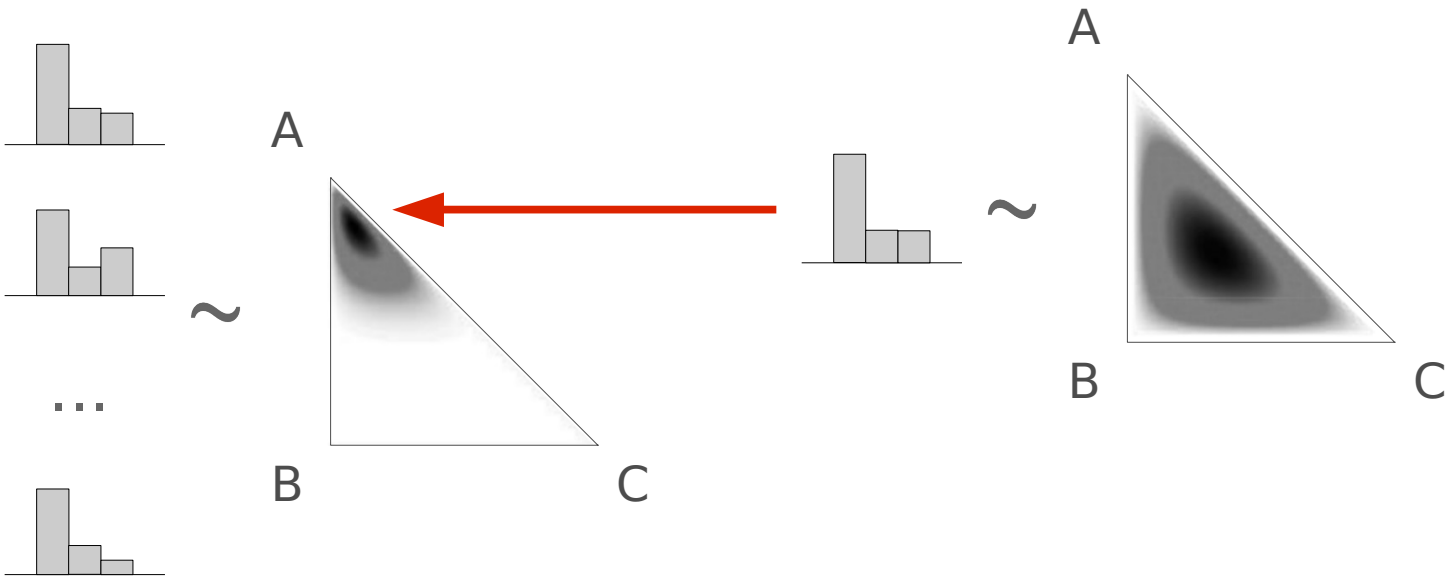

Putting Everything Together

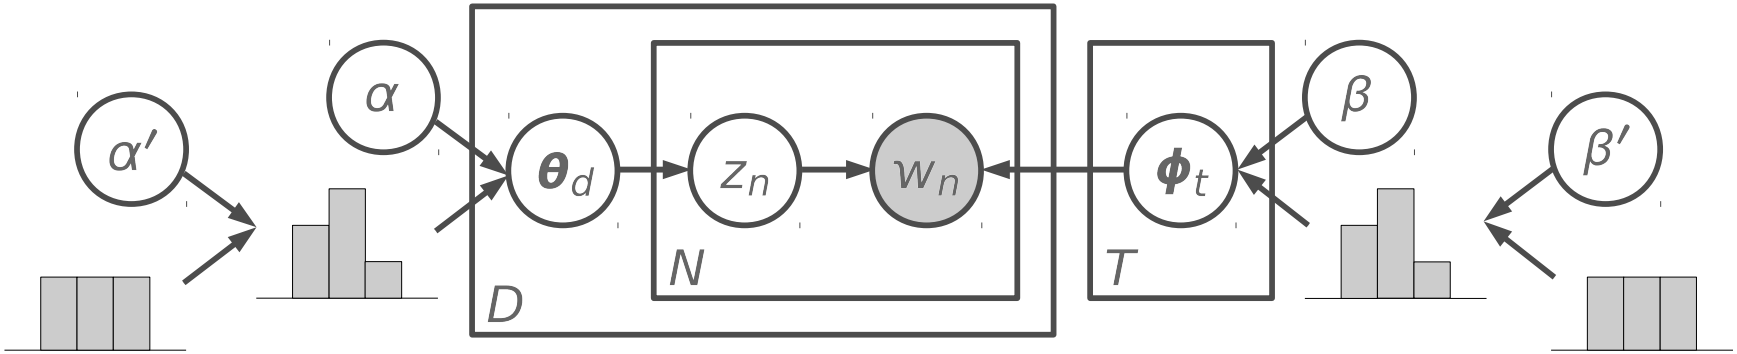

- Asymmetric hierarchical Dirichlet priors
- Integrate out Θ, Φ and base measures
- Learn $\boldsymbol{z}$ and concentration parameters from data

Data Sets

- Carbon nanotechnology patents:
  - Ultimate goal: track innovation and emergence
  - Fullerene and carbon nanotube patents
  - 1,016 abstracts (~100 words each)
  - 103,499 total words; 6,068 unique words
- 20 Newsgroups data (80,012 total words)
- New York Times articles (477,465 total words)

Inferred Topics

before →

a	a	the	the
field	the	of	**invention**
emission	**carbon**	a	of
an	and	to	to
electron	**gas**	and	**present**
...	...	...	...

after →

the	**carbon**	**metal**	**composite**
a	**nanotubes**	**catalytic**	**polymer**
of	**nanotube**	**transition**	**matrix**
to	**catalyst**	**catalyst**	**weight**
and	**substrate**	from	**fiber**
...	...	...	...

Sampled Concentration Parameters

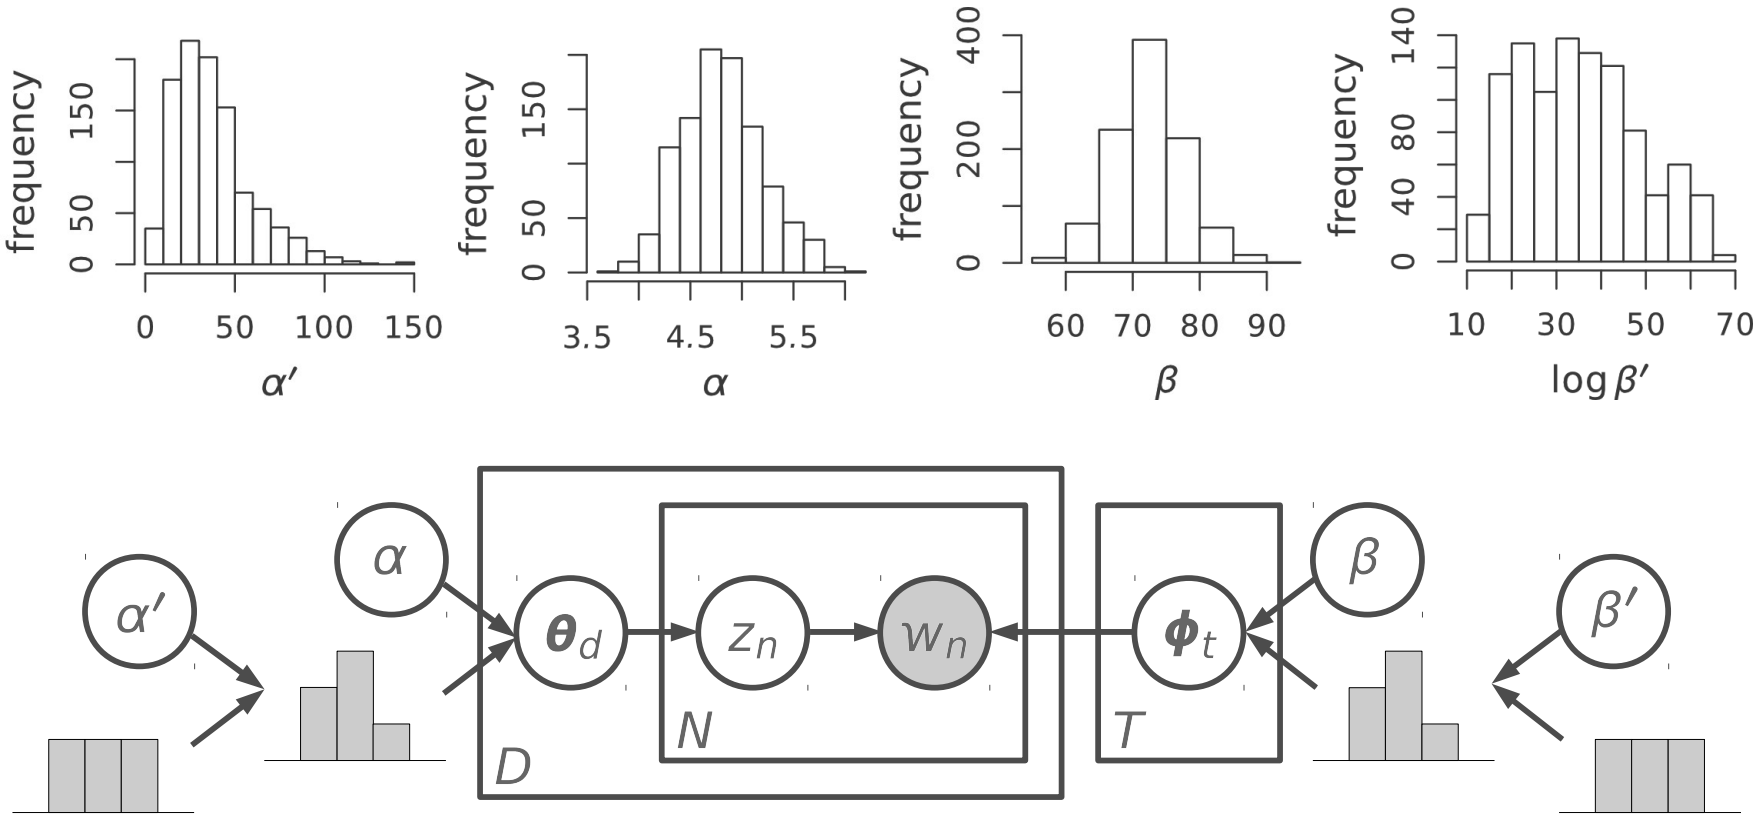

A Theoretical Observation...

- Symmetric Dirichlet is a special case of the hierarchical asymmetric Dirichlet (large concentration parameter)

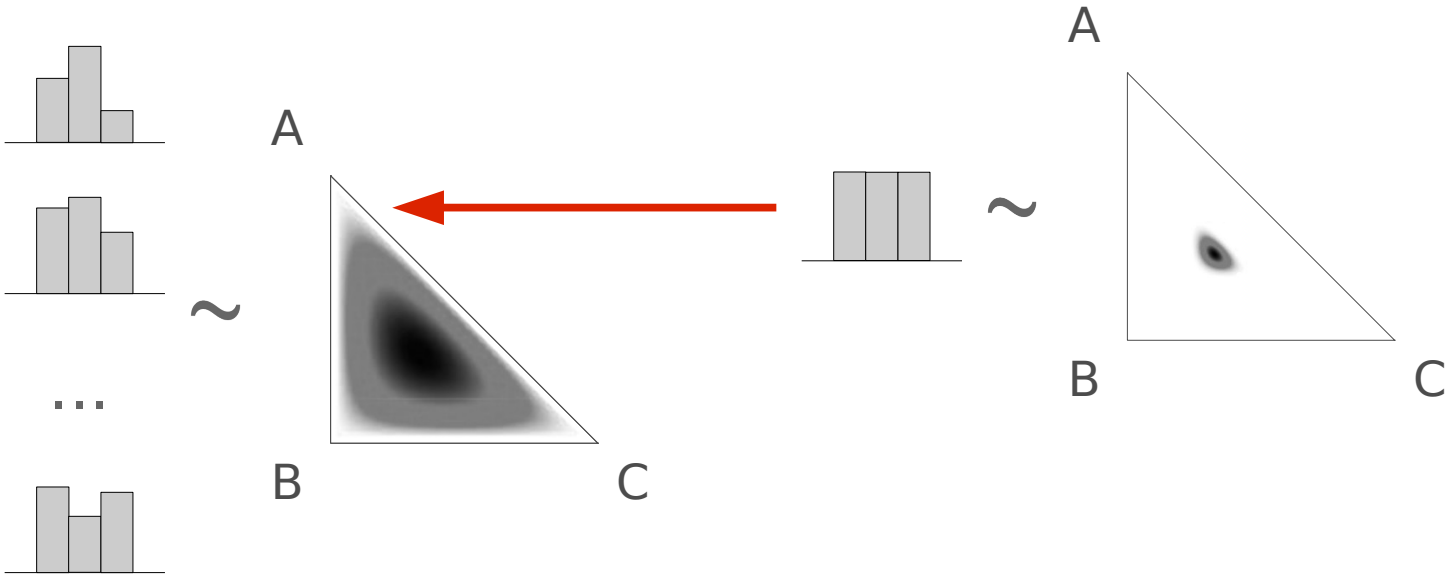

Sampled Concentration Parameters

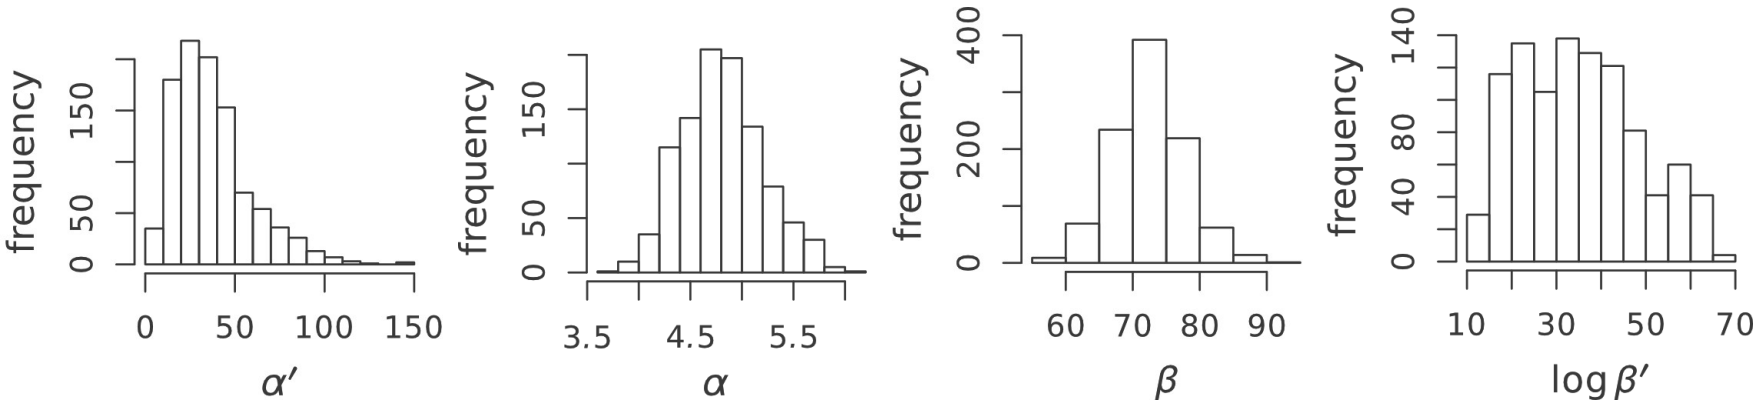

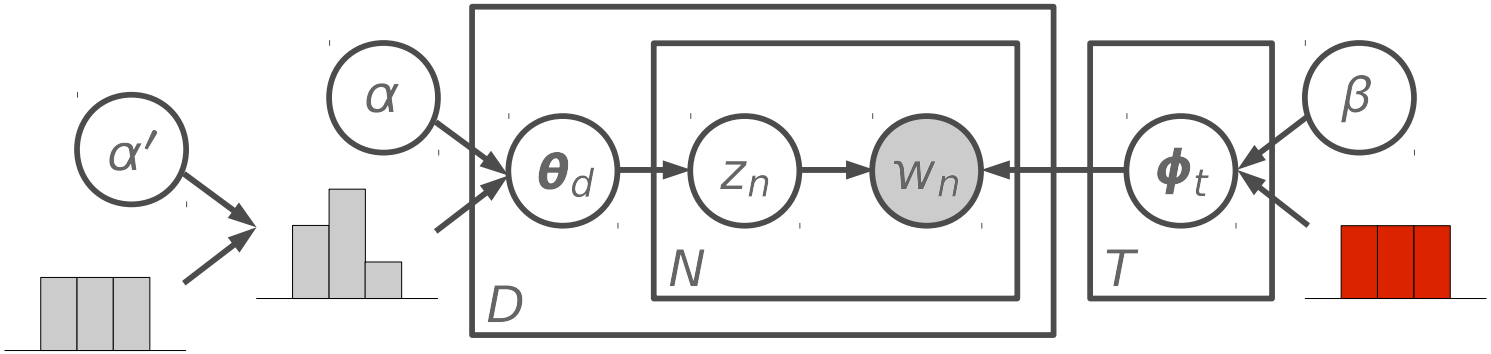

Intuition

- Topics should be distinct from each other:
  - Asymmetric prior over topics makes topics more similar to each other (and to corpus-wide word frequencies)
  - Want a symmetric prior to preserve topic “distinctness”
- Still have to account for power-law word usage:
  - Asymmetric prior over document-specific topic distributions means some topics (e.g., “the, a, of, to ...”) can be used more often than others in all documents

“Off-the-Shelf” Topic Modeling

I can model technology emergence by analyzing patent abstracts!

Great! Let me know if you need any more help!

the	**carbon**	**metal**	**composite**
a	**nanotubes**	**catalytic**	**polymer**
of	**nanotube**	**transition**	**matrix**
to	**catalyst**	**catalyst**	**weight**
and	**substrate**	from	**fiber**
...	...	...	...

Polylingual Topics

CY	sadwrn blaned gallair at lloeren mytholeg
DE	space nasa sojus flug mission
EL	διαστημικό sts nasa αγγλ small
EN	**space mission launch satellite nasa spacecraft**
FA	فضایی ماموریت ناسا مدار فضانورد ماهواره
FI	sojuz nasa apollo ensimmäinen space lento
FR	spatiale mission orbite mars satellite spatial
HE	החלל הארץ חלל כדור א תוכנית
IT	spaziale missione programma space sojuz stazione
PL	misja kosmicznej stacji misji space nasa
RU	космический союз космического спутник станции
TR	uzay soyuz ay uzaya salyut sovyetler

Polylingual Topics

CY	bardd gerddi iaith beirdd fardd gymraeg
DE	dichter schriftsteller literatur gedichte gedicht werk
EL	ποιητής ποίηση ποιητή έργο ποιητές ποιήματα
EN	**poet poetry literature literary poems poem**
FA	شاعر شعر ادبیات فارسی ادبی آثار
FI	runoilija kirjailija kirjallisuuden kirjoitti runo julkaisi
FR	poète écrivain littérature poésie littéraire ses
HE	משורר ספרות שירה סופר שירים המשורר
IT	poeta letteratura poesia opere versi poema
PL	poeta literatury poezji pisarz in jego
RU	поэт его писатель литературы поэзии драматург
TR	şair edebiyat şiir yazar edebiyatı adlı

Polylingual Topic Model

[Mimno et al., '09]

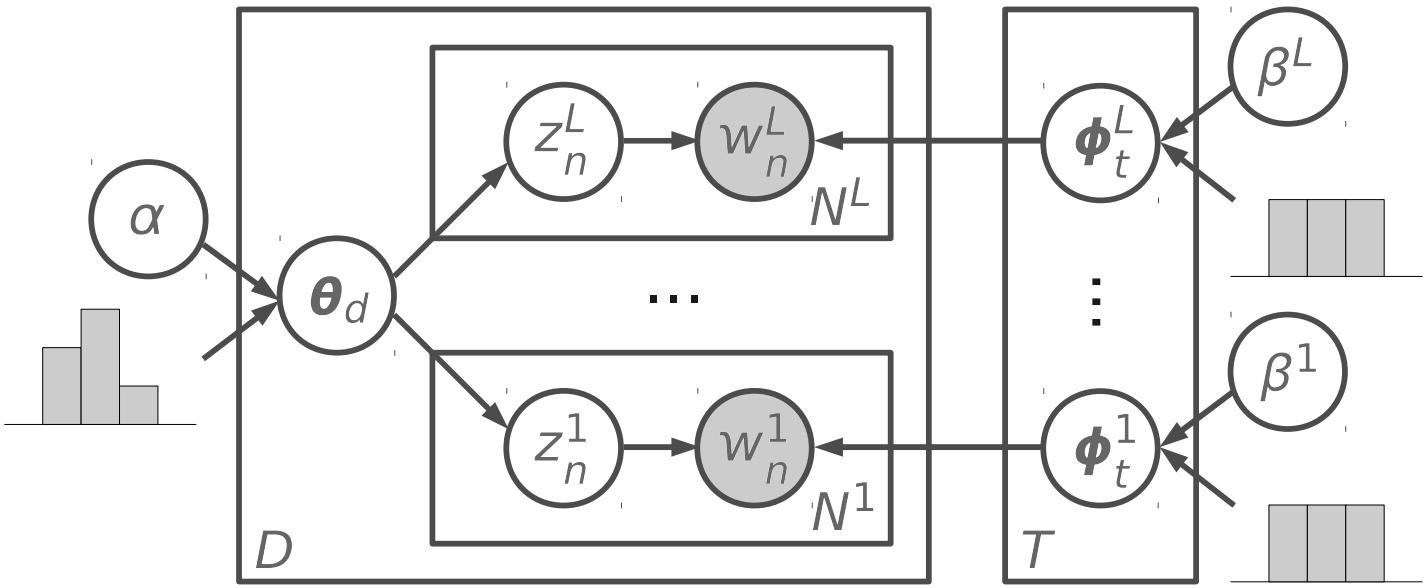

"tuple" of
aligned documents

language-specific
Dirichlet parameters

This Talk

- Statistical topic models for text analysis
- “Off-the-shelf” topic models: priors, stop words
- Studying formerly-classified government documents

In 2009 Alone...

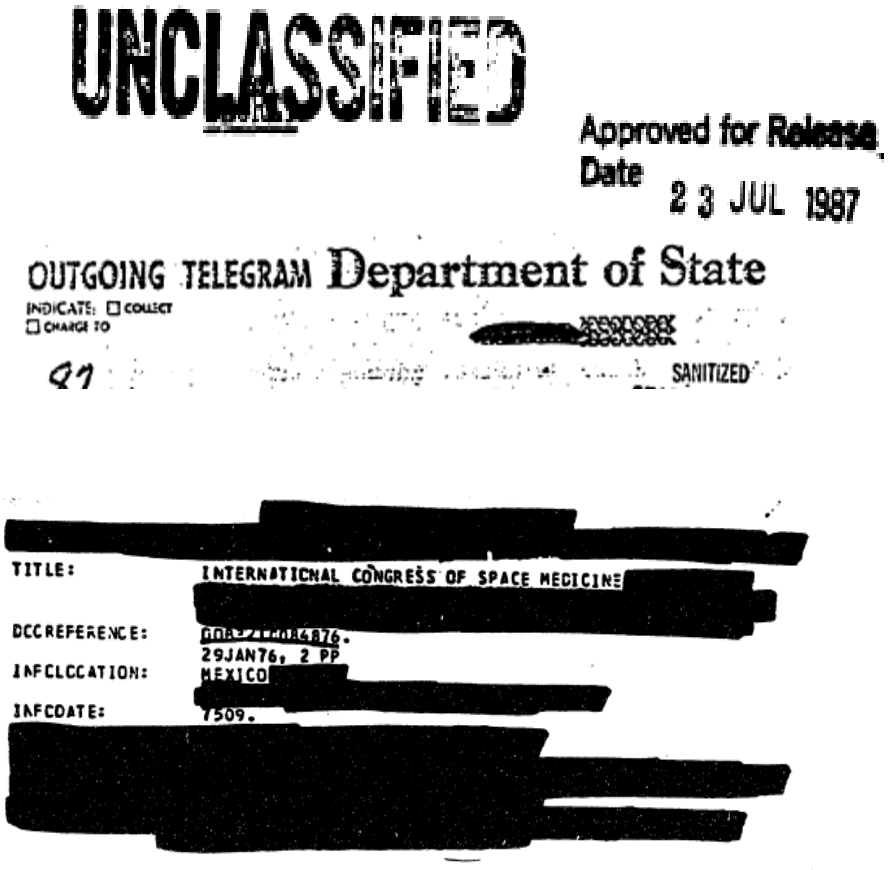
UNCLASSIFIED

Approved for Release
Date 23 JUL 1987

OUTGOING TELEGRAM Department of State

INDICATE: ☐ COLLECT
☐ CHARGE TO

SANITIZED

TITLE: INTERNATIONAL CONGRESS OF SPACE MEDICINE

DOCREFERENCE: 29JAN76, 2 PP

INFLOCATION: MEXICO

INFODATE: 7509.

- 52 million pages reviewed for declassification
- 29 million pages declassified
- $8.8 billion spent on administration of the US government classification system

How Sensitive?

"After a 14-year legal battle by a California history professor, the FBI has released a new cache of material from a 300-page dossier on the late rock star John Lennon, and has agreed to pay $204,000 to cover legal fees incurred in his efforts to open the file. For all the years of challenge, however, the file contains little, if any, new information about Lennon, though it does present some bizarre details, like a description of an antiwar activist trying to train a parrot to speak profanities."

— NYT, 25 September 2007

A Problematic Trade-off

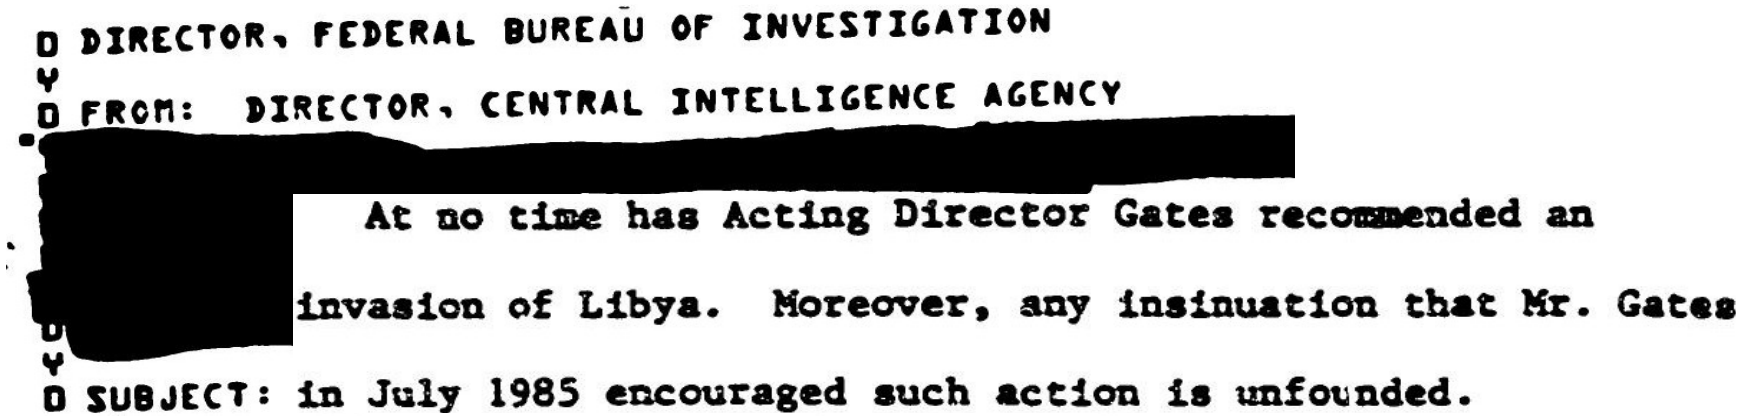

DIRECTOR, FEDERAL BUREAU OF INVESTIGATION

FROM: DIRECTOR, CENTRAL INTELLIGENCE AGENCY

[redacted] At no time has Acting Director Gates recommended an invasion of Libya. Moreover, any insinuation that Mr. Gates in July 1985 encouraged such action is unfounded.

SUBJECT:

- The more data kept secret, the less secure the data:
  - More people need to have access to the data
  - More storage space is required

What We Are *NOT* Studying...

guardian.co.uk

News | Sport | Comment | Culture | Business | Money

News › World news › Afghanistan: The war logs

AFGHANISTAN
THE WAR LOGS

Exploring Declassified Documents

- Declassification goals:
  - Recommend documents for human review
  - Match documents with human reviewers' expertise
- Transparency research goals:
  - High-level characterization of the data
  - Finding specific, known information of interest
  - Finding “interesting” or “unexpected” information

Declassified Documents: DDRS

- ~88,000 formerly-classified government documents
- Created and declassified between 1926 and 2005

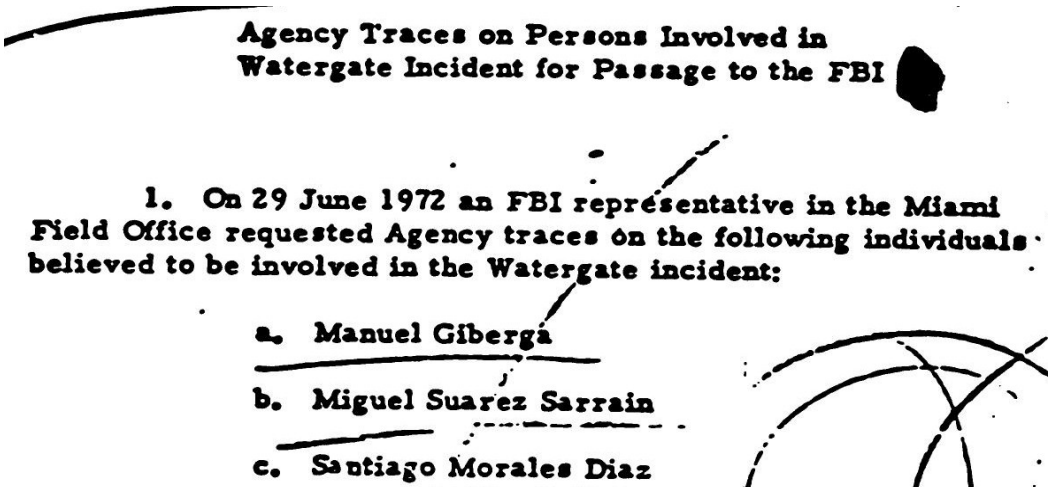

Agency Traces on Persons Involved in
Watergate Incident for Passage to the FBI

1. On 29 June 1972 an FBI representative in the Miami Field Office requested Agency traces on the following individuals believed to be involved in the Watergate incident:

a. Manuel Giberga

b. Miguel Suarez Sarrain

c. Santiago Morales Diaz

Available Information

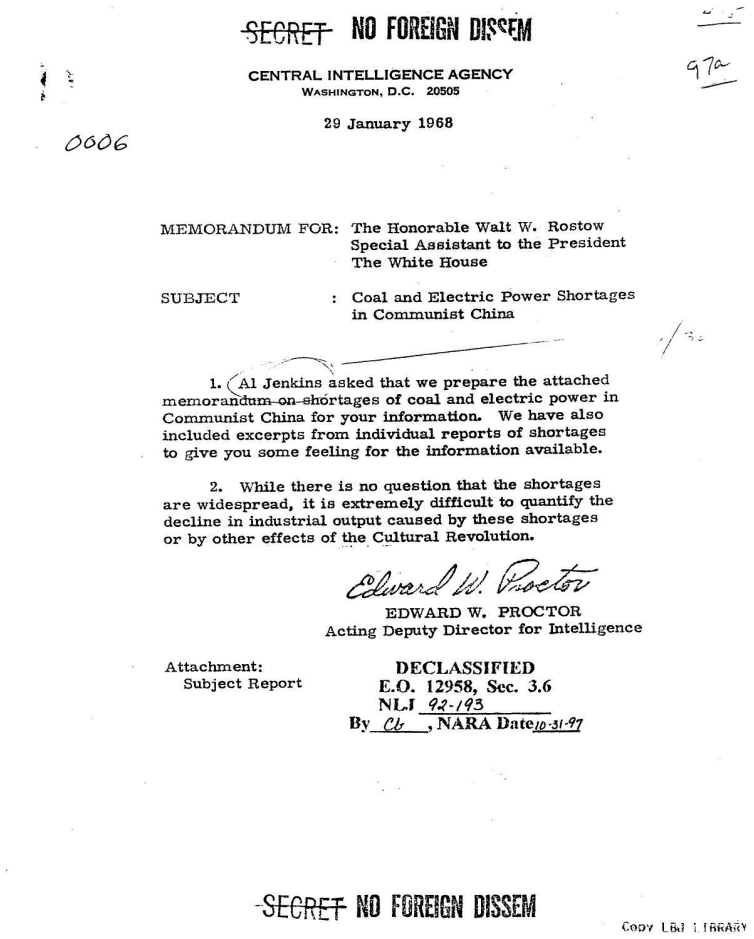

~~SECRET~~ NO FOREIGN DISSEM

CENTRAL INTELLIGENCE AGENCY
WASHINGTON, D.C. 20505

29 January 1968

0006

97a

MEMORANDUM FOR: The Honorable Walt W. Rostow
Special Assistant to the President
The White House

SUBJECT : Coal and Electric Power Shortages
in Communist China

1. Al Jenkins asked that we prepare the attached memorandum on shortages of coal and electric power in Communist China for your information. We have also included excerpts from individual reports of shortages to give you some feeling for the information available.

2. While there is no question that the shortages are widespread, it is extremely difficult to quantify the decline in industrial output caused by these shortages or by other effects of the Cultural Revolution.

Edward W. Proctor

EDWARD W. PROCTOR
Acting Deputy Director for Intelligence

Attachment:
Subject Report

DECLASSIFIED
E.O. 12958, Sec. 3.6
NLJ 92-193
By CL , NARA Date 10-31-97

~~SECRET~~ NO FOREIGN DISSEM

Copy LBJ LIBRARY

- Sanitized?
- Classification level
- Issuer
- Creation date
- Document type
- Declassification date

Available Information

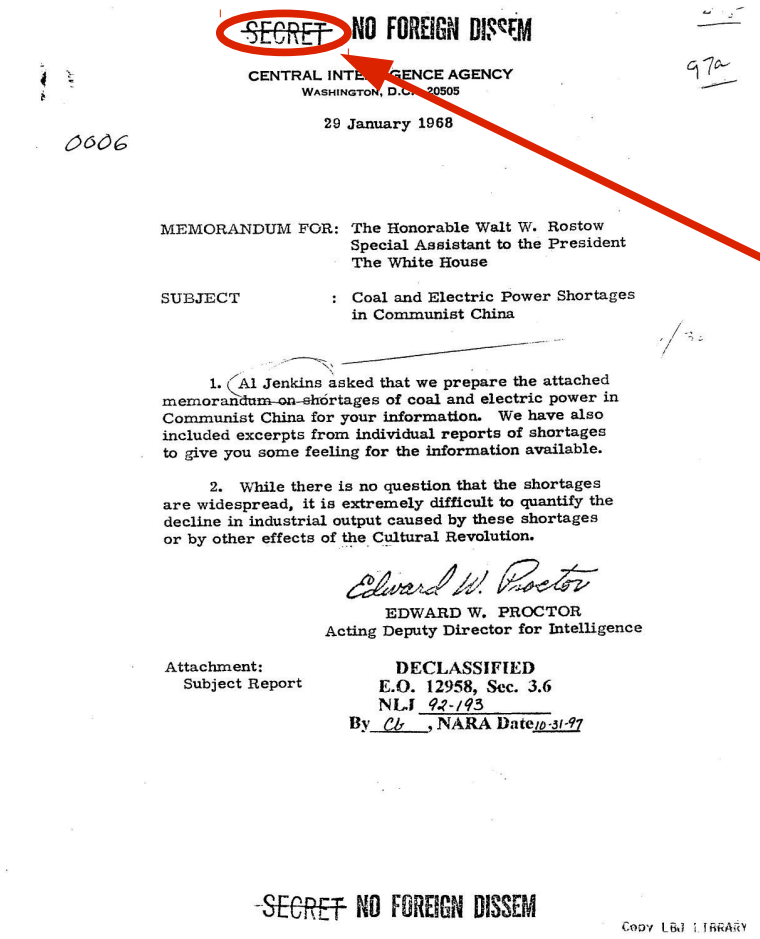

~~SECRET~~ NO FOREIGN DISSEM

CENTRAL INTELLIGENCE AGENCY
WASHINGTON, D.C. 20505

29 January 1968

0006

97a

MEMORANDUM FOR: The Honorable Walt W. Rostow
Special Assistant to the President
The White House

SUBJECT : Coal and Electric Power Shortages
in Communist China

1. Al Jenkins asked that we prepare the attached memorandum on shortages of coal and electric power in Communist China for your information. We have also included excerpts from individual reports of shortages to give you some feeling for the information available.

2. While there is no question that the shortages are widespread, it is extremely difficult to quantify the decline in industrial output caused by these shortages or by other effects of the Cultural Revolution.

Edward W. Proctor

EDWARD W. PROCTOR
Acting Deputy Director for Intelligence

Attachment:
Subject Report

DECLASSIFIED
E.O. 12958, Sec. 3.6
NLJ 92-193
By CL , NARA Date 10-31-97

~~SECRET~~ NO FOREIGN DISSEM

Copy LBJ LIBRARY

- Sanitized?
- Classification level
- Issuer
- Creation date
- Document type
- Declassification date

Available Information

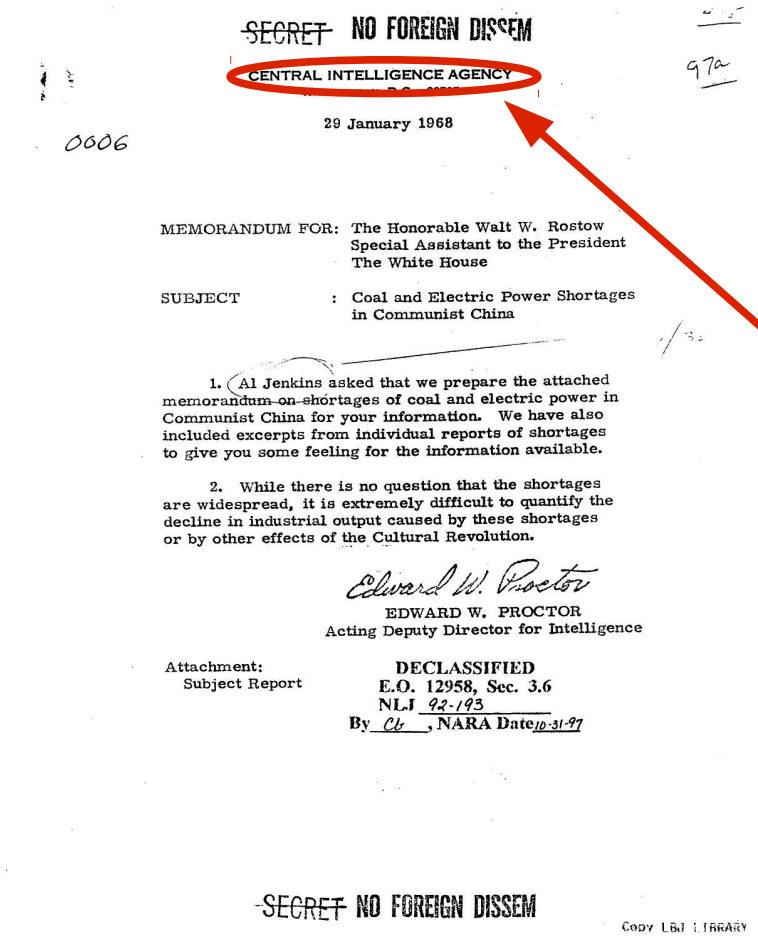

~~SECRET~~ NO FOREIGN DISSEM

CENTRAL INTELLIGENCE AGENCY

29 January 1968

0006

97a

MEMORANDUM FOR: The Honorable Walt W. Rostow
Special Assistant to the President
The White House

SUBJECT : Coal and Electric Power Shortages
in Communist China

1. Al Jenkins asked that we prepare the attached memorandum on shortages of coal and electric power in Communist China for your information. We have also included excerpts from individual reports of shortages to give you some feeling for the information available.

2. While there is no question that the shortages are widespread, it is extremely difficult to quantify the decline in industrial output caused by these shortages or by other effects of the Cultural Revolution.

Edward W. Proctor

EDWARD W. PROCTOR
Acting Deputy Director for Intelligence

Attachment:
Subject Report

DECLASSIFIED
E.O. 12958, Sec. 3.6
NLJ 92-193
By CIA, NARA Date 10-31-97

~~SECRET~~ NO FOREIGN DISSEM

Copy LBJ LIBRARY

- Sanitized?
- Classification level
- Issuer
- Creation date
- Document type
- Declassification date

Available Information

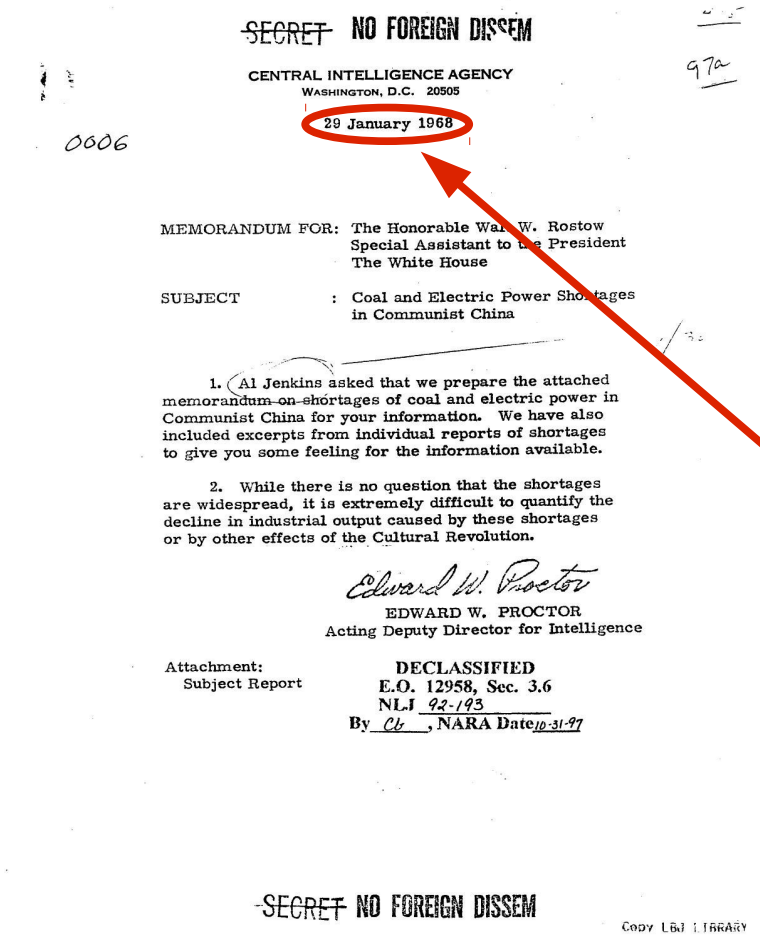

~~SECRET~~ NO FOREIGN DISSEM

CENTRAL INTELLIGENCE AGENCY
WASHINGTON, D.C. 20505

29 January 1968

0006

97a

MEMORANDUM FOR: The Honorable Walt W. Rostow
Special Assistant to the President
The White House

SUBJECT : Coal and Electric Power Shortages in Communist China

1. Al Jenkins asked that we prepare the attached memorandum on shortages of coal and electric power in Communist China for your information. We have also included excerpts from individual reports of shortages to give you some feeling for the information available.

2. While there is no question that the shortages are widespread, it is extremely difficult to quantify the decline in industrial output caused by these shortages or by other effects of the Cultural Revolution.

Edward W. Proctor

EDWARD W. PROCTOR
Acting Deputy Director for Intelligence

Attachment:
Subject Report

DECLASSIFIED
E.O. 12958, Sec. 3.6
NLJ 92-193
By CL , NARA Date 10-31-97

~~SECRET~~ NO FOREIGN DISSEM

COPY LBJ LIBRARY

- Sanitized?
- Classification level
- Issuer
- Creation date
- Document type
- Declassification date

Available Information

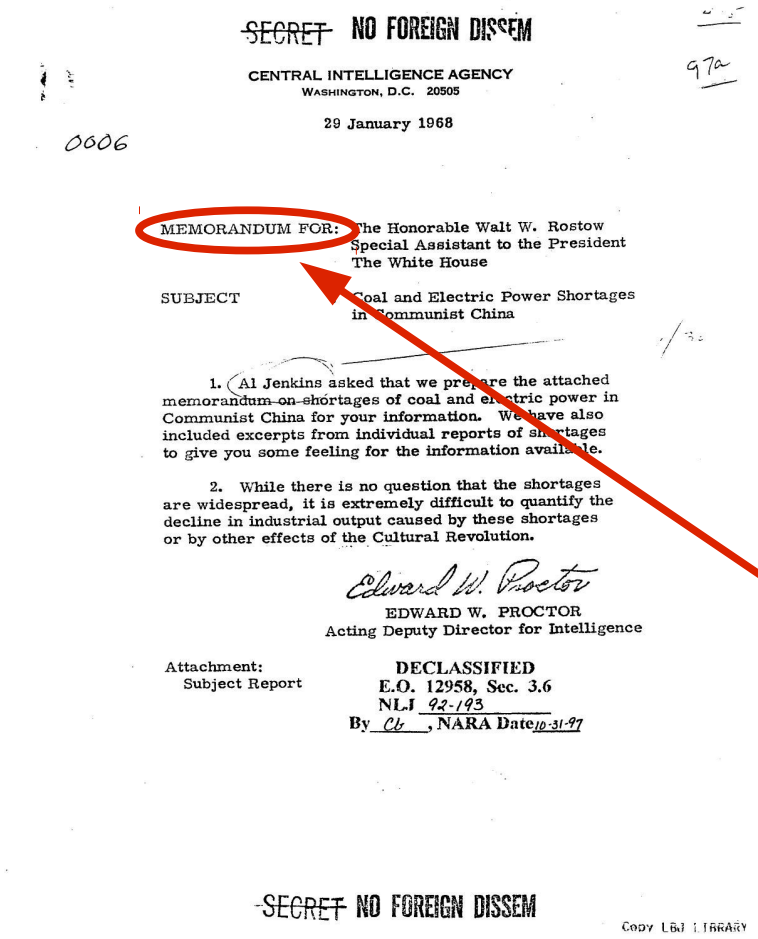

~~SECRET~~ NO FOREIGN DISSEM

CENTRAL INTELLIGENCE AGENCY
WASHINGTON, D.C. 20505

29 January 1968

0006

MEMORANDUM FOR: The Honorable Walt W. Rostow
Special Assistant to the President
The White House

SUBJECT: Coal and Electric Power Shortages in Communist China

1. Al Jenkins asked that we prepare the attached memorandum on shortages of coal and electric power in Communist China for your information. We have also included excerpts from individual reports of shortages to give you some feeling for the information available.

2. While there is no question that the shortages are widespread, it is extremely difficult to quantify the decline in industrial output caused by these shortages or by other effects of the Cultural Revolution.

Edward W. Proctor

EDWARD W. PROCTOR
Acting Deputy Director for Intelligence

Attachment:
Subject Report

DECLASSIFIED
E.O. 12958, Sec. 3.6
NLJ 92-193
By CL, NARA Date 10-31-97

~~SECRET~~ NO FOREIGN DISSEM

COPY LBJ LIBRARY

- Sanitized?
- Classification level
- Issuer
- Creation date
- Document type
- Declassification date

Available Information

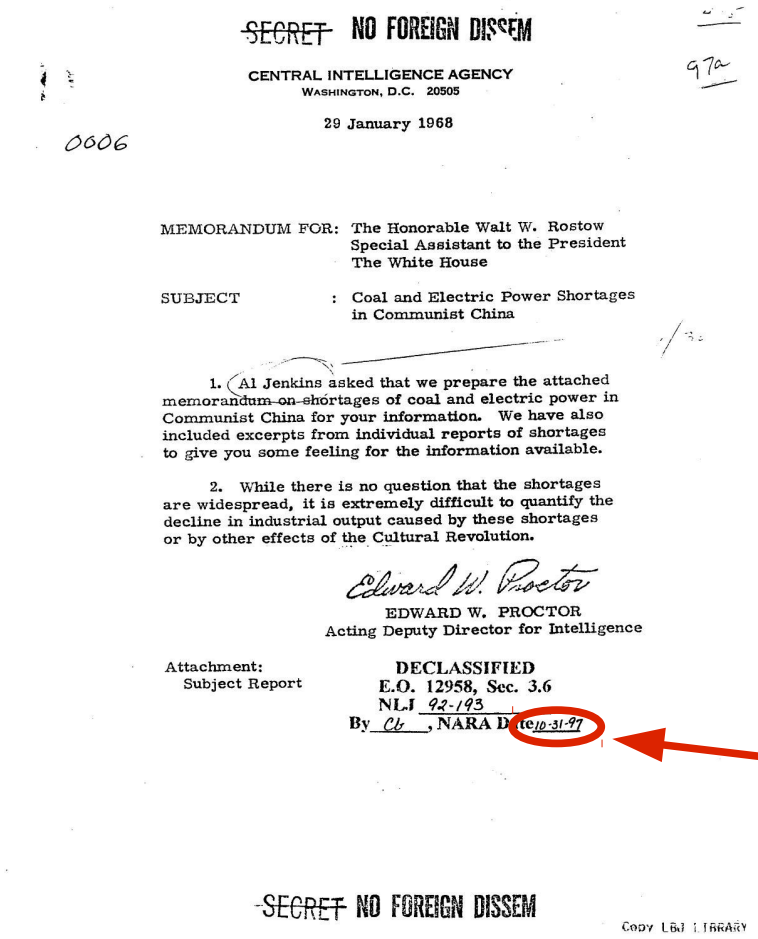

~~SECRET~~ NO FOREIGN DISSEM

CENTRAL INTELLIGENCE AGENCY
WASHINGTON, D.C. 20505

29 January 1968

0006

MEMORANDUM FOR: The Honorable Walt W. Rostow
Special Assistant to the President
The White House

SUBJECT : Coal and Electric Power Shortages
in Communist China

1. Al Jenkins asked that we prepare the attached memorandum on shortages of coal and electric power in Communist China for your information. We have also included excerpts from individual reports of shortages to give you some feeling for the information available.

2. While there is no question that the shortages are widespread, it is extremely difficult to quantify the decline in industrial output caused by these shortages or by other effects of the Cultural Revolution.

Edward W. Proctor

EDWARD W. PROCTOR
Acting Deputy Director for Intelligence

Attachment:
Subject Report

DECLASSIFIED
E.O. 12958, Sec. 3.6
NLJ 92-193
By CL , NARA Date 10-31-97

~~SECRET~~ NO FOREIGN DISSEM

Copy LBJ LIBRARY

- Sanitized?
- Classification level
- Issuer
- Creation date
- Type
- Declassification date

Declassification Durations

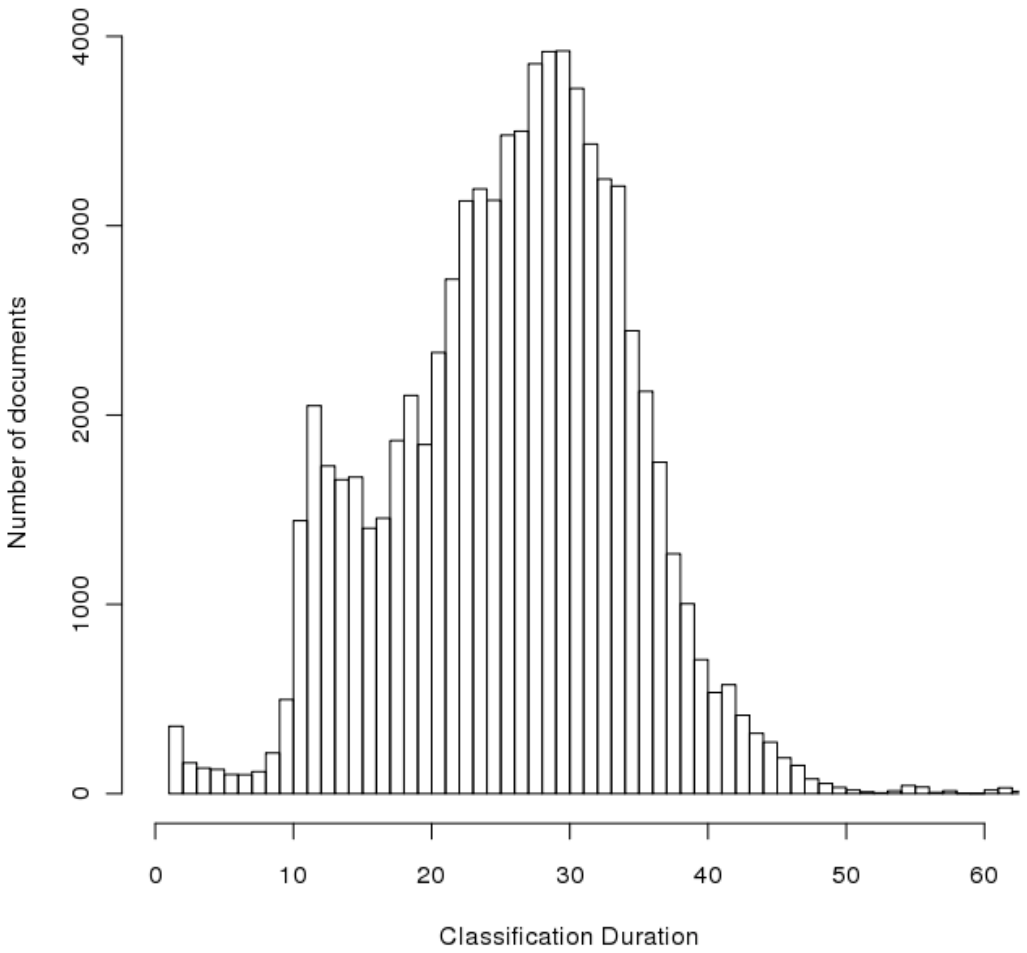

Survival Analysis

- Statistical methods for evaluating “time until death”:
  - Biology/medicine: organism death
  - Engineering: component failure
  - Social science: event durations (e.g., parolee recidivism)
- Goal: model effect on survival time of covariates, e.g.,
  - Vaccine treatments
  - Temperature differences
  - Job placement or education programs

Document “Survival”

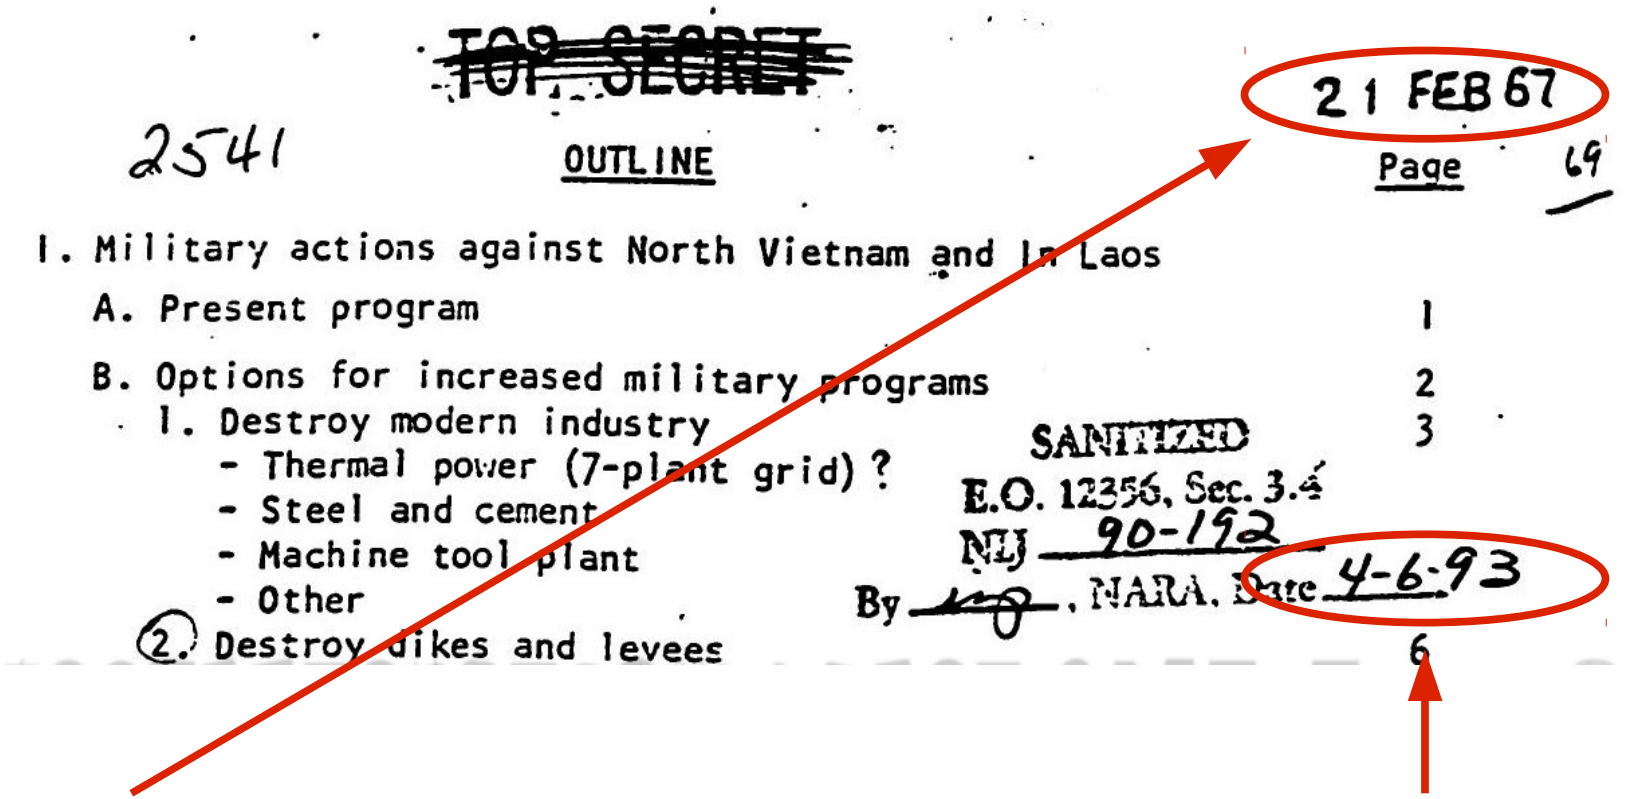

~~TOP SECRET~~

2541

OUTLINE

21 FEB 67

Page 69

I. Military actions against North Vietnam and In Laos

A. Present program 1

B. Options for increased military programs 2

1. Destroy modern industry 3
   - Thermal power (7-plant grid)?
   - Steel and cement
   - Machine tool plant
   - Other

2. Destroy dikes and levees 6

SANITIZED
E.O. 12356, Sec. 3.4
NLJ 90-192
By , NARA, Date 4-6-93

creation date
2/21/67

26 years

declassification date
4/6/93

Survival Distribution of Documents

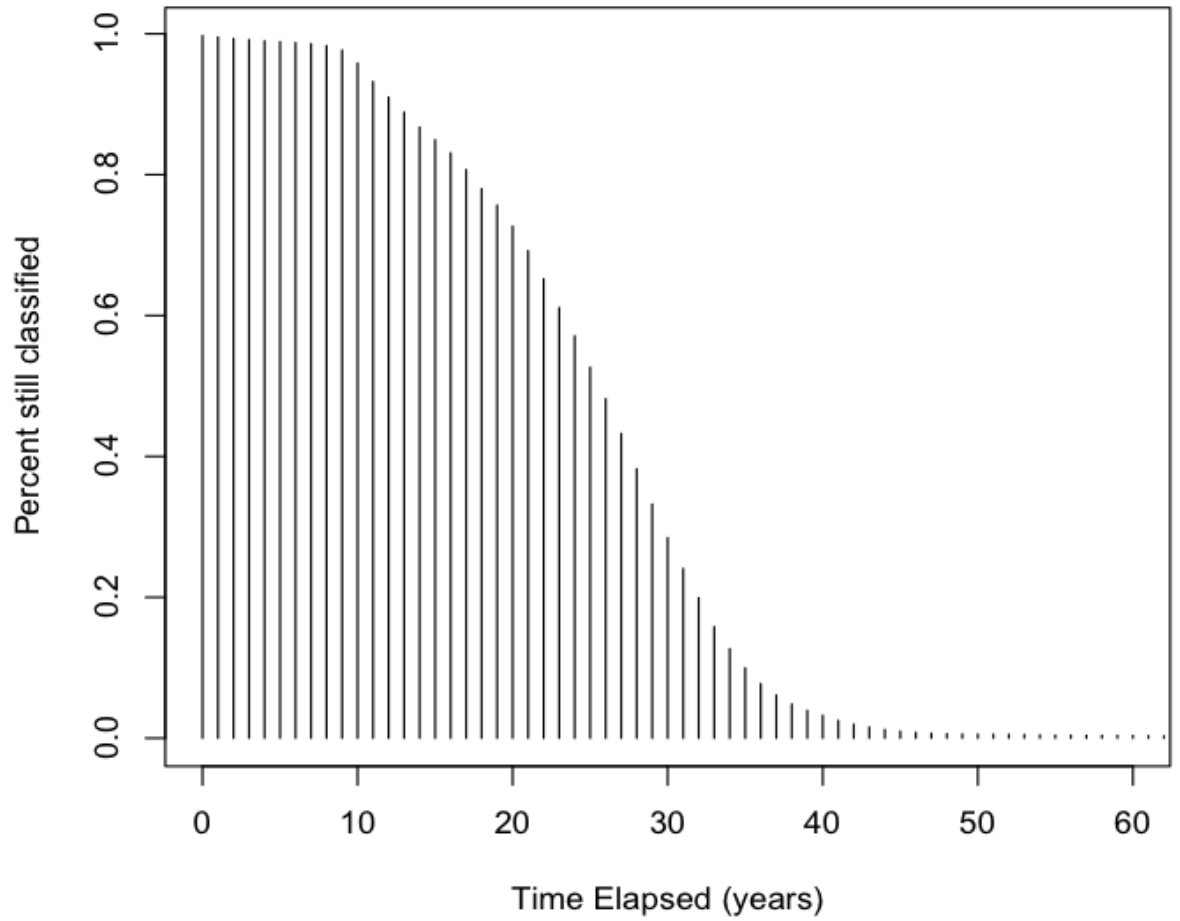

Accelerated Failure Time Models

- Survival analysis with covariates $\boldsymbol{x}_i$
- Linear models for the log of the "duration":

$$\log(t_i) = \boldsymbol{x}_i^\top \boldsymbol{\beta} + \epsilon_i$$

- Parametric: a probability distribution is specified
  - e.g., Weibull, log-normal, gamma, log-logistic...
- Can make predictions for unseen data

Classification and Content

HIS APPROACH WAS, "WELL, OF COURSE, WE KNOW THERE ISN'T ANYTHING TO THIS ALLEGED PHENOMENON (FLYING SAUCERS), BUT ON THE OTHER HAND". DURING HIS TALK SHKLOVSKIY AND OTHER SOVIETS JOKED AND LAUGHED AND OBVIOUSLY DID NOT TAKE THE SPEAKER'S REMARKS SERIOUSLY.

1975 to 1989

1946 to 2003

CENTRAL INTELLIGENCE GROUP

SOVIET CAPABILITIES FOR THE DEVELOPMENT AND PRODUCTION OF CERTAIN TYPES OF WEAPONS AND EQUIPMENT

1. Herein is presented an estimate of Soviet capabilities in the development and production, during the next ten years, of certain weapons and equipment, as follows:

Classification and Content: 1960s

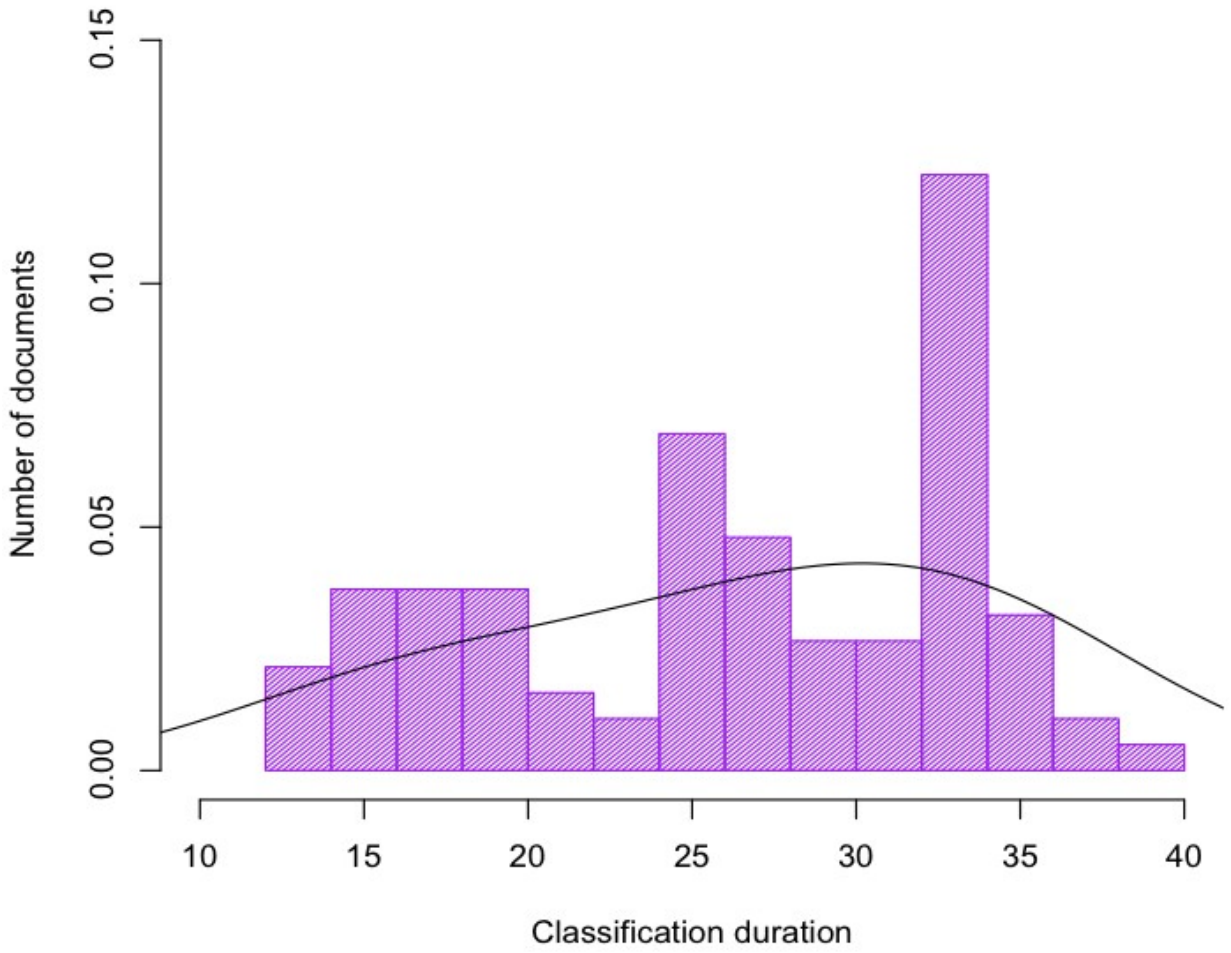

Classification and Content: 1960s

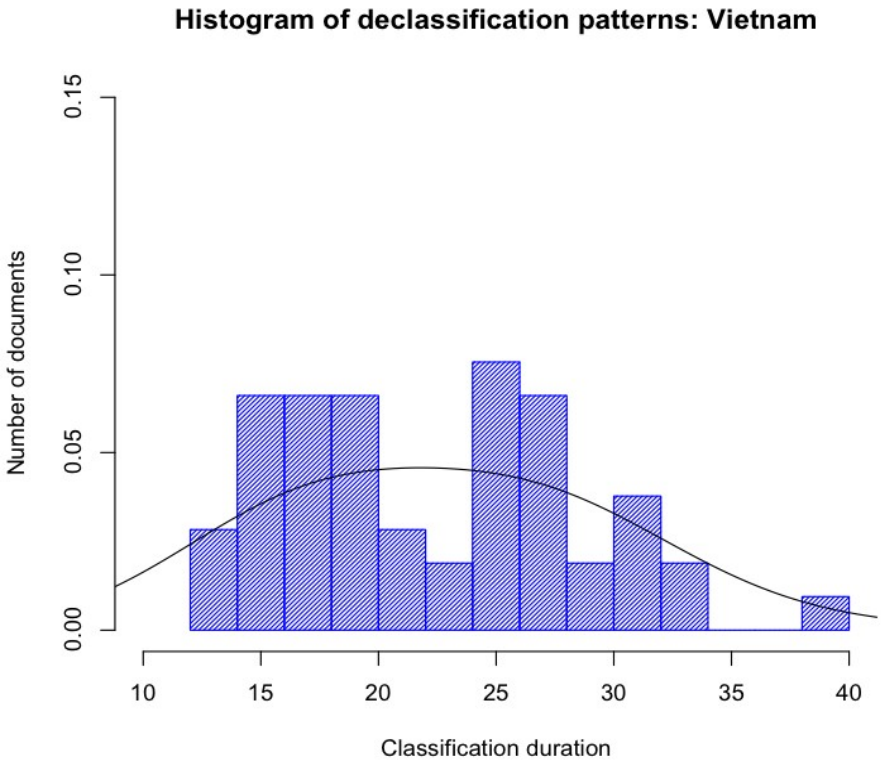

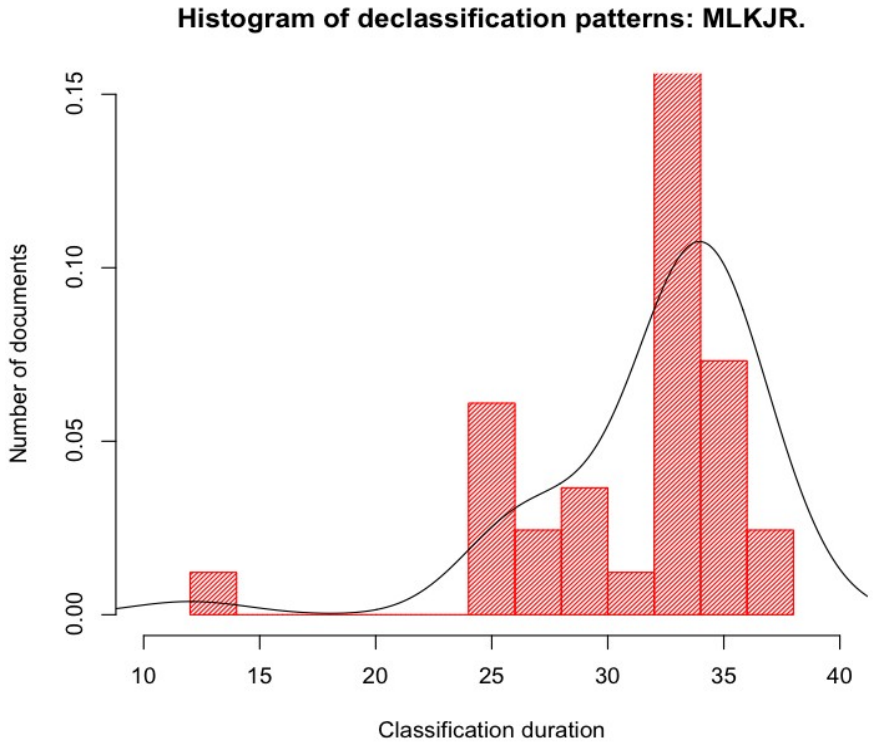

Word Frequencies?

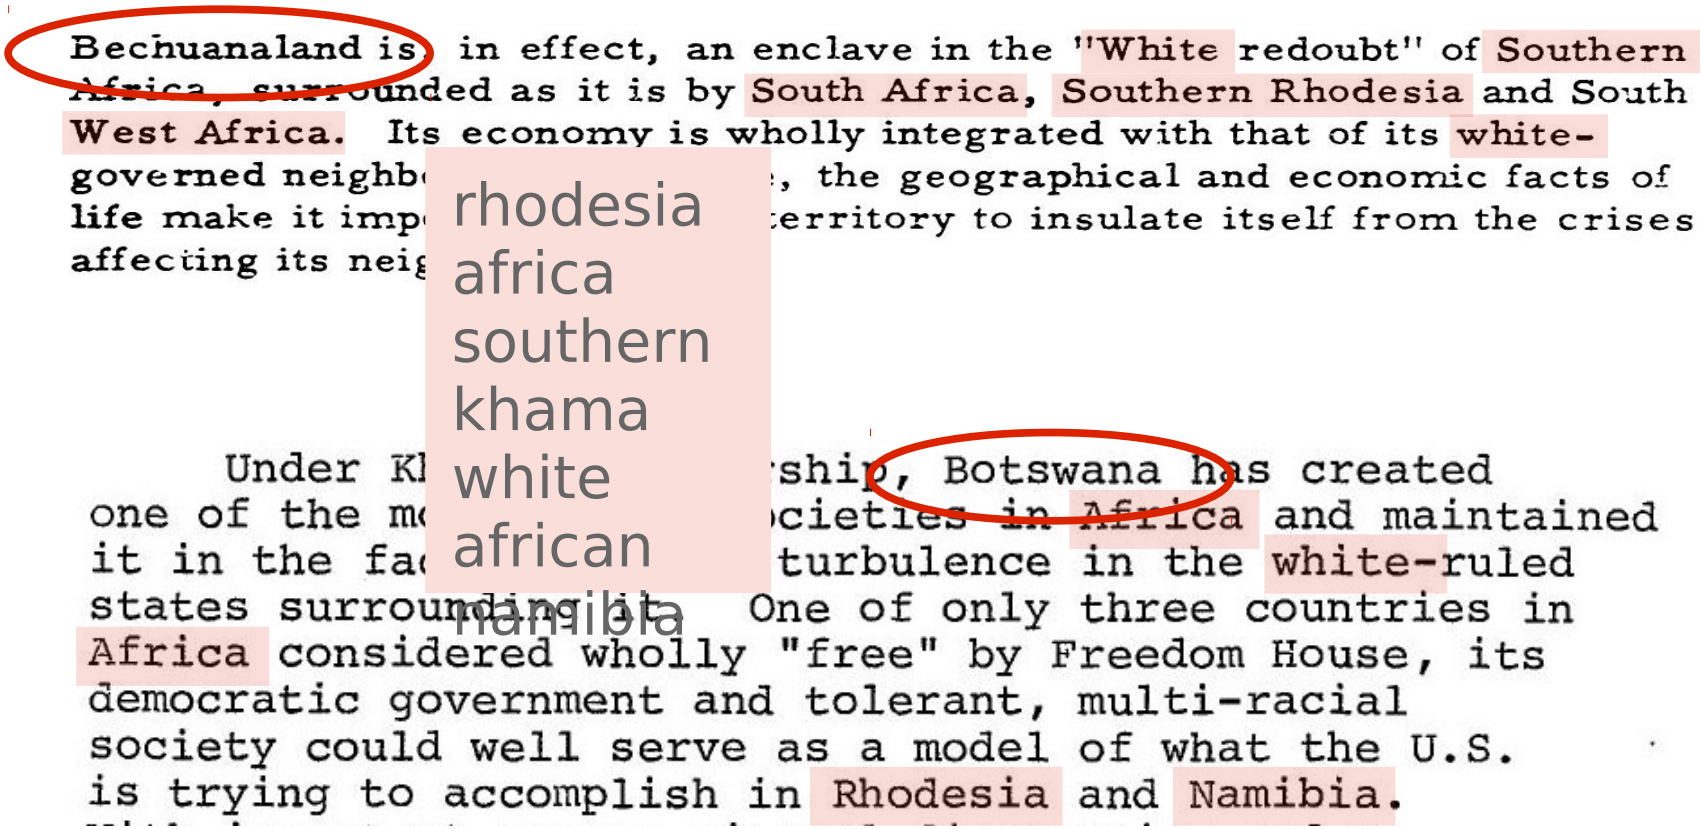

Bechuanaland is, in effect, an enclave in the "White redoubt" of Southern Africa, surrounded as it is by South Africa, Southern Rhodesia and South West Africa. Its economy is wholly integrated with that of its white-governed neighb , the geographical and economic facts of life make it imp territory to insulate itself from the crises affecting its neig

rhodesia
africa
southern
khama
white
african
namibia

Under K ship, Botswana has created one of the m ocieties in Africa and maintained it in the fa turbulence in the white-ruled states surrounding it. One of only three countries in Africa considered wholly "free" by Freedom House, its democratic government and tolerant, multi-racial society could well serve as a model of what the U.S. is trying to accomplish in Rhodesia and Namibia.

Topics in Declassified Documents

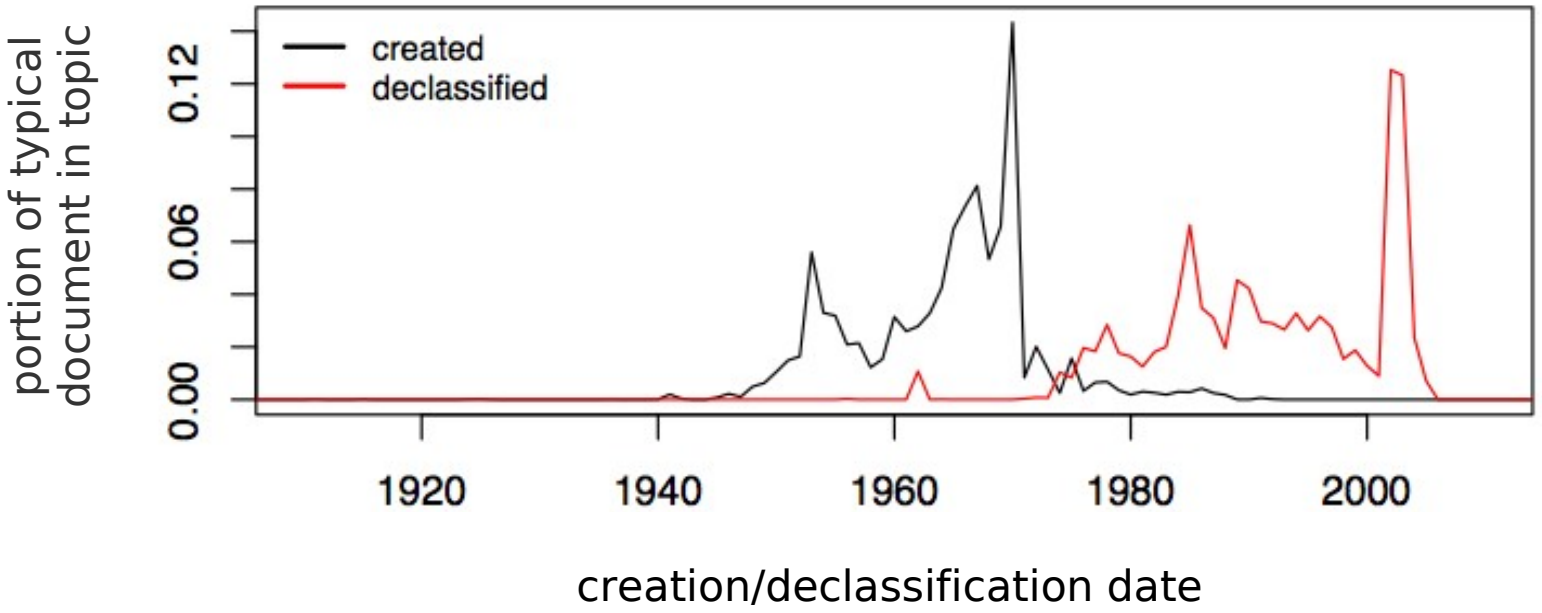

corps, service, volunteers, men, volunteer, age, draft, selective, calls, young, manpower, year, army, deferments, induction, armed, freedom, ...

Topics in Declassified Documents

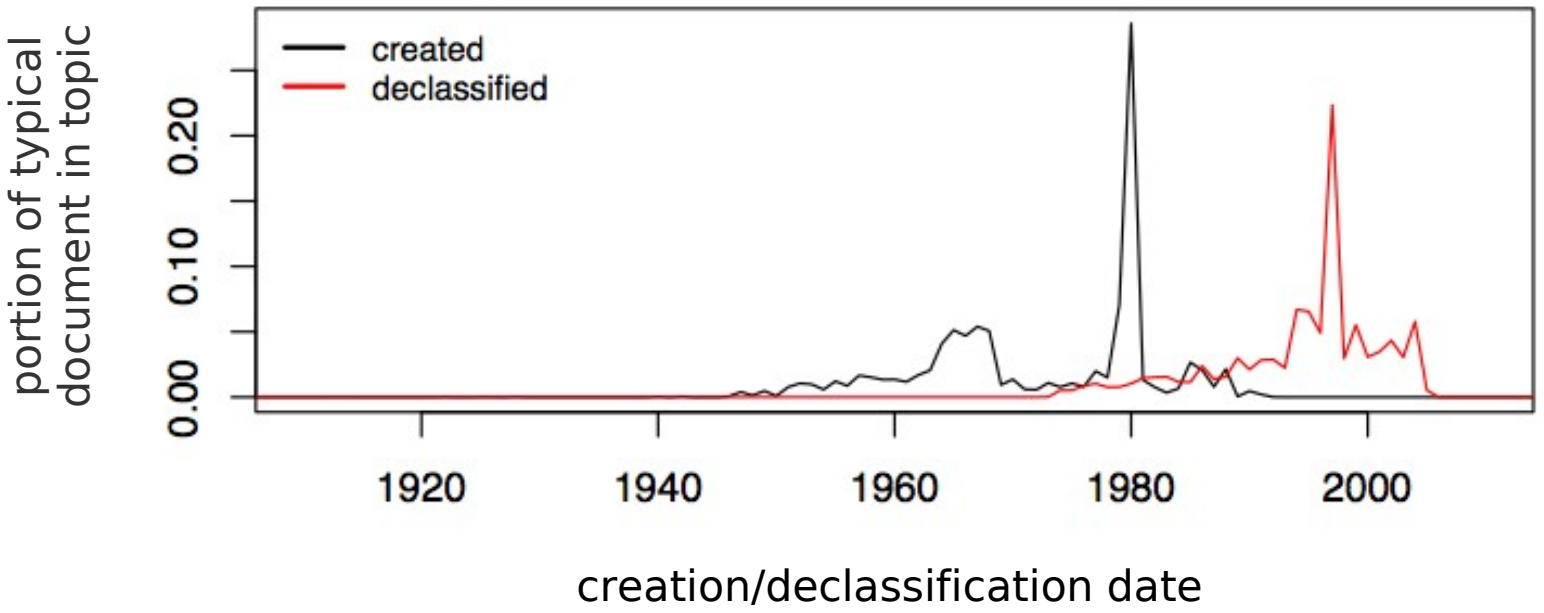

package, hostages, release, hostage, khomeini, packages, ghotbzadeh, held, released, banisadr, revolutionary, debriefing, scenario, family, date, ...

Topics in Declassified Documents

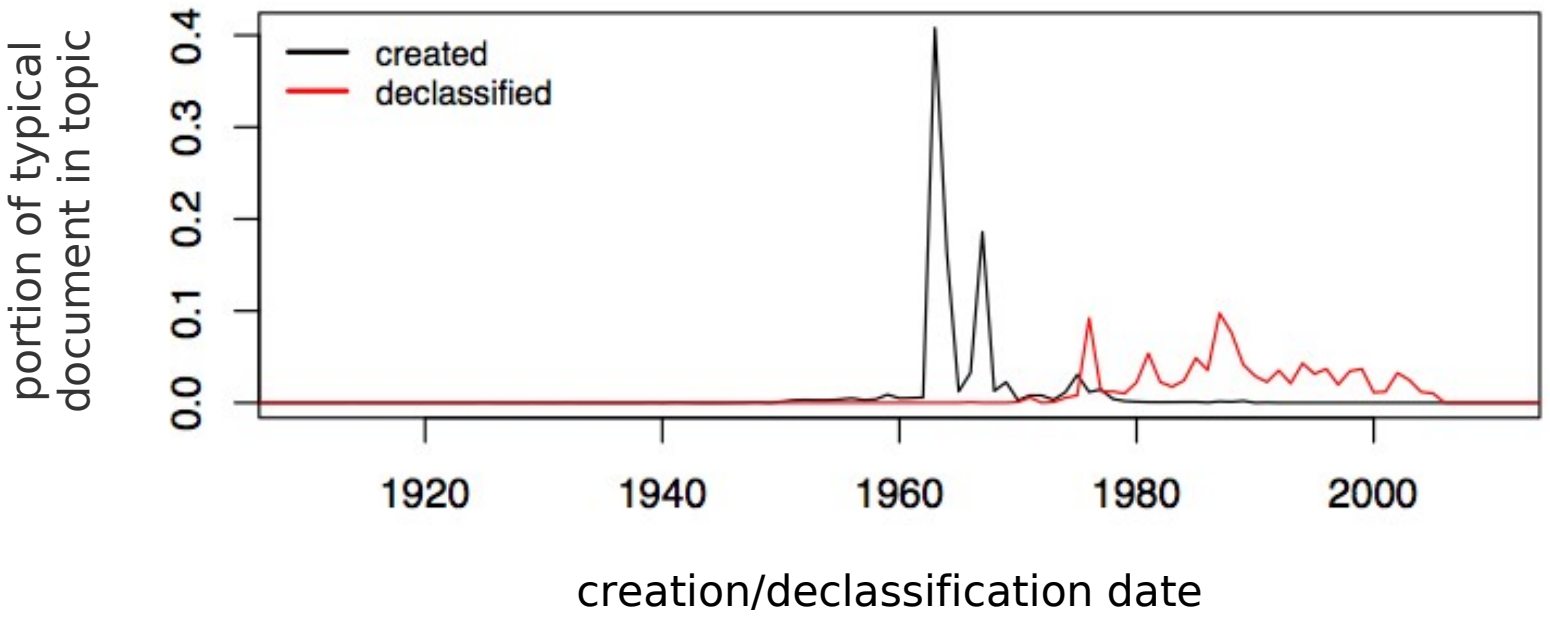

oswald, dallas, assassination, kennedy, texas, fbi, orleans, advised,
lee, president, bureau, started, harvey, john, information, ruby, november, ...

Topics in Declassified Documents

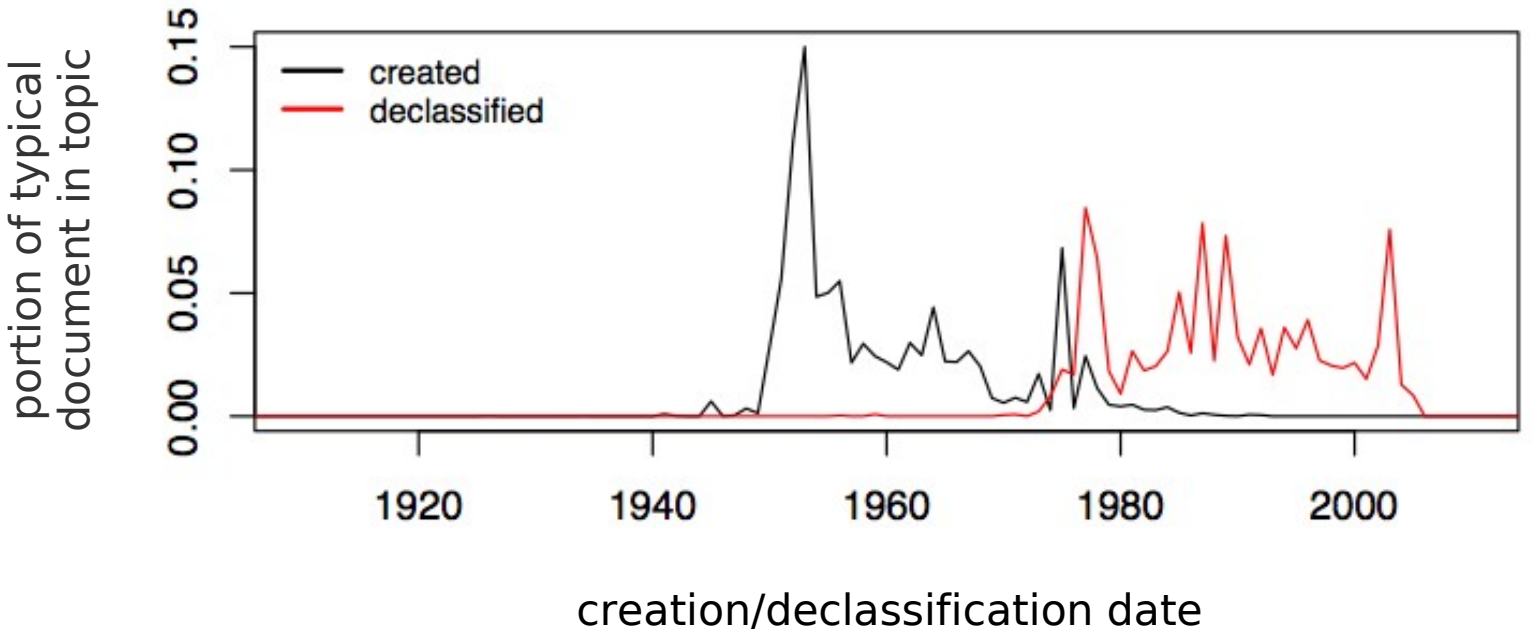

artichoke, subject, drugs, techniques, work, interrogation, writer, drug, lsd, effects, hypnosis, methods, medical, physical, subjects, human, ...

Predicting Duration Using Topics

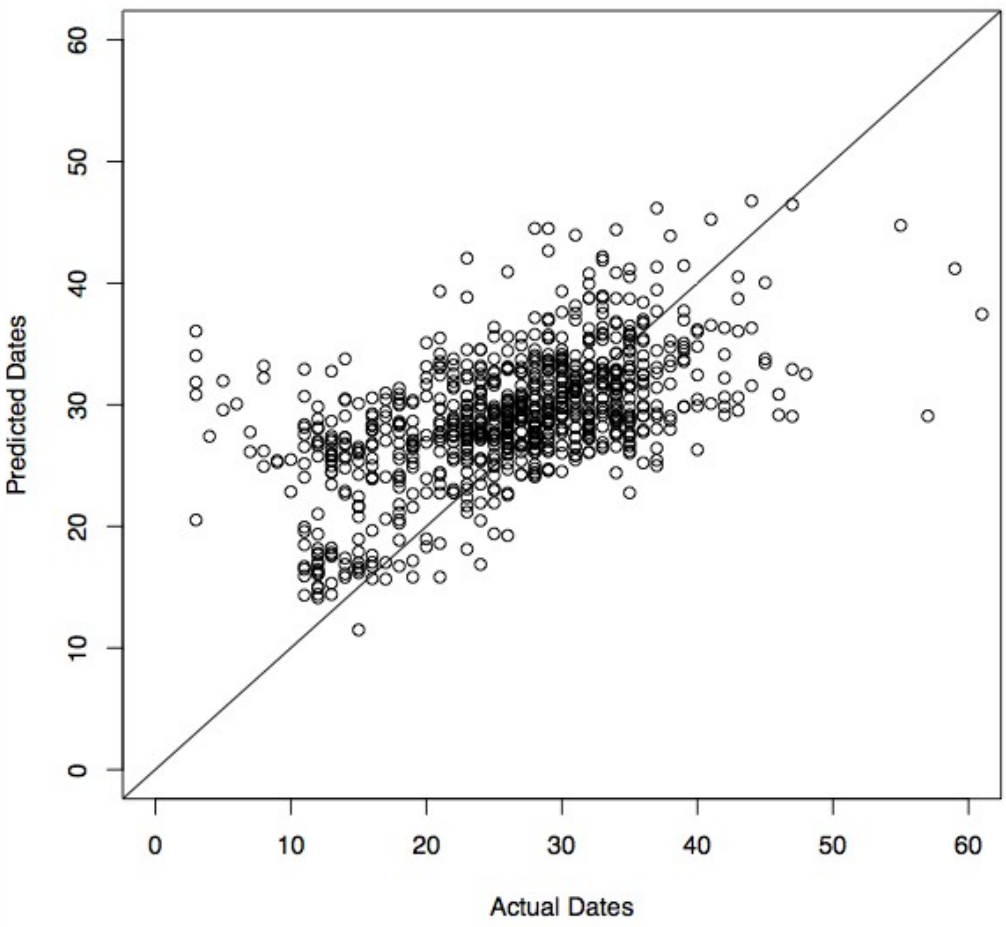

Predicted Duration - Actual Duration

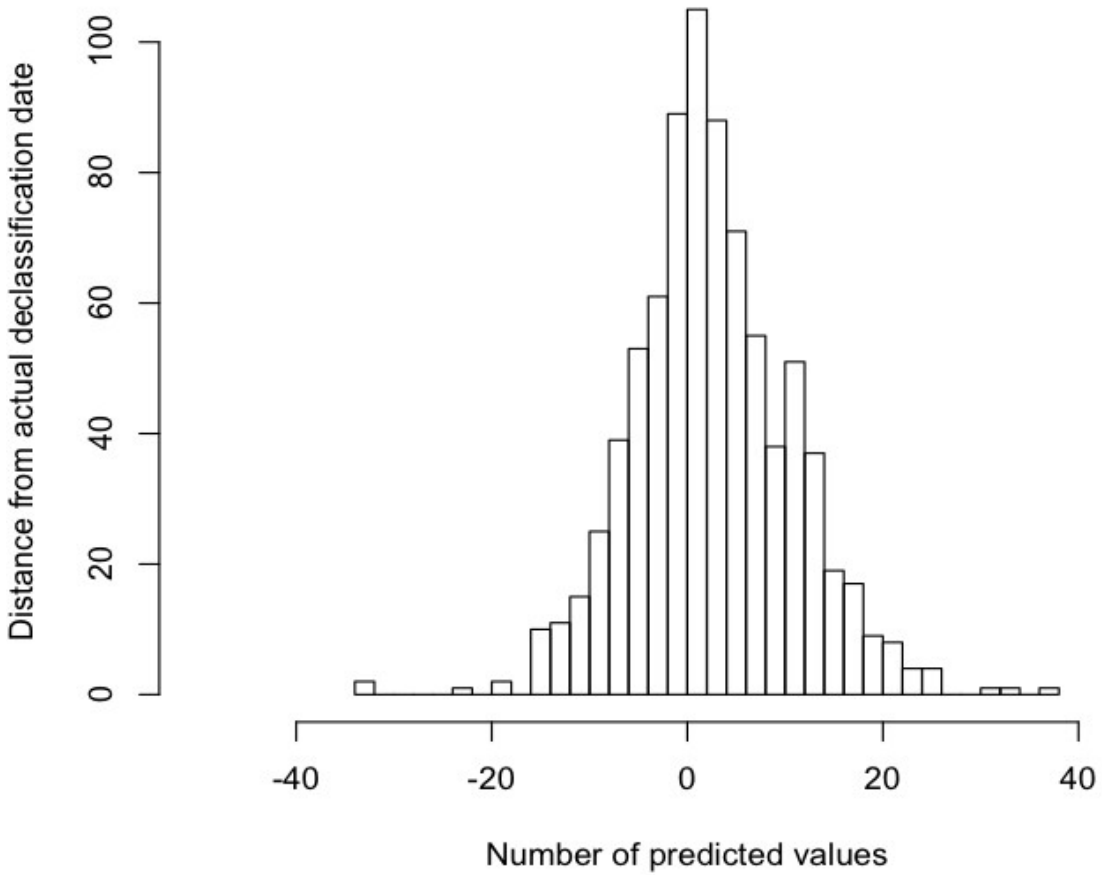

Jointly Modeling Text and Duration

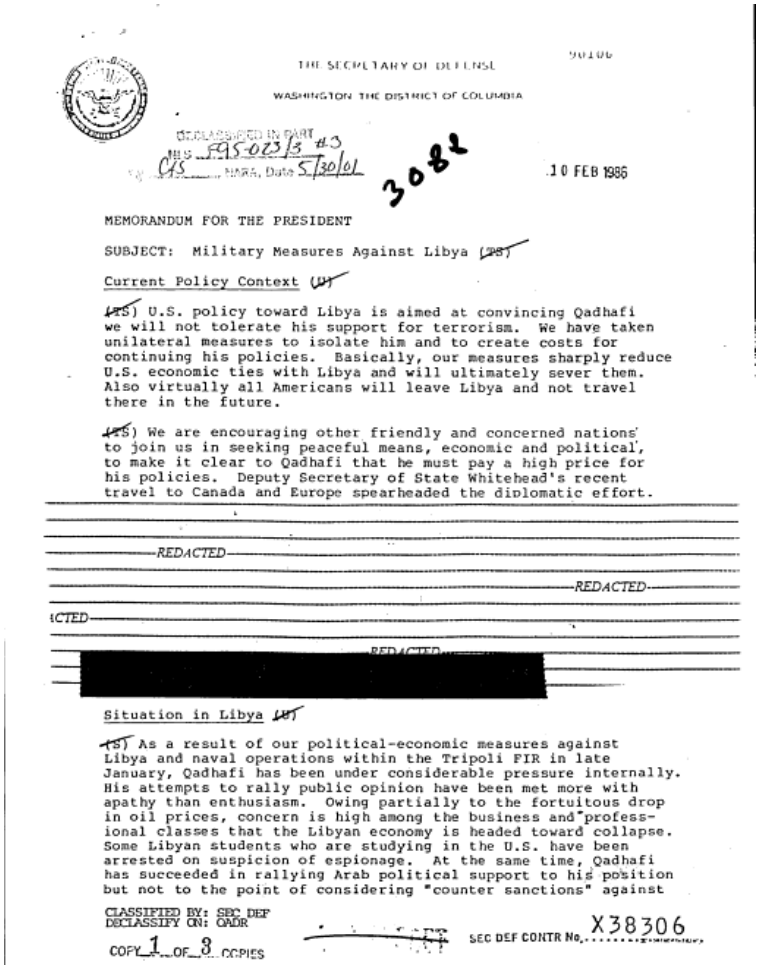

THE SECRETARY OF DEFENSE

WASHINGTON THE DISTRICT OF COLUMBIA

DECLASSIFIED IN PART
Sec. Def. FAS-023/3 #3
By C/S NARA, Date 5/30/01

3082

10 FEB 1986

MEMORANDUM FOR THE PRESIDENT

SUBJECT: Military Measures Against Libya (~~S~~)

Current Policy Context (~~U~~)

(~~S~~) U.S. policy toward Libya is aimed at convincing Qadhafi we will not tolerate his support for terrorism. We have taken unilateral measures to isolate him and to create costs for continuing his policies. Basically, our measures sharply reduce U.S. economic ties with Libya and will ultimately sever them. Also virtually all Americans will leave Libya and not travel there in the future.

(~~S~~) We are encouraging other friendly and concerned nations to join us in seeking peaceful means, economic and political, to make it clear to Qadhafi that he must pay a high price for his policies. Deputy Secretary of State Whitehead's recent travel to Canada and Europe spearheaded the diplomatic effort.

REDACTED

REDACTED

CTED

REDACTED

[redacted]

Situation in Libya (~~U~~)

(~~S~~) As a result of our political-economic measures against Libya and naval operations within the Tripoli FIR in late January, Qadhafi has been under considerable pressure internally. His attempts to rally public opinion have been met more with apathy than enthusiasm. Owing partially to the fortuitous drop in oil prices, concern is high among the business and professional classes that the Libyan economy is headed toward collapse. Some Libyan students who are studying in the U.S. have been arrested on suspicion of espionage. At the same time, Qadhafi has succeeded in rallying Arab political support to his position but not to the point of considering "counter sanctions" against

CLASSIFIED BY: SEC DEF
DECLASSIFY ON: OADR

COPY 1 OF 3 COPIES

SEC DEF CONTR No. X38306

- Topics provide information about classification durations
- Goal: incorporate durations into the generative model
- Infer latent topics using both textual and temporal information

Jointly Modeling Text and Duration

[Shorey et al., '11]

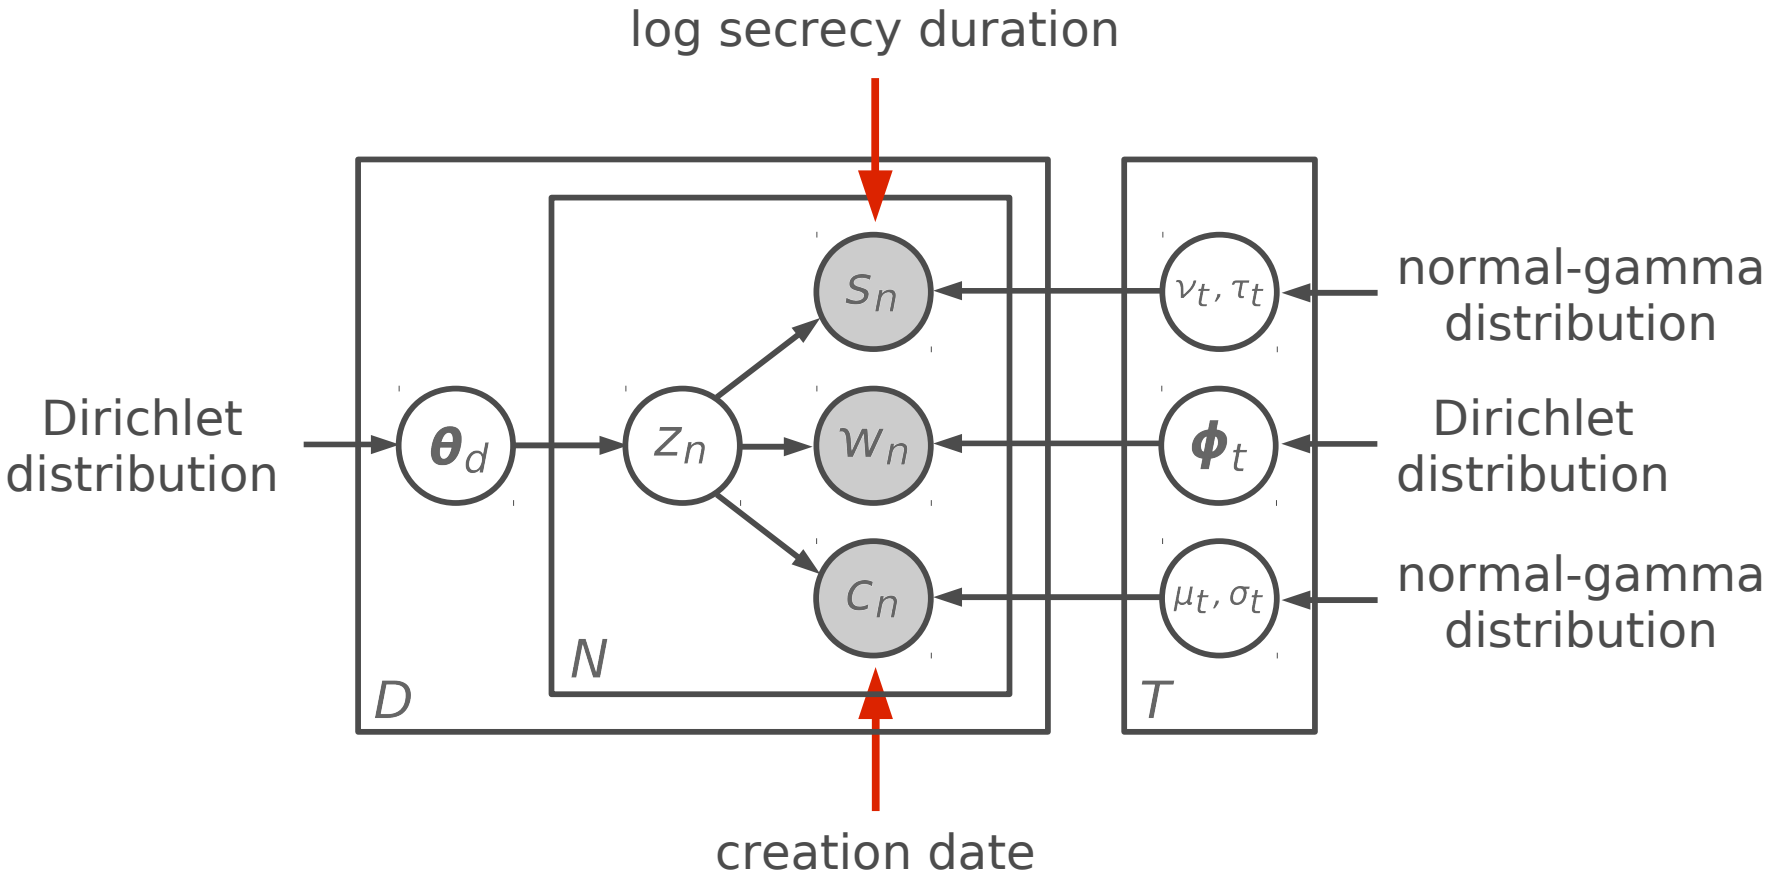

Topic-Specific Duration Distributions

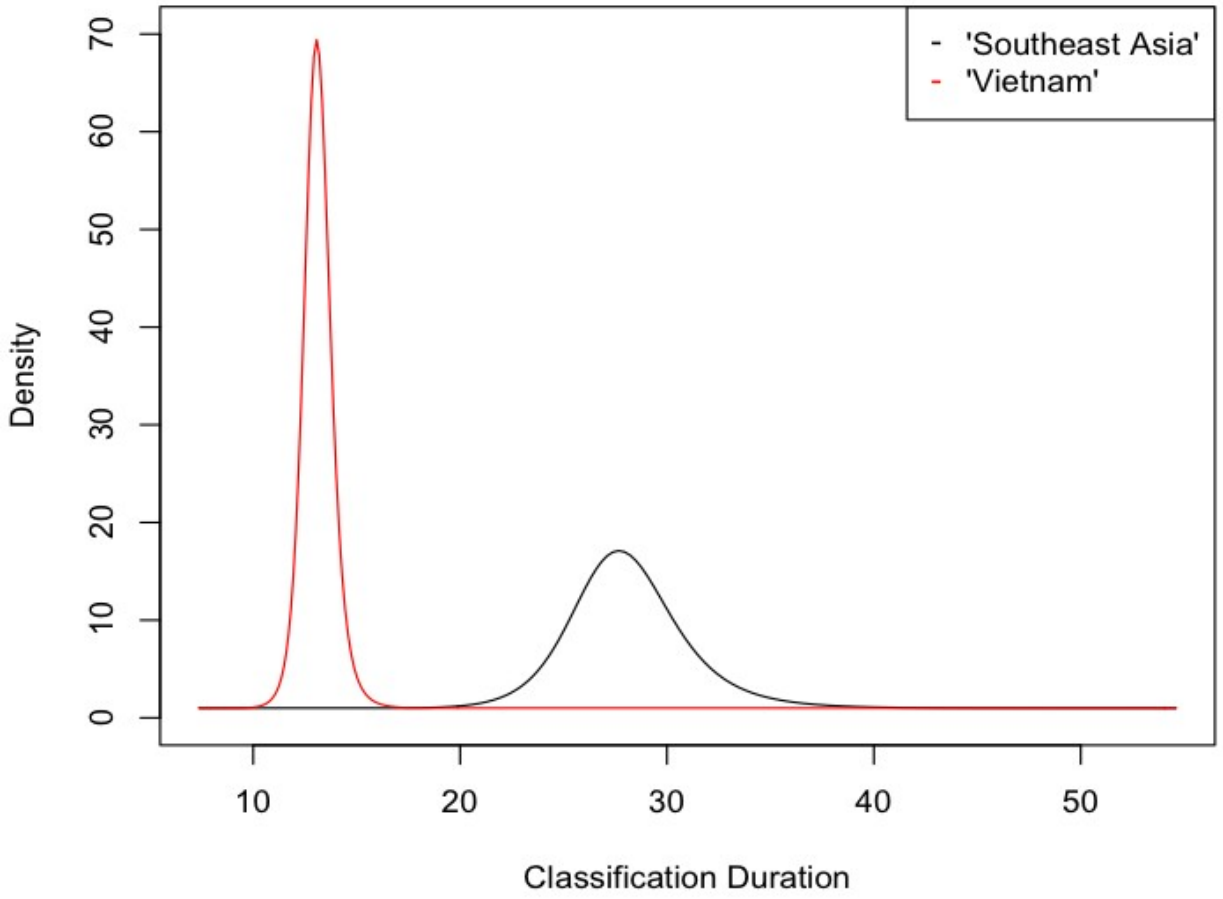

Topic-Specific Duration Distributions

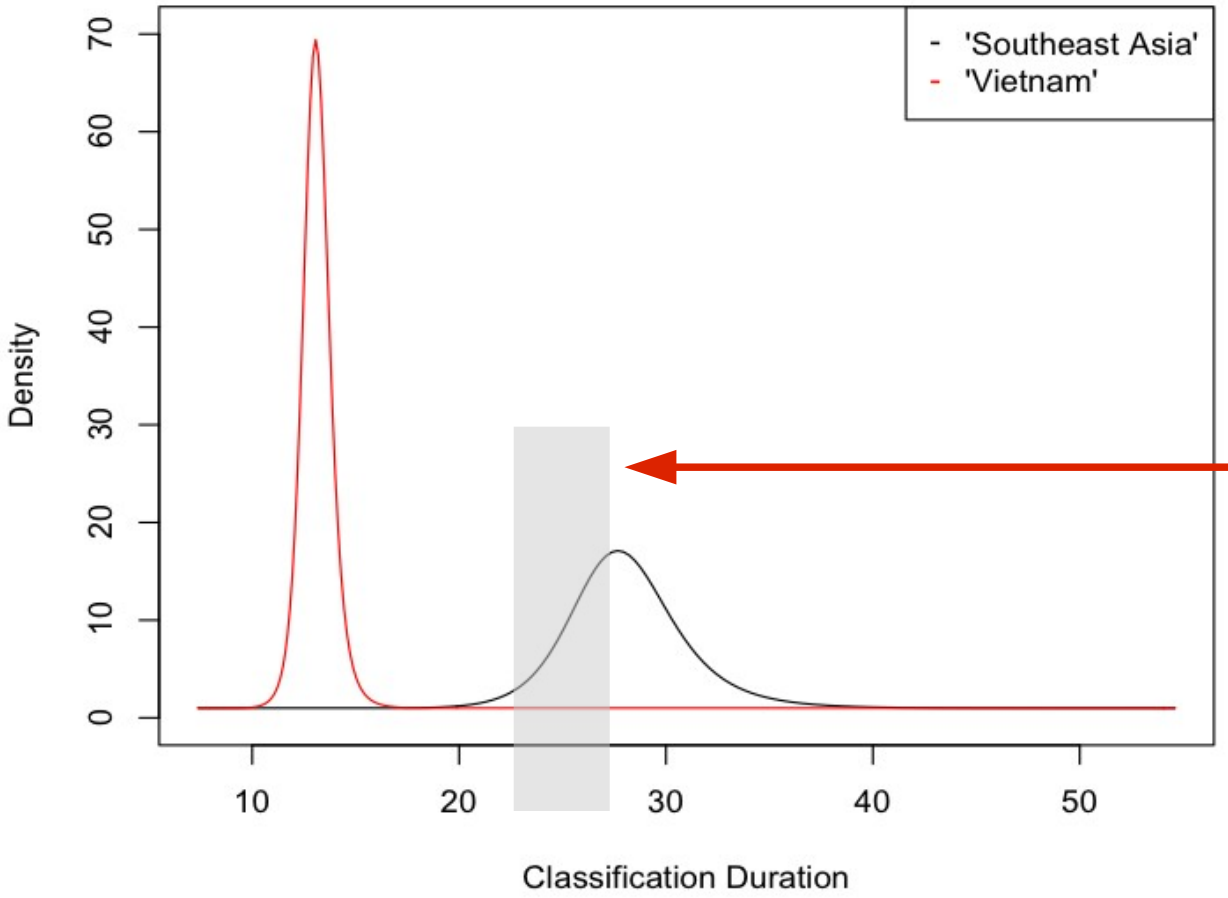

means of all LDA topics relating to Vietnam & Laos

What's Next?

- Predict durations directly from the generative model
  - Mixture vs. admixture topics
  - Supervised topic modeling
  - Unseen content
- Subject matter experts
- Analysis and prediction of redactions

Thanks!

Acknowledgements: B. Desmarais, A. McCallum, D. Mimno, R. Shorey

wallach@cs.umass.edu
http://www.cs.umass.edu/~wallach/