

PROBABILISTIC MACHINE LEARNING  
LECTURE 19  
EXAMPLE: TOPIC MODELS


Philipp Hennig

28 June 2021

EBERHARD KARLS  
UNIVERSITÄT  
TÜBINGEN



FACULTY OF SCIENCE  
DEPARTMENT OF COMPUTER SCIENCE  
CHAIR FOR THE METHODS OF MACHINE LEARNING



#	date	content	Ex	#	date	content	Ex
1	20.04.	Introduction	1	14	09.06.	Generalized Linear Models	
2	21.04.	Reasoning under Uncertainty		15	15.06.	Exponential Families	8
3	27.04.	Continuous Variables	2	16	16.06.	Graphical Models	
4	28.04.	Monte Carlo		17	22.06.	Factor Graphs	9
5	04.05.	Markov Chain Monte Carlo	3	18	23.06.	The Sum-Product Algorithm	
6	05.05.	Gaussian Distributions		19	29.06.	Example: Modelling Topics	10
7	11.05.	Parametric Regression	4	20	30.06.	Mixture Models	
8	12.05.	Learning Representations		21	06.07.	EM	11
9	18.05.	Gaussian Processes	5	22	07.07.	Variational Inference	
10	19.05.	Understanding Kernels		23	13.07.	Fast Variational Inference	12
11	26.05.	Gauss-Markov Models		24	14.07.	Kernel Topic Models	
12	25.05.	An Example for GP Regression	6	25	20.07.	Outlook	
13	08.06.	GP Classification	7	26	21.07.	Revision	

# The Toolbox

---

## Framework:

$$\int p(x_1, x_2) dx_2 = p(x_1) \qquad p(x_1, x_2) = p(x_1 | x_2)p(x_2) \qquad p(x | y) = \frac{p(y | x)p(x)}{p(y)}$$


---

## Modelling:

- ▶ graphical models
- ▶ Gaussian distributions
- ▶ (deep) learnt representations
- ▶ Kernels
- ▶ Markov Chains
- ▶ Exponential Families / Conjugate Priors
- ▶ Factor Graphs & Message Passing

## Computation:

- ▶ Monte Carlo
- ▶ Linear algebra / Gaussian inference
- ▶ maximum likelihood / MAP
- ▶ Laplace approximations
- ▶

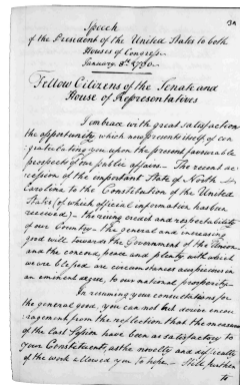
the goal for (most of) the rest of the course:  
Build a Model of History



**[The President] shall from time to time give to the Congress Information of the State of the Union, and recommend to their Consideration such Measures as he shall judge necessary and expedient.**

Article II, §3 of the US Constitution

- ▶ Delivered annually since 1790
- ▶ Summarizes affairs of the US federal government
- ▶ historically delivered in writing, generally spoken since 1982,
- ▶ on radio since 1923, TV since 1947, in the evenings since 1965, webcast since 2002
- ▶ the inaugural SotU of a new president typically has a different tone





**[The President] shall from time to time give to the Congress Information of the State of the Union, and recommend to their Consideration such Measures as he shall judge necessary and expedient.**

Article II, §3 of the US Constitution

- ▶ Delivered annually since 1790
- ▶ Summarizes affairs of the US federal government
- ▶ historically delivered in writing, generally spoken since 1982,
- ▶ on radio since 1923, TV since 1947, in the evenings since 1965, webcast since 2002
- ▶ the inaugural SotU of a new president typically has a different tone



# A Flawed but Useful Summary of (US) History

228 years of history

The SotU Addresses are not a perfect reflection of US history, but they are ...

- ▶ available in their entirety online
- ▶ available without interruption for over 200 years
- ▶ topical
- ▶ given in a reasonably similar setting, annually

Our task: Find **topics** of US history over time.

This is an **unsupervised dimensionality reduction** task.

## Disclaimer:

- ▶ This is not a course in natural language processing!
- ▶ There is an entire toolbox of models for text analysis that will not be discussed here. Some of them have probabilistic interpretation, others don't.
- ▶ The point of this exercise is to try out the tools developed in this course on a practical problem. There is no claim that this is the "best" thing to do

However, the model ultimately developed here is likely unusually expressive in its structure, and more flexible than the standard tools. Key takeaway: It does pay to spend time developing your model!

Our Goal: Build *craftware*: customized, effective and efficient solution to the learning task. Use toolboxes where they help, be willing to write our own solution where necessary.




- ▶  $D = 231$  documents (1790 – 2019; 2 in 1961 (Eisenhower & JFK))
- ▶ individual documents of length  $l_d \sim 10^3$  words
- ▶  $V \sim 10\,000$  words in vocabulary

## A few first simplifications

- ▶ there are many redundant **stop words** required for human understanding but carrying only negligible semantic information
- ▶ since we are looking to *reduce* complexity, we necessarily have to throw out a bit of structure
- ▶ e.g., usage of word is significant, but its position in the text is not crucial. We will model the texts as **Bags of Words**


# Bags of Words

A look at the data



# A Reduced Representation

low-rank decomposition



Consider a dataset  $X \in \mathbb{R}^{D \times V}$ . **Dimensionality Reduction** aims to find an **encoding**  $\phi : \mathbb{R}^V \rightarrow \mathbb{R}^K$  and a **decoding**  $\psi : \mathbb{R}^K \rightarrow \mathbb{R}^V$  with  $K \ll V$  such that the encoded representation

$$Z := \phi(X) \in \mathbb{R}^{D \times K}$$

is a *good approximation* of  $X$  in the sense that some **reconstruction loss** of  $\tilde{X} = \psi(Z)$ ,

$$\mathcal{L}(X, \psi(Z)) = \mathcal{L}(X, \psi \circ \phi(X))$$

is minimized or small. This may be done, e.g., to

- ▶ save memory
- ▶ construct a low-dimensional visualization
- ▶ “find structure”

Data:  $X \in \mathbb{R}^{D \times V} = [\mathbf{x}_1; \dots; \mathbf{x}_D]$ .

- ▶ Consider an orthonormal basis  $\{\mathbf{u}_i\}_{i=1, \dots, V}$ ,  $\mathbf{u}_i^T \mathbf{u}_j = \delta_{ij}$ . Then

$$\mathbf{x}_d = \sum_{i=1}^V (\mathbf{x}_d^T \mathbf{u}_i) \mathbf{u}_i =: \sum_{i=1}^V \alpha_{di} \mathbf{u}_i \quad X = (XU)U^T$$

- ▶ An *approximation* in  $K < D$  degrees of freedom is given by any set  $(A, \mathbf{b}, U)$  as

$$\tilde{\mathbf{x}}_d := \sum_{k=1}^K a_{dk} \mathbf{u}_k + \sum_{\ell=K+1}^V b_{\ell} \mathbf{u}_{\ell}$$

What is the *best* approximation?

Let's find  $(A, \mathbf{b}, U)$  to minimize the *square empirical risk*

$$J = \frac{1}{D} \sum_{d=1}^D \|\mathbf{x}_d - \tilde{\mathbf{x}}_d\|^2 = \frac{1}{D} \sum_{d=1}^D \sum_{v=1}^V \left[ \mathbf{x}_d - \sum_{k=1}^K a_{dk} \mathbf{u}_k - \sum_{j=K+1}^V b_j \mathbf{u}_j \right]_v^2$$

First, let's find  $a_{dk}$  and  $b_j$ : Recall  $\sum_j u_{ij} u_{kj} = \delta_{ik}$ , use  $\bar{\mathbf{x}} := \frac{1}{D} \sum_d \mathbf{x}_d$ , to find

$$\frac{\partial J}{\partial a_{de}} = \frac{2}{D} \sum_{v=1}^V \left[ \mathbf{x}_d - \sum_{k=1}^K a_{dk} \mathbf{u}_k - \sum_{j=K+1}^V b_j \mathbf{u}_j \right]_v (-u_{ev}) = \frac{2}{D} (-\mathbf{x}_d^\top \mathbf{u}_e) + \frac{2}{D} a_{de} \stackrel{!}{=} 0$$

$$\frac{\partial J}{\partial b_e} = \frac{2}{D} \sum_{d=1}^D \sum_{v=1}^V \left[ \mathbf{x}_d - \sum_{k=1}^K a_{dk} \mathbf{u}_k - \sum_{j=K+1}^V b_j \mathbf{u}_j \right]_v (-u_{ev}) = \frac{2}{D} \sum_{d=1}^D (-\mathbf{x}_d^\top \mathbf{u}_e) + 2b_e \stackrel{!}{=} 0$$

Thus  $a_{dk} = \mathbf{x}_d^\top \mathbf{u}_k$ , and  $b_j = \bar{\mathbf{x}}^\top \mathbf{u}_j$ .

With  $a_{dk} = \mathbf{x}_d^T \mathbf{u}_k$ ,  $b_j = \bar{\mathbf{x}}^T \mathbf{u}_j$ , things simplify:

$$\begin{aligned} \mathbf{x}_d - \tilde{\mathbf{x}}_d &= \mathbf{x}_d - \sum_{k=1}^K a_{dk} \mathbf{u}_k - \sum_{j=K+1}^V b_j \mathbf{u}_j = \sum_{\ell=1}^V (\mathbf{x}_d^T \mathbf{u}_\ell) \mathbf{u}_\ell - \sum_{k=1}^K (\mathbf{x}_d^T \mathbf{u}_k) \mathbf{u}_k - \sum_{j=K+1}^V (\bar{\mathbf{x}}^T \mathbf{u}_j) \mathbf{u}_j \\ &= \sum_{\ell=1}^K (\mathbf{x}_d^T \mathbf{u}_\ell) \mathbf{u}_\ell - \sum_{k=1}^K (\mathbf{x}_d^T \mathbf{u}_k) \mathbf{u}_k + \sum_{\ell=K+1}^V (\mathbf{x}_d^T \mathbf{u}_\ell) \mathbf{u}_\ell - \sum_{j=K+1}^V (\bar{\mathbf{x}}^T \mathbf{u}_j) \mathbf{u}_j \\ &= \sum_{j=K+1}^V ((\mathbf{x}_d - \bar{\mathbf{x}})^T \mathbf{u}_j) \mathbf{u}_j, \text{ so, with the sample covariance matrix } S := \frac{1}{D} \sum_{d=1}^D (\mathbf{x}_d - \bar{\mathbf{x}})(\mathbf{x}_d - \bar{\mathbf{x}})^T \\ J &= \frac{1}{D} \sum_{d=1}^D \|\mathbf{x}_d - \tilde{\mathbf{x}}_d\|^2 = \frac{1}{D} \sum_{d=1}^D \sum_{j=K+1}^V ((\mathbf{x}_d - \bar{\mathbf{x}})^T \mathbf{u}_j)^2 = \frac{1}{D} \sum_{j=K+1}^V \sum_{d=1}^D \mathbf{u}_j^T (\mathbf{x}_d - \bar{\mathbf{x}})(\mathbf{x}_d - \bar{\mathbf{x}})^T \mathbf{u}_j \\ &= \sum_{j=K+1}^V \mathbf{u}_j^T S \mathbf{u}_j \end{aligned}$$

# Maybe we can get away with linear algebra?

To find a set of *orthonormal* vectors  $u_j$  to minimize the square reconstruction error

$$J = \frac{1}{D} \sum_{d=1}^D \|\mathbf{x}_d - \tilde{\mathbf{x}}_d\|^2 = \sum_{j=K+1}^V u_j^T S u_j$$

Choose  $U$  as the eigenvectors of the sample covariance  $S := \frac{1}{D} \sum_{d=1}^D (\mathbf{x}_d - \bar{\mathbf{x}})(\mathbf{x}_d - \bar{\mathbf{x}})^T$ , and get the *best rank  $K$  reconstruction*  $\tilde{\mathbf{x}}_d$  by setting

$$\tilde{\mathbf{x}}_d := \sum_{k=1}^K a_{dk} \mathbf{u}_k + \sum_{j=K+1}^V b_j \mathbf{u}_j = \sum_{i=1}^M (\mathbf{x}_d^T \mathbf{u}_i) \mathbf{u}_i + \sum_{i=M+1}^D (\bar{\mathbf{x}}^T \mathbf{u}_i) \mathbf{u}_i$$

This yields  $J = \sum_{j=K+1}^V \lambda_j$  (where  $\lambda_j$  are the eigenvalues of  $S$ , sorted descendingly). If we first center the data  $\hat{X} = X - \mathbf{1}\bar{\mathbf{x}}^T$ , so  $\mathbf{b} = 0$ , the  $U$  are the (right) **singular vectors** of  $\hat{X} = Q\Sigma U^T$ .






Treat the loss, up to scaling, as a non-normalised negative log likelihood:

$$J = -c \cdot \log p(X | \tilde{X}) + \log Z = \frac{1}{D} \sum_{d=1}^D \|\mathbf{x}_d - \tilde{\mathbf{x}}_d\|^2$$

$$\Rightarrow p(X | \tilde{X}) = \prod_{d=1}^D \mathcal{N}(\mathbf{x}_d; \tilde{\mathbf{x}}_d, \sigma^2 I)$$

We also need to encode that we want a *low-dimensional, linear* embedding, and that the embedding should be in terms of *independent* (orthogonal) dimensions.



Thus, consider

$$\mathbf{x}_d = \mathbf{V}\mathbf{a}_d + \boldsymbol{\mu} + \varepsilon \quad \text{with } p(\mathbf{a}_d) = \mathcal{N}(0; I_K), \mathbf{V} \in \mathbb{R}^{V \times K} \text{ and } p(\varepsilon) = \mathcal{N}(0; \sigma^2)$$


with marginal likelihood (where  $C := \mathbf{V}\mathbf{V}^\top + \sigma^2 I$ )

$$p(X) = \int \prod_{d=1}^D p(\mathbf{x}_d | \mathbf{a}_d) p(\mathbf{a}_d) d\mathbf{a}_d = \prod_d \mathcal{N}(\mathbf{x}_d; \boldsymbol{\mu}, C)$$

$$\log p(X) = -\frac{DV}{2} \log(2\pi) - \frac{D}{2} \log |C| - \frac{1}{2} \sum_{d=1}^D (\mathbf{x}_d - \boldsymbol{\mu})^\top C^{-1} (\mathbf{x}_d - \boldsymbol{\mu})$$

$\bar{\mathbf{x}} = \arg \max_{\boldsymbol{\mu}} \log p(X)$ , thus the max. lik. can be written as

$$\log p(X) = -\frac{D}{2} (V \log(2\pi) + \log |C| + \text{tr}(C^{-1} S))$$






$$\log p(X) = -\frac{D}{2} (V \log(2\pi) - \log |C| + \text{tr}(C^{-1}S))$$

yields max. lik. for  $V, \sigma^2$  at [Tipping & Bishop, 1999], with  $RR^T = I_K$  and  $S = U\Lambda U^T$

$$V_{ML} = U_{1:K}(\Lambda_K - \sigma^2 I)^{1/2} R \quad \text{and} \quad \sigma_{ML}^2 = \frac{1}{V-K} \sum_{j=K+1}^V \lambda_j$$

setting  $\sigma^2, \mu, U$  this way, and  $R = I$  w.l.o.g., gives posterior

$$\begin{aligned} p(\mathbf{a}_d | \mathbf{x}_d) &= \mathcal{N}(\mathbf{a}_d; (\mathbf{V}^T \mathbf{V} + \sigma^2 I)^{-1} \mathbf{V}^T (\mathbf{x}_d - \bar{\mathbf{x}}), \sigma^2 (\mathbf{V}^T \mathbf{V} + \sigma^2 I)^{-1}) \\ &= \mathcal{N}(\mathbf{a}_d; \Lambda_K^{-1} (\Lambda_K - \sigma^2 I_K)^{1/2} U_{1:K} (\mathbf{x}_d - \bar{\mathbf{x}}), \sigma^2 \Lambda^{-1}) \end{aligned}$$



So, does it work?

---

```




1 count_vect_lsa = CountVectorizer(max_features=VOCAB_SIZE, stop_words=['000'])
2 X_count = count_vect_lsa.fit_transform(preprocessed).toarray()
3
4 U_, S_, V_T_ = np.linalg.svd(X_count, full_matrices=False)

```

---

# Latent Semantic Indexing / Principal Component Analysis

a first result on our dataset




1. tonight fight taxis faith century today enemy fellow
2. year program world new work need help america
3. dollar war program fiscal year expenditure million united
4. man law dollar business national corporation legislation labor
5. administration policy energy program continue development provide effort
6. war nation power man mexico world peace public
7. united war states american world mexico man nation
8. government people states world free shall dollar constitution
9. year free nation world increase report subject great
10. world free gold government bank note american treasury

- ▶ The singular value decomposition (SVD) minimizes  $\|X - Q\Sigma U'\|_F^2$  for orthonormal matrices  $Q \in \mathbb{R}^{D \times K}$  and  $U \in \mathbb{R}^{V \times K}$ , and a diagonal  $\Sigma \in \mathbb{R}^{K \times K}$  with positive diagonal entries (the *singular values*).
- ▶ We might naively think of  $Q$  as a mapping from documents to topics,  $U'$  from topics to words, and  $\Sigma$  as the relative strength of topics.
- ▶ However, there are several problems:
  - ▶ the matrices  $Q, U$  returned by the SVD are in general *dense*: Every document contains contributions from *every* topic, and *every* topic involves *all* words.
  - ▶ the entries in  $Q, U, \Sigma$  are hard to interpret: They do not correspond to probabilities
  - ▶ the entries of  $Q, U$  can be *negative*! What does it mean to have a negative topic?

# We need Sparsity


How about *one* topic per document?



For PCA, we allowed  $Z \in \mathbb{R}^{D \times K}$ . Maybe we need  $Z \in \{0; 1\}^{D \times K}$  and  $Z \mathbf{1}_K = \mathbf{1}_D$ ?




a supervised problem that can be solved **discriminatively** in a *linear* fashion






a supervised problem that can be solved **discriminatively** in a *nonlinear* fashion





a supervised problem that can be solved **generatively** (in a Gaussian fashion?)






## an unsupervised problem

<https://www.stat.cmu.edu/~larry/all-of-statistics/=data/faithful.dat>



Azzalini, A. and Bowman, A. W. (1990). *A look at some data on the Old Faithful geyser*. Applied Statistics 39, 357-365.








## a Gaussian mixture

$$p(\mathbf{x}_d, Z) = \prod_{d=1}^D p(z_d | \pi) p(x_d | z_d, \boldsymbol{\mu}, \Sigma) = \prod_{d=1}^D \prod_{k=1}^K \pi_k^{z_{dk}} \mathcal{N}(x_d; \boldsymbol{\mu}_k, \Sigma_k)^{z_{dk}}$$




# A Gaussian Mixture isn't quite right yet

word counts aren't real-valued



# A Mixture of Probabilities?

Desiderata



- ▶ topics should be probabilities:  $p(\mathbf{x}_d | k) = \prod_{v=1}^V \theta_{kv}^{x_{dv}}$
- ▶ but documents should have *several* topics! Let  $\pi_{dk}$  be the *probability* to draw a word from topic  $k$