

Partial least squares regression and projection on latent structure regression (PLS Regression)

Hervé Abdi*

Partial least squares (PLS) regression (*a.k.a.* projection on latent structures) is a recent technique that combines features from and generalizes principal component analysis (PCA) and multiple linear regression. Its goal is to predict a set of dependent variables from a set of independent variables or predictors. This prediction is achieved by extracting from the predictors a set of orthogonal factors called *latent* variables which have the best predictive power. These latent variables can be used to create displays akin to PCA displays. The quality of the prediction obtained from a PLS regression model is evaluated with cross-validation techniques such as the bootstrap and jackknife. There are two main variants of PLS regression: The most common one separates the roles of dependent and independent variables; the second one—used mostly to analyze brain imaging data—gives the same roles to dependent and independent variables. © 2010 John Wiley & Sons, Inc. *WIREs Comp Stat* 2010 2 97–106

PLS is an acronym which originally stood for *partial least squares regression*, but, recently, some authors have preferred to develop this acronym as *projection to latent structures*. In any case, PLS regression combines features from and generalizes principal component analysis (PCA) and multiple linear regression. Its goal is to analyze or predict a set of dependent variables from a set of independent variables or predictors. This prediction is achieved by extracting from the predictors a set of orthogonal factors called *latent* variables which have the best predictive power.

PLS regression is particularly useful when we need to predict a set of dependent variables from a (very) large set of independent variables (i.e., predictors). It originated in the social sciences (specifically economics from the seminal work of Herman Wold, see Ref 1) but became popular first in chemometrics (i.e., computational chemistry) due in part to Herman's son Svante,² and in sensory evaluation.³ But PLS regression is also becoming a tool of choice in the social sciences as a multivariate technique for nonexperimental (e.g., Refs 4–6) and

experimental data alike (e.g., neuroimaging, see Refs 7–11). It was first presented as an algorithm akin to the power method (used for computing eigenvectors) but was rapidly interpreted in a statistical framework. (see Refs 12–17).

Recent developments, including, extensions to multiple table analysis, are explored in Ref 18, and in the volume edited by Esposito Vinzi et al. (Ref 19)

PREREQUISITE NOTIONS AND NOTATIONS

The I observations described by K dependent variables are stored in an $I \times K$ matrix denoted \mathbf{Y} , the values of J predictors collected on these I observations are collected in an $I \times J$ matrix \mathbf{X} .

GOAL OF PLS REGRESSION: PREDICT \mathbf{Y} FROM \mathbf{X}

The goal of PLS regression is to predict \mathbf{Y} from \mathbf{X} and to describe their common structure. When \mathbf{Y} is a vector and \mathbf{X} is a full rank matrix, this goal could be accomplished using ordinary multiple regression. When the number of predictors is large compared to the number of observations, \mathbf{X} is likely to be singular and the regression approach is no longer feasible (i.e.,

*Correspondence to: herve@utdallas.edu

School of Behavioral and Brain Sciences, The University of Texas at Dallas, Richardson, TX 75080–3021, USA

DOI: 10.1002/wics.51

because of multicollinearity). This data configuration has been recently often called the ‘*small N large P problem*.’ It is characteristic of recent data analysis domains such as, e.g., bio-informatics, brain imaging, chemometrics, data mining, and genomics.

Principal component regression

Several approaches have been developed to cope with the multicollinearity problem. For example, one approach is to eliminate some predictors (e.g., using stepwise methods, see Ref 20), another one is to use ridge regression.²¹ One method, closely related to PLS regression is called principal component regression (PCR), it performs a principal component analysis (PCA) of the \mathbf{X} matrix and then uses the principal components of \mathbf{X} as the independent variables of a multiple regression model predicting \mathbf{Y} . Technically, in PCA, \mathbf{X} is decomposed using its singular value decomposition (see Refs 22, 23 for more details) as:

$$\mathbf{X} = \mathbf{R}\mathbf{\Delta}\mathbf{V}^T \quad (1)$$

with:

$$\mathbf{R}^T\mathbf{R} = \mathbf{V}^T\mathbf{V} = \mathbf{I}, \quad (2)$$

(where \mathbf{R} and \mathbf{V} are the matrices of the left and right singular vectors), and $\mathbf{\Delta}$ being a diagonal matrix with the singular values as diagonal elements. The singular vectors are ordered according to their corresponding singular value which is the square root of the variance (i.e., eigenvalue) of \mathbf{X} explained by the singular vectors. The columns of \mathbf{V} are called the *loadings*. The columns of $\mathbf{G} = \mathbf{R}\mathbf{\Delta}$ are called the *factor scores* or *principal components* of \mathbf{X} , or simply scores or components. The matrix \mathbf{R} of the left singular vectors of \mathbf{X} (or the matrix \mathbf{G} of the principal components) are then used to predict \mathbf{Y} using standard multiple linear regression. This approach works well because the orthogonality of the singular vectors eliminates the multicollinearity problem. But, the problem of choosing an *optimum* subset of predictors remains. A possible strategy is to keep only a few of the first components. But these components were originally chosen to explain \mathbf{X} rather than \mathbf{Y} , and so, nothing guarantees that the principal components, which ‘explain’ \mathbf{X} optimally, will be relevant for the prediction of \mathbf{Y} .

Simultaneous decomposition of predictors and dependent variables

So, PCA decomposes \mathbf{X} in order to obtain components which best explains \mathbf{X} . By contrast, PLS regression

finds components *from* \mathbf{X} that best predict \mathbf{Y} . Specifically, PLS regression searches for a set of components (called *latent vectors*) that performs a simultaneous decomposition of \mathbf{X} and \mathbf{Y} with the constraint that these components explain as much as possible of the *covariance* between \mathbf{X} and \mathbf{Y} . This step generalizes PCA. It is followed by a regression step where the latent vectors obtained from \mathbf{X} are used to predict \mathbf{Y} .

PLS regression decomposes both \mathbf{X} and \mathbf{Y} as a product of a common set of orthogonal factors and a set of specific loadings. So, the independent variables are *decomposed* as:

$$\mathbf{X} = \mathbf{T}\mathbf{P}^T \quad \text{with} \quad \mathbf{T}^T\mathbf{T} = \mathbf{I}, \quad (3)$$

with \mathbf{I} being the identity matrix (some variations of the technique do not require \mathbf{T} to have unit norms, these variations differ mostly by the choice of the normalization, they do not differ in their final prediction, but the differences in normalization may make delicate the comparisons between different implementations of the technique). By analogy with PCA, \mathbf{T} is called the *score* matrix, and \mathbf{P} the *loading* matrix (in PLS regression the loadings are not orthogonal). Likewise, \mathbf{Y} is *estimated* as:

$$\hat{\mathbf{Y}} = \mathbf{T}\mathbf{B}\mathbf{C}^T, \quad (4)$$

where \mathbf{B} is a diagonal matrix with the ‘regression weights’ as diagonal elements and \mathbf{C} is the ‘weight matrix’ of the dependent variables (see below for more details on the regression weights and the weight matrix). The columns of \mathbf{T} are the *latent vectors*. When their number is equal to the rank of \mathbf{X} , they perform an exact decomposition of \mathbf{X} . Note, however, that the latent vectors provide only an *estimate* of \mathbf{Y} (i.e., in general $\hat{\mathbf{Y}}$ is not equal to \mathbf{Y}).

PLS REGRESSION AND COVARIANCE

The latent vectors could be chosen in a lot of different ways. In fact, in the previous formulation, any set of orthogonal vectors spanning the column space of \mathbf{X} could be used to play the role of \mathbf{T} . In order to specify \mathbf{T} , additional conditions are required. For PLS regression this amounts to finding two sets of weights denoted \mathbf{w} and \mathbf{c} in order to create (respectively) a linear combination of the columns of \mathbf{X} and \mathbf{Y} such that these two linear combinations have maximum covariance. Specifically, the goal is to obtain a first pair of vectors:

$$\mathbf{t} = \mathbf{X}\mathbf{w} \quad \text{and} \quad \mathbf{u} = \mathbf{Y}\mathbf{c} \quad (5)$$

with the constraints that $\mathbf{w}^T \mathbf{w} = 1$, $\mathbf{t}^T \mathbf{t} = 1$ and $\mathbf{t}^T \mathbf{u}$ is maximal. When the first latent vector is found, it is *subtracted* from both \mathbf{X} and \mathbf{Y} and the procedure is re-iterated until \mathbf{X} becomes a null matrix (see the algorithm section for more details).

NIPALS: A PLS ALGORITHM

The properties of PLS regression can be analyzed from a sketch of the original algorithm (called nipals). The first step is to create two matrices: $\mathbf{E} = \mathbf{X}$ and $\mathbf{F} = \mathbf{Y}$. These matrices are then column centered and normalized (i.e., transformed into Z-scores). The sum of squares of these matrices are denoted SS_X and SS_Y . Before starting the iteration process, the vector \mathbf{u} is initialized with random values. The NIPALS algorithm then performs the following steps (in what follows the symbol \propto means ‘to normalize the result of the operation’):

- Step 1. $\mathbf{w} \propto \mathbf{E}^T \mathbf{u}$ (estimate \mathbf{X} weights).
- Step 2. $\mathbf{t} \propto \mathbf{E} \mathbf{w}$ (estimate \mathbf{X} factor scores).
- Step 3. $\mathbf{c} \propto \mathbf{F}^T \mathbf{t}$ (estimate \mathbf{Y} weights).
- Step 4. $\mathbf{u} = \mathbf{F} \mathbf{c}$ (estimate \mathbf{Y} scores).

If \mathbf{t} has not converged, then go to Step 1, if \mathbf{t} has converged, then compute the value of b which is used to predict \mathbf{Y} from \mathbf{t} as $b = \mathbf{t}^T \mathbf{u}$, and compute the factor loadings for \mathbf{X} as $\mathbf{p} = \mathbf{E}^T \mathbf{t}$. Now subtract (i.e., partial out) the effect of \mathbf{t} from both \mathbf{E} and \mathbf{F} as follows $\mathbf{E} = \mathbf{E} - \mathbf{t} \mathbf{p}^T$ and $\mathbf{F} = \mathbf{F} - b \mathbf{t} \mathbf{c}^T$. This subtraction is called a *deflation* of the matrices \mathbf{E} and \mathbf{F} . The vectors \mathbf{t} , \mathbf{u} , \mathbf{w} , \mathbf{c} , and \mathbf{p} are then stored in the corresponding matrices, and the scalar b is stored as a diagonal element of \mathbf{B} . The sum of squares of \mathbf{X} (respectively \mathbf{Y}) explained by the latent vector is computed as $\mathbf{p}^T \mathbf{p}$ (respectively b^2), and the proportion of variance explained is obtained by dividing the explained sum of squares by the corresponding total sum of squares (i.e., SS_X and SS_Y).

If \mathbf{E} is a null matrix, then the whole set of latent vectors has been found, otherwise the procedure can be re-iterated from Step 1 on.

PLS REGRESSION AND THE SINGULAR VALUE DECOMPOSITION

The nipals algorithm is obviously similar to the power method (for a description, see, e.g., Ref 24) which finds eigenvectors. So PLS regression is likely to be closely related to the eigen- and singular value decompositions (see Refs 22,23 for an introduction to these notions)

and this is indeed the case. For example, if we start from Step 1 of the algorithm, which computes: $\mathbf{w} \propto \mathbf{E}^T \mathbf{u}$, and substitute the rightmost term iteratively, we find the following series of equations:

$$\mathbf{w} \propto \mathbf{E}^T \mathbf{u} \propto \mathbf{E}^T \mathbf{F} \mathbf{c} \propto \mathbf{E}^T \mathbf{F} \mathbf{F}^T \mathbf{t} \propto \mathbf{E}^T \mathbf{F} \mathbf{F}^T \mathbf{E} \mathbf{w}. \quad (6)$$

This shows that the weight vector \mathbf{w} is the first right singular vector of the matrix

$$\mathbf{S} = \mathbf{E}^T \mathbf{F}. \quad (7)$$

Similarly, the first weight vector \mathbf{c} is the left singular vector of \mathbf{S} . The same argument shows that the first vectors \mathbf{t} and \mathbf{u} are the first eigenvectors of $\mathbf{E} \mathbf{E}^T \mathbf{F} \mathbf{F}^T$ and $\mathbf{F} \mathbf{F}^T \mathbf{E} \mathbf{E}^T$. This last observation is important from a computational point of view because it shows that the weight vectors can also be obtained from matrices of size I by I .²⁵ This is useful when the number of variables is much larger than the number of observations (e.g., as in the ‘small N , large P problem’).

PREDICTION OF THE DEPENDENT VARIABLES

The dependent variables are predicted using the multivariate regression formula as:

$$\hat{\mathbf{Y}} = \mathbf{T} \mathbf{B} \mathbf{C}^T = \mathbf{X} \mathbf{B}_{\text{PLS}} \text{ with } \mathbf{B}_{\text{PLS}} = (\mathbf{P}^{T+}) \mathbf{B} \mathbf{C}^T \quad (8)$$

(where \mathbf{P}^{T+} is the Moore–Penrose pseudo-inverse of \mathbf{P}^T , see Ref 26). This last equation assumes that both \mathbf{X} and \mathbf{Y} have been standardized prior to the prediction. In order to predict a nonstandardized matrix \mathbf{Y} from a nonstandardized matrix \mathbf{X} , we use $\mathbf{B}_{\text{PLS}}^*$ which is obtained by reintroducing the original units into \mathbf{B}_{PLS} and adding a first column corresponding to the intercept (when using the original units, \mathbf{X} needs to be augmented with a first columns of 1, as in multiple regression).

If all the latent variables of \mathbf{X} are used, this regression is equivalent to PCR. When only a subset of the latent variables is used, the prediction of \mathbf{Y} is optimal for this number of predictors.

The interpretation of the latent variables is often facilitated by examining graphs akin to PCA graphs (e.g., by plotting observations in a $t_1 \times t_2$ space, see Figure 1).

STATISTICAL INFERENCE: EVALUATING THE QUALITY OF THE PREDICTION

Fixed effect model

The quality of the prediction obtained from PLS regression described so far corresponds to a fixed effect model (i.e., the set of observations is considered as the population of interest, and the conclusions of the analysis are restricted to this set). In this case, the analysis is *descriptive* and the amount of variance (of \mathbf{X} and \mathbf{Y}) explained by a latent vector indicates its importance for the set of data under scrutiny. In this context, latent variables are worth considering if their interpretation is meaningful within the research context.

For a fixed effect model, the overall quality of a PLS regression model using L latent variables is evaluated by first computing the predicted matrix of dependent variables denoted $\hat{\mathbf{Y}}^{[L]}$ and then measuring the similarity between $\hat{\mathbf{Y}}^{[L]}$ and \mathbf{Y} . Several coefficients are available for the task. The squared coefficient of correlation is sometimes used as well as its matrix specific cousin the R_V coefficient.²⁷ The most popular coefficient, however, is the *residual sum of squares*, abbreviated as RESS. It is computed as:

$$\text{RESS} = \|\mathbf{Y} - \hat{\mathbf{Y}}^{[L]}\|^2, \quad (9)$$

(where $\|\cdot\|$ is the norm of \mathbf{Y} , i.e., the square root of the sum of squares of the elements of \mathbf{Y}). The smaller

the value of RESS, the better the prediction, with a value of 0 indicating perfect prediction. For a fixed effect model, the larger L (i.e., the number of latent variables used), the better the prediction.

Random effect model

In most applications, however, the set of observations is a *sample* from some population of interest. In this context, the goal is to *predict* the value of the dependent variables for *new* observations originating from the same population as the sample. This corresponds to a *random* model. In this case, the amount of variance explained by a latent variable indicates its importance in the prediction of \mathbf{Y} . In this context, a latent variable is relevant only if it improves the prediction of \mathbf{Y} for new observations. And this, in turn, opens the problem of which and how many latent variables should be kept in the PLS regression model in order to achieve optimal generalization (i.e., optimal prediction for new observations). In order to estimate the generalization capacity of PLS regression, standard parametric approaches cannot be used, and therefore the performance of a PLS regression model is evaluated with computer-based resampling techniques such as the bootstrap and cross-validation techniques where the data are separated into *learning* set (to build the model) and *testing* set (to test the model). A popular example of this last approach is the jackknife (sometimes called the 'leave-one-out' approach). In the jackknife,^{28,29} each

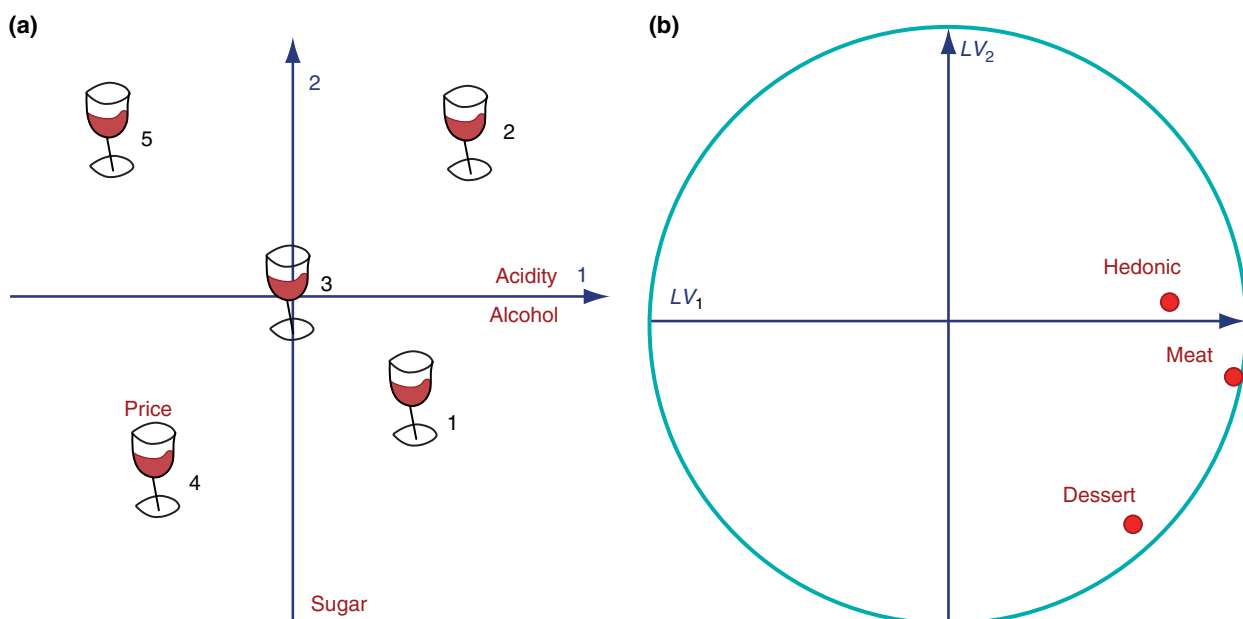


FIGURE 1 | PLS regression regression. (a) Projection of the wines and the predictors on the first two latent vectors (respectively matrices \mathbf{T} and \mathbf{W}). (b) Circle of correlation showing the correlation between the original dependent variables (matrix \mathbf{Y}) and the latent vectors (matrix \mathbf{T}).

observation is, in turn, dropped from the data set, the remaining observations constitute the learning set and are used to build a PLS regression model that is applied to predict the left-out observation which then constitutes the testing set. With this procedure, each observation is predicted according to a random effect model. These predicted observations are then stored in a matrix denoted $\tilde{\mathbf{Y}}$.

For a random effect model, the overall quality of a PLS regression model using L latent variables is evaluated by using L variables to compute—according to the random model—the matrix denoted $\tilde{\mathbf{Y}}^{[L]}$ which stores the predicted values of the observations for the dependent variables. The quality of the prediction is then evaluated as the similarity between $\tilde{\mathbf{Y}}^{[L]}$ and \mathbf{Y} . As for the fixed effect model, this can be done with the squared coefficient of correlation (sometimes called, in this context, the ‘cross-validated r ,’³⁰) as well as the R_V coefficient. By analogy with the RESS coefficient, one can also use the predicted residual sum of squares, abbreviated PRESS. It is computed as:

$$\text{PRESS} = \|\mathbf{Y} - \tilde{\mathbf{Y}}^{[L]}\|^2. \quad (10)$$

The smaller the value of PRESS, the better the prediction for a random effect model, with a value of 0 indicating perfect prediction.

How many latent variables?

By contrast with the fixed effect model, the quality of prediction for a random model does not always increase with the number of latent variables used in the model. Typically, the quality first increases and then decreases. If the quality of the prediction decreases when the number of latent variables increases this indicates that the model is *overfitting* the data (i.e., the information useful to fit the observations from the learning set is not useful to fit *new* observations). Therefore, for a random model, it is critical to determine the optimal number of latent variables to keep for building the model. A straightforward approach is to stop adding latent variables as soon as the PRESS decreases. A more elaborated approach (see, e.g., Ref 16) starts by computing for the ℓ th latent variable the ratio Q_ℓ^2 defined as:

$$Q_\ell^2 = 1 - \frac{\text{PRESS}_\ell}{\text{RESS}_{\ell-1}}, \quad (11)$$

with PRESS_ℓ (resp. $\text{RESS}_{\ell-1}$) being the value of PRESS (resp. RESS) for the ℓ th (resp. $\ell - 1$) latent variable [where $\text{RESS}_0 = K \times (I - 1)$]. A latent variable is kept if its value of Q_ℓ^2 is larger than some arbitrary value generally set equal to $(1 - 95^2) = 0.0975$ (an

alternative set of values sets the threshold to .05 when $I \leq 100$ and to 0 when $I > 100$, see Refs 16,31). Obviously, the choice of the threshold is important from a theoretical point of view, but, from a practical point of view, the values indicated above seem satisfactory.

Bootstrap confidence intervals for the dependent variables

When the number of latent variables of the model has been decided, confidence intervals for the predicted values can be derived using the bootstrap.³² When using the bootstrap, a large number of samples is obtained by drawing, for each sample, observations with replacement from the learning set. Each sample provides a value of \mathbf{B}_{PLS} which is used to estimate the values of the observations in the testing set. The distribution of the values of these observations is then used to estimate the sampling distribution and to derive confidence intervals.

A SMALL EXAMPLE

We want to predict the subjective evaluation of a set of five wines. The dependent variables that we want to predict for each wine are its likeability, and how well it goes with meat or dessert (as rated by a panel of experts, see Table 1). The predictors are the price, sugar, alcohol, and acidity content of each wine (see Table 2).

The different matrices created by PLS regression are given in Tables 3–13. From Table 9, one can find that two latent vectors explain 98% of the variance of \mathbf{X} and 85% of \mathbf{Y} . This suggests that these two dimensions should be kept for the final solution as a fixed effect model. The examination of the two-dimensional regression coefficients (i.e., \mathbf{B}_{PLS} , see Table 10) shows that sugar is mainly responsible for choosing a dessert wine, and that price is negatively correlated with the perceived quality of the wine (at least in this example ...), whereas alcohol is positively correlated with it. Looking at the latent vectors shows that \mathbf{t}_1 expresses price and \mathbf{t}_2 reflects sugar content. This interpretation is confirmed and illustrated in Figure 1a and b which display in (a) the projections on the latent vectors of the wines (matrix \mathbf{T}) and the predictors (matrix \mathbf{W}), and in (b) the correlation between the original dependent variables and the projection of the wines on the latent vectors.

From Table 9, we find that PRESS reaches its minimum value for a model including only the first latent variable and that Q^2 is larger than .0975 only for the first latent variable. So, both PRESS

TABLE 1 | The Matrix Y of the Dependent Variables

Wine	Hedonic	Goes with Meat	Goes with Dessert
1	14	7	8
2	10	7	6
3	8	5	5
4	2	4	7
5	6	2	4

TABLE 2 | The X Matrix of Predictors

Wine	Price	Sugar	Alcohol	Acidity
1	7	7	13	7
2	4	3	14	7
3	10	5	12	5
4	16	7	11	3
5	13	3	10	3

TABLE 3 | The Matrix T

Wine	t_1	t_2	t_3
1	0.4538	-0.4662	0.5716
2	0.5399	0.4940	-0.4631
3	0	0	0
4	-0.4304	-0.5327	-0.5301
5	-0.5633	0.5049	0.4217

TABLE 4 | The Matrix U

Wine	u_1	u_2	u_3
1	1.9451	-0.7611	0.6191
2	0.9347	0.5305	-0.5388
3	-0.2327	0.6084	0.0823
4	-0.9158	-1.1575	-0.6139
5	-1.7313	0.7797	0.4513

TABLE 5 | The Matrix P

	p_1	p_2	p_3
Price	-1.8706	-0.6845	-0.1796
Sugar	0.0468	-1.9977	0.0829
Alcohol	1.9547	0.0283	-0.4224
Acidity	1.9874	0.0556	0.2170

and Q^2 suggest that a model including only the first latent variable is optimal for generalization to new observations. Consequently, we decided to keep one latent variable for the random PLS regression model.

Tables 12 and 13 display the predicted value of \hat{Y} and \tilde{Y} when the prediction uses one latent vector.

SYMMETRIC PLS REGRESSION: BPLS REGRESSION

Interestingly, two different, but closely related, techniques exist under the name of PLS regression. The technique described so far originated from the work of Wold and Martens. In this version of PLS regression, the latent variables are computed from a succession of singular value decompositions followed by deflation of both X and Y . The goal of the analysis is to *predict* Y from X and therefore the roles of X and Y are asymmetric. As a consequence, the latent variables computed to predict Y from X are different from the latent variables computed to predict X from Y .

A related technique, also called PLS regression, originated from the work of Bookstein (Ref 44; see also Ref 33 for early related ideas; Ref 8, and Ref 45 for later applications). To distinguish this version of PLS regression from the previous one, we will call it BPLS regression.

This technique is particularly popular for the analysis of brain imaging data (probably because it requires much less computational time, which is critical taking into account the very large size of brain imaging data sets). Just like standard PLS regression (cf. Eqs. (6) and (7)), BPLS regression starts with the matrix

$$S = X^T Y. \quad (12)$$

The matrix S is then decomposed using its singular value decomposition as:

$$S = W \Theta C^T \text{ with } W^T W = C^T C = I, \quad (13)$$

(where W and C are the matrices of the left and right singular vectors of S and Θ is the diagonal matrix of the singular values, cf. Eq. (1)). In BPLS regression, the latent variables for X and Y are obtained as (cf. Eq. (5)):

$$T = XW \text{ and } U = YC. \quad (14)$$

Because BPLS regression uses a single singular value decomposition to compute the latent variables, they will be identical if the roles of X and Y are reversed: BPLS regression treats X and Y symmetrically. So, while standard PLS regression is akin to multiple regression, BPLS regression is akin to correlation or canonical correlation.³⁴ BPLS regression, however, differs from canonical correlation because

TABLE 6 | The Matrix **W**

	w ₁	w ₂	w ₃
Price	−0.5137	−0.3379	−0.3492
Sugar	0.2010	−0.9400	0.1612
Alcohol	0.5705	−0.0188	−0.8211
Acidity	0.6085	0.0429	0.4218

TABLE 7 | The Matrix **C**

	c ₁	c ₂	c ₃
Hedonic	0.6093	0.0518	0.9672
Goes with meat	0.7024	−0.2684	−0.2181
Goes with dessert	0.3680	−0.9619	−0.1301

BPLS regression extracts the *variance* common to **X** and **Y** whereas canonical correlation seeks linear combinations of **X** and **Y** having the *largest correlation*. In fact, the name of partial least squares *covariance* analysis or canonical *covariance* analysis would probably be more appropriate for BPLS regression.

Varieties of BPLS regression

BPLS regression exists in three main varieties, one of which being specific to brain imaging. The first variety of BPLS regression is used to analyze experimental results, it is called *behavior* BPLS regression if the **Y** matrix consists of measures or *Task* BPLS regression if the **Y** matrix consists of contrasts or describes the experimental conditions with dummy coding.

The second variety is called *mean centered task* BPLS regression and is closely related to barycentric discriminant analysis (e.g., discriminant

TABLE 8 | The **b** Vector

b ₁	b ₂	b ₃
2.7568	1.6272	1.1191

TABLE 10 | The Matrix **B**_{PLS} When Two Latent Vectors Are Used

	Hedonic	Goes with meat	Goes with dessert
Price	−0.2662	−0.2498	0.0121
Sugar	0.0616	0.3197	0.7900
Alcohol	0.2969	0.3679	0.2568
Acidity	0.3011	0.3699	0.2506

TABLE 11 | The Matrix **B**^{*}_{PLS} When Two Latent Vectors Are Used

	Hedonic	Goes with meat	Goes with dessert
Intercept	−3.2809	−3.3770	−1.3909
Price	−0.2559	−0.1129	0.0063
Sugar	0.1418	0.3401	0.6227
Alcohol	0.8080	0.4856	0.2713
Acidity	0.6870	0.3957	0.1919

TABLE 12 | The Matrix **Ŷ** When One Latent Vector Is Used

Wine	Hedonic	Goes with meat	Goes with dessert
1	11.4088	6.8641	6.7278
2	12.0556	7.2178	6.8659
3	8.0000	5.0000	6.0000
4	4.7670	3.2320	5.3097
5	3.7686	2.6860	5.0965

TABLE 13 | The Matrix **Ŷ** When One Latent Vector Is Used

Wine	Hedonic	Goes with meat	Goes with dessert
1	8.5877	5.7044	5.5293
2	12.7531	7.0394	7.6005
3	8.0000	5.0000	6.2500
4	6.8500	3.1670	4.4250
5	3.9871	4.1910	6.5748

TABLE 9 | Variance of **X** and **Y** Explained by the Latent Vectors, RESS, PRESS, and **Q**²

Latent Vector	Cumulative		Cumulative		RESS	PRESS	Q ²
	Percentage of Explained Variance for X	Percentage of Explained Variance for X	Percentage of Explained Variance for Y	Percentage of Explained Variance for Y			
1	70	70	63	63	32.11	95.11	7.93
2	28	98	22	85	25.00	254.86	−280
3	2	100	10	95	1.25	101.56	−202.89

correspondence analysis, see Ref 35). Like discriminant analysis, this approach is suited for data in which the observations originate from groups defined *a priori*, but, unlike discriminant analysis, it can be used for small N , large P problems. The X matrix contains the deviations of the observations to the average vector of all the observations, and the Y matrix uses a dummy code to identify the group to which each observation belongs (i.e., Y has as many columns as there are groups, with a value of 1 at the intersection of the i th row and the k th column indicating that the i th row belongs to the k th group, whereas a value of 0 indicates that it does not). With this coding scheme, the S matrix contains the group barycenters and the BPLS regression analysis of this matrix is equivalent to a PCA of the matrix of the barycenters (which is the first step of barycentric discriminant analysis).

The third variety, which is specific to brain imaging, is called *seed* PLS regression. It is used to study patterns of connectivity between brain regions. Here, the columns of a matrix of brain measurements (where rows are scans and columns are voxels) are partitioned into two sets: a small one called the *seed* and a larger one representing the rest of the brain. In this context, the S matrix contains the correlation between the columns of the seed and the rest of the brain. The analysis of the S matrix reveals the pattern of connectivity between the seed and the rest of the brain.

RELATIONSHIP WITH OTHER TECHNIQUES

PLS regression is obviously related to canonical correlation (see Ref 34), statis, and multiple factor analysis (see Refs 36, 37 for an introduction to these techniques). These relationships are explored in detail in Refs 16, 38–40, and in the volume edited by Esposito Vinzi et al.¹⁹ The main original goal of PLS regression is to preserve the asymmetry of the relationship between predictors and dependent variables, whereas these other techniques treat them symmetrically.

By contrast, BPLS regression is a symmetric technique and therefore is closely related to canonical correlation, but BPLS regression seeks to extract

the variance common to X and Y whereas canonical correlation seeks linear combinations of X and Y having the largest correlation (some connections between BPLS regression and other multivariate techniques relevant for brain imaging are explored in Refs 41–43). The relationships between BPLS regression, and STATIS or multiple factor analysis have not been analyzed formally, but these techniques are likely to provide similar conclusions.

SOFTWARE

PLS regression necessitates sophisticated computations and therefore its application depends on the availability of software. For chemistry, two main programs are used: the first one called SIMCA-P was developed originally by Wold, the second one called the UNSCRAMBLER was first developed by Martens. For brain imaging, SPM, which is one of the most widely used programs in this field, has recently (2002) integrated a PLS regression module. Outside these domains, several standard commercial statistical packages (e.g., SAS, SPSS, Statistica), include PLS regression. The public domain R language also includes PLS regression. A dedicated public domain called SmartPLS is also available.

In addition, interested readers can download a set of matlab programs from the author's home page (www.utdallas.edu/~herve). Also, a public domain set of MATLAB programs is available from the home page of the N-Way project (www.models.kvl.dk/source/nwaytoolbox/) along with tutorials and examples. Staying with matlab, the statistical toolbox includes a PLS regression routine.

For brain imaging (a domain where the Bookstein approach is, by far, the most popular PLS regression approach), a special toolbox written in matlab (by McIntosh, Chau, Lobaugh, and Chen) is freely available from www.rotman-baycrest.on.ca:8080. And, finally, a commercial MATLAB toolbox has also been developed by Eigenresearch.

REFERENCES

1. Wold, H. Estimation of principal components and related models by iterative least squares. In: Krishnaiah PR, ed. *Multivariate Analysis*. New York: Academic Press; 1966, 391–420.
2. Wold S. Personal memories of the early PLS development. *Chemometrics Intell Lab Syst* 2001, 58:83–84.
3. Martens H, Naes T. *Multivariate Calibration*. London: John Wiley & Sons, 1989.

4. Fornell C, Lorange P, Roos J. The cooperative venture formation process: a latent variable structural modeling approach. *Manage Sci* 1990, 36:1246–1255.
5. Hulland J. Use of partial least square in strategic management research: a review of four recent studies. *Strategic Manage J* 1999, 20:195–204.
6. Graham JL, Evenko LI, Rajan MN. An empirical comparison of soviet and american business negotiations. *J Int Business Studies* 1992, 5:387–418.
7. Worsley KJ. An overview and some new developments in the statistical analysis of PET and fMRI data. *Hum Brain Mapp* 1997, 5:254–258.
8. McIntosh AR, Lobaugh NJ. Partial least squares analysis of neuroimaging data: applications and advances. *Neuroimage* 2004, 23:250–263.
9. Giessing C, Fink GR, Roßler F, Thiel CM. fMRI data predict individual differences of behavioral effects of nicotine: a partial least square analysis. *J Cogn Neurosci* 2007, 19:658–6670.
10. Kovacevic N, McIntosh R. Groupwise independent component decomposition of EEG data and partial least square analysis. *NeuroImage* 2007, 35:1103–1112.
11. Wang JY, Bakhadirov K, Devous MD Sr, Abdi H, McColl R, et al. Diffusion tensor tractography of traumatic diffuse axonal injury. *Arch Neurol* 2008, 65:619–626.
12. Burnham AJ, Viveros R, MacGregor JF. Frameworks for latent variable multivariate regression. *J Chemom* 1996, 10:31–45.
13. Garthwaite P. An interpretation of partial least squares. *J Am Stat Assoc* 1994, 89:122–127.
14. Höskuldsson A. Weighting schemes in multivariate data analysis. *J Chemom* 2001, 15:371–396.
15. Phatak A, de Jong S. The geometry of partial least squares. *J Chemom* 1997, 11:311–338.
16. Tenenhaus M. *La régression PLS*. Paris: Technip; 1998.
17. Ter Braak CJF, de Jong S. The objective function of partial least squares regression. *J Chemom* 1998, 12:41–54.
18. Höskuldsson A. Modelling procedures for directed network of data blocks. *Chemometrics Intell Lab Syst* 2009, 97:3–10.
19. Esposito Vinzi V, Chin WW, Henseler J, Wang H, eds. *Handbook of Partial Least Squares Concepts, Methods and Applications in Marketing and Related Fields*. New York: Springer Verlag; 2009.
20. Draper NR, Smith H. *Applied Regression Analysis*. 3rd ed. New York: John Wiley & Sons; 1998.
21. Hoerl AE, Kennard RW. Ridge regression: Biased estimation for nonorthogonal problem. *Technometrics* 1970, 12:55–67.
22. Abdi H. Singular Value Decomposition (SVD) and Generalized Singular Value Decomposition (GSVD). In: Salkind NJ, ed. *Encyclopedia of Measurement and Statistics*. Thousand Oaks, CA: Sage; 2007a, 907–912.
23. Abdi H. Eigen-decomposition: eigenvalues and eigenvectors. In: Salkind NJ, ed. *Encyclopedia of Measurement and Statistics*. Thousand Oaks: Sage; 2007b, 304–308.
24. Abdi H, Valentin D, Edelman B. *Neural Networks*. Thousand Oaks, CA: Sage; 1999.
25. Rännar S, Lindgren F, Geladi P, Wold D. A PLS kernel algorithm for data sets with many variables and fewer objects. Part 1: theory and algorithms. *J Chemom*, 1994, 8:111–125.
26. Abdi H. Linear algebra for neural networks. In: Smelser NJ, Baltes PB, eds. *International Encyclopedia of the Social and Behavioral Sciences*. Oxford, UK: Elsevier; 2001.
27. Abdi H. RV coefficient and Congruence coefficient. In: Salkind NJ, ed. *Encyclopedia of Measurement and Statistics*. Thousand Oaks, CA: Sage; 2007c, 849–853.
28. Quenouille M. Notes on bias and estimation. *Biometrika* 1956, 43:353–360.
29. Efron B. *The Jackknife, the Bootstrap and Other Resampling Plans*. Philadelphia, PA: SIAM; 1982.
30. Wakeling IN, Morris J. A test for significance for partial least squares regression. *J Chemometrics* 1993, 7:291–304.
31. Wold S. PLS for multivariate linear modelling. In: van de Waterbeem H, ed. *QSAR: Chemometric Methods in Molecular Design, Methods and Principles in Medicinal Chemistry, Vol. 2*, Weinheim, Germany: Verlag Chemie; 1995.
32. Efron B, Tibshirani RJ. *An Introduction to the Bootstrap*. New York: Chapman and Hall; 1993.
33. Tucker LR. Determination of parameters of a functional relation by factor analysis. *Psychometrika* 1958, 23:19–23.
34. Gittins R. *Canonical Analysis: A Review with Applications in Ecology*. New York: Springer; 1985.
35. Abdi H. Discriminant correspondence analysis. In: Salkind NJ, ed. *Encyclopedia of Measurement and Statistics*. Thousand Oaks, CA: Sage; 2007d, 270–275.
36. Abdi H, Valentin D. STATIS. In: Salkind NJ, ed. *Encyclopedia of Measurement and Statistics*. Thousand Oaks, CA: Sage; 2007a, 955–962.
37. Abdi H, Valentin D. Multiple factor analysis. In: Salkind NJ, ed. *Encyclopedia of Measurement and Statistics*. Thousand Oaks, CA: Sage; 2007b, 651–657.
38. Pagès J, Tenenhaus M. Multiple factor analysis combined with PLS regression path modeling. Application to the analysis of relationships between physicochemical variables, sensory profiles and hedonic judgments. *Chemometrics Intell Lab Syst* 2001, 58:261–273.
39. Abdi H. Multivariate analysis. In: Lewis-Beck M, Bryman A, Futing T, eds. *Encyclopedia for Research*

- Methods for the Social Sciences*. Thousand Oaks, CA: Sage; 2003b, 699–702.
40. Rosipal R, Krämer N. Overview and recent advances in partial least squares. In: Craig Saunders C, Grobelnik M, Gunn J, Shawe-Taylor J, eds. *Subspace, Latent Structure and Feature Selection: Statistical and Optimization Perspectives Workshop (SLSFS 2005)*. New York: Springer-Verlag; 2006, 34–51.
41. Kherif F, Poline JB, Flandin G, Benali H, Simon O, et al. Multivariate model specification for fMRI data. *NeuroImage* 2002, 16:1068–1083.
42. Friston K, Büchel C. Functional integration. In: Frackowiak RSJ, Friston KJ, Frith CD, Dolan RJ, Price CJ, et al., eds. *Human Brain Function*. New York: Elsevier; 2004, 999–1019.
43. Lazar NA. *The Statistical Analysis of Functional MRI Data*. New York: Springer; 2008.
44. (a) Bookstein FL, Streissguth AP, Sampson PD, Connor PD, Barr HM. Corpus callosum shape and neuropsychological deficits in adult males with heavy fetal alcohol exposure. *Neuro Image* 2002, 15:233–25145. (b) Bookstein FL. Partial least squares: a dose–response model for measurement in the behavioral and brain sciences. [Revised] *Psychology* [an electronic journal] 1994, 5:(23).

FURTHER READING

- Abdi H. PLS regression. In: Lewis-Beck M, Bryman A, Futing T, eds. *Encyclopedia for Research Methods for the Social Sciences*. Thousand Oaks, CA : Sage ; 2003a, 792–795.
- Abdi H, Williams LJ. Principal component analysis. *WIREs Comp Stat*.
- Escofier B, Pagès J. *Analyses Factorielles Multiples*. Paris : Dunod; 1988.
- Frank IE, Friedman JH. A statistical view of chemometrics regression tools. *Technometrics* 1993, 35:109–148.
- Helland IS. PLS and statistical models. *Scand J Stat* 1990, 17:97–114.
- Höskuldson A. PLS methods. *J Chemom* 1988, 2:211–228.
- Geladi P, Kowalski B. Partial least square regression: a tutorial. *Anal Chem Acta* 1986, 35:1–17.
- Martens H, Martens M. *Multivariate Analysis of Quality : An Introduction*. London : John Wiley & Sons; 2001.
- McIntosh AR, Bookstein FL, Haxby JV, Grady CL. Spatial pattern analysis of functional brain images using partial least squares. *Neuroimage* 1996, 3:143–157.