

## Application of feature selection for predicting leaf chlorophyll content in oats (*Avena sativa* L.) from hyperspectral imagery

W.R. Želazny<sup>1,2,\*</sup>

<sup>1</sup>Crop Research Institute, Division of Crop Management Systems, Drnovská 507/73, CZ161 06 Praha 6 Ruzyně, Czech Republic

<sup>2</sup>Czech University of Life Sciences Prague, Faculty of Engineering, Department of Agricultural Machines, Kamýcká 129, CZ165 00 Praha 6 Suchbátka, Czech Republic

\*Correspondence: [wzelazny@vurv.cz](mailto:wzelazny@vurv.cz)

**Abstract.** Feature selection can improve predictions generated by *partial least squares* models. In the context of hyperspectral imaging, it can also enable the development of affordable devices with specialized applications. The feasibility of feature selection for oat leaf chlorophyll estimation from hyperspectral imagery was assessed using a public domain dataset. A wrapper approach resulted in a simplistic model with poor predictive performance. The number of model inputs decreased from 94 to 3 bands when a filter approach based on the *minimum redundancy*, *maximum relevance* criterion was attempted. The filtering led to improved prediction quality, with the *root mean square error* decreasing from 0.17 to 0.16 g m<sup>-2</sup> and *R*<sup>2</sup> increasing from 0.57 to 0.62. Accurate predictions were obtained especially for low chlorophyll levels. The obtained model estimated leaf chlorophyll concentration from near infra-red reflectance, canopy darkness, and its blueness. The prediction robustness needs to be investigated, which can be done by employing an ensemble methodology and testing the model on a new dataset with improved ground-truth measurements and additional crop species.

**Key words:** remote sensing, imaging spectroscopy, unmanned aerial vehicles, partial least squares, reproducibility.

### INTRODUCTION

The indispensability of chlorophyll for plant photosynthesis (Sims & Gamon, 2002; Main et al., 2011) and its contribution to crop optical properties (Ollinger, 2011) make the estimation of leaf chlorophyll concentration an important remote sensing application. In large-scale assessments, leaf chlorophyll remote sensing is useful for yield prediction (Moharana & Dutta, 2016). At finer spatial scales, it can be used for the delineation of management zones for precision agriculture (Miao et al., 2009). As chlorophyll breaks down under stress, its monitoring provides information about the crop status, and enables a timely intervention to prevent the yield loss (Peñuelas et al., 1995; Sims & Gamon, 2002).

Traditional broad-band optical remote sensing relies on vegetation indexes for assessing crop status (e.g., Basso et al., 2016; Domínguez et al., 2017; Barbosa et al., 2019). Consequently, it is of limited use for estimating the concentrations of individual pigments, such as leaf chlorophyll. Many of these indexes have been adapted for use

with hyperspectral imaging products (Miao et al., 2009; Verrelst et al., 2019, often leading to improved results (e.g., Miao et al., 2009). Moharana & Dutta (2016) evaluated ten indexes in terms of rice chlorophyll prediction from proximal spectroradiometric data. Some of the band combinations gave unsatisfactory estimates despite their high performance in other experimental settings, which is a common problem for vegetation indexes. On the other hand, the formulations that excelled during the screening provided realistic maps of rice chlorophyll concentration when applied to EO-1 Hyperion imagery. The limited index transferability across crops can be in part related to differences between plant architectures (Ollinger, 2011). A study involving six crop species evaluated the robustness of relationships between vegetation indexes and leaf chlorophyll with respect to canopy structural parameters. A total of 58 formulations were tested; of this number, only 2 were considered truly robust when applied to both measured and simulated spectra (Zou et al., 2015). Corti et al. (2018) published a meta-analysis intended to identify factors that foster accurate estimation of maize biochemical parameters from optical measurements. Their results suggest that satisfactory predictions can be obtained by avoiding certain families of vegetation indexes—regardless of sensor type, acquisition model, and crop developmental stage. The article indicates that only statistically significant relationships were included in the study, which means that this finding needs to be approached with caution. A recent review by Hatfield et al. (2019) cites additional studies devoted to vegetation indexes suitable for chlorophyll estimation. According to the authors, vegetation indexes should be a first choice in remote sensing applications, as they avoid computational challenges of more sophisticated approaches.

Yet, the main advantage of hyperspectral imagery lies in the possibility of applying ‘full-spectrum’ methods borrowed from chemometrics and machine learning (Corti et al., 2018; Verrelst et al., 2019). *Partial least squares (PLS) regression* was employed to diagnose chlorophyll levels in winter wheat leaf laboratory samples (Zhang et al., 2012). Scanning of single leaves under controlled illumination allowed the authors to evade the challenges inherent to canopy-level imaging in outdoor conditions, and without doubt contributed to extremely accurate ( $R$  statistics up to 0.99) predictions. Unfortunately, unclear study design description undermines the trustworthiness of the findings. Kanning et al. (2018) tested a pushbroom system as a way to overcome some limitations of 2D frame hyperspectral cameras. An experimental winter wheat field was scanned using a UAV, and the measurements subjected to *PLS* modelling. When the model was applied to the pixels of the field orthoimage, the individual nitrogen fertilization treatment levels could be discerned. The estimation quality was sufficient to fit a model for predicting grain yield from the obtained values. Meij et al. (2017) employed *PLS* to predict chlorophyll content in oats from unmanned aerial vehicle (UAV) campaign data. The study also included 25 published vegetation indexes. The *PLS* approach yielded validation predictions inferior to the estimates obtained by using the best of the indexes. Still, according to Verrelst et al. (2019), chemometric methods are in principle more powerful than vegetation indexes for estimating canopy biophysical parameters. The chemometric approach tends, in turn, to be surpassed by machine learning methods, capable of modelling non-linear relationships. A comparison of selected algorithms from both groups demonstrated substantial performance variability within the machine learning family. Robust leaf chlorophyll content predictions for multiple crops were obtained with *kernel ridge* and *Gaussian process regression*. On the other hand, *artificial neural networks*, an approach with a comparable level of sophistication, failed to provide

consistently reliable estimates (Caicedo et al., 2014). By applying *support vector machines (SVM)* to maize hyperspectra, Karimi et al. (2008) obtained very good validation estimates for the tasseling stage. The prediction quality was worse, but still satisfactory, for the early growth stage, which the authors attributed to the soil showing through the crop canopy.

Despite its potential, the adoption of imaging spectroscopy remains hindered, in part by the high investment costs involved (Corti et al., 2018). Scene acquisition using a modern 2D camera tends to be slow due to sequential capture of a large number of bands. As a consequence, the speeds of airborne platforms become constrained (Honkavaara et al., 2017) and band registration needs to be performed during the imagery post-processing (Jakob et al., 2017). The voluminous data contained in hyperspectral data cubes require substantial computational capacities and specialized knowledge to process (Yang et al., 2017; Aasen et al., 2018). In the realm of field point spectrometry, similar challenges have been overcome by the development and commercialization of specialized proximity sensors, such as chlorophyll meters (Govender et al., 2009; Miao et al., 2009). These sensors exploit information from limited numbers of pre-selected bands, and have a predictive model embedded in the firmware to perform the computations. A similar route could be taken for imaging spectrometers in order to make the technology more accessible (Govender et al., 2009). One can envision an affordable specialized device capable of capturing narrow-band imagery, as hyperspectral cameras do, comprising bands that were pre-selected to optimize for accurate remote chlorophyll content estimation.

Feature selection methods have proven to be useful for the screening of spectral bands for a variety of applications. In addition to reducing the number of required model inputs, they were shown to improve the prediction accuracy (Ding & Peng, 2005; Mehmood et al., 2012). Fewer computations are required to process data subjected to feature selection, and model interpretation is facilitated (Ding & Peng, 2005). Band pre-selection prior to data acquisition can also address the problem of slow operation of hyperspectral cameras (Yang et al., 2013; Zhang & He, 2013). As demonstrated by the Zhang & He (2013) oilseed rape yield study, substantial reduction of data volume can be attained without impairing model performance. Discarding of 98% of hyperspectral bands had a minimal effect on the quality of nitrogen content prediction in pepper plants, while significantly simplifying the obtained model (Yu et al., 2014). Behmann et al. (2014) proposed an *SVM* model for detecting water stress in barley. The model inputs comprised vegetation indexes, the combinations of which were determined using wrapper feature selection. Increased detection sensitivity was obtained, allowing for earlier drought detection relative to the raw indexes. The aim of the present study is to investigate the effect of two feature selection approaches on the prediction of leaf chlorophyll concentration in oats from hyperspectral imaging data.

## **MATERIALS AND METHODS**

### **Experimental data**

The present study partially replicates and extends the results of Meij et al. (2017), using the same experimental data. Their experiment evaluated the soil-mediated carry-over effects of preceding and cover crops on crop-of-interest status. The data collection took place in summer 2015, which was the second year of the study, and was focused on experimental plots with oats in the grain-filling developmental stage.

The dataset includes narrow-band reflectance spectra of the experimental plots (one averaged spectrum per plot) obtained from UAV imagery. The spectra cover the range of wavelengths from 450 to 915 nm, i.e., between visible blue and near infra-red. The spectral resolution is 5 nm, thus yielding 94 bands. The spectra are accompanied by ground-truth measurements describing the crop's physiological status. They include, among others, SPAD-estimates of leaf chlorophyll concentrations (one averaged estimate per plot), which are the focus of the present study. There are 56 data points in total, labelled as either calibration or validation data in 1:1 proportion. The dataset is in the public domain, and for the purpose of this study, it was downloaded from the Dryad repository (Meij et al., 2018).

### **Reproduction of the original analysis**

In order to obtain a baseline for the assessment of feature selection performance, a reproduction of the Meij et al. (2017) result was prepared. The original study employed vegetation indexes and *PLS* modelling for predicting leaf chlorophyll from the imaging spectra. This paper focuses on the latter approach.

The data partitioning from the original dataset was preserved, and a *PLS regression* model was fitted to the calibration subset. Leaf chlorophyll concentration was modelled as the dependent variable, and the reflectance values for the whole range of the wavelengths as the independent variables. The number of latent variables was tuned using *leave-one-out cross-validation* by calculating the *cross-validation root mean square error (RMSE)* for each value from between 1 and 20. The validation spectra were then fed to the model exhibiting the lowest error, and the generated predictions compared with the SPAD chlorophyll estimates to obtain *validation RMSE*, *normalized RMSE (NRMSE)*, and the  $R^2$  statistics. To reproduce the original validation results, *RMSE* had to be normalized by dividing it by the mean chlorophyll concentration, rather than the standard deviation or range. Likewise,  $R^2$  had to be calculated as the square of the *correlation* coefficient between the predicted and observed values, rather than derived from the sums of squares.

### **Application of feature selection**

Next, the fitting of the *PLS* model to the calibration dataset was repeated, but in addition to the tuning of the latent variable number, feature selection was performed. Two approaches to feature selection were tested: a filter method based on the *minimum redundancy, maximum relevance (MRMR)* criterion, and a forward selection wrapper method.

Under the filtering approach, variables are evaluated independently of model fitting, according to a measure the value of which determines which of them will be discarded (Mehmood et al., 2012). In the *MRMR* method, this measure is the mutual information shared by the candidate feature and the predicted variable, reduced by the average mutual information shared by the candidate feature and the features already accepted for inclusion into the model. The mutual information is a function of the *correlation* coefficient (De Jay et al., 2013).

With wrapping, models are fitted to multiple pre-selected feature subsets, and the fit quality itself serves as the selection performance criterion, making it a computationally more demanding approach (Mehmood et al., 2012). The wrapper forward selection method is analogous to the forward selection in the *stepwise regression*: candidate

features are picked one by one from the feature pool, and their influence on the performance of the refitted model is assessed. The variable associated with the highest performance increase is kept in the model, and the process continues iteratively, until there is no further improvement.

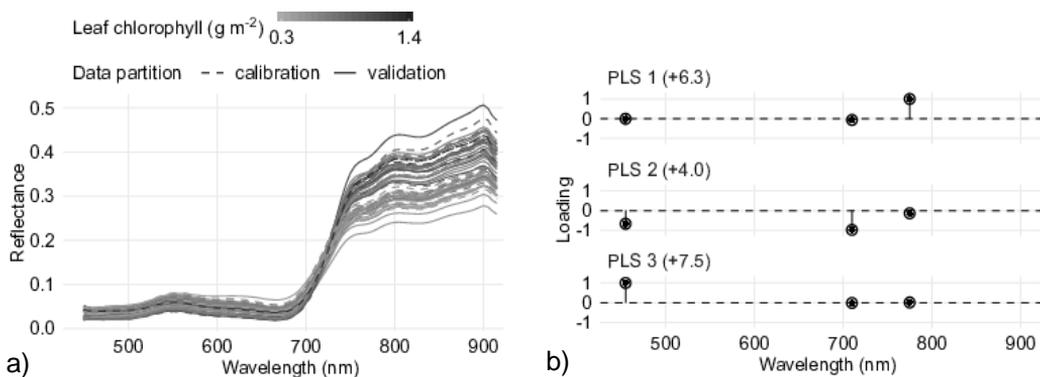
For each method, the present study aimed to obtain a series of models with the input feature number ranging from 2 to all 94 bands (i.e., no selection). In this way, the influence of feature selection intensity on the prediction quality could be investigated.

### Computational reproducibility

The analysis was prepared with reproducibility in mind (Piccolo & Frampton, 2016). It was programmed in the R language (R Core Team, 2019), using the packages *pls* (Mevik et al., 2019) for model fitting, *mRMRe* (De Jay et al., 2013) for assessing the *MRMR* criterion, and *mlr* (Bischl et al., 2016) for model tuning. GNU Make (Stallman et al., 2016) was used as the build tool, and GNU Guix enabled isolation and reproducibility of the software environment for performing the analysis (Courtès & Wurmus, 2015). The computational scripts are available from a Zenodo repository (Želazny, 2020). On an IA-64 machine, the analysis took approximately 100 minutes without parallelization and excluding the time needed to set up the environment. The latter can last hours on the first run, depending on the state of the Guix store (Courtès & Wurmus, 2015) and availability of pre-compiled package substitutes. It is reduced to minutes on subsequent runs.

## RESULTS AND DISCUSSION

### Visual data assessment



**Figure 1.** a) Narrow-band spectra of experimental oat plots in the calibration and validation data subsets acquired using an unmanned aerial vehicle. Line hues reflect the differences in SPAD-estimated leaf chlorophyll concentrations. The figure can be rendered in color by running the computational scripts that accompany the article; b) Loadings in the *partial least squares* model for predicting leaf chlorophyll concentrations from the narrow-band spectra. The model is based on three bands obtained from *minimum redundancy, maximum relevance* filtering. Latent variable loadings are given in the parentheses, wavelength loadings are given on the y axis.

Fig. 1, a depicts the experimental plot spectra matched to the ground-truth data, analogously to Fig. 4 in Meij et al. (2017). High leaf chlorophyll concentration appears

to be associated with increased near infra-red reflectance and a steep red edge—both regions repeatedly considered important for chlorophyll prediction by earlier studies (Govender et al., 2009; Main et al., 2011). On the other hand, contrary to expectation, no apparent red-edge shift can be discerned. The calibration and validation spectra are well mixed in terms of the chlorophyll measurements, as can be expected from the stratified random partitioning, employed by the original study. Regarding the reflectance, the validation subset seems to cover a wider range of values than the calibration subset, but the difference is too small to raise concerns about a mismatch between the partitions.

### Reproduction of Meij et al. (2017)

Despite the variety of existing *PLS* flavours and implementations, the attempt to reproduce the validation results of the Meij et al. (2017) paper turned out to be successful, with only *NRMSE* showing a slight deviation (Table 1, row ‘Reproduction’). However, as discussed above, the high number of bands contributing to the model make the ‘full-spectrum’ approach infeasible for practical application – at least until hyperspectral imagers become affordable (Aasen et al., 2018). In addition, the result of model tuning, which set the number of the latent variables to five, makes an insight into its workings challenging.

**Table 1.** Tuning parameters and validation statistics of the *partial least squares* models. Each model was calibrated using 28 spectra and validated using another set of 28 spectra

Study	Input bands	Latent variables	RMSE (g m <sup>-2</sup> )	NRMSE (%)	R <sup>2</sup>
Meij et al. (2017)	94	5	0.17	23.82	0.57
Reproduction	94	5	0.17	23.75	0.57
Filter feature selection	19	7	0.21	28.36	0.52
Filter feature selection (truncated)	3	3	0.16	21.84	0.62
Wrapper feature selection	1	1	0.20	28.23	0.43

*RMSE* = root mean square error, *NRMSE* = normalized root mean square error.

### Feature selection

The *cross-validation* results of models employing filter feature selection exhibit two local error minima (Fig. 2). The absolute minimum corresponds to 19 input bands, a much lower number than for the reference model, but still too high for developing reasonably priced specialized device. What is more, the model shows higher validation error and involves even more latent variables (seven) than the reproduction model (Table 1, ‘Filter feature selection’).

Conversely, three wavelengths, as in the second minimum, seem a good middle-ground between technical feasibility and expected estimation error. The fact that the number of latent variables in *PLS regression* cannot exceed the number of inputs contributes to the model interpretability. Notable is the improvement of the validation statistics (Table 1, ‘Filter feature selection (truncated)’), which corroborates the positive influence of feature selection on prediction accuracy (Mehmood et al., 2012). Although the obtained gains may seem modest, one should consider other advantages offered by feature selection, such as the reduced cost of a specialized imager (Govender et al., 2009), more efficient data acquisition (Yang et al., 2013; Zhang & He, 2013), and

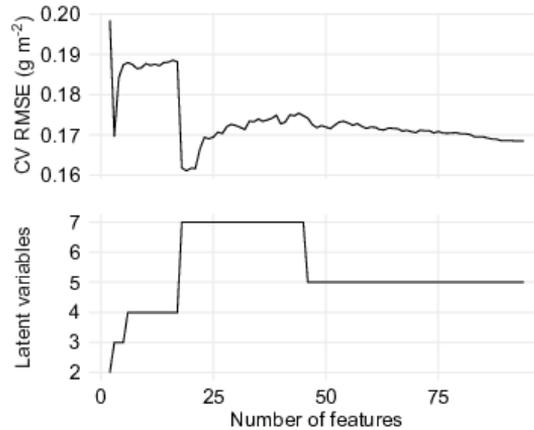
smaller volumes of the collected data (Zhang & He, 2013). On a closer examination, the model appears to give accurate predictions for low levels of chlorophyll, but its performance deteriorates above the level of about  $0.75 \text{ g m}^{-2}$  (Fig. 3). A similar pattern occurred in the Kanning et al. (2018) pushbroom imager study. An attempt to further improve the prediction quality could be made by log-transforming the chlorophyll content values prior to modelling.

Fig. 1, b depicts the band loadings for each latent variable and the latent variable loadings of this model. The chlorophyll content is, thus, predicted as  $LCC = 6.3 PLS_1 + 4.0 PLS_2 + 7.5 PLS_3$ . The value of the first component  $PLS_1 = 0.0 r_{455} - 0.1 r_{710} + 1.0 r_{775}$  corresponds to the near infra-red reflectance, in accordance with the visual assessment, above. The second component  $PLS_2 = -0.7 r_{455} - 1.0 r_{710} - 0.1 r_{775}$  includes the bottom part of the red edge and, interestingly, a blue band., it can be interpreted as canopy darkness (low visible albedo), and linked to the absorbance in the photosynthetically-active spectral region. The third component value  $PLS_3 = 1.0 r_{455} + 0.0 r_{710} + 0.0 r_{775}$  is determined by canopy blueness (blue hue intensity).

Wavelength combinations similar to the one picked by the filtering algorithm seldom occur in vegetation index formulations. They can be found in the Enhanced Vegetation Index (Gao et al., 2000), the Structure Insensitive Pigment Index (Peñuelas et al., 1995), the Modified Simple Ratio, and the Modified Normalized Difference ( $mND_{705}$ ) (Sims & Gamon, 2002). No such index was investigated by Meij et al. (2017). In the study by Main et al. (2011), the first three indexes fared poorly when used for predicting chlorophyll content in maize leaves at various developmental stages. The authors attribute this to the weak relationship between the blue spectral region and the leaf chlorophyll concentration.

Regarding  $mND_{705}$ , it was among the best-performing indexes in Main et al. (2011), and in Miao et al. (2009) – also a maize study. On the other hand, it occurred to be a poor predictor of chlorophyll content in rice (Moharana & Dutta, 2016). The  $mND_{705}$  index formula includes blue reflectance as a way to account for specular reflectance (Sims & Gamon, 2002). The third latent variable of the discussed *PLS* model may play the same role.

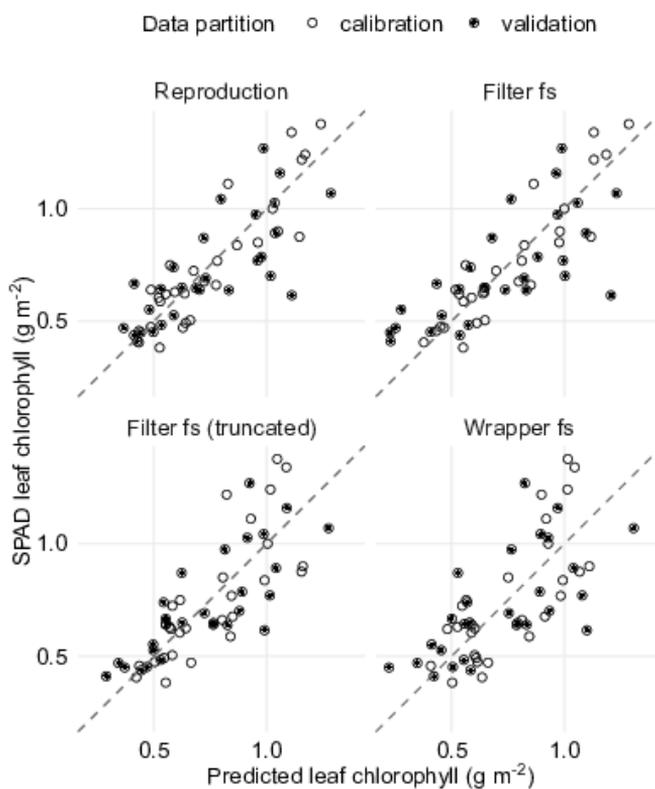
Alternatively, it may adjust for Rayleigh scattering. According to Beisl et al. (2008), atmospheric effects occur even in low-altitude airborne remote sensing applications. Although the analysed dataset has been subjected to atmospheric correction, it was based on a single reference panel measurement (Meij et al., 2017). The weakness of this approach is the assumption of constant illumination conditions as



**Figure 2.** Cross-validation prediction performance and tuning results of the oat leaf chlorophyll prediction models according to the number of features selected using the *minimum redundancy, maximum relevance* filter. *CV RMSE* = cross-validation root mean square error.

individual images are acquired. The blue band information may account for the residual error that still remained after the correction.

The forward selection within the wrapper approach stopped after picking one band (775 nm), thus reducing the *PLS* model to a classical *regression* model with a single independent variable. The selected wavelength lies in the near infra-red spectral region, which agrees with the observation from the visual assessment, above. According to the validation statistics (Table 1 ‘Wrapper feature selection’), despite its extreme simplicity, the model performs surprisingly well in terms of *RMSE*. However, the low  $R^2$  value puts in question the feasibility of its practical use. Moreover, like the preceding model, it exhibits uneven prediction quality for various levels of chlorophyll (Fig. 3).



**Figure 3.** Prediction error patterns of the studied models with respect to the ground-truth data.

In the light of this finding, it can be recommended to avoid wrapper selection for chlorophyll content prediction, especially considering the substantial computational demands of this approach (Ding & Peng, 2005; Mehmood et al., 2012). Conversely, the encouraging results attained with *MRMR* suggest high potential of the filter strategy towards picking highly predictive spectral bands. The *MRMR* criterion seems particularly well-suited to data acquired using optical remote sensing methods. As reflectance measurements exhibit substantial spectral autocorrelation (Karimi et al., 2008; Verrelst et al., 2019), a naive algorithm could pick a set of neighbouring bands, with information content barely exceeding that of a single band. The ‘minimum redundancy’ aspect of *MRMR* avoids this issue by taking correlations between features

into consideration (Ding & Peng, 2005). Still future research might consider examination of feature selection methods from the filter family. The performance of the three classes of methods reviewed by Mehmood et al. (2012): based on loading weights, *regression* coefficients, and variable importance in projection; could be compared, for instance.

### **Possibilities of assessing and improving study generalizability**

The present study illustrates the application of feature selection for obtaining a parsimonious predictive model with high interpretability. Just as omitting model *cross-validation* can lead to over-fitting, a model that performs well on a single validation dataset does not necessarily generalize to new circumstances. This is especially true for unstable models, whose parameters change radically in response to even slight modification of the training data.

In the present study, an improvement of validation statistics was obtained after filtering the spectral bands using the *MRMR* algorithm. As highlighted by De Jay et al. (2013), the algorithm in its original form produces results that are unstable with respect to data modifications. The cited authors proposed an ensemble extension of the filter to stabilize its output.

Ensemble modelling has been shown to improve prediction accuracy, as exemplified by *random forests* (Breiman, 2001), and enable interval estimation, as exemplified by *bootstrap* methods (Wood, 2005). Its obvious application in the discussed study would be to abandon the fixed data partitioning, which was inherited from Meij et al. (2017), in favour of multiple analyses, each based on a different assignment of the data points to the calibration and validation subsets. By the subsequent aggregation of the obtained partial results, the stability of the best performing models could be assessed – not only with respect to the selected wavelengths, but also to their loadings and validation statistics.

Two candidate models fitted to filtered bands were elected by hand for further evaluation based on *CV RMSE* and feature selection intensity as an auxiliary criterion. Repeated data partitioning would result in proliferation of models, making the manual approach unfeasible. Replacing it with an algorithm would necessitate taking both optimization criteria into account, which can be accomplished with aid of *model-based multi-objective optimization* (Horn et al., 2015).

These avenues could not have been taken due to high computational complexity involved, especially if wrapper feature selection were also included. In the future, an adaptation of the analysis for an execution in a high-performance computing environment might be attempted. At that point, an extension of the study to include ensemble modelling would become feasible.

An evident weakness of both the present and the original Meij et al. (2017) study is the fact that the ground-truth data were obtained using a SPAD chlorophyll meter, and thus include spectroscopic estimation errors (Uddling et al., 2007). It is possible that similar errors present in the discussed *PLS* results become masked in the consequence, leading to overoptimistic validation statistics. Therefore, it would be desirable to replicate the study using laboratory analyses for the ground truth, instead.

Spectral responses of leaf pigments differ across plant genotypes. Although the chlorophyll signal is readily discernible in a leaf or canopy spectrum (Ollinger, 2011), the reflectance is modified by additional factors. They include leaf and canopy anatomy and morphology (Asner, 1998; Jacquemoud & Ustin, 2001; Ollinger, 2011) and spectral

properties of additional foliar pigments present in the tissues (Jacquemoud & Ustin, 2001; Ollinger, 2011). Research is needed to establish whether feature selection can yield a set of bands that enable calibration of models for chlorophyll content estimation in multiple crops, and how big this set needs to be for the models to be accurate.

## CONCLUSIONS

Filtering of bands according to the *minimum redundancy, maximum relevance* criterion can improve the performance of a *partial least squares* model aimed at oat leaf chlorophyll prediction from airborne hyperspectral imagery. Chlorophyll concentration can be estimated from near infra-red reflectance, canopy darkness, and its blueness. The obtained size of the feature space (three bands in the present study) is sufficiently small for the development of affordable single-purpose imagers. Although a wrapper approach based on forward feature selection can yield an even more parsimonious model, the resulting prediction quality is not satisfactory. The robustness of the findings remains to be investigated using an ensemble of dataset partitionings and ground truth obtained from laboratory analyses based on samples collected from multiple crops.

ACKNOWLEDGEMENTS. The work was conducted with ~~the~~ financial support from the Ministry of Agriculture of the Czech Republic, institutional support MZE-RO0418; and from the European Union Erasmus+ Programme. Lammert Kooistra and Jan Clevers read the paper draft and provided valuable remarks. Suggestions from three anonymous reviewers were of immense help to further improve the manuscript.

## REFERENCES

- Aasen, H., Honkavaara, E., Lucieer, A. & Zarco-Tejada, P. 2018. Quantitative remote sensing at ultra-high resolution with UAV spectroscopy: A review of sensor technology, measurement procedures, and data correction workflows. *Remote Sensing* **10**, 1091.
- Asner, G.P. 1998. Biophysical and biochemical sources of variability in canopy reflectance. *Remote Sensing of Environment* **64**, 234–253.
- Barbosa, B., Ferraz, G., Gonçalves, L., Marin, D., Maciel, D., Ferraz, P. & Rossi, G. 2019. RGB vegetation indices applied to grass monitoring: A qualitative analysis. *Agronomy Research* **17**, 349–357.
- Basso, B., Fiorentino, C., Cammarano, D. & Schulthess, U. 2016. Variable rate nitrogen fertilizer response in wheat using remote sensing. *Precision Agriculture* **17**, 168–182.
- Behmann, J., Steinrücken, J. & Plümer, L. 2014. Detection of early plant stress responses in hyperspectral images. *ISPRS Journal of Photogrammetry and Remote Sensing* **93**, 98–111.
- Beisl, U., Telaar, J. & Schönemark, M. 2008. Atmospheric correction, reflectance calibration and BRDF correction for ADS40 image data. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **37**, 7–12.
- Bischl, B., Lang, M., Kotthoff, L., Schiffner, J., Richter, J., Studerus, E., Casalicchio, G. & Jones, Z.M. 2016. mlr: Machine learning in R. *Journal of Machine Learning Research* **17**, 1–5.
- Breiman, L. 2001. Random forests. *Machine Learning* **45**, 5–32.
- Caicedo, J.P.R., Verrelst, J., Muñoz-Mari, J., Moreno, J. & Camps-Valls, G. 2014. Toward a semiautomatic machine learning retrieval of biophysical parameters. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* **7**, 1249–1259.

- Corti, M., Cavalli, D., Cabassi, G., Gallina, P.M. & Bechini, L. 2018. Does remote and proximal optical sensing successfully estimate maize variables? A review. *European Journal of Agronomy* **99**, 37–50.
- Courtès, L. & Wurmus, R. 2015. Reproducible and user-controlled software environments in HPC with Guix. In: *Euro-Par 2015: Parallel Processing Workshops*. Vienna University of Technology, Vienna, pp. 579–591.
- De Jay, N., Papillon-Cavanagh, S., Olsen, C., El-Hachem, N., Bontempi, G. & Haibe-Kains, B. 2013. MRMRe: An R package for parallelized mRMR ensemble feature selection. *Bioinformatics* **29**, 2365–2368.
- Ding, C. & Peng, H. 2005. Minimum redundancy feature selection from microarray gene expression data. *Journal of Bioinformatics and Computational Biology* **3**, 185–205.
- Domínguez, J., Kumhálová, J. & Novák, P. 2017. Assessment of the relationship between spectral indices from satellite remote sensing and winter oilseed rape yield. *Agronomy Research* **15**, 55–68.
- Gao, X., Huete, A.R., Ni, W. & Miura, T. 2000. Optical–biophysical relationships of vegetation spectra without background contamination. *Remote Sensing of Environment* **74**, 609–620.
- Govender, M., Govender, P.J., Weiersbye, I.M., Witkowski, E.T.F. & Ahmed, F. 2009. Review of commonly used remote sensing and ground-based technologies to measure plant water stress. *Water SA* **35**, 741–752.
- Hatfield, J.L., Prueger, J.H., Sauer, T.J., Dold, C., O’Brien, P. & Wacha, K. 2019. Applications of vegetative indices from remote sensing to agriculture: Past and future. *Inventions* **4**, 71.
- Honkavaara, E., Rosnell, T., Oliveira, R. & Tommaselli, A. 2017. Band registration of tuneable frame format hyperspectral UAV imagers in complex scenes. *ISPRS Journal of Photogrammetry and Remote Sensing* **134**, 96–109.
- Horn, D., Wagner, T., Biermann, D., Weihs, C. & Bischl, B. 2015. Model-based multi-objective optimization: Taxonomy, multi-point proposal, toolbox and benchmark. In: *International Conference on Evolutionary Multi-Criterion Optimization*. Springer, Guimarães, Portugal, pp. 64–78.
- Jacquemoud, S. & Ustin, S.L. 2001. Leaf optical properties: A state of the art. In: *8th International Symposium of Physical Measurements & Signatures in Remote Sensing*. CNES, Aussois, pp. 223–332.
- Jakob, S., Zimmermann, R. & Gloaguen, R. 2017. The need for accurate geometric and radiometric corrections of drone-borne hyperspectral data for mineral exploration: MEPHySTo – a toolbox for pre-processing drone-borne hyperspectral data. *Remote Sensing* **9**, 88.
- Kanning, M., Kühling, I., Trautz, D. & Jarmer, T. 2018. High-resolution UAV-based hyperspectral imagery for LAI and chlorophyll estimations from wheat for yield prediction. *Remote Sensing* **10**, 2000.
- Karimi, Y., Prasher, S., Madani, A. & Kim, S. 2008. Application of support vector machine technology for the estimation of crop biophysical parameters using aerial hyperspectral observations. *Canadian Biosystems Engineering* **50**, 13–20.
- Main, R., Cho, M.A., Mathieu, R., O’Kennedy, M.M., Ramoelo, A. & Koch, S. 2011. An investigation into robust spectral indices for leaf chlorophyll estimation. *ISPRS Journal of Photogrammetry and Remote Sensing* **66**, 751–761.
- Mehmood, T., Liland, K.H., Snipen, L. & Sæbø, S. 2012. A review of variable selection methods in partial least squares regression. *Chemometrics and Intelligent Laboratory Systems* **118**, 62–69.
- Meij, B. van der, Kooistra, L., Suomalainen, J., Barel, J.M. & Deyn, G.B.D. 2017. Remote sensing of plant trait responses to field-based plant–soil feedback using UAV-based optical sensors. *Biogeosciences* **14**, 733–749.
- Meij, B. van der, Kooistra, L., Suomalainen, J., Barel, J.M. & De Deyn, G.B. 2018. Data from: Remote sensing of plant trait responses to field-based plant–soil feedback using UAV-based optical sensors. doi:10.5061/dryad.75k1d

- Mevik, B.-H., Wehrens, R. & Liland, K.H. 2019. *pls: Partial least squares and principal component regression*. <https://CRAN.R-project.org/package=pls>
- Miao, Y., Mulla, D.J., Randall, G.W., Vetsch, J.A. & Vintila, R. 2009. Combining chlorophyll meter readings and high spatial resolution remote sensing images for in-season site-specific nitrogen management of corn. *Precision Agriculture* **10**, 45–62.
- Moharana, S. & Dutta, S. 2016. Spatial variability of chlorophyll and nitrogen content of rice from hyperspectral imagery. *ISPRS Journal of Photogrammetry and Remote Sensing* **122**, 17–29.
- Ollinger, S.V. 2011. Sources of variability in canopy reflectance and the convergent properties of plants. *New Phytologist* **189**, 375–394.
- Peñuelas, J., Baret, F. & Filella, I. 1995. Semi-empirical indices to assess carotenoids/chlorophyll a ratio from leaf spectral reflectance. *Photosynthetica* **31**, 221–230.
- Piccolo, S.R. & Frampton, M.B. 2016. Tools and techniques for computational reproducibility. *GigaScience* **5**, 30.
- R Core Team. 2019. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna. <https://www.R-project.org/>
- Sims, D.A. & Gamon, J.A. 2002. Relationships between leaf pigment content and spectral reflectance across a wide range of species, leaf structures and developmental stages. *Remote Sensing of Environment* **81**, 337–354.
- Stallman, R.M., McGrath, R. & Smith, P.D. 2016. *GNU Make. A program for directing recompilation*. Free Software Foundation, Boston. 220 pp.
- Uddling, J., Gelang-Alfredsson, J., Piikki, K. & Pleijel, H. 2007. Evaluating the relationship between leaf chlorophyll concentration and SPAD-502 chlorophyll meter readings. *Photosynthesis Research* **91**, 37–46.
- Verrelst, J., Malenovsky, Z., Van der Tol, C., Camps-Valls, G., Gastellu-Etchegorry, J.-P., Lewis, P., North, P. & Moreno, J. 2019. Quantifying vegetation biophysical variables from imaging spectroscopy data: A review on retrieval methods. *Surveys in Geophysics* **40**, 589–629.
- Wood, M. 2005. Bootstrapped confidence intervals as an approach to statistical inference. *Organizational Research Methods* **8**, 454–470.
- Yang, G., Liu, J., Zhao, C., Li, Z., Huang, Y., Yu, H., Xu, B., Yang, X., Zhu, D., Zhang, X., Zhang, R., Feng, H., Zhao, X., Li, Z., Li, H. & Yang, H. 2017. Unmanned aerial vehicle remote sensing for field-based crop phenotyping: Current status and perspectives. *Frontiers in Plant Science* **8**, 1–26.
- Yang, W., Duan, L., Chen, G., Xiong, L. & Liu, Q. 2013. Plant phenomics and high-throughput phenotyping: Accelerating rice functional genomics using multidisciplinary technologies. *Current Opinion in Plant Biology* **16**, 180–187.
- Yu, K.-Q., Zhao, Y.-R., Li, X.-L., Shao, Y.-N., Liu, F. & He, Y. 2014. Hyperspectral imaging for mapping of total nitrogen spatial distribution in pepper plant. *PLoS ONE* **9**, e116205.
- Zhang, D., Wang, X., Ma, W. & Zhao, C. 2012. Research vertical distribution of chlorophyll content of wheat leaves using imaging hyperspectra. *Intelligent Automation & Soft Computing* **18**, 1111–1120.
- Zhang, X. & He, Y. 2013. Rapid estimation of seed yield using hyperspectral images of oilseed rape leaves. *Industrial Crops Products* **42**, 416–420.
- Zou, X., Hernández-Clemente, R., Tammeorg, P., Lizarazo Torres, C., Stoddard, F.L., Mäkelä, P., Pellikka, P. & Möttus, M. 2015. Retrieval of leaf chlorophyll content in field crops using narrow-band indices: Effects of leaf area index and leaf mean tilt angle. *International Journal of Remote Sensing* **36**, 6031–6055.
- Želazny, W.R. 2020. Application of feature selection for predicting leaf chlorophyll content in oats (*Avena sativa* L.) from hyperspectral imagery (computational scripts). doi:10.5281/zenodo.3573320