

MATH REVIEW

PSD $A \geq 0$
 if $\forall v \neq 0: v^T A v \geq 0$

SVD $M = U \Sigma V^T$

U cols = eigenvectors of $M M^T$
 V cols = eigenvectors of $M^T M$
 $\lambda_i = \sum_{j=1}^n u_{ij}^2 = \sum_{j=1}^n v_{ij}^2$

Spectral thm $A = Q \Lambda Q^T$

for square symmetric matrices
 Λ = eigenvalues of A
 Q = eigenvectors of A

ISO contours = ellipses centered around dist's mean

to find direction of major axis:
 1) Find eigenvalues / eigenvectors
 2) Find eigenvector corresponding to largest eigenvalue

covariance matrix $\Sigma = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)(x_i - \mu)^T$
 $X = AZ + \mu$ where $Z \sim N(0, I_n)$
 $\Sigma = A A^T$ * zero mean
 * diagonal entries = variances
 * off-diagonal = covariances
 * uncorrelated: $E\{x_i x_j\} = E\{x_i\} E\{x_j\}$

lagrangian multiplier
 $\mathcal{L}(x, y, \lambda) = f(x, y) - \lambda g(x, y)$
 example $\|w\|_2 = 1$ constrained optimization
 $w_1 - w_1^2, w_2, \lambda(w_1^2 - 1)$

SIGMOID FUNCTION
 $\sigma(x) = \frac{1}{1 + \exp(-x)}$
 * prone to vanishing gradients

LOG-LIKELIHOOD
 TAKE LOG OF $L(\theta)$
 since log is monotonically increasing
 $\ell(\mu, \Sigma) = \log\left(\prod_{i=1}^n p(x_i | \mu, \Sigma)\right)$
 * exponents \rightarrow multiplication
 * products \rightarrow sums

MAP = maximum a posteriori
 * point estimate of param to maximize posterior * uniform prior:
 $\hat{\theta}_{MAP} = \underset{\theta}{\operatorname{argmax}} P(\theta | X)$
 $= \underset{\theta}{\operatorname{argmax}} \frac{P(X | \theta) P(\theta)}{P(X)}$
 $\hat{\theta}_{MAP} = \underset{\theta}{\operatorname{argmax}} P(X | \theta) P(\theta) = \underset{\theta}{\operatorname{argmax}} \left[\prod_{i=1}^n P(x_i | \theta) \right] P(\theta)$

Bayes Rule!!!

$$P(Y|X) = \frac{P(X|Y)P(Y)}{P(X)}$$

posterior

EXPECTATION
 $E\{f\} = \sum_x p(x) f(x)$

variance
 $\operatorname{var}\{X\} = E\{f(x)^2\} - E\{f(x)\}^2$
 covariance

orthonormal vectors
 $\operatorname{cov}\{x_i, y_j\} = E\{x_i y_j\} - E\{x_i\} E\{y_j\}$

L1 NORM

$$\|x\|_1 = \sum_{i=1}^n |x_i|$$

L2 NORM

$$\|x\|_2 = \sqrt{\sum_{i=1}^n x_i^2}$$

$$\|x\|_2^2 = x^T x$$

singular values
 $\sigma_i = \sqrt{\lambda_i}$ where λ_i = eigenvals of $A^T A$

Frobenius Norm

$$\|A\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^n |A_{ij}|^2}$$

$$= \sqrt{\operatorname{tr}(A^T A)} = \sqrt{\sum_{i=1}^n \sigma_i^2}$$

OUTER PRODUCT FORM

$$X^T X = \sum_j x_j x_j^T$$

$$S = \frac{1}{N} \sum_{n=1}^N x_n x_n^T$$

sample covariance if data centered

Normal Gaussian PDF

$$N(x | \mu, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

MULTIVARIATE GAUSSIAN

$$p(x) = \frac{1}{(2\pi)^{\frac{n}{2}} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)\right)$$

variance

examples:
 $L(p) = P(\operatorname{coin} | A)^2 \cdot P(\operatorname{coin} | B) \cdot P(\operatorname{coin} | B)^3$
 $L(\mu, \Sigma) = \prod_{i=1}^n p(x_i | \mu, \Sigma)$
 $L(w) = \prod_{i=1}^n p(y_i | x_i)$

LIKELIHOOD FUNCTION

$$P(X | \mu, \sigma^2) = \prod_{n=1}^N N(x_n | \mu, \sigma^2)$$

count

$$L(\theta) = P(x_1 | \theta)^{c_1} \cdot P(x_2 | \theta)^{c_2} \dots P(x_m | \theta)^{c_m}$$

MLE = max likelihood estimation = finding param vals to maximize likelihood function

$$\underset{\mu, \sigma}{\operatorname{argmax}} p(x_1, x_2, \dots, x_n | \mu, \sigma)$$

* $X^T X$ invertible when full rank
 * features > num data points = underdetermined
 * use regularization
 * data pts > features = overdetermined
 * least squares

LINEAR REGRESSION

$$y = Xw$$

$$\mathcal{L} = \underset{w}{\operatorname{argmin}} \|y - Xw\|_2^2$$

$$w^* = (X^T X)^{-1} X^T y$$

$$\hat{y} = X^T (X^T X)^{-1} X^T y$$

Matrix Derivatives

$$\frac{\partial x^T a}{\partial x} = \frac{\partial a^T x}{\partial x} = a$$

$$\frac{\partial a^T X b}{\partial X} = a b^T$$

$$\frac{\partial a^T X^T b}{\partial X} = b a^T$$

$$\frac{\partial X^T X a}{\partial X} = \frac{\partial a^T X^T a}{\partial X} = a a^T$$

$$\frac{\partial X}{\partial X} = I^N$$

$$\frac{\partial (X A)}{\partial X_{ij}} = \delta_{im} (A)_{mj} = (J^m A)_{ij}$$

$$\frac{\partial (X^T A)}{\partial X_{ij}} = \delta_{im} (A)_{mj} = (J^m A)_{ij}$$

$$\frac{\partial}{\partial X_{mn}} \sum_k X_{km} X_{kn} = 2 \sum_k X_{ki}$$

$$\frac{\partial (Bx + b)^T C (Dx + d)}{\partial X} = B^T C (Dx + d) + D^T C^T (Bx + b)$$

$$\frac{\partial (X^T B X)}{\partial X} = \delta_j (X^T B)_j + \delta_j (B X)_j$$

$$\frac{\partial (X^T B X)}{\partial X} = X^T B^T + B^T X \quad (B^T)_j = \delta_j B_j$$

$$\frac{\partial X^T B X}{\partial X} = (B + B^T) X$$

$$\frac{\partial (X^T C^T D X e)}{\partial X} = D^T X b e^T + D X e b^T$$

$$\frac{\partial}{\partial X} (Xb + e)^T D (Xb + e) = (D + D^T) (Xb + e) b^T$$

Jacobian

$$J = \begin{bmatrix} \frac{\partial z_1}{\partial x} & \dots & \frac{\partial z_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial z_m}{\partial x} & \dots & \frac{\partial z_m}{\partial x_n} \end{bmatrix}$$

Hessian

$$H = \nabla^2 f(x) = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2} & \dots & \frac{\partial^2 f}{\partial x_1 \partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 f}{\partial x_n \partial x_1} & \dots & \frac{\partial^2 f}{\partial x_n^2} \end{bmatrix}$$

ReLU

$$\operatorname{ReLU}(x) = \begin{cases} x & x > 0 \\ 0 & x \leq 0 \end{cases}$$

$$\nabla_x (xw - y)^T (xw - y) = \nabla_x ((Xw)^T (Xw - y)) = \nabla_x (w^T X^T X w - y^T X w + y^T y) = (X^T X + X^T X) w - 2X^T y = 2X^T (Xw - y)$$

$$0 = 2X^T X w - 2X^T y$$

$$2X^T y = 2X^T X w \quad \text{invertible}$$

$$w^* = (X^T X)^{-1} X^T y$$

RIDGE REGRESSION

* underdetermined models

$$\mathcal{L} = (y - Xw)^T (y - Xw) + \lambda \|w\|_2^2$$


invertible when $\lambda > 0$ penalty term

$$w^* = (X^T X + \lambda I)^{-1} X^T y$$

LASSO REGRESSION

* induces sparsity w/ L1 norm
 $w^* = \underset{w}{\operatorname{argmin}} (y - Aw)^T (y - Aw) + \lambda \|w\|_1$
 * sharper corners = more sparse

LINEAR CLASSIFIER



FEATURES = X, LABELS = Y
 Model conditional probability for both classes & learning params that best fit this distribution
 1 Model each class P(X|Y) and P(Y)
 2 Use Bayes to compute P(Y|X)

DISCRIMINATIVE CLASSIFIER

Model conditional probability dist P(Y|X) directly
 uses sigmoid function for binary
 $P(Y=1|X) = \sigma(w^T X) = \frac{1}{1 + e^{-w^T X}}$
 softmax for multiclass classification

LOGISTIC REGRESSION

Goal: find probability that input x belongs to a certain class P(Y=1|X)
 BINARY
 $P(Y=1|X) = \frac{e^{-w^T X}}{1 + e^{-w^T X}}$
 $P(Y=0|X) = \frac{e^{-w^T X}}{1 + e^{-w^T X}}$
 log-odds = $\log\left(\frac{P(Y=1|X)}{P(Y=0|X)}\right)$

NEURAL NETWORK

1 $a_j^{(l)} = \sum_{i \in I^{(l-1)}} \theta_{ji}^{(l)} x_i^{(l-1)}$ (pre activation value for jth neuron in layer l)
 2 $x_j^{(l)} = g_{\sigma}(a_j^{(l)})$ weights neurons from prev layer
 3 $g_{\sigma}: \mathbb{R} \rightarrow \mathbb{R}$ activation function

CROSS ENTROPY LOSS

want to minimize loss minimize using SGD for 1 layer NN $\rightarrow \frac{\partial L}{\partial w_{ij}} = \frac{\partial L}{\partial a_{ij}} \frac{\partial a_{ij}}{\partial w_{ij}}$
 $L = -Y \ln(\hat{Y})$ batch n of m samples
 $J = -\frac{1}{m} \sum_{i=1}^m y_i \ln(\hat{y}_i)$
 chain rule: partial of loss wrt weight = partial of pre-activation value wrt weight

$$P(Y=k|X) = \frac{e^{z_k}}{\sum_{j=1}^K e^{z_j}}$$

negative log likelihood $Z(\theta) = -Y \log p(\hat{z}) - (C-Y) \log(1-p(\hat{z}))$

STOCHASTIC GRADIENT DESCENT

$$\theta_{k+1} = \theta_k - \epsilon_t \nabla_{\theta} Z(x_i, y_i; \theta)$$

BACKPROP:



- Forward pass: calculate preactivation value apply activation function
- Backward pass: Take gradients of loss wrt params of each layer - Use chain rule to propagate backwards - update weights w/ gradient descent after gradients calculated

loss $\frac{\partial L}{\partial z_{ij}} = \frac{\partial L}{\partial a_{ij}} = \frac{\partial L}{\partial y_{ij}} \frac{\partial y_{ij}}{\partial z_{ij}} = \frac{\partial L}{\partial y_{ij}} \mathbb{1}_{z_{ij} > 0} \rightarrow \frac{\partial L}{\partial z} = \frac{\partial L}{\partial y} \circ \mathbb{1}_{z > 0}$
 element wise indicator matrix

CONVOLUTIONAL NEURAL NETWORKS

- convolutional filter: same weights applied across many different locations
 $Z[d_1, d_2, n] = (X * W)[d_1, d_2, n] = \sum_i \sum_j W[i, j, c, n] X[d_1 + i, d_2 + j, c] + b[n]$
 # of channels

- pooling layer: downsample image \rightarrow introduces translational invariance



SKIP CONNECTIONS TO DEAL w/ VANISHING GRADIENTS





Figure 2. Residual learning: a building block.

SEQ2SEQ MODELS



ATTENTION


- Transformer encoder activation to key
- Transformer decoder activation to query
- attention score = dot product of query w/keys
- softmax to normalize attention scores
- send all of $h_1 \dots h_n$ through linear combo dictated by softmax of attn scores

PCA = principal component analysis

- STEP BY STEP:
- center the data around the mean
 - compute covariance matrix $Z = X^T X$
 - EigenDecomp: $X^T X = Q \Lambda Q^T$ or use SVD $X = U \Sigma V^T$ $V =$ eigenvectors $X^T X$
 - keep top k eigenvectors $Q_k = Q_{:,1:k}$ $V =$ eigenvectors $X^T X$
 - Project points down to subspace \rightarrow principal component scores
- * final dim reduced data: $\tilde{X}_k = \tilde{X}_k \Sigma_k^{-1/2} Q_k^T$
 \tilde{X}_k to reconstruct: $\tilde{X}_{reconst} = \tilde{X}_k \Sigma_k^{1/2} Q_k^T = X_k \tilde{X}_k^T$
 * recon loss: $\| \tilde{X}_{reconst} - \tilde{X} \|_F$ \rightarrow want smallest error possible

t-SNE = computed shortest pairwise distance between points

t-distribution = distribution heavier tail
 $Q_{ij} = \frac{(1 + \|x_i - x_j\|)^{-1}}{\sum_k (1 + \|x_i - x_k\|)^{-1}}$ change in computing stochastic neighbors




MULTIHEAD ATTN

- * multiple keys, queries, values for each time step
- * nonlinearities
- * masked decoding \rightarrow not allowed to look at future values


TRANSFORMER

- * stacked self attn layers w/ pos wise nonlinearities
- * easily parallelizable



SELF ATTENTION

want K, Q, V for all positions




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

```

Data: data set X = {x_1, x_2, ..., x_n}
cost function parameters: entropy constant C to set alpha
optimization parameters: number of iterations T, learning rate eta
Result: low-dimensional data representation Q^{(T)} = {q_1, q_2, ..., q_n}
begin
  compute pairwise affinities p_{ij} with entropy constant C to set alpha
  set P_{ij} = \frac{e^{-C p_{ij}}}{\sum_k e^{-C p_{ik}}}
  sample initial solution Q^{(0)} = {q_1, q_2, ..., q_n} from \mathcal{A}(0, 10^{-4})
  for t=1 to T do
    compute low-dimensional affinities q_{ij} (using Equation 4)
    compute gradient \frac{\partial L}{\partial q_{ij}} (using Equation 5)
    set Q^{(t)} = Q^{(t-1)} + \eta \frac{\partial L}{\partial q_{ij}}
  end
  
```

POSITIONAL ENCODING

- * KEEP OVERFITTING \rightarrow use periodicity



K MEANS CLUSTERING

$$\text{argmin}_{c_1, \dots, c_k} \sum_{k=1}^K \sum_{x \in C_k} \|x - c_k\|^2$$

c_k : cluster centroid
 C_k : cluster centroids

- steps:**
1. Compute partition by choosing closest centroid
 2. Compute centers by averaging over partition
 3. Continue until centers do not change
- soft k-means:** use softness of soft partition & distances
 $Y_{ik} = \text{softmax}(\frac{1}{2} \|x_i - c_k\|^2)$
 $\hat{c}_k = \sum_{i=1}^n Y_{ik} x_i$
 $\hat{c}_k = \sum_{i=1}^n r_{ik} x_i$

Voronoi tessellation for euclidean metric all boundaries are linear
 $\|x - x^A\|_2, \|x - x^B\|_2$

Mixture of Gaussians: model cluster as non-spherical gaussian

likelihood: $L_k = P(X_i) = \sum_{k=1}^K P(X_i | Z_i = k)$


$L = \prod_{i=1}^n P(X_i | Z_i = k) = \prod_{k=1}^K \prod_{i: Z_i = k} P(X_i | Z_i = k)$
 (learn these params)

Kleinberg Impossibility Thm

1. Scale invariance: stretching data = same cluster
2. consistency: stretching space between cluster = same cluster
3. richness: allowing to produce any arbitrary partition

Classifier decision outcomes



		MODEL PREDICTIONS	
		Positive	Negative
GROUND TRUTH	Positive	TP	FN
	Negative	FP	TN



sensitivity = $\frac{TP}{TP + FN}$
 specificity = $\frac{TN}{TN + FP}$
 miss rate = $\frac{FN}{TP + FN}$
 fall out = $\frac{FP}{TN + FP}$

Bias-Variance

Bias = avg diff between model output & truth
 low overfitting = high bias
 overfitting = low bias
 variance = variance over all possible train sets



FOR REGULARIZATION: increasing λ
 less flexible model prevents overfitting & higher bias lower variance

Bias = $E[\hat{y} - y] = E[E[\hat{y} - y | X]]$
 Variance = $\text{Var}(\hat{y})$

KNN Algorithm (supervised)

1. for each training sample $(x_i, y_i) \in D$ compute distance between x and x_i
2. choose set of training samples with k smallest distance
3. return majority label of samples in N

Pros: No training, learns complex nonlinear functions

Cons: High storage cost, slow inference, curse of dimensionality (worse in higher dim)

Decision Trees

At each node: split by feature \rightarrow traverse until you hit a leaf node = output

Greedy Algorithm: Next best attribute = feature + split that MAY increase gain / MIN entropy

Entropy: $H(Y) = -\sum_{k=1}^K P(Y=k) \log P(Y=k)$

surprise = $-\log P(Y=k)$

conditional entropy: $H(Y|X) = \sum_{x \in X} P(x) H(Y|X=x)$

Info gain: $I(S, A) = \text{Entropy}(S) - \sum_{s \in S} \frac{|s|}{|S|} \text{Entropy}(s)$

Binary (left/right) Info gain: $H(L|R) - \left(\frac{|L|}{|L|+|R|} H(L) + \frac{|R|}{|L|+|R|} H(R) \right)$

Bayes Optimal Classifier

$P(Y=c|X) = \frac{P(X|c)P(c)}{\sum_{c \in C} P(X|c)P(c)}$
 all $(b, c) \text{ id} \rightarrow \text{all } b \text{ id}$
 $P(a|b, c|d) = P(a|d)P(b|c|d)$
 $\sum P(a, b, c|d) = \sum P(a|d)P(b, c|d)$
 $P(a, b|d) = P(a|d)P(b|d)$

Bayes Decision Boundary

when 2 probabilities weighted w/ cost are equal

$L(c_1, 0) P(Y=0|X) = L(c_1, 1) P(Y=1|X)$

TRIPLET LOSS: minimize distance between reference & pos sample, maximize distance between reference & neg sample

$L_{triplet} = \max(0, d(X_{anchor}, X_{pos}) - d(X_{anchor}, X_{neg}) + \text{margin})$

N-Pair Loss: compare w/ multiple neg samples at same time

$L_{n-pair} = \frac{1}{n} \sum_{i=1}^n \log(1 + \sum_{j \neq i} \exp(-\langle x_i, x_j \rangle - \langle x_i, x_j \rangle))$

MARKOV CHAINS

$A = a_{1,1} \dots a_{1,n}$ \rightarrow N states
 $A = a_{1,1} \dots a_{n,n}$ \rightarrow transition probability matrix A where a_{ij} = Prob of moving from state i to j
 $\pi = \pi_1 \dots \pi_n$ \rightarrow initial probability dist over states

MARKOV ASSUMPTION:
 $P(q_t = a | q_{1:t-1}) = P(q_t = a | q_{t-1})$




for states x_1, \dots, x_n : JOINT DIST: $P(X_1, x_2, \dots, x_n) = P(X_1) \prod_{i=2}^n P(X_i | X_{i-1})$

w/ no independence assumptions: min # params = $K^2 - 1$ # of states

HIDDEN MARKOV MODELS

= events interested in cannot be observed directly

- HMM has
1. set of K states in $Q = \{1, \dots, K\}$
 2. transition matrix A where rows sum to 1
 3. seq of observations $Y_{1:T} = (y_1, \dots, y_T)$
 4. emission prob: $B_{ij} = P(y_t = j | X_t = i)$
- prob of generation $Y = 1$ generated from state $q_0 = k$



5. initial probability distribution (π)


Viterbi Algorithm = find the most likely sequence of hidden states in HMM

Viterbi Pseudocode

```

function VITERBI(O, S, H, Y, A, B): X
for each state i = 1, 2, ..., K do
  T[1][i] = -inf
  T[1][i] = 0
end for
for each observation j = 2, 3, ..., T do
  for each state i = 1, 2, ..., K do
    T[j][i] = max_k (T[j-1][k] * A_ki * B_ki)
    T[j][i] = argmax_k (T[j-1][k] * A_ki * B_ki)
  end for
  and for
  and for
  x_j = argmax_k (T[j][k])
  x_{j-1} = x_j
  for j = T, T-1, ..., 2 do
    x_{j-1} = T[j-1][x_j]
    x_j = x_{j-1}
  end for
  return X
end function
    
```

$\delta_t(i) = \max_k \{ \delta_{t-1}(k) a_{ki} \} b_{ki}$
 * DYNAMIC PROGRAMMING + backtracking



Bagging: Train M models with n' samples, sample with replacement \rightarrow REDUCES VARIANCE

Random forests: same as bagging but at each split, choose only random subset $p' = \sqrt{p}$ features, trained in parallel

Avg model performance: $\hat{y}_{bagging} = \frac{1}{M} \sum_{m=1}^M \hat{y}_m(x)$

weighted avg model: $\hat{y}_{weighted} = \sum_{m=1}^M \alpha_m \hat{y}_m(x)$

Boosting: 1. train next model conditioned on all prev models weights, 2. reweight models to minimize loss, 3. repeat


trained sequentially, reduces bias by weighting misclassified points more, stumps better for boosting

Probabilistic Graph Models
 Node = random variable
 Edge = dependence relationship

Joint Dist:

$P(a,b,c,d) = P(a)P(b|a)P(c|a,b)P(d|c)$
 # of params = 1 + 1 + 2 + 1 = 5
 $P(c_1, s) = P(c_1)P(c_2|c_1)P(c_3|c_1)P(c_4|c_2, c_3)P(s, c_4)$
 # of params = 1 + 2 + 2 + 4 + 1 = 10

MARKOV DECISION PROCESS (MDP)



transition dynamic: $P(S_t, R_t | S_{t-1}, A_{t-1})$
 conditioning on all history = conditioning on just previous state
 maximize sum of discounted rewards / return

Return: $G_t = R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots$
 $= \sum_{k=0}^{\infty} \gamma^k R_{t+k+1} = R_{t+1} + \gamma G_{t+1}$

State Value: $V_{\pi}(s) = E_{\pi}[G_t | S_t = s]$
 Action Value: $Q_{\pi}(s, a) = E_{\pi}[G_t | S_t = s, A_t = a]$
BELLMAN EQUATION


Value function: $V_{\pi}(s) = E_{\pi}[\sum_{t=0}^{\infty} \gamma^t r_{t+1} | S_0 = s]$
 For $R(s)$
 $V^*(s) = \max_a \sum_{s' \in S} P(s'|s, a) V^*(s')$
 For $R(s, a)$
 $V^*(s) = \max_a [R(s, a) + \gamma \sum_{s' \in S} P(s'|s, a) V^*(s')]$
 For $Q(s, a, a')$
 $Q^* = \max_a \sum_{s' \in S} P(s'|s, a) [R(s, a, s') + \gamma V^*(s')]$

POLICY ITERATION:
 1. Init value func & policy randomly
 2. Policy evaluation
 3. Policy improvement
 repeat until converged

Q function: $Q(s, a) = R(s, a) + \gamma \sum_{s'} P(s'|s, a) V(s')$
 Policy Eval: $V(s) = \sum_{a \in A} \pi(a|s) Q(s, a)$
 Policy Improvement: $\pi(s) = \arg \max_a Q(s, a)$
 Value Iteration: $V(s) = \max_a Q(s, a)$
 for each state s

Graph Neural Networks

Msg passing w/ either convolutional / attentional mechanisms



$N^i = \{j | (i, j) \in E\}$
 $N^i = \{j | (i, j) \in E, \alpha_j > 0\}$

Aggregation function: permutation invariant
 will it change if you permute input?
 No: $f(A) = f(A)$
 Yes: $f(A) \neq f(A)$

Neural network: permutation equivariant
 permutation of input = same permutation output
 $f(PA) = P f(A)$

translational equivariance: $[x] \rightarrow [x + c]$ has diff result
 rotational invariance: $[x] \rightarrow [R \cdot x]$ has same result

Langevin MCMC

score-based generative models: $S_0(x) \propto \nabla_x \log p_{data}(x)$
 Langevin dynamics to sample from dist: $x_{t+1} = x_t + \eta \nabla_x \log p_{data}(x_t) + \sqrt{2\eta} z_t$ where $z_t \sim N(0, I)$
 learning score-based models:

- Max likelihood: $\min_{\theta} E_{p_{data}} [\log p_{\theta}(x)]$
- Score matching: $\min_{\theta} E_{p_{data}} [\|\nabla_x \log p_{\theta}(x) - S_0(x)\|^2]$
- Denoising approaches

KERNELS = fun model of high dim set of features w/o blowing up complexity

- Project features to higher dim space ($x \rightarrow \beta(x)$)
- Rewrite all training / inference w/ only inner products between features $\beta(x_i)^T \beta(x_j)$
- Write kernel function that computes inner products between high-dim features $k(x_i, x_j) = \beta(x_i)^T \beta(x_j)$

2 conditions for kernel

- k has inner product rep: $\exists \beta: \mathbb{R}^D \rightarrow \mathcal{H}$ s.t. $\forall x_i, x_j \in \mathbb{R}^D, k(x_i, x_j) = \langle \beta(x_i), \beta(x_j) \rangle$
- for every sample $x_1, \dots, x_n \in \mathbb{R}^D$
 $K = \begin{bmatrix} k(x_1, x_1) & \dots & k(x_1, x_n) \\ \vdots & \ddots & \vdots \\ k(x_n, x_1) & \dots & k(x_n, x_n) \end{bmatrix}$ is PSD ($a^T K a \geq 0$)
 cond 1 implies cond 2

conv layer output: $H^i = \left[\frac{H \cdot K + \beta}{S} \right] + 1$
 $S(z) = \frac{1}{1 + e^{-z}} \rightarrow \frac{\partial S(z)}{\partial z} = S(z)(1 - S(z))$
 $\frac{d}{dx} f(x)$ where $f(x)$ is PDF: $-x f(x)$

least squares denoising

- easier to sample if we
- Add gaussian noise
 - sample from noisy distribution
 - denoise noisy to clean
- $\mathcal{L}(q(y)) = \int \|x - E(y)\|^2 p(x) p(y) dx$

$\nabla_y \log P(y) = \frac{\nabla_y P(y)}{P(y)} = \int \nabla_y P(y|x) \frac{p(x)}{P(y)}$
 $y = x + z \rightarrow P(y|x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(y-x)^2}{2\sigma^2}\right]$
 $q^*(y) = y + \sigma^2 \nabla \log p(y)$
 score func: $\nabla_y \log p(y) = \frac{1}{\sigma^2} [q^*(y) - y]$

$q^*(y) = E(x|y) = \int x P(x|y) dx = \int x P(y|x) P(x) dx = \int x P(y|x) p(x) dx = \frac{\int x P(y|x) p(x) dx}{P(y)}$
 $y_{t+1} = y_t + \frac{\sigma^2}{2} [q^*(y_t) - y_t] + \sqrt{2\sigma^2} z_t$
 Langevin to sample: step size \uparrow Normal