Δelta 2026

Day 1 - Latent Spaces & Neural Optimization

Machine Learning Foundations for Traffic Systems

TIME: 3 HOURS • TOTAL: 50 POINTS

Each problem is scored independently. Justify all answers with complete proofs. Partial credit is awarded for significant progress. No calculators or electronic devices are permitted.

Problem 1 — 7 points

(a) In a continuous latent state traffic model, the cost of routing a real-valued volume of traffic $x_i \geq 0$ along edge i is modeled by a strictly convex function $C(x_i)$. Suppose n temporally identical intervals process a combined total traffic volume V. Prove that the total congestion cost $\sum_{i=1}^{n} C(x_i)$ is strictly minimized when the traffic is uniformly distributed across all intervals, i.e., $x_i = V/n$ for all i.

(b) A machine learning load-balancer distributes a continuous total traffic mass T across two parallel highways. Let the allocated traffic be $t_1, t_2 \geq 0$ such that $t_1 + t_2 = T$. The highways have latent capacity constants $c_1, c_2 > 0$. The total systemic energy cost is modeled as $E = \frac{t_1^2}{c_1} + \frac{t_2^2}{c_2}$. Prove that the lowest possible system energy is exactly $\frac{T^2}{c_1 + c_2}$, and identify the unique traffic assignment (t_1, t_2) that achieves this.

Problem 2 — 9 points

The Principal Traffic Corridors. A city's traffic authority monitors vehicle flow across n intersections over m days ($m \geq n$). Let $A \in \mathbb{R}^{m \times n}$ be the traffic matrix, where $A_{i,j}$ represents the volume of vehicles at intersection j on day i. To simplify the traffic control system, the authority wishes to approximate A with a rank-k matrix B representing k "principal traffic corridors," where $k < n$.

Prove that the minimum possible sum of squared errors $\sum_{i,j}(A_{i,j} - B_{i,j})^2$ over all rank-k matrices B is exactly equal to $\sum_{i=k+1}^{n} \sigma_i^2$, where $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_n \geq 0$ are the singular values of the traffic matrix A.

Problem 3 — 10 points

The Automated Routing Optimizer. An AI regulates traffic light timings across a grid. The total delay experienced by commuters for a given configuration of timings $x \in \mathbb{R}^d$ is modeled by a strictly convex and twice-differentiable delay function $f(x)$. Suppose there exist constants $0 < \mu \leq L$ such that the Hessian satisfies $\mu I \preceq \nabla^2 f(x) \preceq LI$ for all $x \in \mathbb{R}^d$.

The AI updates the timings iteratively using the rule $x_{t+1} = x_t - \frac{1}{L}\nabla f(x_t)$. Let x^* be the unique optimal timing configuration that minimizes f.

Prove that the delay difference converges exponentially:

$$f(x_{t+1}) - f(x^*) \;\leq\; \left(1 - \frac{\mu}{L}\right)\left(f(x_t) - f(x^*)\right).$$

Problem 4 — 11 points

The Piecewise Linear Zoning Boundaries. A centralized routing system divides a continuous city space $\mathbb{R}^d$ into distinct "autonomous traffic zones." The system uses m linear sensors. The output of the i-th sensor at position $x \in \mathbb{R}^d$ is activated if $w_i^T x + b_i > 0$, where $w_i \in \mathbb{R}^d$ and $b_i \in \mathbb{R}$.

A traffic zone is defined as a region of space where the activation pattern (which sensors are strictly positive and which are zero or negative) is constant. Assuming the hyperplanes defined by $w_i^T x + b_i = 0$ are in **general position** (no two are parallel, no $d+1$ intersect at a point, etc.), prove that the maximum number of distinct autonomous traffic zones formed is exactly

$$\sum_{i=0}^{d} \binom{m}{i}.$$

Problem 5 — 13 points

The Latent Delivery Relocation. A fleet of autonomous bikes starts distributed across a city according to a continuous Gaussian density $\rho_0 = \mathcal{N}(m_0, \Sigma_0)$ in $\mathbb{R}^d$. By the end of the day, they must be repositioned to match a target demand distribution $\rho_1 = \mathcal{N}(m_1, \Sigma_1)$. The cost of moving a bike from x to y is proportional to the squared Euclidean distance $\|x - y\|^2$.

The dispatcher seeks the optimal transport map $T : \mathbb{R}^d \to \mathbb{R}^d$ that pushes ρ_0 forward to ρ_1 while minimizing the total expected transportation cost.

Prove that the optimal movement map is an affine transformation $T(x) = A(x - m_0) + m_1$, and explicitly find the unique symmetric positive-definite matrix A in terms of Σ_0 and Σ_1.

End of Round 3