

It's a jungle out there

semi-supervised learning    **overfitting**    stochastic gradient descent    **SVM**    *Q learning*  
Gaussian processes    **deterministic noise**    data snooping    learning curves  
*distribution-free*    *linear regression*    VC dimension    mixture of experts  
collaborative filtering    nonlinear transformation    **sampling bias**    *neural networks*    *no free lunch*  
**decision trees**    *RBF*    *training versus testing*    noisy targets    *Bayesian prior*  
active learning    **linear models**    bias-variance tradeoff    weak learners  
*ordinal regression*    cross validation    logistic regression    **data contamination**  
**ensemble learning**    types of learning    perceptrons    **hidden Markov models**  
exploration versus exploitation    error measures    *kernel methods*    graphical models  
**is learning feasible?**    **soft-order constraint**  
*clustering*    regularization    weight decay    *Occam's razor*    *Boltzmann machines*

# The map

