WEEKLY PARTICIPATION 4

Assume that our population follows the model where the log-odds of y are linear given $\mathbf{x}$ with bias and coefficients ω_0 and $\boldsymbol{\omega}$ as before. Assume also that we expect that the true parameter $\boldsymbol{\omega}$ is likely to be sparse. The targets are in ± 1.

State:

- the appropriate loss function $\ell(u, v)$ for fitting the data,

- the appropriate regularizer $\Omega(\boldsymbol{\omega})$,

- and the objective function for learning the parameters ω_0 and $\boldsymbol{\omega}$ using RERM, given that we have n training examples.