

LMP 1210 - Basic principles of machine learning in biomedical research

Rahul G. Krishnan

Canada CIFAR AI Chair

Tier II Canada Research Chair
in Computational Medicine

Unsupervised Learning vs Supervised Learning

	<i>Supervised Learning</i>	<i>Unsupervised Learning</i>
<i>Discrete</i>	classification or categorization	clustering
<i>Continuous</i>	regression	dimensionality reduction

Why dimension reduction?

High-dimensional data is everywhere!


- High-Dimensions = Lot of Features

Document classification

Features per document =
thousands of words/unigrams



Why dimension reduction?



Curse of Dimensionality


Why dimension reduction?

Useful for:

- Visualization
- More efficient use of resources (e.g., time, memory, communication)
- Statistical: fewer dimensions → better generalization
- Noise removal (improving data quality)
- Further processing by machine learning algorithms

PCA: Principle Component Analysis


What is PCA: Unsupervised technique for extracting variance structure from high dimensional datasets.



- PCA is an orthogonal projection or transformation of the data into a (possibly lower dimensional) subspace so that the variance of the projected data is maximized.


PCA: Principle Component Analysis

Intrinsically lower dimensional than the dimension of the ambient space.



Only one relevant feature


If we rotate data, again only one coordinate is more important.



Both features are relevant

Question: Can we transform the features so that we only need to preserve one latent feature?


PCA: Principle Component Analysis




In case where data lies on or near a low d -dimensional linear subspace, axes of this subspace are an effective representation of the data.

Identifying the axes is known as [Principal Components Analysis](#), and can be obtained by using classic matrix computation tools (Eigen or Singular Value Decomposition).


Example: 2D Gaussian



Example: 2D Gaussian





Example: 2D Gaussian



How PCA?

Intuition: Maximizing the Variances

- Consider the two projections below
- Which maximizes the variance?



How PCA?

Intuition: Maximizing the Variances

- Given data $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$, compute covariance matrix Σ

$$\Sigma = \frac{1}{m} \sum_{i=1}^m (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T \quad \text{where} \quad \bar{\mathbf{x}} = \frac{1}{m} \sum_{i=1}^m \mathbf{x}_i$$

- PCA** basis vectors = the eigenvectors of Σ

We get the eigenvectors using an eigendecomposition.

Power iteration (Von Mises iteration is a standard algorithm for this)

- Larger eigenvalue \Rightarrow more important eigenvectors

More on Co-Variance

- Covariance as a measure of how much **each of the dimensions** vary from the mean **with respect to each other**.
- What is the interpretation of covariance calculations?
e.g.: 2 dimensional data set
x: number of hours studied for a subject
y: marks obtained in that subject
covariance value is say: 104.53
what does this value mean?

More on Co-Variance

- Covariance as a measure of how much **each of the dimensions** vary from the mean **with respect to each other**.
- A positive value of covariance indicates **both dimensions increase or decrease together** e.g. as the number of hours studied increases, the marks in that subject increase.
- A negative value indicates **while one increases the other decreases**, or vice-versa
- If covariance is zero: the two dimensions are **independent** of each other e.g. heights of students vs the marks obtained in a subject

More on Co-Variance

- By finding the **eigenvalues and eigenvectors** of the covariance matrix, we find that the eigenvectors with the largest eigenvalues correspond to the dimensions that have the strongest correlation in the dataset.
- This is the principal component.
- PCA is a useful statistical technique that has found application in:
 - fields such as face recognition and image compression
 - finding patterns in data of high dimension.

PCA Summary

Goal: Find r -dim projection that best preserves variance

1. Compute mean vector μ and covariance matrix Σ of original points
2. Compute eigenvectors and eigenvalues of Σ
3. Select top r eigenvectors
4. Project points onto subspace spanned by them:


$$y = A(x - \mu)$$

where y is the new point, x is the old one,
and the rows of A are the eigenvectors

How many PCs

- For n original dimensions, sample covariance matrix is $n \times n$, and has up to n eigenvectors. So n PCs.
- Where does dimensionality reduction come from?

Can *ignore* the components of lesser significance.



You do lose some information, but if the eigenvalues are small, you don't lose much

- n dimensions in original data
- calculate n eigenvectors and eigenvalues
- choose only the first p eigenvectors, based on their eigenvalues

Visualization of data using t-SNE


SNE - Stochastic Neighbor Embedding

Van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).


- t-SNE is an alternative dimensionality reduction algorithm.
- PCA tries to find a **global** structure
 - ▶ Low dimensional subspace
 - ▶ Can lead to local inconsistencies
 - ▶ Far away point can become nearest neighbors
- t-SNE tries to preserve **local** structure
 - ▶ Low dimensional neighborhood should be the same as original neighborhood.
- Unlike PCA almost only used for visualization
 - ▶ No easy way to embed new points

Visualization of MNIST using PCA

3 6 8 1 7 9 6 6 4 1
6 7 5 7 8 6 3 4 8 5
2 1 7 9 7 1 2 4 4 5
4 8 1 9 0 1 8 8 9 4
7 6 1 8 6 4 1 5 6 0
7 5 9 2 6 5 8 1 9 7
1 2 2 2 2 3 4 4 8 0
0 2 3 8 0 7 3 8 5 7
0 1 4 6 4 6 0 2 4 3
7 1 2 8 1 6 9 8 6 1



Visualization of MNIST using t-SNE



SNE -- Stochastic Neighboring Embedding

SNE basic idea:

- "Encode" high dimensional neighborhood information as a distribution
- Intuition: Random walk between data points.
 - ▶ High probability to jump to a close point
- Find low dimensional points such that their neighborhood distribution is similar.
- How do you measure distance between distributions?
 - ▶ Most common measure: KL divergence

SNE -- Stochastic Neighboring Embedding

- Consider the neighborhood around an input data point $\mathbf{x}_i \in \mathbb{R}^d$
- Imagine that we have a Gaussian distribution centered around \mathbf{x}_i
- Then the probability that \mathbf{x}_i chooses some other datapoint \mathbf{x}_j as its neighbor is in proportion with the density under this Gaussian
- A point closer to \mathbf{x}_i will be more likely than one further away

SNE -- Stochastic Neighboring Embedding

The $i \rightarrow j$ probability (should be familiar from A1Q2), is the probability that point \mathbf{x}_i chooses \mathbf{x}_j as its neighbor

$$P_{j|i} = \frac{\exp(-\|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\|^2 / 2\sigma_i^2)}{\sum_{k \neq i} \exp(-\|\mathbf{x}^{(i)} - \mathbf{x}^{(k)}\|^2 / 2\sigma_i^2)}$$

With $P_{i|i} = 0$

- The parameter σ_i sets the size of the neighborhood
 - ▶ Very low σ_i - all the probability is in the nearest neighbor.
 - ▶ Very high σ_i - Uniform weights.
- Here we set σ_i differently for each data point
- Results depend heavily on σ_i - it defines the neighborhoods we are trying to preserve.
- Final distribution over pairs is symmetrized: $P_{ij} = \frac{1}{2N}(P_{i|j} + P_{j|i})$

SNE -- Stochastic Neighboring Embedding

Given $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(N)} \in \mathbb{R}^D$ we define the distribution P_{ij}

Goal: Find good embedding $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)} \in \mathbb{R}^d$ for some $d < D$ (normally or 3)

How do we measure an embedding quality?

For points $\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(N)} \in \mathbb{R}^d$ we can define distribution Q similarly the same (notice no σ_i^2 and not symmetric)

$$Q_{ij} = \frac{\exp(-\|\mathbf{y}^{(i)} - \mathbf{y}^{(j)}\|^2)}{\sum_k \sum_{l \neq k} \exp(-\|\mathbf{y}^{(l)} - \mathbf{y}^{(k)}\|^2)}$$

Optimize Q to be close to P


- ▶ Minimize KL-divergence

SNE pitfall: Crowding problem


- In high dimension we have more room, points can have a lot of different neighbors
- In 2D a point can have a few neighbors at distance one all far from each other - what happens when we embed in 1D?
- This is the "crowding problem" - we don't have enough room to accommodate all neighbors.
- This is one of the biggest problems with SNE.
- t-SNE solution: Change the Gaussian in Q to a heavy tailed distribution.
 - ▶ if Q changes slower, we have more "wiggle room" to place points at.

SNE pitfall: Crowding problem

Similarity in high dimension



Similarity in low dimension



There is much **more space** in high dimensions.

t-Distributed SNE

t-Distributed Stochastic Neighbor Embedding

- Student-t Probability density $p(x) \propto (1 + \frac{x^2}{\nu})^{-(\nu+1)/2}$
 - ▶ for $\nu = 1$ we get $p(x) \propto \frac{1}{1+x^2}$
- Probability goes to zero much slower than a Gaussian.
- Can show it is equivalent to averaging Gaussians with some prior over σ^2
- We can now redefine Q_{ij} as

$$Q_{ij} = \frac{(1 + \|\mathbf{y}_i - \mathbf{y}_j\|^2)^{-1}}{\sum_k \sum_{l \neq k} (1 + \|\mathbf{y}_k - \mathbf{y}_l\|^2)^{-1}}$$

- We leave P_{ij} as is!

How t-SNE? (Optional)

Algorithm 1: Simple version of t-Distributed Stochastic Neighbor Embedding.

Data: data set $\mathcal{X} = \{x_1, x_2, \dots, x_n\}$,

cost function parameters: perplexity $Perp$,

optimization parameters: number of iterations T , learning rate η , momentum $\alpha(t)$.

Result: low-dimensional data representation $\mathcal{Y}^{(T)} = \{y_1, y_2, \dots, y_n\}$.

begin

 compute pairwise affinities $p_{j|i}$ with perplexity $Perp$ (using Equation 1)

 set $p_{ij} = \frac{p_{j|i} + p_{i|j}}{2n}$

 sample initial solution $\mathcal{Y}^{(0)} = \{y_1, y_2, \dots, y_n\}$ from $\mathcal{N}(0, 10^{-4}I)$

for $t=1$ **to** T **do**

 compute low-dimensional affinities q_{ij} (using Equation 4)

 compute gradient $\frac{\delta C}{\delta \mathcal{Y}}$ (using Equation 5)

 set $\mathcal{Y}^{(t)} = \mathcal{Y}^{(t-1)} + \eta \frac{\delta C}{\delta \mathcal{Y}} + \alpha(t) (\mathcal{Y}^{(t-1)} - \mathcal{Y}^{(t-2)})$

end

end

Gradient Descent

Breaks

UMAP: Uniform Manifold Approximation and Projection

<https://umap-learn.readthedocs.io/en/latest/>

- UMAP is a replacement for tSNE to fulfil the same role
- Conceptually very similar to tSNE, but with a couple of relevant (and somewhat technical) changes
- Practical outcome is:
 - UMAP is quite a bit quicker than tSNE
 - UMAP can preserve more global structure than tSNE*
 - UMAP can run on raw data without PCA preprocessing*
 - UMAP can allow new data to be added to an existing projection

* In theory, but possibly not in practice

UMAP: Uniform Manifold Approximation and Projection

<https://umap-learn.readthedocs.io/en/latest/>

- Instead of the single perplexity value in tSNE, UMAP defines
 - **Nearest neighbours**: the number of expected nearest neighbours – basically the same concept as perplexity
 - **Minimum distance**: how tightly UMAP packs points which are close together
- Nearest neighbours will affect the influence given to global vs local information. Min dist will affect how compactly packed the local parts of the plot are.

UMAP: Uniform Manifold Approximation and Projection

<https://umap-learn.readthedocs.io/en/latest/>


- Speed – mostly a level of maths I'm not going to get into!
 - UMAP skips a normalisation step in the calculation of high dimensional distances which speeds it up
 - In the 2D projection UMAP uses a more efficient method to shuffle the cells into their final position
 - Doesn't have to measure every cell to decide on what to move
 - Uses an algorithm which can be multi-threaded
 - Algorithm is more deterministic, allowing more data to be projected later

UMAP is better than tSNE?


2D t-SNE projection



2D UMAP projection



Original 3D Data




3D mammoth skeleton projected into 2D

tSNE:	Perplexity 2000	2h 5min
UMAP:	Nneigh 200, mindist 0.25,	3min

<https://pair-code.github.io/understanding-umap/>


UMAP is better than tSNE?



- It may perform better on more complex datasets
- It's certainly quicker


<https://pair-code.github.io/understanding-umap/>

Dimension Reduction using Deep Learning: Auto-Encoder



- An autoencoder is a special type of feed forward neural network which does the following


Dimension Reduction using Deep Learning: Auto-Encoder



$$\mathbf{h} = g(W\mathbf{x}_i + \mathbf{b})$$

- An autoencoder is a special type of feed forward neural network which does the following
- Encodes its input \mathbf{x}_i into a hidden representation \mathbf{h}

Dimension Reduction using Deep Learning: Auto-Encoder




$$\mathbf{h} = g(W\mathbf{x}_i + \mathbf{b})$$

$$\hat{\mathbf{x}}_i = f(W^*\mathbf{h} + \mathbf{c})$$

- An autoencoder is a special type of feed forward neural network which does the following
- Encodes its input \mathbf{x}_i into a hidden representation \mathbf{h}
- Decodes the input again from this hidden representation


Dimension Reduction using Deep Learning: Auto-Encoder



$$\mathbf{h} = g(W\mathbf{x}_i + \mathbf{b})$$
$$\hat{\mathbf{x}}_i = f(W^*\mathbf{h} + \mathbf{c})$$

- An autoencoder is a special type of feed forward neural network which does the following
- Encodes its input \mathbf{x}_i into a hidden representation \mathbf{h}
- Decodes the input again from this hidden representation
- The model is trained to minimize a certain loss function which will ensure that $\hat{\mathbf{x}}_i$ is close to \mathbf{x}_i (we will see some such loss functions soon)

How about a fat auto-encoder?




$$\mathbf{h} = g(W\mathbf{x}_i + \mathbf{b})$$

$$\hat{\mathbf{x}}_i = f(W^*\mathbf{h} + \mathbf{c})$$

- Let us consider the case when $\dim(\mathbf{h}) \geq \dim(\mathbf{x}_i)$
- In such a case the autoencoder could learn a trivial encoding by simply copying \mathbf{x}_i into \mathbf{h} and then copying \mathbf{h} into $\hat{\mathbf{x}}_i$

How about a fat auto-encoder?

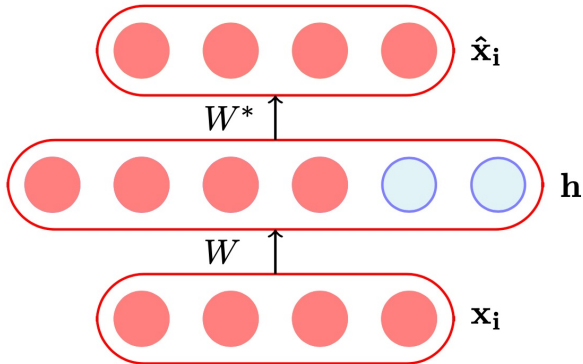


$$\mathbf{h} = g(W\mathbf{x}_i + \mathbf{b})$$

$$\hat{\mathbf{x}}_i = f(W^*\mathbf{h} + \mathbf{c})$$

- Let us consider the case when $\dim(\mathbf{h}) \geq \dim(\mathbf{x}_i)$
- In such a case the autoencoder could learn a trivial encoding by simply copying \mathbf{x}_i into \mathbf{h} and then copying \mathbf{h} into $\hat{\mathbf{x}}_i$

How about a fat auto-encoder?




$$\mathbf{h} = g(W\mathbf{x}_i + \mathbf{b})$$

$$\hat{\mathbf{x}}_i = f(W^*\mathbf{h} + \mathbf{c})$$

- Let us consider the case when $\dim(\mathbf{h}) \geq \dim(\mathbf{x}_i)$
- In such a case the autoencoder could learn a trivial encoding by simply copying \mathbf{x}_i into \mathbf{h} and then copying \mathbf{h} into $\hat{\mathbf{x}}_i$
- Such an identity encoding is useless in practice as it does not really tell us anything about the important characteristics of the data

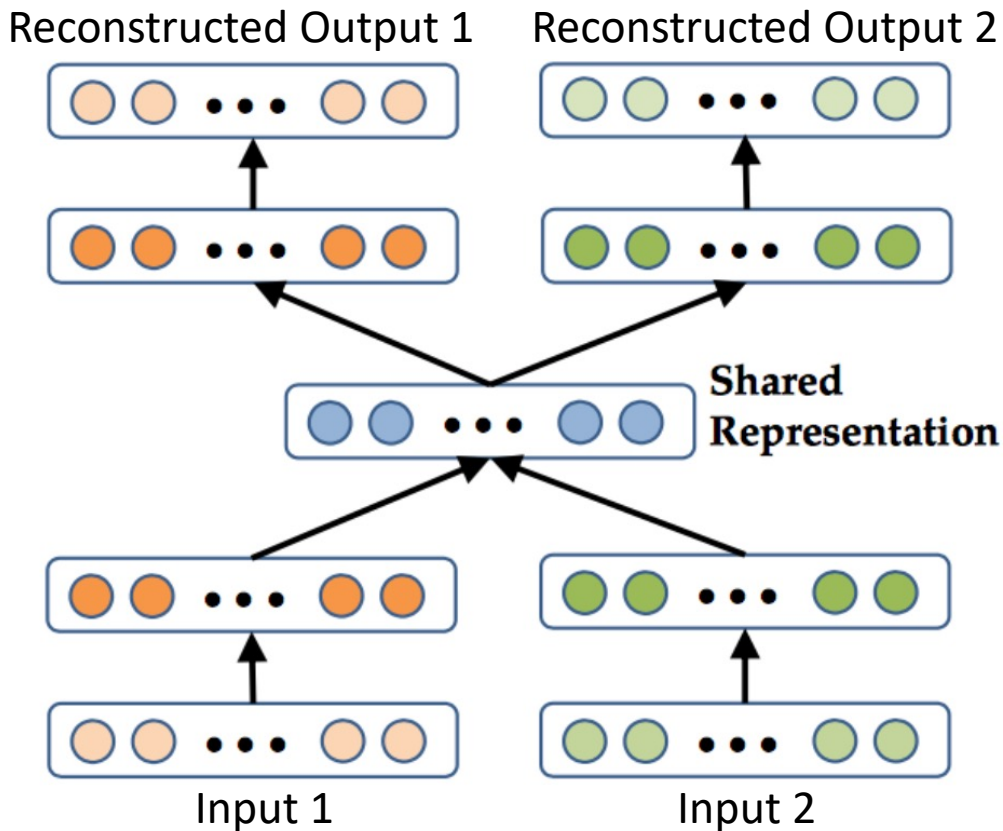
Going Deeper: Stacked Auto-encoder



How to choose encoder/decoder

- Tabular Data (e.g., clinical variables, gene expression profiles etc.) : Multi-layer Perceptron/Fully Connected Layers
- Imaging Data (e.g., MRI, CT scans etc.) : Convolutional Neural Networks
- Sequence Data (e.g., Texts, ECG etc.) : Recurrent Neural Networks

Multi-Modal Auto-encoder



Take-Aways

- Dimension Reduction is useful for removing noise, visualization, reducing computational cost, data compression.
- PCA is a linear dimension reduction method that tries to maximize the variances in low-dimensional space.
- T-SNE and UMAP are non-linear visualization approaches that aim to preserve neighboring similarities.
- Deep Auto-Encoder is a non-linear representation learning approach that aims to reconstruct the inputs.