Source task and conventions

Examine the limit in Corollary 14.6 and construct a counterexample at $\lambda = 0$.

Conventional solution and proof

(a) With the source constants held fixed, the additive terms are

$$\frac{4Mr^2}{\lambda m} + \left(\frac{8Mr^2}{\lambda} + M\right)\sqrt{\frac{\log(1/\delta)}{2m}},$$

which diverge as $\lambda \downarrow 0$ when the relevant constants are nonzero. Thus this corollary gives no finite unregularized guarantee. Allowing a loss bound M to depend on λ cannot improve this conclusion without additional control. (b) Consider one-dimensional least squares with m examples, $m-1$ copies of $(x, y) = (0, 0)$, and one example $(\varepsilon, 1)$, $0 < \varepsilon \leq 1$. The unique empirical minimizer is $w = 1/\varepsilon$. Replace the last example by $(\varepsilon, 0)$; the unique minimizer is $w' = 0$. Both datasets have bounded inputs and labels. At the test point $(1, 0)$ their squared losses differ by $1/\varepsilon^2$, which is arbitrarily large. Therefore no finite uniform stability bound holds over this input/label class, even for fixed m. The failure is caused by arbitrarily small design singular values, not by choosing among nonunique minimizers.

Structural or analytic derivation

Regularization supplies the lower spectral bound $\lambda I \preceq XX^\top + \lambda I$. At zero regularization a bound is possible only on a restricted class with a uniform positive design-eigenvalue bound, or another hypothesis that controls prediction changes. The counterexample identifies the missing assumption exactly.

Comparative cost and preferred method