Abstract

David Kepplinger, Peter Filzmoser and Kurt Varmuza:

Variable selection with genetic algorithms using repeated cross-validation of PLS regression models as fitness measure

Genetic algorithms are a widely used method in chemometrics for extracting variable subsets with high prediction power. Most fitness measures used by these genetic algorithms are based on the ordinary least-squares fit of the resulting model to the entire data or a subset thereof. Due to multicollinearity, partial least squares regression is often more appropriate, but rarely considered in genetic algorithms due to the additional cost for estimating the optimal number of components.

We introduce two novel fitness measures for genetic algorithms, explicitly designed to estimate the internal prediction performance of partial least squares regression models built from the variable subsets. Both measures estimate the optimal number of components using cross-validation and subsequently estimate the prediction performance by predicting the response of observations not included in model-fitting. This is repeated multiple times to estimate the measures' variations due to different random splits. Moreover, one measure was optimized for speed and more accurate estimation of the prediction performance for observations not included during variable selection. This leads to variable subsets with high internal and external prediction power.

Results on high-dimensional chemical-analytical data show that the variable subsets acquired by this approach have competitive internal prediction power and superior external prediction power compared to variable subsets extracted with other fitness measures.