

REAL WORLD EXAMPLES

Cancer prediction

- Trained model to predict "patient has cancer" from medical records. Excellent results on holdout test data!
- Turns out one of the features was hospital name, and these can be highly correlated w/ having cancer (e.g. cancer specialty hospitals)
- But not helpful when hospital not yet assigned.

↳ Label leakage

Literature

- Predict author political affiliations based on the "mind metaphors" they used
- Model trained with (input sentence, ^{author} political label) and high accuracy
- But researchers, in splitting up sentences into training / test / validation put all authors in each. This was too much context. When assigning only one set for a given author, results were worse.
- Know what the data represents!