

Optimization Algorithms Overview

Method	Details
Simplex Method	<p>Description: Solves LPs by moving along the edges of the feasible region to find the optimal vertex.</p> <p>Convergence: Exponential in the worst case but generally polynomial time for practical problems.</p> <p>When to Use: Suitable for small to medium-sized LPs, especially when a vertex solution is desired.</p>
Gradient Descent (GD)	<p>Description: Iteratively updates variables in the direction of the negative gradient.</p> <p>Convergence Rates: Convex and L-smooth: $O(1/k)$. μ-strongly convex and L-smooth: $O((1 - \mu/L)^k)$. γ-PL condition and L-smooth: $O((1 - \gamma/L)^k)$.</p> <p>When to Use: For smooth and (strongly) convex problems; careful step size selection is crucial.</p>
Newton's Method	<p>Description: Utilizes second-order derivative (Hessian) to find stationary points.</p> <p>Convergence: μ-strongly convex, L-smooth, Lipschitz continuous Hessian: Globally linear with fast local quadratic convergence, once it is close enough to optimum.</p> <p>When to Use: For smooth problems requiring high precision and solving linear systems with the Hessian isn't too expensive.</p>
BFGS	<p>Description: A quasi-Newton method approximating the Hessian matrix iteratively using the secant condition.</p> <p>Convergence: μ-strongly convex, L-smooth, Lipschitz continuous Hessian: Globally linear. Exhibits fast superlinear convergence once it is close enough to optimum.</p> <p>When to Use: When exact Hessians in Newton's method are too costly to compute or invert, but faster convergence than GD is desired.</p>
L-BFGS	<p>Description: Memory-efficient version of BFGS, storing only a limited amount of information to approximate the Hessian.</p> <p>Convergence: Linear convergence, performs better than (A)GD in practice but no supporting theory.</p> <p>When to Use: Ill-conditioned large-scale optimization problems with memory constraints.</p>
Interior-Point Methods (IPM)	<p>Description: Solves constrained optimization problems using barrier functions.</p> <p>Convergence: Polynomial time for LPs, QPs, and conic optimization problems.</p> <p>When to Use: For small to medium-scale problems with complex constraints. When a reliable, high-accuracy solution is required.</p>
Accelerated Gradient Descent (Nesterov's)	<p>Description: Enhances basic GD by incorporating a momentum term.</p> <p>Convergence Rates: Convex and L-smooth: Optimal $O(1/k^2)$. μ-strongly convex and L-smooth: $O((1 - \sqrt{\mu/L})^k)$.</p> <p>When to Use: For smooth, convex problems needing faster convergence than standard GD and provable rates.</p>
Stochastic Gradient Descent (SGD)	<p>Description: Computes the gradient based on a subset of data at each iteration.</p> <p>Convergence Rates: Convex and Lipschitz continuous: $O(1/\sqrt{k})$ with decaying step size $\eta_k = O(1/\sqrt{k})$. μ-strongly convex, L-smooth:</p> <ol style="list-style-type: none"> linear to a ball centered at the optimum of radius $O(\epsilon)$, when using appropriate <i>fixed</i> step size equal to $O(\epsilon)$. $O(1/k)$ with decaying step size $\eta_k = O(1/k)$. <p>When to Use: Large-scale or online learning problems where full gradient computation is intractable.</p>
Stochastic Variance Reduced Gradient (SVRG)	<p>Description: Mitigates gradient variance in SGD by periodically computing a full gradient.</p> <p>Convergence Rates: L-smooth and convex: $O(1/k)$. μ-strongly convex and L-smooth: Linear, $O((1 - \mu/L)^k)$.</p> <p>When to Use: Large-scale learning with variance reduction needs, feasible when full gradient computation intermittently is affordable.</p>
Projected Gradient Descent	<p>Description: Extends GD by projecting onto the feasible set after each iteration.</p> <p>Convergence Rates: Convex and L-smooth: $O(1/k)$. μ-strongly convex and L-smooth: Linear $O((1 - \mu/L)^k)$.</p> <p>When to Use: Constrained optimization within a convex feasible set with an easy-to-compute projection operator.</p>