

## Lecture 6

### Problem 1

Consider the case of (regularized) linear regression. Which type of regularization is traditionally associated with sparsity, that is, producing a model where coefficients corresponding to non-predictive features are 0?

- (a) Ridge (i.e.,  $\ell_2$ -regularization)
- (b) LASSO (i.e.,  $\ell_1$ -regularization)

### Problem 2

Alex trains a neural network to predict life expectancy based on weekly alcohol consumption and weekly exercise. He uses observational data from the Federal Office of Statistics. For a weekly consumption of two liters of beer and six hours of exercise the neural network predicts an average life expectancy of 80 years. His network also predicts that the average life expectancy would increase to 85 years if the weekly exercise hours are increased from six to seven. Does this mean that the network is able to reason about counterfactual questions?

- (a) Yes, the network can answer questions of the type “What would happen if...?”, which means that it can reason about counterfactuals.
- (b) No, the network is not necessarily capturing the causal relationships that determine life expectancy.
- (c) It depends on the architecture of the network.