Roll. No:	Name:

Total Marks: 33, Total Time: 150 mins

Instructions

1. The exam is divided into three sections: short answer questions in the first section, followed by two sections with problems that require a detailed solution. For the first section, provide the final answer ONLY. For the second and third sections, provide detailed answers showing all the necessary steps.
2. Use rough sheets for any calculations *if necessary,* and do not submit the rough sheets. Do not use a pencil for writing the answers.
3. Assume standard data whenever you feel that the given data is insufficient. However, do quote your assumptions explicitly.

I Short answer questions (Answer any six)

Note: If more than six questions are answered, then the first six answers will be considered for evaluation.

$\boxed{1.5}$ 1. For a twice continuously differentiable function f, let

$$\theta^* = \arg\min_{\theta} f(\eta) + \nabla f(\eta)^{\intercal}(\theta - \eta) + \frac{1}{2\alpha}\left\|\eta - \theta\right\|^2, \text{ and}$$
$$\tilde{\theta}^* = \arg\min_{\theta} f(\eta) + \nabla f(\eta)^{\intercal}(\theta - \eta) + \frac{1}{2\alpha}(\theta - \eta)^{\intercal}\nabla^2 f(\eta)(\theta - \eta).$$

Provide explicit expressions for θ^* and $\tilde{\theta}^*$.

$\boxed{1.5}$ 2. Consider the following function:

$$f(x_1, x_2) = x_1^2 x_2 + x_2^3 x_1.$$

Provide a descent direction for f at $(2, 1)$.

$\boxed{1.5}$ 3. True or false: Let $f : \mathbb{R}^n \to \mathbb{R}^n$ be a contraction mapping with modulus α under the Euclidean norm. Let x^* denote the fixed point of f. Then, $f(x) \leq x$ implies $x^* \leq x$. Here $\leq$ is element-wise.

$\boxed{1.5}$ 4. Consider the function $f(x) = \frac{x}{2} + \sqrt{1 - x^2}$ on the domain $[-1, 1]$. Find the maximum and minimum of this function on the given domain.

$\boxed{1.5}$ 5. Suppose x^* minimizes the convex and L-smooth function $f : \mathbb{R}^d \to \mathbb{R}$. If $f(x_n) \stackrel{\text{wp1}}{\to} f(x^*)$, what can be said about x_n?
(A) $x_n \stackrel{\text{wp1}}{\to} x^*$ (B) Nothing (C) $x_n \stackrel{\text{p}}{\to} x^*$ (D) $x_n \stackrel{\text{d}}{\to} x^*$

Here $\stackrel{\text{p}}{\to}$ and $\stackrel{\text{d}}{\to}$ denote convergence in probability and distribution, respectively.

$\boxed{1.5}$ 6. Say whether each of the statements below is true or false for a discounted MDP.

(a) There exists a policy π such that $\lim_{k\to\infty} T_\pi^k J < \lim_{k\to\infty} T^k J$ for some vector J.

(b) If the single stage cost $g(x,a,x')$ is replaced by $g'(x,a,x') = g(x,a,x') + 10, \forall (x,a,x')$, then the optimal policy remains unaffected.

1.5 7. Consider a optimization of a convex and smooth function f using zeroth-order information. Consider the following two settings: In the first setting, the function f is observable exacty (i.e., without noise) at any input x. On the other hand, in the second setting, for any input x, one can obtain observations $\hat{f}(x) = f(x) + \epsilon$, where ϵ is standard normal. Suppose we form a zeroth-order gradient estimate, say $\widehat{\nabla} f(\cdot)$ using the simultaneous perturbation method, and perform gradient descent, i.e., the following update (with obvious notation):

$$x_{t+1} = x_t - a_t \widehat{\nabla} f(x_t).$$

As a function of t specify bounds on $\mathbb{E}(f(x_t)) - f(x^*)$, where x^* is a minimum. Use the oh-notation, and specify bounds for both settings mentioned above.

II. Medium answer problems (Answer any two)

Note: Provide detailed answers with proper justification. Further, if more than two questions are answered, then the first two answers will be considered for evaluation.

1. A discounted MDP is specified below.

   States $\{1,2\}$, actions $\{a,b\}$ in state 1, and $\{c,d\}$ in state 2. The transition probabilities are

$$p_{11}(a) = p_{12}(a) = 0.5; \quad p_{11}(b) = 0.8, p_{12}(b) = 0.2;$$
$$p_{21}(c) = 0.4, p_{22}(c) = 0.6; \quad p_{21}(d) = 0.7, p_{22}(d) = 0.3;$$

   The discount factor $\alpha = 0.9$.

   The time-invariant single-stage costs are as follows:

$$g(1,a,1) = -9, g(1,a,2) = -3, g(1,b,1) = -4, g(1,b,2) = -4,$$
$$g(2,c,1) = -3, g(2,c,2) = 7, g(2,d,1) = -1, g(2,d,2) = 10.$$

   For each of the policies given below, find the expected discounted cumulative cost with start state 1.

1.5 (a) Policy π: $\pi(1) = a, \pi(2) = c$.

1.5 (b) Policy $\widetilde{\pi}$: $\widetilde{\pi}(1) = b, \widetilde{\pi}(2) = d$.

2 (c) Is the optimal policy different from the two policies listed above? Why or why not?

2. Consider a machine that can be in one of the following two states: 'good' and 'bad'. If the machine is in a good state in the current period, then it will transition to a bad state in the next period with probability (w.p.) p_1. On the other hand, a machine in 'bad' state in the current period has to undergo maintenance, and transitions to a 'good' state in the next period w.p. p_2, and remains in the 'bad' state w.p. $(1-p_2)$. Suppose the machine in 'good' state earns A INR (i.e., a cost of $-A$ INR) per-period, while the per-period maintenance cost for a machine in 'bad' state is B INR.

   Answer the following:

3 (a) Formulate this problem in the infinite horizon discounted cost framework, with $\alpha \in (0,1)$ denoting the discount factor. What is the expected discounted cost of the machine for each possible initial state, i.e., 'good' and 'bad'?

2 (b) Suppose a new machine starts in the 'good' state, and costs M INR. Compare the purchase cost to the expected discounted cost to infer when it is optimal to a buy a new machine.

5 3. Consider a linear stochastic approximation algorithm with the following update iteration:

$$\theta_{n+1} = \theta_n + a(n)\left(A_{n+1}\theta_n + b_{n+1}\right),$$

where $a(n)$ is the step size, while A_n and b_n are matrices and vectors that satisfy

$$\mathbb{E}\left[A_{n+1} \mid \theta_1, \ldots, \theta_n\right] = A, E\left[b_{n+1} \mid \theta_1, \ldots, \theta_n\right] = b,$$

where A is a negative-definite matrix. Moreover, $\mathbb{E}\left[\|(A_n - A)\|^2\right] \leq C_1$ and $\mathbb{E}\left[\|b_n - b\|^2\right] \leq C_2$. Use the Kushner-Clark lemma to establish asymptotic convergence of θ_n.

III. Long answer problems (Answer any two)

Note: Provide detailed answers with proper justification.
If more than two questions are answered, then the first two answers will be considered for evaluation.

1. Let $f(x) = \frac{1}{p}\sum_{i=1}^{p} f_i(x)$, where f_i is a L-smooth function, for $i = 1, \ldots, p$ and f is μ-strongly convex.

Suppose you are given gradient inputs from a zeroth-order oracle, i.e., for any $i = 1, \ldots, p$ and any x, an optimization algorithm can obtain a (random) gradient estimate $\widehat{\nabla} f_i(x)$ that satisfies

$$\mathbb{E}\left\|\widehat{\nabla} f_i(x) - \nabla f_i(x)\right\| \leq C_1\delta^2, \text{ and } \mathbb{E}\left\|\widehat{\nabla} f_i(x) - \nabla f_i(x)\right\|^2 \leq \frac{C_2}{\delta^2}.$$

In the above, δ is a bias-variance tradeoff parameter, which an optimization algorithm gets to choose before querying the oracle.

Answer the following:

1 (a) For minimizing f, write the update iteration of a SGD-type algorithm with stepsize denoted by α_k and iterate by x_k. The gradient inputs are from the oracle defined above, and the SGD-type algorithm is required to use only one of the f_is in each update iteration.

4 (b) Analyze the non-asymptotic behavior of the algorithm from the part above. In particular, provide a bound on $\mathbb{E}\left[f(x_{k+1}) - f(x^*)\right]$. Specify the SGD algorithm parameters α_k and δ_k precisely. Here δ_k is the bias-variance tradeoff parameter used for obtaining the gradient estimate.

Show your work in arriving at the bound.

2 (c) What would be the corresponding rate for a GD algorithm that uses gradients for each f_i.

2. Consider the following problem, which is a variant of mean estimation. For a continuous random variable (r.v.) X with cumulative distribution function F and for a given $\alpha \in (0,1)$, define

$$q_\alpha(X) = F^{-1}(\alpha).$$

Notice that $q_\alpha(X)$ is the median of the distribution of X when $\alpha = 0.5$.

Let $\{X_n\}_{n\geq 1}$ be a independent sequence of r.v.s with common distribution F.

Answer the following:

$\boxed{3}$ (a) Derive a stochastic approximation algorithm for estimating $q_\alpha(X)$ for a pre-specified α. Let q_n denote the iterate after observing samples $X_1, \ldots, X_n$. The algorithm should be iterative, i.e., given an estimate q_n at time instant n and a new sample X_{n+1}, the algorithm should perform an incremental update using q_n, X_{n+1} to arrive at q_{n+1}.

$\boxed{3}$ (b) Provide a sketch of the convergence analysis of the algorithm from part (a) above, in particular, to specify suitable assumptions so that q_n converges almost surely to $q_\alpha(X)$?

$\boxed{2}$ (c) Consider the following alternative observation model is as follows: At time instant n, the stochastic approximation algorithm picks a threshold, say T, and the environment returns a boolean that indicates whether $X_{n+1} < T$ or not. Discuss the changes to the stochastic approximation from the part above to handle this threshold-based model.

3. Given a dataset $D_n = \{(a_i, y_i); i = 1, .., n\}$ with $a_i \in \mathbb{R}^d$ and $y_i \in \mathbb{R}$, consider the linear regression problem of finding the minimizer x^* of the following objective:

$$J(x) = \frac{1}{2n} \sum_{i=1}^{n} (y_i - x^T a_i)^2. \tag{1}$$

Let Φ be the $n \times d$ matrix whose i^{th} row is a_i^T. Assume Φ has full column rank. Let $A = \frac{1}{n}\Phi^\intercal\Phi$.

Answer the following:

$\boxed{1}$ (a) Find the gradient and Hessian of J at a given point x.

$\boxed{1}$ (b) Show that J is strongly-convex.

$\boxed{1}$ (c) Write down the update rule for a gradient descent (GD) algorithm to find the minimizer x^* of J.

$\boxed{3}$ (d) Show that the gradient descent iterate, say x_t, after t iterations, satisfies the following bounds:

$$\|x_t - x^*\|^2 \leq (x_0 - x^*)^\intercal (I - \alpha A)^{2t} (x_0 - x^*), \tag{2}$$
$$J(x_t) - J(x^*) \leq (x_0 - x^*)^\intercal (I - \alpha A)^{2t} A (x_0 - x^*), \tag{3}$$

where α is the constant stepsize used by the GD algorithm.

$\boxed{2}$ (e) Let μ and L denote the smallest and largest eigenvalues of A. Using (2), establish the following bound for GD:

$$\|x_t - x^*\|^2 \leq \left(1 - \frac{\mu}{L}\right)^{2t} \|x_0 - x^*\|^2 .$$