

Lecture 15. Self-attention Mechanism and Transformers

Bao Wang

Department of Mathematics

Scientific Computing and Imaging Institute

University of Utah

Math 5750/6880, Fall 2021

Embeddings and Language Models

Converting words to vectors

- In natural language processing (NLP), our inputs are sequences of words, but deep learning needs vectors.
- How to convert words to vectors?
- Simplest idea: one-hot encoding.

One-hot encoding

Vocabulary

index:	Word:
0	aardvark
1	able
...	...
2409	black
2410	bling
...	...
3202	candid
3203	cast
3204	cat
...	...
5281	is
5282	island
...	...
8676	the
8677	thing
...	...
9999	zombie

the

cat

is

black

Problems with one-hot encoding

- Scales poorly with vocabulary size.
- Very high-dimensional sparse vectors: neural network operations work poorly.
- Violates what we know about word similarity (e.g. “run” is as far away from “running” as from “poetry”).

Map one-hot to dense vectors

sparse one-hot
encoding of words

aardvark	1	0	0	...	0	0	0
black	0	0	...	1	...	0	0
cat	0	0	...	1	...	0	0
duvet	0	0	...	1	...	0	0
zombie	0	0	0	...	0	0	1

VxV matrix



**VxE
embedding
matrix**

	animal	fluffiness	dangerous	spooky
aardvark	0.97	0.03	0.15	0.04
black	0.07	0.01	0.20	0.95
cat	0.98	0.98	0.45	0.35
duvet	0.01	0.84	0.12	0.02
zombie	0.74	0.05	0.98	0.93

VxE matrix

Problem: how do we find the values of the embedding matrix?

Solution: Learn as part of the task


```
CLASS torch.nn.Embedding(num_embeddings: int, embedding_dim: int,  
padding_idx: Optional[int] = None, max_norm: Optional[float] = None,  
norm_type: float = 2.0, scale_grad_by_freq: bool = False, sparse: bool  
= False, _weight: Optional[torch.Tensor] = None)
```

[SOURCE]

- A simple lookup table that stores embeddings of a fixed dictionary and size.
- This module is often used to store word embeddings and retrieve them using indices. The input to the module is a list of indices, and the output is the corresponding word embeddings.

Word2Vec – 1

"king"



“king”




“Man”




“Woman”



king - man + woman



king - man + woman \approx queen



Some NLP Applications

Q&A: SQuAD

- 100K question-answer pairs
- Answers are always spans in the question

In meteorology, precipitation is any product of the condensation of atmospheric water vapor that falls under **gravity**. The main forms of precipitation include drizzle, rain, sleet, snow, **graupel** and hail... Precipitation forms as smaller droplets coalesce via collision with other rain drops or ice crystals **within a cloud**. Short, intense periods of rain in scattered locations are called "showers".

What causes precipitation to fall?

gravity

What is another main form of precipitation besides drizzle, rain, snow, sleet and hail?

graupel

Where do water droplets collide with ice crystals to form precipitation?

within a cloud

Natural Language Inference: SNLI

- What relation exists between piece of text and hypothesis?
- 570K pairs

A man inspects the uniform of a figure in some East Asian country.	contradiction C C C C C	The man is sleeping
An older and younger man smiling.	neutral N N E N N	Two men are smiling and laughing at the cats playing on the floor.
A black race car starts up in front of a crowd of people.	contradiction C C C C C	A man is driving down a lonely road.
A soccer game with multiple males playing.	entailment E E E E E	Some men are playing a sport.
A smiling costumed woman is holding an umbrella.	neutral N N E C N	A happy woman in a fairy costume holds an umbrella.

GLUE – 1

Description	Data example
Is the sentence grammatical or ungrammatical?	"This building is than that one." = Ungrammatical
Is the movie review positive, negative, or neutral?	"The movie is funny , smart , visually inventive , and most of all , alive ." = .93056 (Very Positive)
Is the sentence B a paraphrase of sentence A?	A) "Yesterday , Taiwan reported 35 new infections , bringing the total number of cases to 418 ." B) "The island reported another 35 probable cases yesterday , taking its total to 418 ." = A Paraphrase
How similar are sentences A and B?	A) "Elephants are walking down a trail." B) "A herd of elephants are walking along a trail." = 4.6 (Very Similar)
Are the two questions similar?	A) "How can I increase the speed of my internet connection while using a VPN?" B) "How can Internet speed be increased by hacking through DNS?" = Not Similar
Does sentence A entail or contradict sentence B?	A) "Tourist Information offices can be very helpful." B) "Tourist Information offices are never of any help." = Contradiction


GLUE – 2

Does sentence B contain the answer to the question in sentence A?	A) "What is essential for the mating of the elements that create radio waves?" B) "Antennas are required by any radio receiver or transmitter to couple its electrical connection to the electromagnetic field." = Answerable
Does sentence A entail sentence B?	A) "In 2003, Yunus brought the microcredit revolution to the streets of Bangladesh to support more than 50,000 beggars, whom the Grameen Bank respectfully calls Struggling Members." B) "Yunus supported more than 50,000 Struggling Members." = Entailed
Sentence B replaces sentence A's ambiguous pronoun with one of the nouns - is this the correct noun?	A) "Lily spoke to Donna, breaking her concentration." B) "Lily spoke to Donna, breaking Lily's concentration." = Incorrect Referent

- 9 tasks, model score is averaged across them.

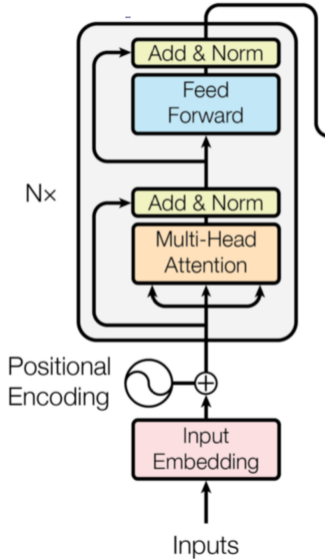
Transformers

Attention is all you need



- Encoder-decoder with only attention and fully connected layers (no recurrent or convolutions).
- Set new SOTA on translation datasets and many more.

Attention is all you need



- For simplicity, we can focus just on the encoder. For instance, BERT is just the encoder.
- The components:
 - (Masked) Self-attention
 - Positional encoding
 - Layer normalization

Basic self-attention


- **Input:** sequence of tensors $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t$.
- **Output:** sequence of tensors, each one a weighted sum of the input sequence:

$$\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_t, \quad \text{where } \mathbf{y}_i = \sum_j w_{ij} \mathbf{x}_j.$$

– w_{ij} is a function of \mathbf{x}_i and \mathbf{x}_j

$$w_{ij} = \frac{\exp(w'_{ij})}{\sum_j \exp(w'_{ij})}, \quad \text{where } w'_{ij} = \mathbf{x}_i^\top \mathbf{x}_j.$$


Basic self-attention



Basic self-attention – Problems


- **No learned weights** → **Let us learn some weights!**
- Order of the sequence does not affect result of computations.

Query, key, value



- Every input vector x_i is used in 3 ways:
 - Compared to every other vector to compute attention weights for its own output y_i (query).
 - Compared to every other vector to compute attention weight w_{ij} for output y_j (key).
 - Summed with other vectors to form the result of the attention weighted sum (value).

Query, key, value



- We can process each input vector to fulfill the three roles with matrix multiplication.
- Learning the matrices, i.e., **learning attention**


$$q_i = W_q x_i, \quad k_i = W_k x_i, \quad v_i = W_v x_i$$

$$w'_{ij} = q_i^\top k_j,$$

$$w_{ij} = \text{softmax}(w'_{ij})$$

$$y_i = \sum_j w_{ij} v_j.$$


Multi-head attention




- Multiple “heads” of attention just means learning different sets of W_q , W_k , and W_v matrices simultaneously.
- Implemented as just a single matrix anyway ...

Transformer

- Self-attention layer → Layer normalization → Dense layer




Layer normalization



Gradient of larger parameter dominates the update

Both parameters can be updated in equal proportions



- Neural net layers work best when input vectors have uniform mean and std in each dimension.
- As inputs flow through the network, means and std's get blown out.
- Layer normalization is a hack to reset things to when we want them in between layers.

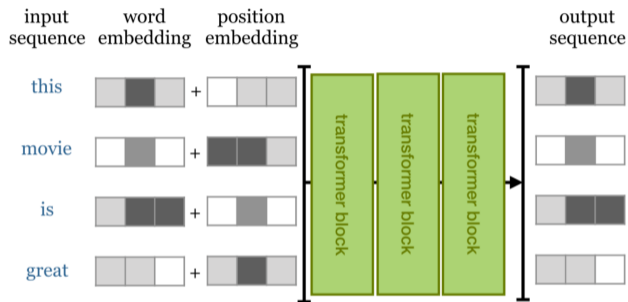
Transformer

So far:

- Learned query, key, value weights.
- Multiple heads.
- **Order of the sequence does not affect result of computations:** Let us encode each vector with position.

Transformer

- Position embedding: just what it sounds!




For instance, let d be the dimension of the word embedding, then one particular position embedding scheme is

$$PE_{(pos,2i)} = \sin\left(\frac{pos}{10000^{2i/d}}\right); \quad PE_{(pos,2i+1)} = \cos\left(\frac{pos}{10000^{2i/d}}\right)$$


Transformer: last trick

- Since the Transformer sees all inputs at once, to predict next vector in sequence (e.g. generate text), we need to mask the future.



Transformer: last trick

- Since the Transformer sees all inputs at once, to predict next vector in sequence (e.g. generate text), we need to mask the future.




raw attention weights

⊗




mask



y_i cannot see x_j for $j > i$.


Transformers-based AI Models

Transformers: Recap



- Encoder-decoder for translation.


Transformers: Recap





- Encoder-decoder for translation.
- Later models made it mostly just the encoder or just the decoder.
- ... but then the latest models are back to encoder-decoder.

GPT/GPT-2

- Generative Pre-trained Transformer
- GPT learns to predict the next word in the sequence
- Since it conditions only on preceding words, it uses masked self-attention



GPT/GPT-2




Trained on 8M web pages




Model Dimensionality: 768
117M params



Model Dimensionality: 1024
345M params



Model Dimensionality: 1280
762M params

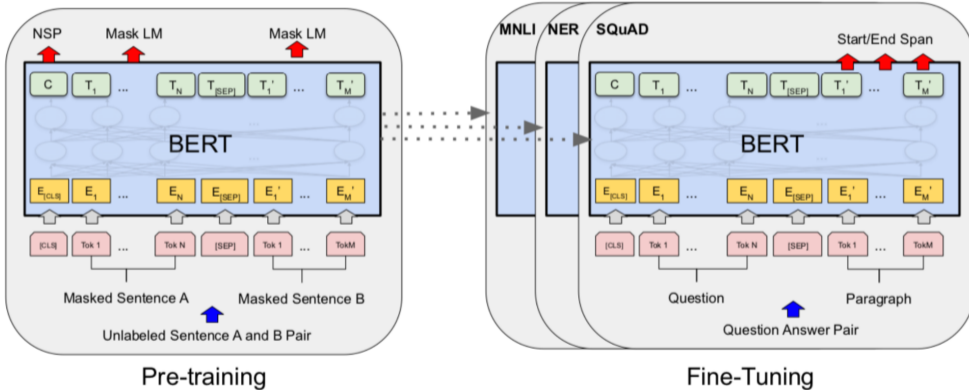


Model Dimensionality: 1600
1.5B params


BERT

- Bidirectional Encoder Representations from Transformers
- Encoder blocks only (no masking)
- BERT involves pre-training on a lot of text with 15% of all words masked out. Also, sometimes predicting whether one sentence follows another.
- 340M parameters: 24 transformer blocks, embedding dim of 1024, 16 attention heads.

BERT



More transformer-based AI models



Linearized Self-Attention

Katharopoulos et al. Transformers are RNNs: Fast Autoregressive Transformers with Linear Attention, ICML 2020.

Self-attention mechanism: Recap

Self-attention mechanism: Transforms sequences $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_N]^\top \in \mathbb{R}^{N \times D_x}$ as follows:

Step 1. Project \mathbf{X} into \mathbf{Q} , \mathbf{K} , and \mathbf{V} :

$$\mathbf{Q} = \mathbf{X}\mathbf{W}_Q^\top; \mathbf{K} = \mathbf{X}\mathbf{W}_K^\top; \mathbf{V} = \mathbf{X}\mathbf{W}_V^\top,$$

where $\mathbf{W}_Q, \mathbf{W}_K \in \mathbb{R}^{D \times D_x}$, and $\mathbf{W}_V \in \mathbb{R}^{D_v \times D_x}$ are learnable. $\mathbf{Q} := [\mathbf{q}_1, \dots, \mathbf{q}_N]^\top$, similar for \mathbf{K} and \mathbf{V} .

Step 2. Output sequence $\hat{\mathbf{V}} := [\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_N]$, where

$$\hat{\mathbf{v}}_i = \sum_{j=1}^N \text{softmax}\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{D}}\right) \mathbf{v}_j.$$

Self-attention mechanism: Bottlenecks

$$\hat{\mathbf{v}}_i = \sum_{j=1}^N \text{softmax}\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{D}}\right) \mathbf{v}_j, \iff \hat{\mathbf{V}} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{D}}\right) \mathbf{V} := \mathbf{A}\mathbf{V}.$$

Bottleneck: Store \mathbf{A} takes $\mathcal{O}(N^2)$ memory and compute $\mathbf{A}\mathbf{V}$ costs $\mathcal{O}(N^2)$ time.

Relaxation

Note that

$$\hat{\mathbf{v}}_i = \sum_{j=1}^N \text{softmax}\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{D}}\right) \mathbf{v}_j,$$

can be written as

$$\hat{\mathbf{v}}_i = \frac{\sum_{j=1}^N \text{sim}(\mathbf{q}_i, \mathbf{k}_j) \mathbf{v}_j}{\sum_{j=1}^N \text{sim}(\mathbf{q}_i, \mathbf{k}_j)},$$

where $\text{sim}(\mathbf{q}_i, \mathbf{k}_j) = \exp\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{D}}\right)$.

Relaxation

Note that

$$\hat{\mathbf{v}}_i = \sum_{j=1}^N \text{softmax}\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{D}}\right) \mathbf{v}_j,$$

can be written as

$$\hat{\mathbf{v}}_i = \frac{\sum_{j=1}^N \text{sim}(\mathbf{q}_i, \mathbf{k}_j) \mathbf{v}_j}{\sum_{j=1}^N \text{sim}(\mathbf{q}_i, \mathbf{k}_j)},$$

where $\text{sim}(\mathbf{q}_i, \mathbf{k}_j) = \exp\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{D}}\right)$.

Key idea: Replace $\text{sim}(\mathbf{q}_i, \mathbf{k}_j)$ with a kernel $k(\mathbf{q}_i, \mathbf{k}_j)$ that can be represented as the inner product of a feature map on \mathbf{q}_i and \mathbf{k}_j , i.e., $k(\mathbf{q}_i, \mathbf{k}_j) = \phi(\mathbf{q}_i)^\top \phi(\mathbf{k}_j)$.

Linearized self-attention

$$\hat{\mathbf{v}}_i = \frac{\sum_{j=1}^N \text{sim}(\mathbf{q}_i, \mathbf{k}_j) \mathbf{v}_j}{\sum_{j=1}^N \text{sim}(\mathbf{q}_i, \mathbf{k}_j)} \approx \frac{\sum_{j=1}^N k(\mathbf{q}_i, \mathbf{k}_j) \mathbf{v}_j}{\sum_{j=1}^N k(\mathbf{q}_i, \mathbf{k}_j)} = \frac{\phi(\mathbf{q}_i)^\top \sum_{j=1}^N \phi(\mathbf{k}_j) \mathbf{v}_j^\top}{\phi(\mathbf{q}_i)^\top \sum_{j=1}^N \phi(\mathbf{k}_j)}.$$

Advantages of linearized self-attention

$$\hat{\mathbf{V}} = \mathbf{AV} \approx (\mathbf{ab}^\top)\mathbf{V} = \mathbf{a}(\mathbf{b}^\top\mathbf{V}),$$

– \mathbf{ab}^\top is the rank-one approximation of the matrix \mathbf{A} .

Remark. Computing \mathbf{AV} requires $\mathcal{O}(N^2)$ complexity in computational time and memory footage. However, compute $\mathbf{a}(\mathbf{b}^\top\mathbf{V})$ only requires $\mathcal{O}(N)$ complexity in computational time and memory footage.

Remark. Rank-one approximation of \mathbf{A} is a quite poor approximation.

FMMformer: Efficient and Flexible Transformers with Decomposed Near-field and Far-field Attention

T. Nguyen, V. Suliafu, S. Osher, L. Chen, and B. Wang, FMMformer: Efficient and Flexible Transformers with Decomposed Near-field and Far-field Attention, NeurIPS 2021.

$$\hat{\mathbf{v}}_i = \sum_{j=1}^N \text{softmax}\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{D}}\right) \mathbf{v}_j, \iff \hat{\mathbf{v}}_i = \underbrace{\sum_{j=1}^N k(\mathbf{q}_i, \mathbf{k}_j) \mathbf{v}_j}_{\text{Force calculation?}}$$

Simplest idea: cutoff \Rightarrow sparse (local) attention. **Problematic** if $k(\mathbf{q}_i, \mathbf{k}_j) \gtrsim \frac{1}{\|\mathbf{q}_i - \mathbf{k}_j\|}$.

Physics analogue: gravitational and electrostatics force calculation. **Long-range force** where the potential decays at the rate $1/\|\mathbf{q}_i - \mathbf{k}_j\|$.

Computational math toolbox: particle mesh Ewald (PME) (Ewald, Ann. Phys. 1921.); **fast multipole method (FMM)** (Greengard and Rokhlin, JCP, 1987.)

PME: calculate near-field interaction in real-space and calculate far-field interaction in the k -space.
FMM: direct calculation of near-field interaction and coarse-grain far-field interaction.

Fast multipole method & its algebraic interpretation

Key idea of FMM: *far-field interaction can be well-approximated by separable low-rank matrices* while the near-field interaction can be calculated directly.

Let $\mathbf{A} \in \mathbb{R}^{N \times N}$ encodes interaction among all particles, where $\mathbf{A}(i, j) = g(\|\mathbf{q}_i - \mathbf{k}_j\|)$ with g be smooth except at 0 and $g(st) = g(s)g(t)$. E.g., $g(\|\mathbf{q}_i - \mathbf{k}_j\|) = 1/\|\mathbf{q}_i - \mathbf{k}_j\|$.

Def. [Well-separation] Let us partition $\{1, \dots, N\}$ into two groups $\{T_1, T_2\}$, then T_1 is called well-separated from T_2 if $\exists \mathbf{k}^*$ and a number δ s.t.

$$\|\mathbf{k}_j - \mathbf{k}^*\| \leq \delta \|\mathbf{q}_i - \mathbf{k}^*\| \quad \forall i \in T_1, j \in T_2, \quad \text{e.g., } \mathbf{k}^* \text{ is the center of vectors in } T_2.$$

In this case, $\|\mathbf{q}_i - \mathbf{k}_j\| \approx \|\mathbf{q}_i - \mathbf{k}^*\|$ for any $j \in T_2$. The i^{th} row of \mathbf{A} will be a constant after excluding a banded matrix \mathbf{D} from \mathbf{A} , i.e., $\mathbf{A} - \mathbf{D}$ is low-rank.


Gravitational force calculation



Proposition. [Informal] Let $\{T_1, T_2\}$ be two well-separated index sets. If g satisfies certain conditions, for any $\epsilon > 0$, the sub-matrix $A(T_1, T_2)$ can be approximated by a rank p matrix to a relative tolerance $\epsilon > 0$ in the sense that: there exists rank p matrices $\mathbf{U} \in \mathbb{R}^{|T_1| \times p}$, $\mathbf{V} \in \mathbb{R}^{|T_2| \times p}$, with $p \geq C |\log_\delta \epsilon|$ for some constant C , such that

$$|\mathbf{A}(i, j) - (\mathbf{UV}^\top)(i, j)| \leq \epsilon, \quad \forall i \in T_1, j \in T_2.$$

Taking above well-separated set partitioning recursively, we get the following \mathcal{H} -matrix.



Low-rank approximation via kernel tricks

Benefits of low-rank approximation: $\mathbf{L}\mathbf{V}$ requires $\mathcal{O}(N^2)$ computational time and memory costs if $\mathbf{L} \in \mathbb{R}^{N \times N}$. However, if \mathbf{L} is rank r with $r \ll N$, then

$$\mathbf{L}\mathbf{V} = \underbrace{(\mathbf{a}_1\mathbf{b}_1^\top + \mathbf{a}_2\mathbf{b}_2^\top + \cdots + \mathbf{a}_r\mathbf{b}_r^\top)}_{\mathcal{O}(N^2)}\mathbf{V} = \underbrace{\mathbf{a}_1(\mathbf{b}_1^\top\mathbf{V}) + \mathbf{a}_2(\mathbf{b}_2^\top\mathbf{V}) + \cdots + \mathbf{a}_r(\mathbf{b}_r^\top\mathbf{V})}_{\mathcal{O}(N)},$$

Practical low-rank attention:

$$\hat{\mathbf{v}}_i = \frac{\sum_{j=1}^N k(\mathbf{q}_i, \mathbf{k}_j)\mathbf{v}_j}{\sum_{j=1}^N k(\mathbf{q}_i, \mathbf{k}_j)} = \frac{\sum_{j=1}^N \phi(\mathbf{q}_i)^\top \phi(\mathbf{k}_j)\mathbf{v}_j}{\sum_{j=1}^N \phi(\mathbf{q}_i)^\top \phi(\mathbf{k}_j)} = \frac{\phi(\mathbf{q}_i)^\top \sum_{j=1}^N \phi(\mathbf{k}_j)\mathbf{v}_j}{\phi(\mathbf{q}_i)^\top \sum_{j=1}^N \phi(\mathbf{k}_j)},$$

i.e.,

$$\hat{\mathbf{V}} = \frac{\phi(\mathbf{Q})(\phi(\mathbf{K})^\top \mathbf{V})}{\phi(\mathbf{Q})\phi(\mathbf{K})^\top}.$$

Select a set of linearly independent feature maps $\{\phi_l(\cdot)\}_{l=1}^r \Rightarrow$ **rank- r attention**.

Banded matrix modeling of near-field attention

We model the near-field attention with the following banded matrix

$$\mathbf{D} = \text{softmax} \left(\text{band}_k \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{D}} \right) \right), \quad (1)$$

with $k \ll N$.

FMMformer

Original self-attention mechanism:

$$\hat{\mathbf{v}}_i = \sum_{j=1}^N \text{softmax}\left(\frac{\mathbf{q}_i^\top \mathbf{k}_j}{\sqrt{D}}\right) \mathbf{v}_j, \iff \hat{\mathbf{V}} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{D}}\right) \mathbf{V} := \mathbf{A}\mathbf{V}.$$

FMMformer:

$$\hat{\mathbf{V}} = \underbrace{w_1 (\mathbf{D}\mathbf{V})}_{\text{banded: near-field attention}} + w_2 \underbrace{\left(\sum_{l=1}^r \frac{\phi_l(\mathbf{Q})(\phi_l(\mathbf{K})^\top \mathbf{V})}{\phi_l(\mathbf{Q})\phi_l(\mathbf{K})^\top} \right)}_{\text{rank } r: \text{ far-field attention}},$$

where w_1 and w_2 are two learnable weights with positivity constraints.

Random Fourier Features

Rahimi and Recht, Random Features for Large-Scale Kernel Machines, NeurIPS, 2007.

Question

Given a kernel $k(\mathbf{x}, \mathbf{y}) = \phi(\mathbf{x})^\top \phi(\mathbf{y})$, where $\phi(\cdot)$ is a feature map which can be infinite-dimensional. How to construct a finite-dimensional explicit feature map $\mathbf{z}(\cdot)$ such that $k(\mathbf{x}, \mathbf{y}) \approx \mathbf{z}(\mathbf{x})^\top \mathbf{z}(\mathbf{y})$. In particular, consider the following kernels:

$$\exp(\mathbf{x}^\top \mathbf{y}),$$

and

$$\exp(-\|\mathbf{x} - \mathbf{y}\|_2^2).$$

Linear models

Consider a learning problem with data and targets $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$ where $\mathbf{x}_n \in \mathcal{X}$ and $y_n \in \mathcal{Y}$. Ignoring the bias, a linear model finds a hyperplane \mathbf{w} s.t. the decision function

$$f^*(\mathbf{x}) = \mathbf{w}^\top \mathbf{x}, \quad (2)$$

is optimal for some loss function.

For example, in logistic regression, we compute the logistic function of $f(\mathbf{x})$, and then threshold the output probability to produce a binary classifier with $\mathcal{Y} = \{0, 1\}$.

Kernel machines

- Kernel machine (kernel method): the input domain \mathcal{X} is mapped into another space \mathcal{V} in which the targets may be a linear function of the data.
- The dimension of \mathcal{V} may be high or even infinite, but kernel methods avoid operating explicitly in this space using the kernel trick.
- Kernel trick: if $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a positive definite kernel, then by Mercer's theorem there exists a basis function or feature map $\phi : \mathcal{X} \rightarrow \mathcal{V}$ such that

$$k(\mathbf{x}, \mathbf{y}) = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle_{\mathcal{V}}, \quad (3)$$

where $\langle \cdot, \cdot \rangle_{\mathcal{V}}$ is an inner product in \mathcal{V} .

Dual form of linear regression

- Linear regression model: $y(\mathbf{x}) = \mathbf{w}^\top \phi(\mathbf{x})$.
- Consider the following regularized sum-of-squares error function for linear regression

$$J(\mathbf{w}) = \frac{1}{2} \sum_{n=1}^N \{ \mathbf{w}^\top \phi(\mathbf{x}_n) - y_n \}^2 + \frac{\lambda}{2} \mathbf{w}^\top \mathbf{w}, \quad \lambda \geq 0. \quad (4)$$

- $$\nabla_{\mathbf{w}} J(\mathbf{w}) = 0 \Rightarrow \mathbf{w} = -\frac{1}{\lambda} \sum_{n=1}^N \{ \mathbf{w}^\top \phi(\mathbf{x}_n) - y_n \} \phi(\mathbf{x}_n) = \sum_{n=1}^N a_n \phi(\mathbf{x}_n) = \Phi^\top \mathbf{a},$$

where Φ is the design matrix, whose n th row is given by $\phi(\mathbf{x}_n)^\top$. The vector $\mathbf{a} = (a_1, \dots, a_N)^\top$ with

$$a_n = -\frac{1}{\lambda} \{ \mathbf{w}^\top \phi(\mathbf{x}_n) - y_n \}. \quad (5)$$

Dual form of linear regression

- We can reformulate the least squares algorithm in terms of the parameter vector \mathbf{a} instead of \mathbf{w} , resulting in a *dual representation*. Substitute $\mathbf{w} = \Phi^\top \mathbf{a}$ into $J(\mathbf{w})$ gives

$$J(\mathbf{a}) = \frac{1}{2} \mathbf{a}^\top \Phi \Phi^\top \Phi \Phi^\top \mathbf{a} - \mathbf{a}^\top \Phi \Phi^\top \mathbf{y} + \frac{1}{2} \mathbf{y}^\top \mathbf{y} + \frac{\lambda}{2} \mathbf{a}^\top \Phi \Phi^\top \mathbf{a}, \quad (6)$$

where $\mathbf{y} = (y_1, \dots, y_N)^\top$.

Dual form of linear regression

- Define the *Gram* matrix $\mathbf{K} = \Phi\Phi^\top \in \mathbb{R}^{N \times N}$ with $K_{nm} = \phi(\mathbf{x}_n)^\top \phi(\mathbf{x}_m) = k(\mathbf{x}_n, \mathbf{x}_m)$. In terms of the Gram matrix, the sum-of-squares error function can be written as

$$J(\mathbf{a}) = \frac{1}{2} \mathbf{a}^\top \mathbf{K} \mathbf{K} \mathbf{a} - \mathbf{a}^\top \mathbf{K} \mathbf{y} + \frac{1}{2} \mathbf{y}^\top \mathbf{y} + \frac{\lambda}{2} \mathbf{a}^\top \mathbf{K} \mathbf{a}. \quad (7)$$

$$\nabla_{\mathbf{a}} J(\mathbf{a}) = 0 \Rightarrow \mathbf{a} = (\mathbf{K} + \lambda \mathbf{I}_N)^{-1} \mathbf{y}.$$

Dual form of linear regression

- Substitute this back into the linear regression model, we obtain the following prediction for a new input \mathbf{x}

$$y(\mathbf{x}) = \mathbf{w}^\top \phi(\mathbf{x}) = \mathbf{a}^\top \Phi \phi(\mathbf{x}) = \mathbf{k}(\mathbf{x})^\top (\mathbf{K} + \lambda \mathbf{I}_N)^{-1} \mathbf{y}, \quad \text{Dual formulation!} \quad (8)$$

where we have defined the vector $\mathbf{k}(\mathbf{x})$ with elements $k_n(\mathbf{x}) = k(\mathbf{x}_n, \mathbf{x})$.

- The dual formulation allows the solution to the least-squares problem to be expressed entirely in terms of the kernel function $k(\mathbf{x}, \mathbf{x}')$.
- Note that the prediction at \mathbf{x} is given by a linear combination of the target values from the training set.

Dual form of linear regression

$$y(\mathbf{x}) = \sum_{n=1}^N \alpha_n k(\mathbf{x}, \mathbf{x}_n),$$

where $\boldsymbol{\beta} = (\alpha_1, \dots, \alpha_N)^\top = (\mathbf{K} + \lambda \mathbf{I}_N)^{-1} \mathbf{y}$.

Kernel machines

- Using the kernel trick and a representer theorem, kernel methods construct nonlinear models of \mathcal{X} that are linear in $k(\cdot, \cdot)$,

$$f^*(\mathbf{x}) = \underbrace{\sum_{n=1}^N \alpha_n k(\mathbf{x}, \mathbf{x}_n)}_{\text{dual form}} = \underbrace{\langle \mathbf{w}, \phi(\mathbf{x}) \rangle_{\mathcal{V}}}_{\text{original form}} . \quad (9)$$

- Provided we have a positive definite kernel function $k(\cdot, \cdot)$, we can avoid operating in the possibly infinite-dimensional space \mathcal{V} and instead only compute over N data points.

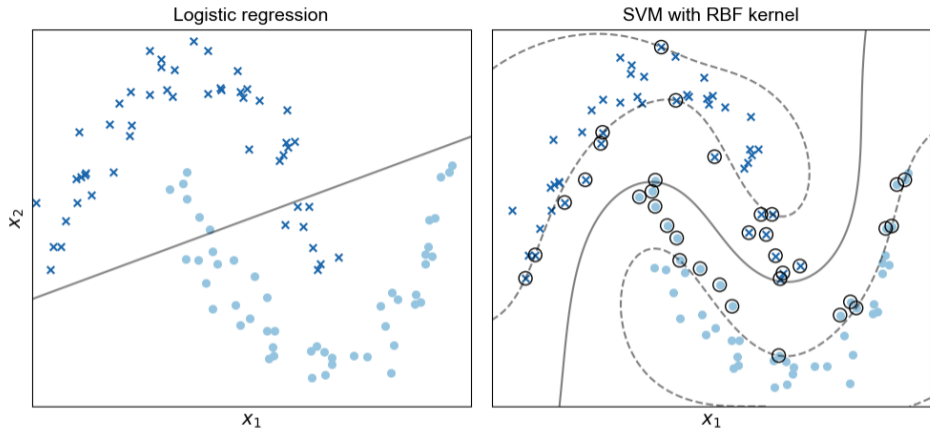


Figure: Logistic regression (left) with decision boundary denoted with a solid line and SVM with radial basis function (RBF) kernel (right) on the two moon dataset. Support vectors are denoted with circles, and the margins are denoted with dashed lines.

Kernel methods: Gaussian process, kernel PCA, kernel regression, etc.

Remarks on kernel methods

- While the kernel trick is a beautiful idea and the conceptual backbone of kernel machines, the problem is that for large datasets (for huge N), the machine must operate on a covariance matrix \mathbf{K}_X , induced by a particular kernel function, that is $N \times N$,

$$\mathbf{K}_X = \begin{pmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & k(\mathbf{x}_1, \mathbf{x}_2) & \cdots & k(\mathbf{x}_1, \mathbf{x}_N) \\ k(\mathbf{x}_2, \mathbf{x}_1) & k(\mathbf{x}_2, \mathbf{x}_2) & \cdots & k(\mathbf{x}_2, \mathbf{x}_N) \\ \vdots & \vdots & \ddots & \vdots \\ k(\mathbf{x}_N, \mathbf{x}_1) & k(\mathbf{x}_N, \mathbf{x}_2) & \cdots & k(\mathbf{x}_N, \mathbf{x}_N) \end{pmatrix} \quad (10)$$

- $\beta = (\mathbf{K} + \lambda \mathbf{I}_N)^{-1} \mathbf{y}$, which requires inverting an $N \times N$ matrix.
- Furthermore, to evaluate a test data point, the model must evaluate the sum in (9). While we're avoiding computing in \mathcal{V} , we are stuck operating in \mathbb{R}^N . In the area of big data, kernel methods do not necessarily scale.

Random features

Key idea: approximate the above inner product $\langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle_{\mathcal{V}}$ with a randomized map $\mathbf{z} : \mathbb{R}^D \rightarrow \mathbb{R}^R$ where ideally $R \ll N$,

$$k(\mathbf{x}, \mathbf{y}) = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle_{\mathcal{V}} \approx \mathbf{z}(\mathbf{x})^\top \mathbf{z}(\mathbf{y}). \quad (11)$$

Why does this work and why is it a good idea?

Random features

- The representer theorem tells us that the optimal solution is a weighted sum of kernel evaluated at our observations. If we have a good approximation of $\phi(\cdot)$, then

$$f^*(\mathbf{x}) = \sum_{n=1}^N \alpha_n k(\mathbf{x}_n, \mathbf{x}) = \sum_{n=1}^N \alpha_n \langle \phi(\mathbf{x}_n), \phi(\mathbf{x}) \rangle_{\mathcal{V}} \approx \sum_{n=1}^N \alpha_n \mathbf{z}(\mathbf{x}_n)^\top \mathbf{z}(\mathbf{x}) = \boldsymbol{\beta}^\top \mathbf{z}(\mathbf{x}). \quad (12)$$

- If $\mathbf{z}(\cdot)$ is a good approximation of $\phi(\cdot)$, then we can project our data using $\mathbf{z}(\cdot)$ and then use fast linear models in \mathbb{R}^R rather than \mathbb{R}^N because both $\boldsymbol{\beta}$ and $\mathbf{z}(\cdot)$ are in \mathbb{R}^R .
- So the task at hand is to find a random projection $\mathbf{z}(\cdot)$ such that it well-approximates the corresponding nonlinear kernel machine.

Random features

Key idea: $k(\mathbf{x}, \mathbf{y}) = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle_{\mathcal{V}} \approx \mathbf{z}(\mathbf{x})^{\top} \mathbf{z}(\mathbf{y})$.

- The idea was inspired by the following observation. Let $\mathbf{w} \in \mathbb{R}^D$ be a random vector such that

$$\mathbf{w} \sim \mathcal{N}_D(0, \mathbf{I}). \quad (13)$$

Now define h as

$$h : \mathbf{x} \rightarrow \exp(i\mathbf{w}^{\top} \mathbf{x}). \quad (14)$$

Above, i is the imaginary unit. Let the superscript $*$ denote the complex conjugate.

- Importantly, recall that the complex conjugate of e^{ix} is e^{-ix} . Then note

$$\begin{aligned}\mathbb{E}_{\mathbf{w}}[h(\mathbf{x})h(\mathbf{y})^*] &= \mathbb{E}_{\mathbf{w}}[\exp(i\mathbf{w}^\top(\mathbf{x} - \mathbf{y}))] = \int_{\mathbb{R}^D} p(\mathbf{w}) \exp(i\mathbf{w}^\top(\mathbf{x} - \mathbf{y})) d\mathbf{w} \\ &= \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{y})^\top(\mathbf{x} - \mathbf{y})\right).\end{aligned}\tag{15}$$

In other words, the expected value of $h(\mathbf{x})h(\mathbf{y}^*)$ is the Gaussian kernel.

Gaussian kernel derivation, i.e., Equation (15)

Let $\boldsymbol{\delta} = \mathbf{x} - \mathbf{y}$:

$$\begin{aligned}\mathbb{E}_{\mathbf{w}}[h(\mathbf{x})h(\mathbf{y})^*] &= \mathbb{E}_{\mathbf{w}}[\exp(i\mathbf{w}^\top \mathbf{x}) \exp(-i\mathbf{w}^\top \mathbf{y})] \\ &= \mathbb{E}_{\mathbf{w}}[\exp(i\mathbf{w}^\top \boldsymbol{\delta})] = \int_{\mathbb{R}^D} \rho(\mathbf{w}) \exp(i\mathbf{w}^\top \boldsymbol{\delta}) d\mathbf{w} \\ &= (2\pi)^{-D/2} \int_{\mathbb{R}^D} \exp\left(-\frac{1}{2}\mathbf{w}^\top \mathbf{w}\right) \exp(i\mathbf{w}^\top \boldsymbol{\delta}) d\mathbf{w} \\ &= (2\pi)^{-D/2} \int_{\mathbb{R}^D} \exp\left(-\left(\frac{1}{2}\mathbf{w}^\top \mathbf{w} - i\mathbf{w}^\top \boldsymbol{\delta}\right)\right) d\mathbf{w} \\ &= (2\pi)^{-D/2} \int_{\mathbb{R}^D} \exp\left(-\frac{1}{2}(\mathbf{w}^\top \mathbf{w} - 2i\mathbf{w}^\top \boldsymbol{\delta} - \boldsymbol{\delta}^\top \boldsymbol{\delta}) - \frac{1}{2}\boldsymbol{\delta}^\top \boldsymbol{\delta}\right) d\mathbf{w} \\ &= (2\pi)^{-D/2} \exp\left(-\frac{1}{2}\boldsymbol{\delta}^\top \boldsymbol{\delta}\right) \underbrace{\int_{\mathbb{R}^D} \exp\left(-\frac{1}{2}(\mathbf{w} - i\boldsymbol{\delta})^\top (\mathbf{w} - i\boldsymbol{\delta})\right) d\mathbf{w}}_{(2\pi)^{D/2}} \\ &= \exp\left(-\frac{1}{2}\boldsymbol{\delta}^\top \boldsymbol{\delta}\right) = k(\boldsymbol{\delta}).\end{aligned}\tag{16}$$

In other words, $k(\cdot)$ is the Gaussian kernel with $\rho(\mathbf{w})$ be a spherical Gaussian.

Bochner's theorem

The above result is a specific instance of Bochner's theorem.

Bochner's theorem. A continuous kernel $k(\mathbf{x}, \mathbf{y}) = k(\mathbf{x} - \mathbf{y})$ on \mathbb{R}^D is positive definite if and only if $k(\Delta)$ is the Fourier transform of a non-negative measure.

The Fourier transform of a non-negative measure, call it $p(\mathbf{w})$, is

$$k(\Delta) = \int p(\mathbf{w}) \exp(i\mathbf{w}\Delta) d\mathbf{w}. \quad (17)$$

Remark. Bochner's theorem gives us a general framework to approximate any shift invariant kernel (Gaussian, Laplace, and Cauchy kernels) by re-defining $h(\cdot)$ in (14) to depend on \mathbf{w} from any non-negative measure $p(\mathbf{w})$, not just the spherical Gaussian in (13). Furthermore, if we sample R i.i.d. realizations $\{\mathbf{w}_r\}_{r=1}^R$, we can lower the variance of this approximation:

$$\begin{aligned}
k(\mathbf{x}, \mathbf{y}) &= k(\mathbf{x} - \mathbf{y}) = \int p(\mathbf{w}) \exp(i\mathbf{w}^\top (\mathbf{x} - \mathbf{y})) d\mathbf{w} \\
&= \mathbb{E}_{\mathbf{w}} [\exp(i\mathbf{w}^\top (\mathbf{x} - \mathbf{y}))] \\
&\stackrel{\underbrace{1}}{\approx} \frac{1}{R} \sum_{r=1}^R \exp(i\mathbf{w}_r^\top (\mathbf{x} - \mathbf{y})) \\
&= \begin{pmatrix} \frac{1}{\sqrt{R}} \exp(i\mathbf{w}_1^\top \mathbf{x}) \\ \frac{1}{\sqrt{R}} \exp(i\mathbf{w}_2^\top \mathbf{x}) \\ \vdots \\ \frac{1}{\sqrt{R}} \exp(i\mathbf{w}_R^\top \mathbf{x}) \end{pmatrix}^\top \begin{pmatrix} \frac{1}{\sqrt{R}} \exp(-i\mathbf{w}_1^\top \mathbf{y}) \\ \frac{1}{\sqrt{R}} \exp(-i\mathbf{w}_2^\top \mathbf{y}) \\ \vdots \\ \frac{1}{\sqrt{R}} \exp(-i\mathbf{w}_R^\top \mathbf{y}) \end{pmatrix} \\
&\stackrel{\underbrace{2}}{=} \mathbf{h}(\mathbf{x})\mathbf{h}(\mathbf{y})^*.
\end{aligned} \tag{18}$$

- Step 1 is a Monte Carlo approximation of the expectation (\mathbf{w}_r s are sampled from the distribution $p(\mathbf{w})$).

- Step 2 is the definition of a random map $\mathbf{h} : \mathbb{R}^D \rightarrow \mathbb{R}^R$, so an R -vector of normalized $h(\cdot)$ transformations.

$$\begin{aligned}
k(\mathbf{x}, \mathbf{y}) &= k(\mathbf{x} - \mathbf{y}) \\
&= \left(\begin{array}{c} \frac{1}{\sqrt{R}} \exp(i\mathbf{w}_1^\top \mathbf{x}) \\ \frac{1}{\sqrt{R}} \exp(i\mathbf{w}_2^\top \mathbf{x}) \\ \vdots \\ \frac{1}{\sqrt{R}} \exp(i\mathbf{w}_R^\top \mathbf{x}) \end{array} \right)^\top \left(\begin{array}{c} \frac{1}{\sqrt{R}} \exp(-i\mathbf{w}_1^\top \mathbf{y}) \\ \frac{1}{\sqrt{R}} \exp(-i\mathbf{w}_2^\top \mathbf{y}) \\ \vdots \\ \frac{1}{\sqrt{R}} \exp(-i\mathbf{w}_R^\top \mathbf{y}) \end{array} \right) \\
&\underbrace{=}_2 \mathbf{h}(\mathbf{x})\mathbf{h}(\mathbf{y})^*.
\end{aligned}$$

Remark. Note that we've talked about the dot product $\mathbf{z}(\mathbf{x})^\top \mathbf{z}(\mathbf{y})$, but above we have $\mathbf{h}(\mathbf{x})\mathbf{h}(\mathbf{y})^*$. As we will see next, the imaginary part of our random map will disappear, and the new transform is what we used in machine learning.

Fine tuning

We have discussed the big idea of a low-dimensional, randomized map and why it might work, let us get into the weeds.

- First, note that since both our distribution $\mathcal{N}_D(0, \mathbf{I})$ and the kernel $k(\Delta)$ are real-valued, we can write

$$\exp(i\mathbf{w}^\top(\mathbf{x} - \mathbf{y})) \underbrace{=} \cos(\mathbf{w}^\top(\mathbf{x} - \mathbf{y})) - \cancel{i \sin(\mathbf{w}^\top(\mathbf{x} - \mathbf{y}))} = \cos(\mathbf{w}^\top(\mathbf{x} - \mathbf{y})) \quad (19)$$

Euler's formula

- We can then define $z_{\mathbf{w}}(\mathbf{x})$ —note that this is still not yet the bolded \mathbf{z} —without the imaginary unit as

$$\begin{aligned} \mathbf{w} &\sim \rho(\mathbf{w}) \\ b &\sim \text{Uniform}(0, 2\pi) \\ z_{\mathbf{w}}(\mathbf{x}) &= \sqrt{2} \cos(\mathbf{w}^\top \mathbf{x} + b). \end{aligned} \quad (20)$$

This works because

$$\begin{aligned}
\mathbb{E}_{\mathbf{w}}[z_{\mathbf{w}}(\mathbf{x})z_{\mathbf{w}}(\mathbf{y})] &= \mathbb{E}_{\mathbf{w}}[\sqrt{2} \cos(\mathbf{w}^{\top} \mathbf{x} + b)\sqrt{2} \cos(\mathbf{w}^{\top} \mathbf{y} + b)] \\
&\stackrel{(1)}{=} \mathbb{E}_{\mathbf{w}}[\cos(\mathbf{w}^{\top}(\mathbf{x} + \mathbf{y}) + 2b)] + \mathbb{E}_{\mathbf{w}}[\cos(\mathbf{w}^{\top}(\mathbf{x} - \mathbf{y}))] \\
&\stackrel{(2)}{=} \mathbb{E}_{\mathbf{w}}[\cos(\mathbf{w}^{\top}(\mathbf{x} - \mathbf{y}))]
\end{aligned} \tag{21}$$

- (1) is because of the following trigonometry identity

$$\cos(x + y) = \cos(x) \cos(y) - \sin(x) \sin(y).$$

- (2) uses the fact that since $b \sim \text{Uniform}(0, 2\pi)$, the expectation w.r.t. b is zero:

Lemma.

$$\mathbb{E}_{\mathbf{w}}[\cos(\mathbf{w}^{\top}(\mathbf{x} + \mathbf{y}) + 2b)] = \mathbb{E}_{\mathbf{w}}[\mathbb{E}_b[\cos(\mathbf{w}^{\top}(\mathbf{x} + \mathbf{y}) + 2b)|\mathbf{w}]] = 0. \tag{22}$$

Proof of the Lemma. Note that

$$\mathbb{E}_{\mathbf{w}}[\cos(\mathbf{w}^\top(\mathbf{x} + \mathbf{y}) + 2b)] = \mathbb{E}_{\mathbf{w}}[\mathbb{E}_b[\cos(\mathbf{w}^\top(\mathbf{x} + \mathbf{y}) + 2b)|\mathbf{w}]]$$

holds by the law of total expectation. We claim the inner conditional expectation is zero. To ease notation, let $\mathbf{t} = \mathbf{w}^\top(\mathbf{x} - \mathbf{y})$. Then

$$\begin{aligned}\mathbb{E}_b[\cos(\mathbf{t} + 2b)|\mathbf{w}] &= \int_0^{2\pi} \frac{\cos(\mathbf{t} + 2b)}{2\pi} db \\ &= \frac{1}{2\pi} \int_0^{2\pi} \cos(\mathbf{t} + 2b) db \\ &= \frac{1}{2\pi} \left[\sin(\mathbf{t} + 2b) \Big|_0^{2\pi} \right] \\ &= \frac{1}{2\pi} \left[\sin(\mathbf{t}) - \sin(\mathbf{t} + 4\pi) \right] \\ &= 0\end{aligned}$$

The last step holds because $\sin(\mathbf{t}) = \sin(\mathbf{t} \pm 2\pi k)$ for any integer k .

Fine tuning

We are now ready to define the random map $\mathbf{z} : \mathbb{R}^D \rightarrow \mathbb{R}^R$ such that (11) holds. Let

$$\mathbf{z}(\mathbf{x}) = \begin{pmatrix} \frac{1}{\sqrt{R}} z_{\mathbf{w}_1}(\mathbf{x}) \\ \frac{1}{\sqrt{R}} z_{\mathbf{w}_2}(\mathbf{x}) \\ \vdots \\ \frac{1}{\sqrt{R}} z_{\mathbf{w}_R}(\mathbf{x}) \end{pmatrix} \quad (23)$$

and therefore

$$\begin{aligned} \mathbf{z}(\mathbf{x})^\top \mathbf{z}(\mathbf{y}) &= \frac{1}{R} \sum_{r=1}^R z_{\mathbf{w}_r}(\mathbf{x}) z_{\mathbf{w}_r}(\mathbf{y}) = \frac{1}{R} \sum_{r=1}^R 2 \cos(\mathbf{w}_r^\top \mathbf{x} + b_r) \cos(\mathbf{w}_r^\top \mathbf{y} + b_r) \\ &= \frac{1}{R} \sum_{r=1}^R \cos(\mathbf{w}_r^\top (\mathbf{x} - \mathbf{y})) \approx \mathbb{E}_{\mathbf{w}}[\cos(\mathbf{w}^\top (\mathbf{x} - \mathbf{y}))] = k(\mathbf{x}, \mathbf{y}). \end{aligned} \quad (24)$$

We now have a simple algorithm to estimate a shift invariant, positive definite kernel. Draw R samples of $\mathbf{w} \sim p(\mathbf{w})$ and $b \sim \text{Uniform}(0, 2\pi)$ and then compute $\mathbf{z}(\mathbf{x})^\top \mathbf{z}(\mathbf{x})$.

Alternative random Fourier features

An alternative version of random Fourier feature is

$$z_{\mathbf{w}_r}(\mathbf{x}) = \begin{pmatrix} \cos(\mathbf{w}_r^\top \mathbf{x}) \\ \sin(\mathbf{w}_r^\top \mathbf{x}) \end{pmatrix} \quad (25)$$

Draw $R' = R/2$ samples

$$\mathbf{w}_r \sim p(\mathbf{w}). \quad (26)$$

Then

$$\begin{aligned} \frac{1}{R'} \sum_{r=1}^{R'} \sum_{r=1}^{R'} z_{\mathbf{w}_r}(\mathbf{x})^\top z_{\mathbf{w}_r}(\mathbf{y}) &\equiv \frac{2}{R} \sum_{r=1}^{R/2} \left(\begin{pmatrix} \cos(\mathbf{w}_r^\top \mathbf{x}) \\ \sin(\mathbf{w}_r^\top \mathbf{x}) \end{pmatrix}^\top \begin{pmatrix} \cos(\mathbf{w}_r^\top \mathbf{y}) \\ \sin(\mathbf{w}_r^\top \mathbf{y}) \end{pmatrix} \right) \\ &= \frac{2}{R} \sum_{r=1}^{R/2} \cos(\mathbf{w}_r^\top \mathbf{x}) \cos(\mathbf{w}_r^\top \mathbf{y}) + \sin(\mathbf{w}_r^\top \mathbf{x}) \sin(\mathbf{w}_r^\top \mathbf{y}) \\ &\underbrace{=}^* \frac{2}{R} \sum_{r=1}^{R/2} \cos(\mathbf{w}_r^\top \mathbf{x} - \mathbf{w}_r^\top \mathbf{y}) \approx \mathbb{E}_{\mathbf{w}}[\cos(\mathbf{w}^\top (\mathbf{x} - \mathbf{y}))] = k(\mathbf{x}, \mathbf{y}). \end{aligned} \quad (27)$$

* in the last equation due to the product identities from trigonometry:

$$2 \sin(x) \sin(y) = \cos(x - y) - \cancel{\cos(x + y)}; \quad 2 \cos(x) \cos(y) = \cos(x - y) + \cancel{\cos(x + y)}. \quad (28)$$

The right-most terms above cancel in (27), and we get $2 \cos(x - y)$.

Example: Gaussian kernel approximation

Let us first approximate a Gaussian kernel using random Fourier features. Sample R i.i.d. \mathbf{w} variables from a spherical Gaussian and then compute

$$\mathbf{z}(\mathbf{x})^\top \mathbf{z}(\mathbf{y}) = \frac{1}{R} \sum_{r=1}^R \mathbf{z}_{\mathbf{w}_r}(\mathbf{x})^\top \mathbf{z}_{\mathbf{w}_r}(\mathbf{y}) = \frac{1}{R} \sum_{r=1}^R \cos(\mathbf{w}_r^\top (\mathbf{x} - \mathbf{y})). \quad (29)$$

for each (\mathbf{x}, \mathbf{y}) pair in the data. The result $N \times N$ matrix is the approximate covariance matrix induced by the Gaussian kernel function. Concretely, let \mathbf{Z}_X denote $\mathbf{z}(\cdot)$ applied to all N samples \mathbf{x}_n . Thus, \mathbf{Z}_X is $N \times R$ and therefore

$$\mathbf{K}_X \approx \begin{bmatrix} \mathbf{z}(\mathbf{x}_1) \\ \vdots \\ \mathbf{z}(\mathbf{x}_N) \end{bmatrix} [\mathbf{z}(\mathbf{x}_1) \quad \cdots \quad \mathbf{z}(\mathbf{x}_N)] = \mathbf{Z}_X \mathbf{Z}_X^\top, \quad (30)$$

because

$$\begin{pmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & \cdots & k(\mathbf{x}_1, \mathbf{x}_N) \\ \vdots & \ddots & \vdots \\ k(\mathbf{x}_N, \mathbf{x}_1) & \cdots & k(\mathbf{x}_N, \mathbf{x}_N) \end{pmatrix} \approx \begin{pmatrix} \mathbf{z}(\mathbf{x}_1)^\top \mathbf{z}(\mathbf{x}_1) & \cdots & \mathbf{z}(\mathbf{x}_1)^\top \mathbf{z}(\mathbf{x}_N) \\ \vdots & \ddots & \vdots \\ \mathbf{z}(\mathbf{x}_N)^\top \mathbf{z}(\mathbf{x}_1) & \cdots & \mathbf{z}(\mathbf{x}_N)^\top \mathbf{z}(\mathbf{x}_N) \end{pmatrix} \quad (31)$$




Figure: As R increases, the covariance matrix approximation improves because each cell value uses more Monte Carlo samples to estimate the basis function $\phi(\cdot)$ associated with $k(\cdot, \cdot)$ for the pair of samples associated with that cell.