

Chapter 2

Non-Negative Matrix Factorization


Part 2: Variations & applications



Geometry of NMF

Geometry of NMF

NMF factors
Data points
Convex cone
Projections



Sparsity in NMF

Sparsity: desiderata

- Sparse factor matrices are often preferred
 - Simpler to interpret (zeroes can be ignored)
 - Agrees with our intuition on parts of whole
 - Faster computations, less space
- NMF is sometimes claimed to automatically yield sparse factors
 - In practice, this is often not the case

Enforcing sparsity

- A common solution: change the target function to minimize

The diagram shows two boxes at the top: 'Parameters' in blue and 'Regularizers' in green. A blue arrow points from 'Parameters' to the coefficient α in the equation. A green arrow points from 'Regularizers' to the coefficient β in the equation. A second blue arrow points from 'Parameters' to the matrix \mathbf{W} in the $\text{sparsity}(\mathbf{W})$ term.

$$\frac{1}{2} \|\mathbf{A} - \mathbf{W}\mathbf{H}\|_F^2 + \alpha \cdot \text{sparsity}(\mathbf{W}) + \beta \cdot \text{sparsity}(\mathbf{H})$$

- How to define sparsity?
- Naïve: $\text{sparsity}(\mathbf{W}) = \text{nnz}(\mathbf{W})/\text{size}(\mathbf{W})$
 - Non-convex, non-nice to optimize directly

Frobenius regularizer

- A.k.a. Tikhonov regularizer
- Minimize $\frac{1}{2} \left(\|\mathbf{A} - \mathbf{W}\mathbf{H}\|_F^2 + \alpha \|\mathbf{W}\|_F^2 + \beta \|\mathbf{H}\|_F^2 \right)$
 - Doesn't help much with sparsity
 - Used to impose smoothness

ALS-NMF with Frobenius regularizers

- The update rules for ALS with Frobenius regularizer are

$$\mathbf{W} \leftarrow [(\mathbf{A}\mathbf{H}^T)(\mathbf{H}\mathbf{H}^T + \alpha\mathbf{I})^{-1}]_+$$

$$\mathbf{H} \leftarrow [(\mathbf{W}^T\mathbf{W} + \beta\mathbf{I})^{-1}(\mathbf{W}^T\mathbf{A})]_+$$

- Uses the fact that $\mathbf{X}^+ = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T$ if \mathbf{X} has full column rank (homework)

L_1 (Lasso) regularizer

- Using L_1 based instead of L_2 based regularizer helps obtaining sparse solutions

$$\frac{1}{2} \|\mathbf{A} - \mathbf{W}\mathbf{H}\|_F^2 + \alpha \sum_{i,j} \mathbf{W}_{ij} + \beta \sum_{i,j} \mathbf{H}_{ij}$$

- Larger values of α and β yield sparse solutions (e.g. $\alpha, \beta \in [0.01, 0.5]$)
- Still no guarantees on sparsity

ALS-NMF with Lasso

- The update rules are

$$\mathbf{W} \leftarrow [(\mathbf{A}\mathbf{H}^T - \alpha \mathbf{1}_{n \times k})(\mathbf{H}\mathbf{H}^T)^{-1}]_+$$

$$\mathbf{H} \leftarrow [(\mathbf{W}^T \mathbf{W})^{-1}(\mathbf{W}^T \mathbf{A} - \beta \mathbf{1}_{k \times m})]_+$$

- $\mathbf{1}_{n \times k}$ is n -by- k matrix of all 1s
- Requires columns of \mathbf{W} to be normalized to unit L_1 after each update

$$\mathbf{W}_{ij} \leftarrow \mathbf{W}_{ij} / \sum_i \mathbf{W}_{ij}$$

Hoyer's sparse NMF

- Hoyer (2004) considers the following sparsity function for n -dimensional vector \mathbf{x}

$$\text{sparsity}(\mathbf{x}) = \frac{\sqrt{n} - (\sum_i |\mathbf{x}_i|) / \sqrt{\sum_i \mathbf{x}_i^2}}{\sqrt{n} - 1}$$

- $\text{sparsity}(\mathbf{x}) = 1$ iff $\text{nnz}(\mathbf{x}) = 1$
- $\text{sparsity}(\mathbf{x}) = 0$ iff $|\mathbf{x}_i| = |\mathbf{x}_j|$ for all i, j

Hoyer's sparse NMF

- Hoyer's algorithm obtains an NMF using factor matrices with user-defined level of sparsity
- After every update, the columns of **W** and rows of **H** are updated s.t. their L_2 is constant and L_1 is set to desired level of sparsity

Getting the desired level of sparsity

- Set $\mathbf{s} = \mathbf{x} + (L_1 - \sum_i \mathbf{x}_i)/\text{dim}(\mathbf{x})$ and $Z = \{\}$ Fix L_1
- **repeat**
 - Set $m_i = L_1/(\text{dim}(\mathbf{x}) - \text{size}(\mathbf{Z}))$ for $i \notin Z$ and $m_i = 0$ o/w Fix L_2
 - Set $\mathbf{s} = \mathbf{m} + \alpha(\mathbf{s} - \mathbf{m})$ where α is s.t. $\|\mathbf{s}\|_2 = L_2$
 - If $\mathbf{s} \geq 0$, **return** \mathbf{s} Are we done?
 - Set $Z = Z \cup \{i : \mathbf{s}_i < 0\}$ and $\mathbf{s}_i = 0$ for all $i \in Z$
 - Set $c = (\sum_i s_i - L_1)/(\text{dim}(\mathbf{x}) - \text{size}(Z))$ Truncate negative values and fix L_1 again
 - Set $s_i = s_i - c$ for all $i \notin Z$

Other forms of NMF

Normalized NMF

- Columns of **W** (and/or rows of **H**) should be normalized to sum to unity
 - Stability of the solution and interpretability
- If only **W** (or **H**) is normalized, the weights can be pushed to the other matrix
- To normalize both, use k -by- k diagonal **Σ** s.t.
$$\sigma_{ii} = ||\mathbf{W}_i||_1 \times ||(\mathbf{H}^T)_i||_1$$
 - Normalized NMF: **$W\Sigma H$**

Semi-orthogonal NMF

- In **semi-orthogonal NMF** we restrict \mathbf{H} to row-orthogonal:
minimize $\|\mathbf{A} - \mathbf{WH}\|_F$ s.t. $\mathbf{HH}^T = \mathbf{I}$ and \mathbf{W} and \mathbf{H} are nonnegative
 - Solutions are unique (up to permutations)
 - The problem is “equivalent” to k -means
 - In the sense that the optimal solutions have the same value

NMF and clustering

- In k -means, we minimize

$$\sum_{j=1}^k \sum_{i \in C_j} \|\mathbf{a}_i - \boldsymbol{\mu}_j\|_2^2 = \sum_{j=1}^k \sum_{i=1}^n \mathbf{G}_{ij} \|\mathbf{a}_i - \boldsymbol{\mu}_j\|_2^2$$

- $\boldsymbol{\mu}_j$ is the centroid of the j th cluster C_j

- \mathbf{G} is n -by- k **cluster assignment matrix**

- $\mathbf{G}_{ij} = 1$ if $i \in C_j$ and 0 otherwise

- Equivalently: $\|\mathbf{A} - \mathbf{GM}\|_F^2$

Type of NMF if
 \mathbf{A} is nonnegative!

- \mathbf{M} is k -by- m containing the centroids as its rows

Orthogonal tri-factor NMF

- We can find NMF where both \mathbf{W} and \mathbf{H} are (column/row) orthogonal
 - Often too restrictive; cannot handle different scales
- In **orthogonal nonnegative tri-factorization** we add third non-negative matrix \mathbf{S} :
minimize $\|\mathbf{A} - \mathbf{WSH}\|_F$ s.t. $\mathbf{W}^T \mathbf{W} = \mathbf{I}$, $\mathbf{H} \mathbf{H}^T = \mathbf{I}$, and all matrices are non-negative

More on tri-factorization

- **S** does not have to be square
 - **W** is n -by- k , **S** is k -by- l , **H** is l -by- m
 - Different number of row and column factors
- If orthogonal NMF “clusters” columns of **A** , this “bi-clusters” rows and columns simultaneously

Computing semi-orthogonal NMF

- If \mathbf{H} has to be orthogonal, either
 - update as usual and set after every iteration
$$\mathbf{H} \leftarrow [\mathbf{H}\mathbf{H}^T]^{-1/2}\mathbf{H} ; \text{ or}$$
 - update $\mathbf{H}_{ij} \leftarrow \mathbf{H}_{ij} \sqrt{\frac{(\mathbf{W}^T \mathbf{A})_{ij}}{(\mathbf{W}^T \mathbf{A} \mathbf{H}^T \mathbf{H})_{ij}}}$
- \mathbf{W} is updated as usual (w/o constraints)
 - If \mathbf{W} needs to be orthogonal, the update rules are changed accordingly

Computing the orthogonal tri-factorization

- The update rules for orthogonal tri-factorization are

$$\mathbf{H}_{ij} \leftarrow \mathbf{H}_{ij} \sqrt{\frac{((\mathbf{WS})^T \mathbf{A})_{ij}}{((\mathbf{WS})^T \mathbf{A} \mathbf{H}^T \mathbf{H})_{ij}}}$$

$$\mathbf{W}_{ij} \leftarrow \mathbf{W}_{ij} \sqrt{\frac{(\mathbf{A}(\mathbf{SH})^T)_{ij}}{(\mathbf{WW}^T \mathbf{A}(\mathbf{SH})^T)_{ij}}}$$

$$\mathbf{S}_{ij} \leftarrow \mathbf{S}_{ij} \sqrt{\frac{(\mathbf{W}^T \mathbf{A} \mathbf{H}^T)_{ij}}{(\mathbf{W}^T \mathbf{W} \mathbf{S} \mathbf{H} \mathbf{H}^T)_{ij}}}$$

Other optimization functions

Kullback–Leibler divergence

- The **Kullback–Leibler divergence** of Q from P , $D_{\text{KL}}(P||Q)$, measures the expected number of **extra** bits required to code samples from P when using a code optimized for Q

$$D_{\text{KL}}(P||Q) = \sum_i P(i) \ln \frac{P(i)}{Q(i)}$$

- P and Q are probability distributions
- Non-negative and **non-symmetric**

Generalized KL-divergence and matrix factorizations

- The standard KL-divergence requires P and Q be probability distributions (e.g. $\sum_i P(i) = 1$)
- The **generalized KL-divergence** (or **I-divergence**) removes this requirement:

$$D_{\text{GKL}}(P||Q) = \sum_i \left(P(i) \ln \frac{P(i)}{Q(i)} - P(i) + Q(i) \right)$$

- In NMF, $P = \mathbf{A}$ and $Q = \mathbf{WH}$:

$$D_{\text{GKL}}(\mathbf{A}||\mathbf{WH}) = \sum_{i,j} \left(\mathbf{A}_{ij} \ln \frac{\mathbf{A}_{ij}}{(\mathbf{WH})_{ij}} - \mathbf{A}_{ij} + (\mathbf{WH})_{ij} \right)$$

KL v.s. GKL in NMF

- KL requires **A** to be considered as a probability distribution
 - $\sum_{i,j} \mathbf{A}_{i,j} = 1$ (or row/column normalization)
 - **WH** should be normalized the same way
- GKL only requires non-negativity
 - But inherently assumes integer data
 - Looses a bit of the probability interpretation

NMF for GKL

- The update rules for multiplicative GKL NMF algorithm are

$$\mathbf{H}_{kj} \leftarrow \mathbf{H}_{kj} \frac{\sum_{i=1}^n \mathbf{W}_{ik} (\mathbf{A}_{ij} / (\mathbf{W}\mathbf{H})_{ij})}{\sum_{i=1}^n \mathbf{W}_{ik}}$$

$$\mathbf{W}_{ik} \leftarrow \mathbf{W}_{ik} \frac{\sum_{j=1}^m (\mathbf{A}_{ij} / (\mathbf{W}\mathbf{H})_{ij}) \mathbf{H}_{kj}}{\sum_{j=1}^m \mathbf{H}_{kj}}$$


- The columns of \mathbf{W} are normalized to sum to unity after every iteration

Applications of NMF


Hand-written digits




(a) PCA - $J = 20$



(b) NMF - $J = 20$



(c) PCA - $J = 50$




(d) NMF - $J = 50$

Separation of various spectra

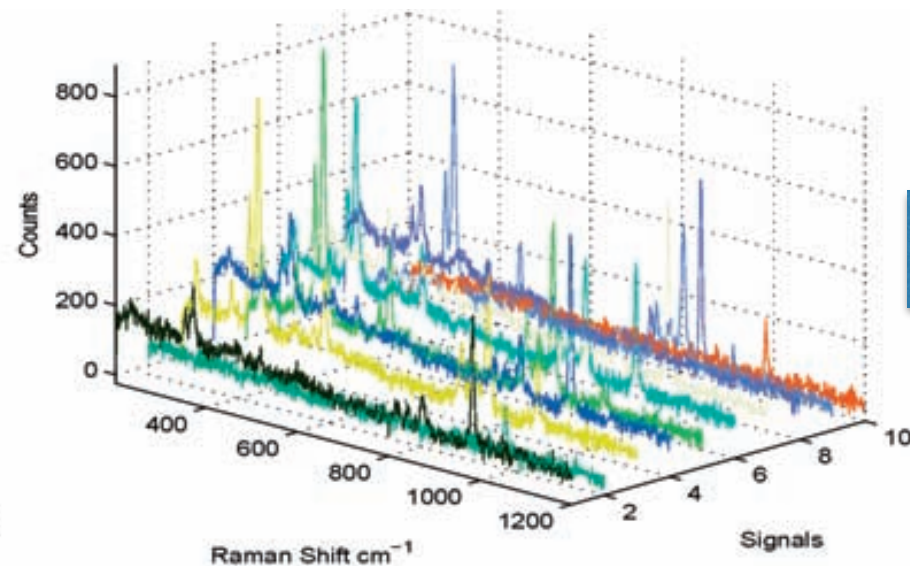
- In spectroscopy imaging we often have multiple observations of signals over some spectrum
- Observations-by-spectrum non-negative matrix
- The signals constitute an additive mixture of “pure” signals

Raman spectroscopy

Ground
truth




(a) Epsomite and Demantoid spectra




(b) Ten components of 256 noisy mixtures

Observation

Estimated



(c) Estimated Raman spectra



(d) Smoothed spectra

Estimated

Text mining and pLSA

- Consider a document-term matrix **A**
 - a_{ij} is the number of times term j appears in document i

Can we find these topics automatically?

A =

	Environmet			democrat	republican
	air	water	pollution		
doc 1	3	2	8	0	0
doc 2	1	4	12	0	0
doc 3	0	0	0	10	11
doc 4	0	0	0	8	5
doc 5	1	1	1	1	1

Politics

The idea

- Normalized \mathbf{A}' that sums to 1 can be considered as a probability distribution

$$P(d, w) = \mathbf{A}'_{d,w}$$

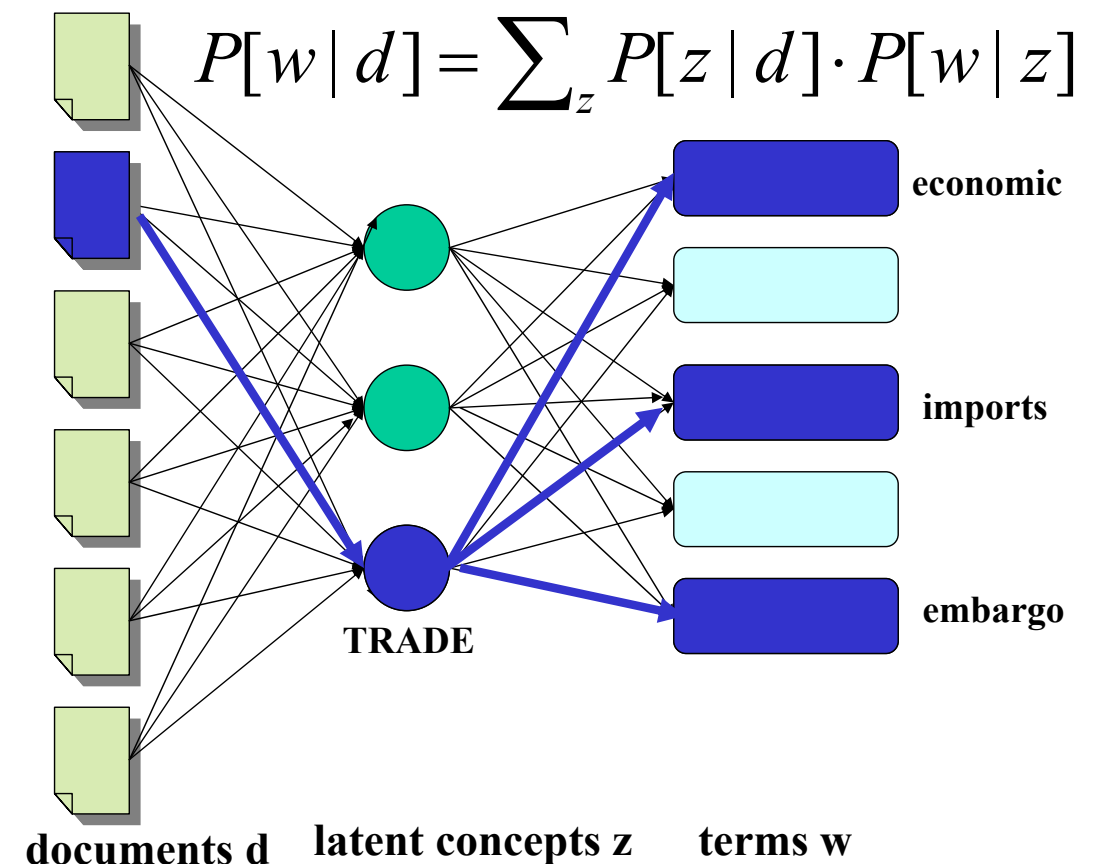
- Model with topics:

$$P(d, w) = \sum_k P(k)P(d | k)P(w | k)$$

$$P(w | d) = \sum_z P(w | z)P(z | d)$$

Generative process


- Pick a document according to $P(d)$
- Select a topic according to $P(z | d)$
- Select a word according to $P(w | z)$



pLSA as NMF

- In the NMF version of the **probabilistic latent semantic analysis** (pLSA) we are given
 - documents-by-terms matrix **A** and rank r
- We have to find
 - n -by- r non-negative **W** (columns sum to unity)
 - r -by- r diagonal non-negative **Σ**
 - r -by- m non-negative **H** (rows sum to unity)
- Minimizing $D_{\text{GKL}}(\mathbf{A} \parallel \mathbf{W}\mathbf{\Sigma}\mathbf{H})$

Geometry of pLSA



pLSA example

air wat pol dem rep

0.04	0.03	0.12	0	0
0.01	0.06	0.17	0	0
0	0	0	0.14	0.16
0	0	0	0.12	0.07
0.01	0.01	0.01	0.01	0.01

A

Here, A is normalized

0.39	0
0.52	0
0	0.58
0	0.36
0.09	0.06

W

How strong
the topic is
in the document?

0.48	0
0	0.52

Σ

Overall
frequency

air wat pol dem rep

0.15	0.21	0.64	0	0
0	0	0	0.53	0.47

H

How strong the
word is in the
topic?

NMF algorithm for pLSA

- Compute \mathbf{W} and \mathbf{H} using GKL NMF algorithm
- Normalize columns (rows) of \mathbf{W} (\mathbf{H}) and put the multipliers to Σ
 - Normalize Σ to sum to unity
- Real implementations would require tempering to avoid over-fitting

pLSA example

- Concepts (10 of 128) extracted from Science Magazine articles (12K)

$P(w z)$	universe	0.0439	drug	0.0672	cells	0.0675	sequence	0.0818	years	0.156
	galaxies	0.0375	patients	0.0493	stem	0.0478	sequences	0.0493	million	0.0556
	clusters	0.0279	drugs	0.0444	human	0.0421	genome	0.033	ago	0.045
	matter	0.0233	clinical	0.0346	cell	0.0309	dna	0.0257	time	0.0317
	galaxy	0.0232	treatment	0.028	gene	0.025	sequencing	0.0172	age	0.0243
	cluster	0.0214	trials	0.0277	tissue	0.0185	map	0.0123	year	0.024
	cosmic	0.0137	therapy	0.0213	cloning	0.0169	genes	0.0122	record	0.0238
	dark	0.0131	trial	0.0164	transfer	0.0155	chromosome	0.0119	early	0.0233
	light	0.0109	disease	0.0157	blood	0.0113	regions	0.0119	billion	0.0177
	density	0.01	medical	0.00997	embryos	0.0111	human	0.0111	history	0.0148
$P(w z)$	bacteria	0.0983	male	0.0558	theory	0.0811	immune	0.0909	stars	0.0524
	bacterial	0.0561	females	0.0541	physics	0.0782	response	0.0375	star	0.0458
	resistance	0.0431	female	0.0529	physicists	0.0146	system	0.0358	astrophys	0.0237
	coli	0.0381	males	0.0477	einstein	0.0142	responses	0.0322	mass	0.021
	strains	0.025	sex	0.0339	university	0.013	antigen	0.0263	disk	0.0173
	microbiol	0.0214	reproductive	0.0172	gravity	0.013	antigens	0.0184	black	0.0161
	microbial	0.0196	offspring	0.0168	black	0.0127	immunity	0.0176	gas	0.0149
	strain	0.0165	sexual	0.0166	theories	0.01	immunology	0.0145	stellar	0.0127
	salmonella	0.0163	reproduction	0.0143	aps	0.00987	antibody	0.014	astron	0.0125
	resistant	0.0145	eggs	0.0138	matter	0.00954	autoimmune	0.0128	hole	0.00824

Source: Thomas Hofmann, Tutorial at ADFOCS 2004

pLSA applications

- Topic modeling
- Clustering documents and terms
- Information retrieval
 - Similar to LSA/LSI
- Generalizes better than LSA
 - But outperformed by **Latent Dirichlet Allocation** (LDA)

NMF summary

- Parts-of-whole interpretation
 - Often easier/more appropriate than SVD
- Hard to compute and non-unique
 - Local updates (multiplicative, gradient, ALS)
- Many applications and specific variations
- Still under very active research

Literature

- Berry, Browne, Langville, Pauca & Plemmons (2007): *Algorithms and applications for approximate nonnegative matrix factorization*. Comput. Stat. Data Anal. 52, pp. 155–173
- Hoyer (2004): *Non-negative matrix factorizations with sparseness constraints*. J. Mach. Learn. Res. 5, pp. 1457–1469
- Ding, Li, Peng & Park (2006): *Orthogonal nonnegative matrix tri-factorizations for clustering*. In KDD '06, pp. 126–135
- Cichocki, Zdunek, Phan & Amari (2009): *Nonnegative matrix and tensor factorizations*. John Wiley & Sons Chichester, UK