The denominator is strictly positive: $I - H = \lambda C \succ 0$. Therefore (c) has the corrected expression

$$\boxed{\widehat{R}_{\mathrm{LOO}} = \frac{1}{m}\sum_i \left(\frac{\widehat{y}_i - y_i}{1 - H_{ii}}\right)^2 = \frac{1}{m}\sum_i \left(\frac{\alpha_i}{C_{ii}}\right)^2, \quad \alpha = Cy.}$$

The source prints H_{ii} in the denominator; with $K = 2I, \lambda = 1, y = \mathbf{1}$, the true LOO risk is 1, while that expression is $1/4$.

(d) If all diagonal entries equal γ, then $\widehat{R}_{\mathrm{LOO}} = \widehat{R}_S/(1-\gamma)^2$. Since $0 \preceq H \prec I$, $0 \le \gamma < 1$. For positive empirical risk the errors coincide exactly at $\gamma = 0$. If the empirical risk is zero they coincide for every admissible γ; with $\lambda > 0$ and this strictly regularized formulation, zero training residual forces $y = 0$. Under averaged-loss conventions, deletion changes the effective coefficient unless λ is adjusted to keep the unnormalised penalty fixed.

Structural or analytic derivation

An independent derivation partitions $A = K + \lambda I$ around coordinate i. Its Schur complement $s = a - b^\top B^{-1} b$ is positive; block inversion gives $C_{ii} = 1/s$ and $\alpha_i = (y_i - b^\top B^{-1} y_{-i})/s$. The deleted prediction is $b^\top B^{-1} y_{-i}$, so the deleted residual is exactly α_i / C_{ii}. This proves the shortcut without solving m separate systems.

Comparative cost and preferred method

Independent dense refitting uses m factorizations of order $m-1$, giving $O(m^4)$ arithmetic. One factorization plus all inverse-diagonal entries costs $O(m^3)$, including triangular solves against the identity; storage remains $O(m^2)$. Solving only for α is insufficient: the diagonal must also be computed. The matched benchmark checks every residual against explicit deletion. The identity does not apply exactly to logistic regression or to a preprocessing step that is re-estimated after each deletion without further analysis.