

*Classification problems in systems biology are often characterized by very large numbers of input features, many of which are irrelevant to the classification task. Unfortunately, which input features are indeed relevant is rarely known a priori. We have been working to develop algorithms in which feature selection is an integral part of the classifier learning process. These algorithms do not adopt a traditional "filter" or "wrapper" strategy for feature selection, but rather jointly optimize feature selection and classifier design through a single objective function. This is accomplished by means of a Bayesian formulation in which a prior promotes sparsity of the learned classifier weights. We present a quick sketch of the details of the approach, along with some results of the application of these algorithms to a few high-dimensional problems in systems biology.*