Abstract

Isotonic regression is a constrained statistical method used to estimate parameters (e.g., population means) under order restrictions, such as monotonicity (e.g., $\mu_1 \leq \mu_2 \leq \cdots \leq \mu_k$). It minimizes a weighted sum of squared deviations, $\sum(Y_i - \mu_i)^2 w_i$, subject to pairwise constraints defined by a quasi-order or matrix inequalities. Key computational tools include the **Pool Adjacent Violators Algorithm (PAVA)**, which iteratively pools violating adjacent blocks, and connections to quadratic programming. The method extends to generalized linear models, exponential families (e.g., Poisson, binomial), and nonparametric settings. Applications span hypothesis testing, ANOVA under order constraints, and shape-restricted function estimation. Theoretical insights link it to projections onto convex sets and greatest convex minorants. Isotonic regression provides a flexible framework for incorporating prior knowledge of parameter orderings, balancing computational efficiency with rigorous statistical inference.