


CS-GY 6923: Lecture 13

Semantic Embeddings, Beyond Autoencoders

NYU Tandon School of Engineering, Prof. Christopher Musco

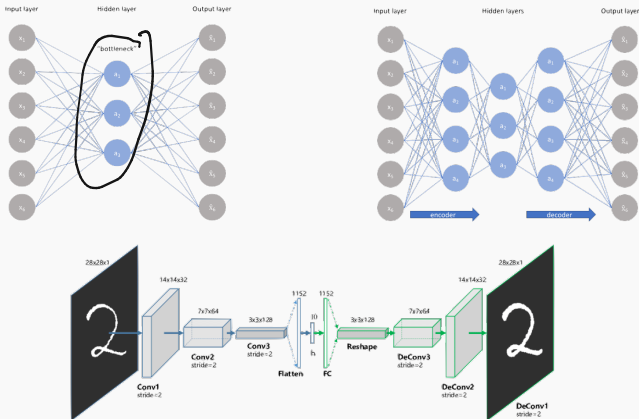
AUTOENCODER



- Let $\underline{f}_{\underline{\theta}} : \underline{\mathbb{R}^d} \rightarrow \underline{\mathbb{R}^d}$ be our model.
- Let L_{θ} be a loss function. E.g. squared loss:
 $L_{\theta}(\mathbf{x}) = (\|\mathbf{x} - \underline{f}_{\theta}(\mathbf{x})\|_2^2)$
- Train model: $\theta^* = \min_{\theta} \sum_{i=1}^n L_{\theta}(\mathbf{x}_i)$.

AUTOENCODER

Important property of autoencoders: no matter what architecture is used, there must always be a bottleneck with fewer parameters than the input. The bottleneck ensures information is “distilled” from low-level features to high-level features.




AUTOENCODER

Separately name mapping from input to bottleneck, and from bottleneck to output.

Encoder: $e : \mathbb{R}^d \rightarrow \mathbb{R}^k$


Decoder: $d : \mathbb{R}^k \rightarrow \mathbb{R}^d$

$$\underline{f(x)} = d(e(x))$$



AUTOENCODER RECONSTRUCTION

Example image reconstructions from autoencoder:



<https://www.biorxiv.org/content/10.1101/214247v1.full.pdf>

Input parameters: $d = 49152$.


Bottleneck “latent” parameters: $(k = 1024.)$

Lots of applications:

- Data compression.
- Data denoising and repair.
- Data synthesis.
- (**Feature learning.**)

PRINCIPAL COMPONENT ANALYSIS

Simple linear autoencoder:



$$f(x)^T = x^T W_1 W_2$$

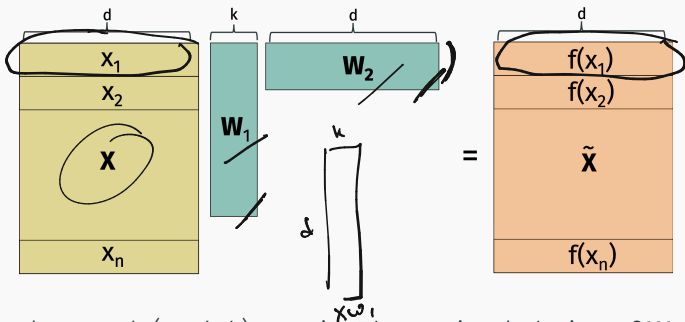
PRINCIPAL COMPONENT ANALYSIS

Given training data set x_1, \dots, x_n , let X denote our data matrix.

Let $\tilde{X} = \underline{XW_1W_2}$ denote the autoencoded data. Want to

minimize $\|X - \tilde{X}\|_F^2$.


$$x_i^T w_1, w_2$$



\tilde{X} is a low-rank (rank k) matrix. The optimal choice of W_1 and W_2 can be found using algorithms for optimal low-rank approximation.

SINGULAR VALUE DECOMPOSITION

Any matrix X can be written using its singular value decomposition:




Where $\underline{U^T U} = I$, $\underline{V^T V} = I$, and $\sigma_1 \geq \sigma_2 \geq \dots \sigma_d \geq 0$. I.e. U and V are orthogonal matrices.

Can be computed in $O(nd^2)$ time (faster with approximation algos).

SINGULAR VALUE DECOMPOSITION

We obtain an optimal autoencoder by setting $\underline{W_1} = \underline{V_k}$, $\underline{W_2} = \underline{V_k^T}$
 $\underline{f(x)}^\top = \underline{x}^\top \underline{V_k} \underline{V_k^T}$




For many natural data sets, get a good approximation even when k is chosen to be relatively small compared to d .

What do principal components and loading vectors look like?

PRINCIPAL COMPONENTS

MNIST principal components:



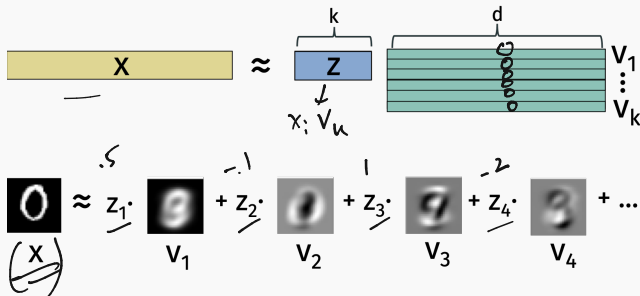
Often principal components are difficult to interpret.

LOADING VECTORS

$$x_i \approx f(x_i) = \underline{x_i V_u V_u^T}$$

What do the **loading vectors** look like?


The loading vector z for an example x contains coefficients which recombine the top k principal components v_1, \dots, v_k to approximately reconstruct x .



Provide a short “finger print” for any image x which can be used to reconstruct that image.

LOADING VECTORS: SIMILARITY VIEW

For any x with loading vector z , z_i is the inner product similarity between x and the i^{th} principal component v_i .




$$z_1 = \langle \begin{matrix} \text{O} \\ \underline{x} \end{matrix}, \begin{matrix} \text{S} \\ v_1 \end{matrix} \rangle \quad z_2 = \langle \begin{matrix} \text{O} \\ x \end{matrix}, \begin{matrix} \text{O} \\ v_2 \end{matrix} \rangle \quad z_3 = \langle \begin{matrix} \text{O} \\ x \end{matrix}, \begin{matrix} \text{O} \\ v_3 \end{matrix} \rangle \dots$$

$$z_i[j] = \langle x_i, v_j \rangle$$

LOADING VECTORS: PROJECTION VIEW

For an example x_i , the loading vector \underline{z}_i contains the coordinates in the projection space:

$$\| \underline{z}_i - \underline{z}_j \|_2 \approx \| x_i - x_j \|_2$$



Visual way of seeing what we argued last time: inner products and distances between loading vectors should approximate inner products and distances in the original space.

TERM DOCUMENT MATRIX

Word-document matrices tend to be low rank.




Documents tend to fall into a relatively small number of different categories, which use similar sets of words:

- **Financial news:** *markets, analysts, dow, rates, stocks*
- **US Politics:** *president, senate, pass, slams, twitter, media*
- **StackOverflow posts:** *python, help, convert, javascript*

LATENT SEMANTIC ANALYSIS

Latent semantic analysis = PCA applied to a word-document matrix (usually from a large corpus). One of the most fundamental techniques in natural language processing (NLP).




Each column of z corresponds to a latent “category” or “topic”. Corresponding row in Y corresponds to the “frequency” with which different words appear in documents on that topic.

LATENT SEMANTIC ANALYSIS

$$\langle \underline{x_i}, \underline{x_j} \rangle \quad \langle \underline{z_i}, \underline{z_j} \rangle$$


Similar documents have similar LSA document vectors. I.e. $\langle z_i, z_j \rangle$ is large.

- z_j provides a more compact “finger print” for documents than the long bag-of-words vectors. Useful for e.g search engines.
- Comparing document vectors is (often more effective) than comparing raw BOW features. Two documents can have $\langle z_i, z_j \rangle$ large even if they have no overlap in words. E.g. because both share a lot of words with words with another document k , or with a bunch of other documents.




EIGENFACES

Same fingerprinting idea was also important in early facial recognition systems based on “eigenfaces”:




Each image above is one of the principal components of a dataset containing images of faces.

WORD EMBEDDINGS




- $\langle y_i, z_a \rangle \approx 1$ when doc_a contains word_i.
- If word_i and word_j both appear in doc_a, then $\langle y_i, z_a \rangle \approx \langle y_j, z_a \rangle \approx 1$, so we expect $\langle y_i, y_j \rangle$ to be large.



If two words appear in the same document their, word vectors tend to point more in the same direction.

SEMANTIC EMBEDDINGS

Result: Map words to numerical vectors in a (semantically) meaningful way. Similar words map to similar vectors. Dissimilar words to dissimilar vectors.



Extremely useful “side-effect” of LSA.

Capture e.g. the fact that “great” and “excellent” are near synonyms. Or that “difficult” and “easy” are antonyms.

Word embeddings are considered a type of semantic embedding.

They can be obtained by training on a very large corpus of text (e.g. Wikipedia, Twitter, news data sets) and then used for many different tasks in the future as an initial way to convert text data to numerical data.

WORD EMBEDDINGS: MOTIVATING PROBLEM

Review 1: Very small and handy for traveling or camping. Excellent quality, operation, and appearance. +


Review 2: So far this thing is great. Well designed, compact, and easy to use. I'll never use another can opener. +

Review 3: Not entirely sure this was worth \$20. Mom couldn't figure out how to use it and it's fairly difficult to turn for someone with arthritis. -

Goal is to classify reviews as “positive” or “negative”.

WORD EMBEDDINGS

Another view on word embeddings from LSA:




We chose Z to equal $\underline{XV}_k = \underline{U}_k \underline{\Sigma}_k$ and $Y = \underline{V}_k^T$.

Could have just as easily set $\underline{Z} = \underline{U}_k$ and $\underline{Y} = \underline{\Sigma}_k \underline{V}_k^T$, so Z has orthonormal columns.

$$\underline{Z} \cdot \underline{Y} = \underline{U}_k \underline{\Sigma}_k \underline{V}_k^T$$

WORD EMBEDDINGS

Another view on word embeddings from LSA:



- $\underline{X} \approx \underline{Z}\underline{Y}$
- $\underline{X^T X} \approx \underline{Y^T Z^T Z Y} = \underline{Y^T Y}$
- So for $word_i$ and $word_j$, $\langle y_i, y_j \rangle \approx \underline{\underline{[X^T X]_{i,j}}}$.


What does the i, j entry of $\underline{\underline{X^T X}}$ represent?

WORD EMBEDDINGS

n

	car	loan	house	...	dog	cat
doc_1	0	0	1	0	1	0
doc_2	0	0	0	1	0	0
...	1	1	0	1	0	0
doc_n	1	0	0	0	0	1

term-document matrix X




What does the i, j entry of $X^T X$ represent?

of docs that both word i and word j appear in.



$=$




$\langle \mathbf{y}_i, \mathbf{y}_j \rangle$ is larger if $word_i$ and $word_j$ appear in more documents together (high value in word-word co-occurrence matrix, $\mathbf{X}^T\mathbf{X}$).
Similarity of word embeddings mirrors similarity of word context.

General word embedding recipe:

1. Choose similarity metric $k(word_i, word_j)$ which can be computed for any pair of words.
2. Construct similarity matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$ with $\mathbf{M}_{i,j} = k(word_i, word_j)$.
3. Find low rank approximation $\mathbf{M} \approx \mathbf{Y}^T\mathbf{Y}$ where $\mathbf{Y} \in \mathbb{R}^{k \times n}$.
4. Columns of \mathbf{Y} are word embedding vectors.

WORD EMBEDDINGS



How do current state-of-the-art methods differ from LSA?

- Similarity based on co-occurrence in smaller chunks of words. E.g. in sentences or in any consecutive sequences of 3, 4, or 10 words.
- Usually transformed in non-linear way. E.g.
 $k(\text{word}_i, \text{word}_j) = \frac{p(i,j)}{p(i)p(j)}$ where $p(i,j)$ is the frequency both i,j appeared together, and $p(i), p(j)$ is the frequency either one appeared.

MODERN WORD EMBEDDINGS

Computing word similarities for “window size” 4:

The girl walks to her dog to the park.
It can take a long time to park your car in NYC.
The dog park is always crowded on Saturdays.

The girl walks to her dog to the park.
It can take a long time to park your car in NYC.
The dog park is always crowded on Saturdays.


The girl walks to her dog to the park.
It can take a long time to park your car in NYC.
The dog park is always crowded on Saturdays.

	<u>dog</u>	<u>park</u> / crowded	<u>the</u>	
<u>dog</u>	0	3	0	3
<u>park</u>	2	0	1	2
<u>crowded</u>	0	1	0	0
<u>the</u>	3	2	0	0

Current state of the art models: GloVe, word2vec.

- ~~word2vec~~ was originally presented as a shallow neural network model, but it is equivalent to matrix factorization method (Levy, Goldberg 2014).
- For word2vec, similarity metric is the “point-wise mutual information”: $\log\left(\frac{p(i,j)}{p(i)p(j)}\right)$.


CAVEAT ABOUT FACTORIZATION



SVD will not return a symmetric factorization in general. In fact, if M is not positive semidefinite¹ then the optimal low-rank approximation does not have this form.

¹i.e., κ is not a positive semidefinite kernel.

CAVEAT ABOUT FACTORIZATION




- For each word i we get a left and right embedding vector w_i and y_i . It's reasonable to just use one or the other.
- If $\langle y_i, y_j \rangle$ is large and positive, we expect that y_i, y_j have similar similarity scores with other words, so they typically are still related words.
- Another option is to use as your features for a word the concatenation $[w_i, y_i]$

If you want to use word embeddings for your project, the easiest approach is to download pre-trained word vectors:

- Original gloVe website:
`https://nlp.stanford.edu/projects/glove/`
- Compilation of many sources:
`https://github.com/3Top/word2vec-api`

WORD EMBEDDINGS MATH

Lots of cool demos online for what can be done with these embeddings. E.g. “vector math” to solve analogies.




Queen - woman



London : England

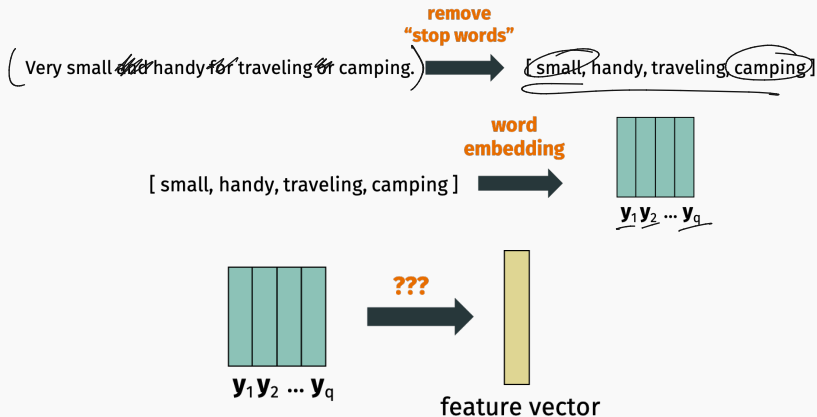
Rome : _____



(England - London + Rome)

USING WORD EMBEDDINGS


How to go from word embeddings to features for a whole sentence or chunk of text?




USING WORD EMBEDDINGS

A few simple options:

Feature vector $\mathbf{x} = \frac{1}{q} \sum_{i=1}^q \mathbf{y}_i$.




Feature vector $\mathbf{x} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_q]$.



USING WORD EMBEDDINGS


To avoid issues with inconsistent sentence length, word ordering, etc., can concatenate a fixed number of top principal components of the matrix of word vectors:



There are much more complicated approaches that account for word position in a sentence. Lots of pretrained libraries available (e.g. Facebook's `InferSent`).

SEMANTIC EMBEDDINGS

The same approach used for word embeddings can be used to obtain meaningful numerical features for any other data where there is a natural notion of similarity.




For example, the items could be nodes in a social network graph. Maybe we want to predict an individual's age, level of interest in a particular topic, political leaning, etc.

NODE EMBEDDINGS

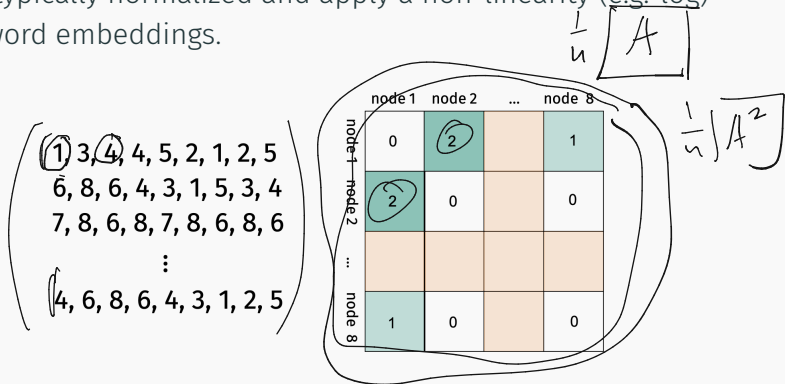


Generate random walks (e.g. “sentences” of nodes) and measure similarity by node co-occurrence frequency.



NODE EMBEDDINGS


Again typically normalized and apply a non-linearity (e.g. log) as in word embeddings.



Popular implementations: DeepWalk, Node2Vec. Again initially derived as simple neural network models, but are equivalent to matrix-factorization (Qiu et al. 2018).

BEYOND AUTOENCODERS


While they have many applications, one of the original goals of autoencoders was to learn “high level” features when we did not have substantial training data.



Typical supervised approach to transfer learning.

BEYOND AUTOENCODERS

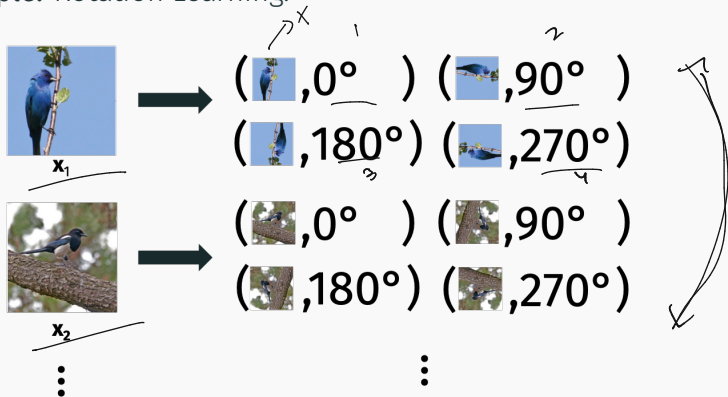
Growing realization: The types of features needed to precisely reconstruct an image (e.g. with small ℓ_2 error) don't exactly match up with the features required to understand an image.



SELF-SUPERVISED LEARNING

Automatically create a supervised learning problem with a “simpler” task than image reconstruction.

Example: Rotation Learning.



4 class learning problem. Train a supervised neural network

SELF-SUPERVISED LEARNING

Example: AlphaCode (released this year by Google/Deepmind).
Can process a description of a task and output correct code to complete the task.

Input

You are given two strings s and t , both consisting of lowercase English letters. You are going to type the string t character by character, from the first character to the last one.

When typing a character, instead of pressing the button corresponding to it, you can press the "Backspace" button. It deletes the last character you have typed among those that aren't deleted yet (or does nothing if there are no characters in the current string). For example, if s is "abcde" and you press Backspace instead of typing the first and the fourth characters, you will get the string "bcd" (the first press of Backspace deletes no character, and the second press deletes the character "a"). Another example, if s is "abcde" and you press Backspace instead of the last two letters, then the resulting text is "a".

Your task is to determine whether you can obtain the string s , if you type the string t and press "Backspace" instead of typing several (maybe zero) characters of s .

Output

For each test case, print "YES" if you can obtain the string s by typing the string t and replacing some characters with presses of "Backspace" button, or "NO" if you cannot.

You may print each letter in any case ("YES", "yes", "Yes" will all be recognized as positive answer, "NO", "no" and "No" will all be recognized as negative answer).

Input
4
ababa
ba
ababa
bb
aaa
aaaa
ababa
ababa

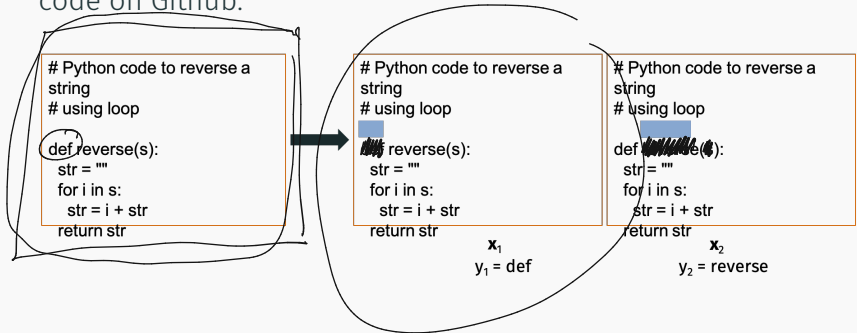
Output
YES
NO
NO
YES



```
t=int(input())
for i in range(t):
    s=input()
    t=input()
    a=[]
    b=[]
    for j in s:
        a.append(j)
    for j in t:
        b.append(j)
    a.reverse()
    b.reverse()
    c=[]
    while len(b)!=0 and len(a)!=0:
        if a[0]==b[0]:
            c.append(b.pop(0))
            a.pop(0)
        elif a[0]!=b[0] and len(a)==1:
            a.pop(0)
            a.pop(0)
        elif a[0]!=b[0] and len(a)==1:
            a.pop(0)
    if len(b)==0:
        print("YES")
    else:
        print("NO")
```


There is not a whole lot of training data for this sort of problem!

Key Component in AlphaCode: Semi-supervised learning using code on Github.



SELF-SUPERVISED LEARNING

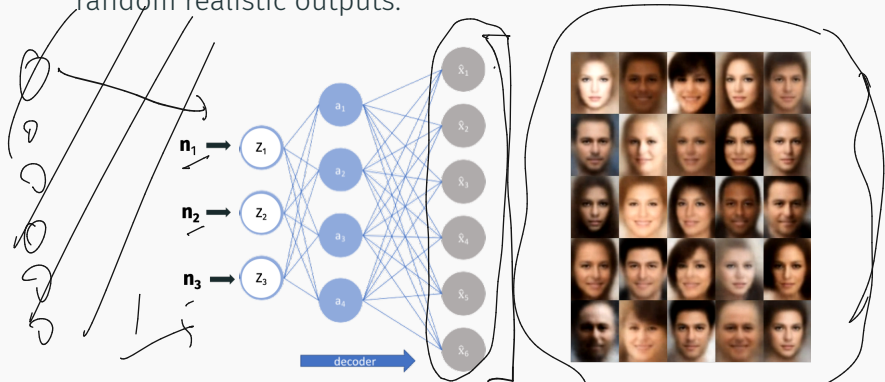
There are a lot of different possibilities!



- **Advantage:** Self-supervised learning tends to outperform autoencoders for feature learning (e.g. better performance in transfer learning tasks).
- **Disadvantage:** There is no “decoder” function, so no natural way to use these techniques for e.g. data compression, super resolution, or synthetic data generation.

Autoencoder approach to generative ML: Pretrain auto-encoder. Feed random inputs into decoder to produce random realistic outputs.

2014



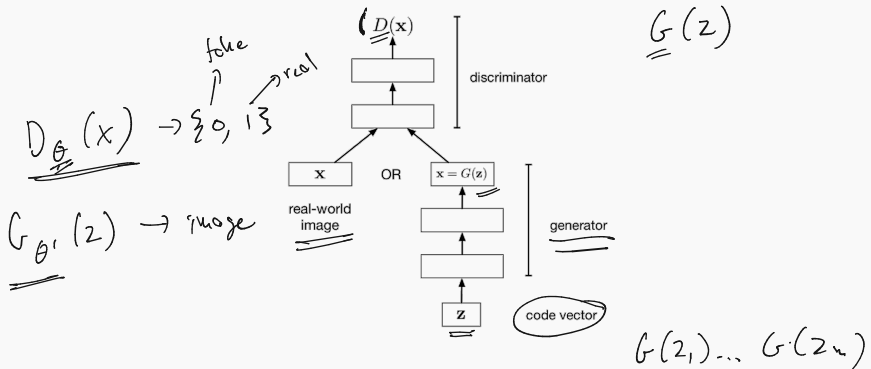
Pretty cool, but tends to produce images with immediately recognizable flaws (e.g. soft edges, high-frequency artifacts).

GENERATIVE ADVERSARIAL NETWORKS (GANS)

Lots of efforts to hand-design regularizers that penalize images that don't look realistic to the human eye.

Main idea behind GANs: Use machine learning to automatically encourage realistic looking images.

GENERATIVE ADVERSARIAL NETWORKS (GANS)

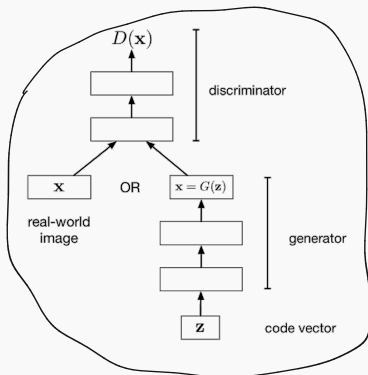


Let $\underline{x_1, \dots, x_n}$ be real images and let $\underline{z_1, \dots, z_m}$ be random code vectors. The goal of the discriminator is to output a number between $[0, 1]$ which is close to 0 if the image is fake, close to 1 if it's real.

Train weights of discriminator D_{θ} to minimize:

$$\min_{\theta} \left(\sum_{i=1}^n -\log(D_{\theta}(x_i)) + \sum_{i=1}^m -\log(1 - D_{\theta}(G_{\theta}(z_i))) \right)$$

GENERATIVE ADVERSARIAL NETWORKS (GANS)



Goal of the generator $G_{\theta'}$ is the opposite. We want to maximize:

$$\left(\max_{\theta'} \sum_{i=1}^n -\log(D_{\theta}(x_i)) + \sum_{i=1}^m -\log(1 - D_{\theta}(G_{\theta'}(z_i))) \right)$$

This is called an “adversarial loss function”. D is playing the role of the adversary.

GENERATIVE ADVERSARIAL NETWORKS (GANS)

$$\theta^*, \theta'^* \text{ solve } \min_{\theta} \max_{\theta'} \sum_{i=1}^n -\log(D_{\theta}(x_i)) + \sum_{i=1}^m -\log(1 - D_{\theta}(G_{\theta'}(z_i)))$$

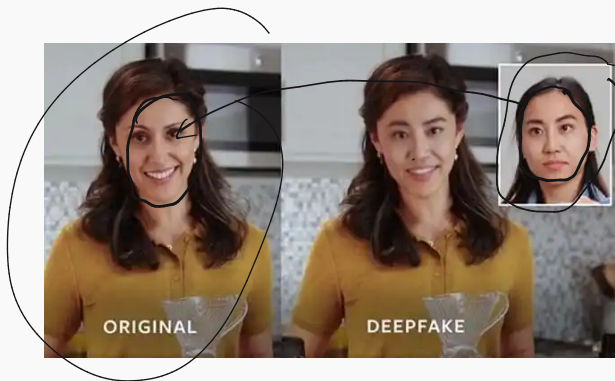
This is called a minimax optimization problem. Really tricky to solve in practice.

- Repeatedly play: Fix one of θ^* or θ'^* , train the other to convergence, repeat.
- Simultaneous gradient descent: Run a single gradient descent step for each of θ^* , θ'^* and update D and G accordingly. Difficult to balance learning rates.
- Lots of tricks (e.g. slight different loss functions) can help.

GENERATIVE ADVERSARIAL NETWORKS (GANS)



GENERATIVE ADVERSARIAL NETWORKS (GANS)



GENERATIVE ADVERSARIAL NETWORKS (GANS)

