

1. Math Prerequisites

- Data matrix: $\mathbf{X} \in \mathbb{R}^{N \times D}$, weights $\mathbf{w} \in \mathbb{R}^D$, targets $\mathbf{y} \in \mathbb{R}^N$.
- Train / test errors: $\mathcal{E}_{\text{train}}, \mathcal{E}_{\text{test}}$ (avoid confusion with expectation $\mathbb{E}[\cdot]$).

- $\mathbf{A} \in \mathbb{R}^{m \times n}, \mathbf{B} \in \mathbb{R}^{n \times p} \Rightarrow (\mathbf{AB})^T = \mathbf{B}^T \mathbf{A}^T$.
- Trace: $\text{Tr}(\mathbf{A}) = \sum_i A_{ii} = \sum_i \lambda_i$. $\text{Tr}(\mathbf{ABC}) = \text{Tr}(\mathbf{BCA})$.
- Eigendecomp: $\mathbf{A} = \mathbf{QAQ}^T$ (Symm). SVD: $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$.
- Gradients: $\nabla_{\mathbf{x}}(\mathbf{b}^T \mathbf{x}) = \mathbf{b}$, $\nabla_{\mathbf{x}} \|\mathbf{x}\|_2^2 = 2\mathbf{x}$.
- Quadratic Form: $f(\mathbf{x}) = \mathbf{x}^T \mathbf{A} \mathbf{x}$.
 - $\nabla_{\mathbf{x}} f = (\mathbf{A} + \mathbf{A}^T)\mathbf{x}$. (If symm: $2\mathbf{A}\mathbf{x}$), $\nabla^2 f = \mathbf{A} + \mathbf{A}^T$.

- Hessian \mathbf{H} : Matrix of 2nd derivatives, $H_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}$.
 - Taylor (2nd order): $f(\mathbf{x} + \mathbf{d}) \approx f(\mathbf{x}) + \nabla f^T \mathbf{d} + \frac{1}{2} \mathbf{d}^T \mathbf{H} \mathbf{d}$.
 - Curvature: Eigenvalues of \mathbf{H} determine shape.
 - $\mathbf{H} \succ 0$ (PD) \Rightarrow Local Min (Bowl), $\mathbf{H} \prec 0$ (ND) \Rightarrow Local Max (Hill).
 - Indefinite (mixed signs) \Rightarrow Saddle Point.

Bayes: $P(Y|X) = \frac{P(X|Y)P(Y)}{P(X)} = \frac{P(X|Y)P(Y)}{\sum_y P(X|y)P(y)}$. If $X \perp\!\!\!\perp Y$,

- Expectation: $\mathbb{E}[aX + b] = a\mathbb{E}[X] + b$, then $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$.
- Variance: $\text{Var}(X) = \mathbb{E}[X^2] - (\mathbb{E}[X])^2$, $\text{Var}(aX + b) = a^2 \text{Var}(X)$.

Gaussian: $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{\exp(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu}))}{(2\pi)^{D/2} |\boldsymbol{\Sigma}|^{1/2}}$.

Linear Algebra Checks (Exam '24)

- Invertibility: For $\mathbf{X} \in \mathbb{R}^{N \times D}$, $\mathbf{X}^T \mathbf{X}$ is $D \times D$.
- If $D > N$ (High dim), $\text{rank}(\mathbf{X}) \leq N < D$. $\mathbf{X}^T \mathbf{X}$ is singular (not invertible). LS solution not unique.
- Variance Transformation: $\text{Var}(\mathbf{A}\mathbf{x} + \mathbf{b}) = \mathbf{A} \text{Var}(\mathbf{x}) \mathbf{A}^T$. (1D: $a^2 \sigma^2$).

2. Statistical Learning Theory

Bias-Variance Decomposition (MSE) True model: $y = f(x) + \epsilon$, $\mathbb{E}[\epsilon] = 0$, $\text{Var}(\epsilon) = \sigma^2$. Estimator $\hat{f}_D(x)$ trained on dataset D . Expectation is over dataset D (and noise), with fixed x : $\mathbb{E}_D[(y - \hat{f}_D(x))^2] = \text{Bias}^2(\hat{f}) + \text{Var}(\hat{f}) + \sigma^2$, $\text{Bias}^2(\hat{f}) = (\mathbb{E}_D[\hat{f}_D(x)] - f(x))^2$, $\text{Var}(\hat{f}) = \mathbb{E}_D[(\hat{f}_D(x) - \mathbb{E}_D[\hat{f}_D(x)])^2]$. **Overfitting (high Var)**: $\mathcal{E}_{\text{train}} \ll \mathcal{E}_{\text{test}}$. **Underfitting (high Bias)**: both errors high. **Regularization**: $\text{Var} \downarrow$, $\text{Bias} \uparrow$. More data: $\text{Var} \downarrow$.

Exact Formula Check (Exam) $\mathbb{E}[(Y - \hat{f})^2] = \text{Bias}[\hat{f}]^2 + \text{Var}[\hat{f}] + \sigma^2$. *Note*: Irreducible error σ^2 cannot be removed.

Train / Test / Validation

Standardization: $\tilde{x}_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j}$.

- Important: compute μ, σ using train only; apply to val/test.

Bayes Risk & Calibration (Theory) Bayes Opt. Classifier: $f^*(x) = \text{sign}(2\eta(x) - 1)$ where $\eta(x) = P(Y = 1|X = x)$. **Bayes Risk**: $\mathcal{L}^* = \mathbb{E}[\min(\eta(X), 1 - \eta(X))]$. **Calibration**: Minimizing surrogate ϕ -risk leads to f^* if ϕ convex, diff. at 0, $\phi'(0) < 0$.

- Square: $g^*(x) = 2\eta(x) - 1$. (Calibrated 已校准) **Logistic**: $g^*(x) = \log \frac{\eta(x)}{1-\eta(x)}$. (Calibrated) **Hinge**: $g^*(x) = \text{sign}(2\eta(x) - 1)$. (Calibrated)

3. Likelihood & MLE Principles

Definitions & Properties Given i.i.d. data $D = \{x_1, \dots, x_N\}$ and model param θ .

- Likelihood: $L(\theta) = P(D|\theta) = \prod_{n=1}^N P(x_n|\theta)$. (Func of θ , not x).
- Log-Likelihood (LL): $\ell(\theta) = \ln L(\theta) = \sum_{n=1}^N \ln P(x_n|\theta)$.
- MLE: $\hat{\theta}_{MLE} = \arg \max_{\theta} L(\theta) \equiv \arg \max_{\theta} \ell(\theta) \equiv \arg \min_{\theta} (-\ell(\theta))$.
- Equivalences: Constants (e.g., $\frac{1}{2}$) or scaling do NOT change arg max. (Exam '23 Q3).

Properties of MLE (Handwrite)

- Consistent: Converges to true θ^* as $N \rightarrow \infty$.
- Asymptotic Normality: $\hat{\theta} \approx \mathcal{N}(\theta^*, I(\theta^*)^{-1})$.
- Efficiency: Achieves Cramér-Rao lower bound (lowest variance).
- $I(\theta)$ is Fisher Info.

Key Connections (Derivations)

- Gaussian \rightarrow MSE: If noise $\epsilon \sim \mathcal{N}(0, \sigma^2)$, then $y_n \sim \mathcal{N}(\mathbf{w}^T \mathbf{x}_n, \sigma^2)$.

$$\max \sum \ln \left(C \cdot e^{-\frac{(y-\hat{y})^2}{2\sigma^2}} \right) \iff \min \sum (y_n - \mathbf{w}^T \mathbf{x}_n)^2$$

Maximizing Gaussian LL is equivalent to Minimizing MSE (Least Squares).

- Bernoulli \rightarrow Logistic Loss: If $y \in \{0, 1\}$, $P(y|\mathbf{x}) = \hat{y}^y (1 - \hat{y})^{1-y}$.

$$-\ln L(\mathbf{w}) = -\sum [y_n \ln \hat{y}_n + (1 - y_n) \ln(1 - \hat{y}_n)]$$

Maximizing Bernoulli LL is equivalent to Minimizing Cross-Entropy.

4. Linear Regression

Data $D = \{(\mathbf{x}_n, y_n)\}_{n=1}^N$, $\mathbf{x}_n \in \mathbb{R}^D$ (optionally augmented with 1). Model: $y_n = \mathbf{w}^T \mathbf{x}_n + \epsilon_n$. **Least Squares (MSE)** $L(\mathbf{w}) = \frac{1}{2N} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2$ **Gradient**: $\nabla_{\mathbf{w}} L = \frac{1}{N} \mathbf{X}^T (\mathbf{X}\mathbf{w} - \mathbf{y})$ **Normal Equation**: $\mathbf{X}^T \mathbf{X} \mathbf{w} = \mathbf{X}^T \mathbf{y}$.

If $\mathbf{X}^T \mathbf{X}$ invertible (full column rank): $\mathbf{w}^* = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$. **Invariance**: Centering/Scaling features does NOT change prediction performance for unregularized OLS.

Ridge Regression (L2) $\hat{\mathbf{w}}_{\text{Ridge}} = \arg \min_{\mathbf{w}} \frac{1}{2N} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \frac{\lambda}{2} \|\mathbf{w}\|_2^2$
 $\mathbf{w}^* = (\mathbf{X}^T \mathbf{X} + N\lambda \mathbf{I})^{-1} \mathbf{X}^T \mathbf{y}$.

- Always unique since $\mathbf{X}^T \mathbf{X} + N\lambda \mathbf{I} \succ 0$ for $\lambda > 0$.
- MAP view: Gaussian prior $\mathbf{w} \sim \mathcal{N}(0, \tau^2 \mathbf{I})$ gives $\lambda \propto 1/\tau^2$ (up to scaling conventions).

- Lasso Regression (L1)** $\hat{\mathbf{w}}_{\text{Lasso}} = \arg \min_{\mathbf{w}} \frac{1}{2N} \|\mathbf{y} - \mathbf{X}\mathbf{w}\|_2^2 + \lambda \|\mathbf{w}\|_1$
- Sparsity: promotes $w_j = 0$ (feature selection).
 - Subgradient: $\partial |w_j| = \{-1\}$ if $w_j < 0$, $\{1\}$ if $w_j > 0$, $[-1, 1]$ if $w_j = 0$.

Perceptron Convergence Proof Assumptions: Data $\|\mathbf{x}_n\| \leq R$, Linear Sep: $y_n \mathbf{w}_*^T \mathbf{x}_n \geq \gamma$. **Algorithm**: If err, $\mathbf{w}_{t+1} = \mathbf{w}_t + y_i \mathbf{x}_i$. (Init $\mathbf{w}_0 = \mathbf{0}$).

- Lower Bound (Dot Prod): $\mathbf{w}_*^T \mathbf{w}_{t+1} = \mathbf{w}_*^T \mathbf{w}_t + \underbrace{y_i \mathbf{w}_*^T \mathbf{x}_i}_{\geq \gamma} \geq \mathbf{w}_*^T \mathbf{w}_t + \gamma \implies \mathbf{w}_*^T \mathbf{w}_t \geq t\gamma$.

- Upper Bound (Norm):
 $\|\mathbf{w}_{t+1}\|_2^2 = \|\mathbf{w}_t\|_2^2 + \underbrace{2y_i \mathbf{w}_t^T \mathbf{x}_i}_{\leq 0(\text{mistake})} + \underbrace{\|\mathbf{x}_i\|_2^2}_{\leq R^2} \leq \|\mathbf{w}_t\|_2^2 + R^2 \implies \|\mathbf{w}_t\|_2^2 \leq tR^2$.

- Conclusion (Steps t): $t^2 \gamma^2 \leq (\mathbf{w}_*^T \mathbf{w}_t)^2 \leq \|\mathbf{w}_*\|_2^2 \|\mathbf{w}_t\|_2^2 \leq \|\mathbf{w}_*\|_2^2 tR^2 \implies t \leq \frac{R^2 \|\mathbf{w}_*\|_2^2}{\gamma^2}$. (Novikoff's Thm).

5. Optimization (优化算法)

Gradient Descent (GD, 梯度下降) 更新公式: $\mathbf{w}_{t+1} = \mathbf{w}_t - \gamma \nabla L(\mathbf{w}_t)$.

- If L convex (凸) and γ small enough, converges to global optimum.
- Cost per step (全量计算代价): $O(ND)$. (N samples, D dims).

Projected Gradient Descent For constrained opt (带约束优化) $\min_{\mathbf{w} \in C} L(\mathbf{w})$: $\mathbf{w}_{t+1} = \Pi_C(\mathbf{w}_t - \gamma \nabla L(\mathbf{w}_t))$

- Π_C : Projection Operator (投影算子). Maps point back to valid set C (e.g., if $\|\mathbf{w}\| > 1$, clip it to boundary).

Stochastic Gradient Descent (SGD, 随机梯度下降)

- Sample n (随机抽样): $\mathbf{w}_{t+1} = \mathbf{w}_t - \gamma \nabla L_n(\mathbf{w}_t)$.
- Cost per step: $O(D)$ (Fast!); noisy gradients (震荡大).
- Typically use decaying LR (学习率衰减) γ_t (or schedules).

Optimization Variants (变种)

- Momentum (动量法): Accumulate velocity (累积速度/惯性).
 $\mathbf{m}_{t+1} = \beta \mathbf{m}_t + (1 - \beta) \mathbf{g}_t$; $\mathbf{w}_{t+1} = \mathbf{w}_t - \gamma \mathbf{m}_{t+1}$
- Adam: Adaptive moments (自适应矩估计).
 $\mathbf{m}_{t+1} = \beta_1 \mathbf{m}_t + (1 - \beta_1) \mathbf{g}_t$ (1st moment/Mean 均值)
 $\mathbf{v}_{t+1} = \beta_2 \mathbf{v}_t + (1 - \beta_2) \mathbf{g}_t^2$ (2nd moment/Var 方差)
 $\mathbf{w}_{t+1} = \mathbf{w}_t - \gamma \hat{\mathbf{m}}_{t+1} / (\sqrt{\hat{\mathbf{v}}_{t+1}} + \epsilon)$ (Scale by std dev).
- Sign-SGD: Direction only (仅用符号). $\mathbf{w}_{t+1} = \mathbf{w}_t - \gamma \text{sign}(\mathbf{g}_t)$.

Subgradient Definition (次梯度 - 解决不可导) For convex non-diff function f (e.g., L1 norm $|x|$), \mathbf{g} is a subgradient at \mathbf{w} if: $f(\mathbf{u}) \geq f(\mathbf{w}) + \mathbf{g}^T (\mathbf{u} - \mathbf{w})$, $\forall \mathbf{u}$. *Meaning*: The tangent plane is always below the function (切线在函数下方).

Newton's Method (牛顿法) $\mathbf{w}_{t+1} = \mathbf{w}_t - \gamma [\mathbf{H}(\mathbf{w}_t)]^{-1} \nabla L(\mathbf{w}_t)$

- Quadratic convergence (二次收敛): Very fast near optimum.
- Cost: Hessian inversion is $O(D^3)$. (Too slow for large dims).

6. Logistic Regression

Model: $\hat{y} = \sigma(z)$, $z = \mathbf{w}^T \mathbf{x}$. **Sigmoid Deriv**: $\sigma' = \sigma(1 - \sigma)$.

Loss (NLL): $L(\mathbf{w}) = -\sum [y_n \log \hat{y}_n + (1 - y_n) \log(1 - \hat{y}_n)]$.

Gradient: $\nabla_{\mathbf{w}} L = \frac{1}{N} \sum (\hat{y}_n - y_n) \mathbf{x}_n = \frac{1}{N} \mathbf{X}^T (\hat{\mathbf{y}} - \mathbf{y})$.

Hessian Derivation: Chain rule on error term $(\hat{y}_n - y_n)$:

$$\frac{\partial}{\partial \mathbf{w}} (\dots) = \frac{\partial \hat{y}_n}{\partial z_n} \mathbf{x}_n = \hat{y}_n (1 - \hat{y}_n) \mathbf{x}_n = s_n \mathbf{x}_n \text{ (Scalar } s_n \cdot \text{ Vector } \mathbf{x}_n)$$

Matrix Form: Summing $s_n \mathbf{x}_n \mathbf{x}_n^T$ over N samples:

$$\mathbf{H} = \frac{1}{N} \mathbf{X}^T \mathbf{S} \mathbf{X} \text{ where } \mathbf{S} = \text{diag}(s_1, \dots, s_N).$$

- Convexity: $\forall \mathbf{v} \neq \mathbf{0}, \mathbf{v}^T \mathbf{H} \mathbf{v} = \frac{1}{N} \sum s_n (\mathbf{v}^T \mathbf{x}_n)^2 \geq 0 \implies$ Global Min.
- Newton Update: $\mathbf{w} \leftarrow \mathbf{w} - (\mathbf{X}^T \mathbf{S} \mathbf{X})^{-1} \mathbf{X}^T (\hat{\mathbf{y}} - \mathbf{y})$.

Formal Convexity Proof Steps (Exam '24) To prove convexity, show Hessian $\mathbf{H} \succeq 0$.

- Derive Grad: $\nabla L = \frac{1}{N} \mathbf{X}^T (\hat{\mathbf{y}} - \mathbf{y})$.
- Derive Hessian: $\mathbf{H} = \frac{1}{N} \mathbf{X}^T \mathbf{S} \mathbf{X}$ with $\mathbf{S} = \text{diag}(\hat{y}_n(1 - \hat{y}_n))$.
- Proof: For any vector \mathbf{v} , $\mathbf{v}^T \mathbf{H} \mathbf{v} = \frac{1}{N} \mathbf{v}^T \mathbf{X}^T \mathbf{S} \mathbf{X} \mathbf{v} = \frac{1}{N} (\mathbf{X} \mathbf{v})^T \mathbf{S} (\mathbf{X} \mathbf{v})$.
- Conc: Since $\hat{y}(1 - \hat{y}) > 0$, \mathbf{S} is positive definite diagonal, so quadratic form ≥ 0 .

7. SVM & Kernels

Soft-Margin SVM (Primal) obj: Minimize *Hinge Loss* + L_2 Regularization:

$$J(\mathbf{w}, b) = \frac{1}{N} \sum_{n=1}^N \underbrace{\max(0, 1 - y_n(\mathbf{w}^T \mathbf{x}_n + b))}_{\text{Hinge Loss}} + \frac{\lambda}{2} \|\mathbf{w}\|_2^2$$

SGD Updates (Pick random (\mathbf{x}_n, y_n)):

- Case 1 ($y_n f(\mathbf{x}_n) < 1$): Violation or inside margin. $\mathbf{w} \leftarrow (1 - \gamma\lambda) \mathbf{w} + \gamma y_n \mathbf{x}_n$.
- Case 2 ($y_n f(\mathbf{x}_n) \geq 1$): Correct & safe. $\mathbf{w} \leftarrow (1 - \gamma\lambda) \mathbf{w}$ (Weight decay only)

Note: $C \propto \frac{1}{\lambda}$. Small λ / Large $C \rightarrow$ Hard Margin (Overfit risk).

Dual Formulation (Kernelized) $\max_{\alpha} \sum_{n=1}^N \alpha_n - \frac{1}{2} \sum_{i,j} \alpha_i \alpha_j y_i y_j k(\mathbf{x}_i, \mathbf{x}_j)$ **Constraints**: $0 \leq \alpha_n \leq C$ AND $\sum \alpha_n y_n = 0$.

KKT Conditions & Support Vectors (SVs)

- $\alpha_n = 0$: Correct ($y_n f(\mathbf{x}_n) > 1$). Not an SV.
- $\alpha_n > 0$: Support Vector.
 - $0 < \alpha_n < C$: On Margin ($y_n f(\mathbf{x}_n) = 1$). Use for b^* !
 - $\alpha_n = C$: Inside Margin/Error ($y_n f(\mathbf{x}_n) < 1$).

- $\mathbf{w}^* = \sum \alpha_n y_n \mathbf{x}_n$ (Exists only if linear kernel).
- Bias b^* : Average over SVs with $0 < \alpha_k < C$: $b^* = \frac{1}{|\mathcal{S}|} \sum_{k \in \mathcal{S}} (y_k - \sum_{m \in \text{SV}} \alpha_m y_m k(\mathbf{x}_m, \mathbf{x}_k))$
- Predict: $y_{\text{new}} = \text{sign}(\sum_{n \in \text{SV}} \alpha_n y_n k(\mathbf{x}_n, \mathbf{x}_{\text{new}}) + b^*)$.

Representer Theorem (Handwrite) For any loss L and regularizer $\Omega(\|\mathbf{w}\|_2)$ (strictly monotonic), the optimal solution \mathbf{w}^* lies in the span of the data:

$$\mathbf{w}^* = \sum_{n=1}^N \alpha_n \mathbf{x}_n = \mathbf{X}^T \boldsymbol{\alpha}$$

This allows kernelizing linear models: $\mathbf{w}^T \mathbf{x} = \sum \alpha_n \langle \mathbf{x}_n, \mathbf{x} \rangle = \sum \alpha_n k(\mathbf{x}_n, \mathbf{x})$.

Kernels & Mercer's Theorem Def: $k(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle$. Valid iff Gram Matrix $\mathbf{K} \succeq 0$. **Common Kernels**:

- Linear: $\mathbf{x}^T \mathbf{x}'$. Poly: $(\gamma \mathbf{x}^T \mathbf{x}' + c)^d$.
- RBF (Gaussian): $k(\mathbf{x}, \mathbf{x}') = \exp(-\gamma \|\mathbf{x} - \mathbf{x}'\|^2)$.
 - $\gamma \approx \frac{1}{2\sigma^2}$. Large γ : Narrow peak \rightarrow Overfitting.
 - Small γ : Flat \rightarrow Underfitting (Linear-like).

Kernel Construction Rules (Closure): If k_1, k_2 are valid kernels, $c > 0$, $f(\cdot)$ polynomial with positive coefficients:

- Sum: $k_1 + k_2$ Product: $k_1 \cdot k_2$ Scale: $c \cdot k_1$
- Mapping: $k(\phi(\mathbf{x}), \phi(\mathbf{x}'))$ Exp: $\exp(k_1)$
- Poly: $f(k_1)$ (e.g., $k_1^2 + 3k_1 + 1$)

Disproving Validity (Exam '24) To show $k(\mathbf{x}, \mathbf{y})$ is invalid, find set $\{\mathbf{x}_1, \mathbf{x}_2\}$ s.t. Gram Matrix \mathbf{K} is not PSD ($\det(\mathbf{K}) < 0$). **Example**: $k(x, y) = (xy + c)^d$ with $c < 0$. Pick $x_1 = 1, x_2 = -1$.

$$\mathbf{K} = \begin{pmatrix} (1+c)^d & (c-1)^d \\ (c-1)^d & (1+c)^d \end{pmatrix}. \text{ Check if } \det(\mathbf{K}) = (1+c)^{2d} - (c-1)^{2d} < 0.$$

8. Unsupervised Learning

PCA Centered data. Cov $\mathbf{S} = \frac{1}{N} \mathbf{X}^T \mathbf{X}$. $\max_{\mathbf{u}} \mathbf{u}^T \mathbf{S} \mathbf{u}$ s.t. $\mathbf{u}^T \mathbf{u} = 1$ Principal components: eigenvectors of \mathbf{S} with largest eigenvalues.

K-Means Objective: $J = \sum_{k=1}^K \sum_{n \in C_k} \|\mathbf{x}_n - \boldsymbol{\mu}_k\|_2^2$

- Initialize $\boldsymbol{\mu}_k$.
- Assign $z_n = \arg \min_k \|\mathbf{x}_n - \boldsymbol{\mu}_k\|_2^2$.
- Update $\boldsymbol{\mu}_k = \frac{1}{|C_k|} \sum_{n \in C_k} \mathbf{x}_n$.

GMM + EM: $p(\mathbf{x}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$

- E-step: $\gamma_{nk} = p(z_n = k|\mathbf{x}_n) = \frac{\pi_k \mathcal{N}(\mathbf{x}_n|\boldsymbol{\mu}_k)}{\sum_j \pi_j \mathcal{N}(\mathbf{x}_n|\boldsymbol{\mu}_j)}$.

M-step: update $\pi_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k$ via weighted MLE.

Weighted GMM Updates Obj: $\max \sum w_n \log \sum \pi_k \mathcal{N}$. E-Step q_{nk} same as std. M-Step: $\boldsymbol{\mu}_k = \frac{\sum w_n q_{nk} \mathbf{x}_n}{\sum w_n q_{nk}}$, $\boldsymbol{\Sigma}_k = \frac{\sum w_n q_{nk} (\mathbf{x}_n - \boldsymbol{\mu}_k)(\mathbf{x}_n - \boldsymbol{\mu}_k)^T}{\sum w_n q_{nk}}$, $\pi_k = \frac{\sum w_n q_{nk}}{\sum w_n}$.

GMM Free Parameters Count (Exam '22) Total params for K clusters in D dim:

- Full Covariance: $K - 1$ (weights) + KD (means) + $K \frac{D(D+1)}{2}$ (cov).
- Spherical ($\sigma^2 I$): $K - 1 + KD + K = KD + 2K - 1$.

9. Neural Networks

1. Backpropagation (The δ Rule) Forward: $\mathbf{z}^{(l)} = \mathbf{W}^{(l)} \mathbf{a}^{(l-1)} + \mathbf{b}^{(l)}$, $\mathbf{a}^{(l)} = \phi(\mathbf{z}^{(l)})$.
 Backward: Propagate error $\delta^{(l)} = \frac{\partial L}{\partial \mathbf{z}^{(l)}}$.

- Output: $\delta^L = \nabla_{\mathbf{a}^L} L \odot \phi'(\mathbf{z}^L)$.
- Hidden: $\delta^l = (\mathbf{W}^{l+1})^T \delta^{l+1} \odot \phi'(\mathbf{z}^l)$.
- Grads: $\frac{\partial L}{\partial \mathbf{W}^l} = \delta^l (\mathbf{a}^{l-1})^T$, $\frac{\partial L}{\partial \mathbf{b}^l} = \delta^l$.

Matrix Shapes (Exam '22) Layer $\mathbf{Y} = \mathbf{XW} + \mathbf{b}$. Given $\delta_Y = \partial L / \partial \mathbf{Y}$.

- $\frac{\partial L}{\partial \mathbf{W}} = \mathbf{X}^T \delta_Y$ $\frac{\partial L}{\partial \mathbf{X}} = \delta_Y \mathbf{W}^T$ $\frac{\partial L}{\partial \mathbf{b}} = \text{sum}_{\text{rows}}(\delta_Y)$.

Universal Approximation (Barron) A single hidden layer NN with sufficiently many neurons and non-polynomial activation can approximate any continuous function on a compact set.

2. Complexity: Width H vs. Depth L , For Fully Connected layers:

- Widening ($H \rightarrow 2H$): Params/Ops $\propto H^2$. (Quad. cost).
- Deepening ($L \rightarrow 2L$): Params/Ops $\propto L$. (Linear cost).
- Result: Deep & Narrow is computationally cheaper than Shallow & Wide for same param count.

3. Softmax Properties Given logits \mathbf{x} , prediction is $\arg \max_k x_k$.

- Shift ($\mathbf{x} \rightarrow \mathbf{x} + \mathbf{b}$): Probs unchanged. Acc & Loss Unchanged.
- Scale ($\mathbf{x} \rightarrow \alpha \mathbf{x}$, $\alpha > 0$): Order preserved.
 - Accuracy: Unchanged.
 - Loss: Changes! $\alpha > 1$ (sharp) \rightarrow Loss \downarrow . $\alpha < 1$ (flat) \rightarrow Loss \uparrow .

4. Regularization & Init

• Xavier: $\text{Var}(W) \approx \frac{2}{n_{in} + n_{out}}$ (Sigmoid). He: $\frac{2}{n_{in}}$ (ReLU).

• Batch Norm: $\frac{\alpha - \mu_B}{\sigma_B} \cdot \gamma + \beta$. (Dep. on batch).

• Layer Norm: $\frac{x_{nk} - \mu_n}{\sigma_n}$. Indep of batch. (For Transf.)

• Dropout: Train: mask w/ p . Test: scale weights by $(1 - p)$.

Activation Derivatives GeLU: $\phi(z) = z \sigma(cz) = z \frac{1}{1 + e^{-cz}}$. ($c \approx 1.702$).

- Deriv: $\phi'(z) = \sigma(cz) + z \cdot \sigma'(cz) (1 - \sigma(cz)) \cdot c$.
- Limits: $z \rightarrow \infty$: $\phi'(z) \rightarrow 1 + 0 = 1$ (Like ReLU). $z \rightarrow -\infty$: $\phi'(z) \rightarrow 0$ (Like ReLU).
- Relation: Smooth approx of ReLU. Non-monotonic.

10. Convolutional Neural Networks (CNNs)

Dimensions & Layers (尺寸计算) Output Size (输出边长): $W_{out} = \lfloor \frac{W_{in} - K + 2P}{S} \rfloor + 1$
 Output Depth (输出深度): C_{out} = number of filters (滤波器数量)

- W_{in} : Input Size (输入边长); K : Kernel Size (卷积核大小)
- P : Padding (填充); S : Stride (步长)
- Stride (S): Step size (滑动步长).
- Pooling (池化): W, H reduced by S (尺寸减小); C unchanged (通道不变).
- Flatten (展平): 3D vol \rightarrow 1D vec. Size: $(W \cdot H \cdot C)$.

Learnable Parameters (参数量 - Exam '24)

• Conv Layer: $\underbrace{(K^2 \cdot C_{in} + 1)}_{\text{weights+bias per filter}} \cdot \underbrace{C_{out}}_{\text{\# filters}}$

- FC Layer: $(N_{in} + 1) \cdot N_{out}$. (+1 for bias 偏置).
- Pool Layer: 0 parameters (无参数).

Architecture Insights

- ResNet: $y = F(x) + x$; identity shortcut (恒等跳跃) eases opt.
- Receptive Field (感受野): Input area "seen" by pixel. \uparrow with depth.
- 1x1 Conv: Changes channel C (调整通道数); keeps W, H .

11. Transformers & NLP

1. Self-Attention Attention ($\mathbf{Q}, \mathbf{K}, \mathbf{V}$) = $\text{softmax}\left(\frac{\mathbf{QK}^T}{\sqrt{d_k}}\right) \mathbf{V}$

- $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ from same source (Self-Attn). $Q \in \mathbb{R}^{S \times d_k}$. 这里是 d 下标 k.
- Scale $\sqrt{d_k}$: Prevents dot product from blowing up \rightarrow prevents softmax saturation \rightarrow prevents vanishing gradients. Complexity: $O(S^2 d + S d^2)$. S^2 is bottleneck (Attn Matrix). vs Linear Attn $O(S)$.

2. Positional Embeddings (PE, 位置编码)

- Transformer w/o PE is Permutation Equivariant (置换等变性) \rightarrow acts like a Set Function.
- 考点: w/o PE, it treats input as a "bag of words".
 - If inputs are identical ($x_i = x_j$), outputs are identical ($z_i = z_j$), regardless of position.
 - Cannot solve order-tasks (e.g., "output 0 at odd pos" 奇数位输出 0).

- PE is added (按元素相加) (not concat) to input embeddings.
- $E_{final} = E_{word} + E_{pos}$. (Dimension stays d_{model}).

3. BERT vs. GPT (Architecture)

- BERT (Encoder): Masked LM. Predicts masks in parallel.
 - Independence Assumption: Predicts $P(w_A, w_B | C) \approx P(w_A | C) P(w_B | C)$. Ignores dependency between masked tokens.
- GPT (Decoder): Autoregressive. Sequential generation.
 - Masked Attn: Enforces causality (cant see future).

4. Theoretical Limits Transformer has bounded computation per token ($O(1)$ depth). Cannot solve problems requiring linear time $\Omega(i)$ w.r.t input index.

12. Generative AI

VAE (Variational Autoencoder) Objective: Maximize ELBO (Evidence Lower Bound). $\mathcal{L} = \mathbb{E}_q[\log p(\mathbf{x} | z)] - D_{KL}(q(z | \mathbf{x}) \parallel p(z))$

- Reconstruction
- Regularization
- Rec. Term: Ensure output \approx input.
- KL Term: Force latent z to be $\mathcal{N}(0, I)$.
- Reparameterization Trick: $z = \mu + \sigma \odot \epsilon$ (allows backprop).

Diffusion Models (DDPM)

- Forward (No train): $q(x_t | x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t I)$. Property: $q(x_t | x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t} x_0, (1 - \bar{\alpha}_t) I)$.
- Reverse (Train): Train NN ϵ_θ to predict the noise. Approx $q(x_{t-1} | x_t)$ with $p_\theta(x_{t-1} | x_t)$.
 $\mu_\theta(x_t, t) = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right)$.
- Loss (MSE): $\mathcal{L} = \|\epsilon - \epsilon_\theta(x_t, t)\|^2$. Score Matching $\approx \nabla \log p(x_t)$.

Score Matching Derivation True Score $\nabla_{x_t} \log p(x_t)$ is intractable. Tractable form: Condition on x_0 . $\nabla_{x_t} \log p(x_t | x_0) = \nabla \log \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t} x_0, (1 - \bar{\alpha}_t) I) = -\frac{x_t - \sqrt{\bar{\alpha}_t} x_0}{1 - \bar{\alpha}_t} = -\frac{\sqrt{1 - \bar{\alpha}_t} \epsilon}{\sqrt{1 - \bar{\alpha}_t}^2} = -\frac{\epsilon}{\sqrt{1 - \bar{\alpha}_t}}$. \implies Matching score is equiv to predicting noise ϵ .

GAN (Generative Adversarial Nets) Min-Max Game (Zero-sum): \uparrow see up

- D (Discriminator): Maximize. Real $x \rightarrow 1$, Fake $G(z) \rightarrow 0$.
- G (Generator): Minimize. Trick D so $D(G(z)) \rightarrow 1$.
- Issue: Mode collapse, unstable training.

13. Metrics & Misc

Metrics: $P = TP / (TP + FP)$, $R = TP / (TP + FN)$, $F1 = 2PR / (P + R)$

Conf Mat: $y \rightarrow \hat{y}$: $1 \rightarrow 1$ (TP), $0 \rightarrow 0$ (TN), $0 \rightarrow 1$ (FP/Type I), $1 \rightarrow 0$ (FN/Type II). 混淆矩阵: 预测 $\hat{y} = 1$: 对是 TP/错是 FP(Type I); 预测 $\hat{y} = 0$: 错是 FN(Type II)/对是 TN.

Fairness Criteria Variables: A (sensitive), Y (true), \hat{Y} (pred). (xxx same).

- 1. Independence: $\hat{Y} \perp\!\!\!\perp A$. $P(\hat{Y} = 1 | A = a) = P(\hat{Y} = 1 | A = b)$. (Rate \hat{Y}).
 - 2. Separation: $\hat{Y} \perp\!\!\!\perp A | Y$. $P(\hat{Y} = 1 | Y = y, A = a) = P(\hat{Y} = 1 | Y = y, A = b)$. (TPR/FPF).
 - 3. Sufficiency: $Y \perp\!\!\!\perp A | \hat{Y}$. $P(Y = 1 | \hat{Y} = s, A = a) = P(Y = 1 | \hat{Y} = s, A = b)$. (PPV).
- Impossibility Thm: If base rates differ ($P(Y|A) \neq P(Y|A')$) and predictor imperfect, cannot satisfy all 3 simultaneously.

Mitigation Strategies (Handwrite)

- Pre-processing: Adjust features x to be uncorrelated with A .
- In-processing: Add regularization term to loss during training.
- Post-processing: Adjust thresholds/outputs of learned classifier.

14. Adversarial Robustness

Goal: Maximization Problem Find perturbation δ to maximize loss (fool the model) under constraint (imperceptible): $\max_{\delta} L(f(\mathbf{x} + \delta), y)$ s.t. $\|\delta\|_p \leq \epsilon$ White-box: Model parameters/gradients known. Black-box: Only outputs known (transfer attack).

Black-Box Variants (Handwrite)

- Score-based: Can query continuous confidence scores (probabilities).
- Decision-based: Can only query discrete hard labels (0/1).
- 1. FGSM (Fast Gradient Sign Method) Constraint: ℓ_∞ norm (Change each pixel $\leq \epsilon$). Linear loss around \mathbf{x} (1st order Taylor): $\mathbf{x}_{adv} = \mathbf{x} + \epsilon \cdot \text{sign}(\nabla_{\mathbf{x}} L(\theta, \mathbf{x}, y))$
- Why Sign? Optimal direction for ℓ_∞ constraint (corners of the hypercube). One-step: Fast but underfits (linear approximation).
- 2. PGD (Projected Gradient Descent) Iterative FGSM. More powerful, harder to defend. $\mathbf{x}_{t+1} = \Pi_{\mathbf{x} + S}(\mathbf{x}_t + \alpha \cdot \text{sign}(\nabla_{\mathbf{x}} L))$
- Project (II): Clip II: range $[x - \epsilon, x + \epsilon] \cap [0, 1]$.
- Universal 1st-order adversary.
- 3. ℓ_2 Attack (Gradient Direction) Total energy limited ($\|\delta\|_2 \leq \epsilon$).

$\delta = \epsilon \cdot \frac{\nabla_{\mathbf{x}} L}{\|\nabla_{\mathbf{x}} L\|_2}$ (No sign function!)

Defense Adversarial Training:

Train on $\{\mathbf{x}_{adv}, y\}$. $\min_{\theta} \mathbb{E}_{(\mathbf{x}, y) \sim D} [\max_{\delta} L(f_{\theta}(\mathbf{x} + \delta), y)]$.

Linear Model Exact Attack (Exam '21) Model $y = \mathbf{w}^T \mathbf{x}$, Loss $(y - \mathbf{w}^T(\mathbf{x} + \delta))^2$, constraint $\|\delta\|_{\infty} \leq \epsilon$.

- Goal: Maximize error. Make $\mathbf{w}^T \delta$ as negative/positive as possible.
- Optimal δ^* : $\delta_i = -\epsilon \cdot \text{sign}(w_i)$ (if trying to decrease pred).
- Max Loss: $(y - \mathbf{w}^T \mathbf{x} - \underbrace{\mathbf{w}^T \delta^*}_{-\epsilon \|\mathbf{w}\|_1})^2 = (y - \mathbf{w}^T \mathbf{x} + \epsilon \|\mathbf{w}\|_1)^2$.

15. Matrix Factorization (RecSys)

Model Definition Approximate rating matrix $\mathbf{R} \approx \mathbf{P} \times \mathbf{Q}^T$. Prediction for user u , item i : $\hat{r}_{ui} = \mathbf{p}_u \cdot \mathbf{q}_i^T = \sum_{f=1}^k p_{u,f} \cdot q_{i,f}$ where $\mathbf{P} \in \mathbb{R}^{m \times k}$, $\mathbf{Q} \in \mathbb{R}^{n \times k}$ (k : latent dim, $k \ll m, n$).

Objective Function (Regularized SE) Minimize loss J over observed set \mathcal{K} : $J = \sum_{(u,i) \in \mathcal{K}} (r_{ui} - \mathbf{p}_u \cdot \mathbf{q}_i^T)^2 + \lambda (\|\mathbf{p}_u\|_2^2 + \|\mathbf{q}_i\|_2^2)$

SGD Update Rules (Calculations) 1. Calculate Error: $e_{ui} = r_{ui} - \hat{r}_{ui}$

2. Update Factors: (Learning rate η). Note: Term $-\lambda \mathbf{p}_u$ is weight decay (prevents overfit). $\mathbf{p}_u \leftarrow \mathbf{p}_u + \eta (e_{ui} \mathbf{q}_i - \lambda \mathbf{p}_u)$; $\mathbf{q}_i \leftarrow \mathbf{q}_i + \eta (e_{ui} \mathbf{p}_u - \lambda \mathbf{q}_i)$

- Biased MF: $\hat{r}_{ui} = \mu + b_u + b_i + \mathbf{p}_u \cdot \mathbf{q}_i^T$.
- ALS (Alternating Least Squares): Fix \mathbf{P} , solve \mathbf{Q} (OLS); Fix \mathbf{Q} , solve \mathbf{P} .
- Good for implicit feedback & parallelization.

SVD vs. MF: SVD: Requires dense matrix (no missing). Slow $O(mn^2)$. MF: Handles sparse data. Fast $O(|\mathcal{K}|k)$. Approx.

16. Contrastive Learning (SimCLR)

Objective: InfoNCE Loss Maximize similarity of views from same image (pos); push away others (neg). $\mathcal{L} = -\log \frac{\exp(\text{sim}(z_i, z_j) / \tau)}{\sum_k \exp(\text{sim}(z_i, z_k) / \tau)}$

Key Components

- Augmentations: NOT random/all. Composition is key.
 - Crop & Color Jitter: Critical. Forces model to learn shape/texture, not color histograms.
- Batch Size: Needs large batch (more negatives).
- Temperature τ :
 - Low τ : Softmax becomes sharp. Model focuses heavily on Hard Negatives (difficult samples).
 - High τ : Gradients uniform across all negatives.

17. K-Nearest Neighbors (KNN)

- Classification: Majority vote of K neighbors. Regression: Average of K neighbors. Hyperparameter K :
 - Small K (e.g., 1): High Variance, Low Bias. Complex boundary. Overfits (Train Error = 0).
 - Large K (e.g., N): Low Variance, High Bias. Simple boundary. Underfits (Predicts majority class).
- Distance: Sensitive to feature scaling. Must standardize data first!
- Curse of Dimensionality: As $D \uparrow$, all points become equidistant; Euclidean distance fails.

18. Advanced Exam Concepts (Tricks)

- Convexity Operations (Crucial) Let f, g be convex functions.
 - Sum: $\alpha f(x) + \beta g(x)$ is convex ($\forall \alpha, \beta \geq 0$).
 - Max: $h(x) = \max(f(x), g(x))$ is Convex.
 - Min: $h(x) = \min(f(x), g(x))$ is NOT guaranteed convex.
 - Composition: $f(g(x))$ convex if f convex non-decreasing & g convex.

2. Loss Minimizers

- Minimize MSE (L_2): $\sum (x_i - \theta)^2 \implies \theta = \text{Mean}(x)$.
- Minimize MAE (L_1): $\sum |x_i - \theta| \implies \theta = \text{Median}(x)$.

3. ReLU Network Construction Identities

- Identity: $x = \text{ReLU}(x) - \text{ReLU}(-x)$. Weights: $\mathbf{w} = [1, -1]^T$, $\mathbf{v} = [1, -1]^T$. Abs: $|x| = \text{ReLU}(x) + \text{ReLU}(-x)$. Weights: $\mathbf{w} = [1, -1]^T$, $\mathbf{v} = [1, 1]^T$. Max: $\max(a, b) = b + \text{ReLU}(a - b)$. $F(x) = \text{ReLU}(x_1 - x_2) + x_2$.

4. Algorithm Nuances

- K-Means+ Init: Pick c_1 random. Sample next c with prob $P(x) \propto D(x)^2$ (Sq distance to nearest center). K-Means Convergence: J is non-increasing ($J_{t+1} \leq J_t$). Converges to local min. EM Algorithm: Likelihood $L(\theta)$ is non-increasing per iter. Batch Norm (Test Mode): Use global running μ, σ (EMA from train), NOT batch stats.
- GAN Optimal: $D^*(x) = 0.5$. If D is perfect (100% acc), gradients vanish \rightarrow training fails.