The Annals of Applied Statistics
2011, Vol. 5, No. 4, 2266–2277
DOI: 10.1214/10-AOAS408

THE DUALITY DIAGRAM IN DATA ANALYSIS: EXAMPLES OF MODERN APPLICATIONS

BY OMAR DE LA CRUZ[1] AND SUSAN HOLMES[2]

Stanford University

Today's data-heavy research environment requires the integration of different sources of information into structured data sets that can not be analyzed as simple matrices. We introduce an old technique, known in the European data analyses circles as the Duality Diagram Approach, put to new uses through the use of a variety of metrics and ways of combining different diagrams together. This issue of the Annals of Applied Statistics contains contemporary examples of how this approach provides solutions to hard problems in data integration. We present here the genesis of the technique and how it can be seen as a precursor of the modern kernel based approaches.

1. Introduction. Multivariate statistical methods have been used for many decades to deal with situations in which two or more variables are measured or recorded for each unit.

A classical example of this situation is Guerry's data set, in which several variables meant to capture "moral qualities" (e.g., literacy, crime rate, suicide rate) were tabulated for each of the departments in which France was divided at the time (1833). This data set suggests that one can be interested in how the variables change as one moves around in France, or one can be interested in how the departments compare to each other based on the measured characteristics. It is of special interest how these two approaches can be combined; this is considered in detail in Dray and Jombart (2011).

Besides having a combination of two essentially multidimensional sources of information, like geographic location plus recorded data, another layer of complexity is added when one more variable like time is added, leading to what essentially are two or more data cubes. A typical example of this is ecological data, where species abundances are measured at different, specified locations, over the course of time. The different approaches used in this setting are reviewed in Thioulouse (2011), using the duality diagram setting as a unifying framework.

Such an approach is not limited to animal or plant species spread over a geographic area; the advent of *metagenomics* has made it possible to study the abundances of bacterial species in locations like ocean or pond water, or even the human

Received August 2010.
[1]Supported in part by NSF Grant DMS-EMSW21-VIGRE-0502385.
[2]Supported in part by NIH Grant R01-GM086884.
Key words and phrases. Duality, gPCA, generalized SVD, kernel methods, RV coefficient.

gut. In this case it becomes important to incorporate information not only about location in space, but also in the phylogenetic landscape, by using established or inferred phylogenetic trees for the bacterial species detected. This problem is addressed, using the duality diagram formalism, in Purdom (2011).

As an outgrowth of methods favored by French statisticians, Cailliez and Pages (1976) proposed a unifying framework capable of including many methods reinvented and used by different groups in different countries as special cases. This framework is based on the analysis of certain linear operators between inner-product spaces which can be naturally associated to a data matrix, in the same way Kernel matrices are used today in machine learning Schölkopf, Smola and Muller (1998). This is explained in detail in works like Escoufier (2006) and Holmes (2006). In this article we present some motivations behind the choices made for this approach in the accompanying papers.

The notion of duality is everywhere in Mathematics, appearing under different guises in the most diverse fields; and it is often remarkably useful. The idea of duality was introduced in the analysis of multivariate data by the French school of data analysts as a way to unify a suite of methods that turned out to be exactly or almost exactly equivalent to methods known by a different name, and the duality-diagram formalism provides a simple way to put all these methods in the same context.

Since this approach is the basis of the special articles presented together here [Dray and Jombart (2011), Purdom (2011), Thioulouse (2011)], this short introduction aims to establish the basic facts and notation. The abstract approach in the duality diagram setup is often intimidating and it possibly turns away some interested readers; we hope we can show here that these notions are actually natural, and that the overhead due in understanding the notation pays off handsomely in the breadth and complexity of applications.

2. The data matrix as an operator between inner-product spaces. Today the distinction between the space of rows of the matrix as a sample from a population and the space of columns as the fixed variables on which the observations were measured has been softened and we often hear the term ‘transposable’ data. The definitions presented here explain this row-column duality.

By dispensing of the traditional probabilistic sample-population interpretation, European data analysts in the 1970s [Benzécri (1973), Cailliez and Pages (1976), Gifi (1990)] can be seen in hindsight as precursors of the current Machine learning schools. It is interesting to remember that all these schools had precursors who spent time at the AT&T laboratories in New Jersey at a time when John Tukey was active there.

Consider an $n \times p$ matrix $\mathbf{X}$ containing data for variables $V_1, \ldots, V_p$ collected from n individuals or units. This matrix defines an operator $L_{\mathbf{X}} : \mathbb{R}^p \to \mathbb{R}^n$ by the rule $\mathbf{v} \mapsto \mathbf{B}\mathbf{v}$. What interpretation can we give to such a map? The vector $\mathbf{v}$ can be considered to contain the coefficients for linearly combining the variables

$V_1, \ldots, V_p$ into a new, synthetic, variable. In that sense, it becomes apparent that actually we should consider this a map from $\mathbb{R}^{p*}$, the dual space of $\mathbb{R}^p$, into $\mathbb{R}^n$. The map $L_{\mathbf{X}}$ provides a way to fill in the n values for the new synthetic variable $V = v_1 V_1 + \cdots + v_p V_p$, which could have been defined even before collecting the data. From now on we will abuse notation and identify the operator $L_{\mathbf{X}}$ and the matrix $\mathbf{X}$ (and will do the same with other similarly defined operators and matrices). We have then the following portion of the diagram:

$$\mathbb{R}^{p*} \xrightarrow{\mathbf{X}} \mathbb{R}^n.$$

2.1. *Adjoint operators as a useful formalism.* Recall that the adjoint of a linear transformation $T : \mathbb{V}_1 \to \mathbb{V}_2$ between inner product spaces is defined as the mapping $T^* : \mathbb{V}_2 \to \mathbb{V}_1$ that satisfies

$$\langle Tu, z\rangle_2 = \langle u, T^* z\rangle_1 \qquad \forall u \in V_1, \forall z \in V_2$$

(for simplicity, we will only consider spaces with scalars in $\mathbb{R}$ in this article; this is enough for most data analyses). This can be seen as just a clever way of extending the notion of matrix transpose to a more general setting, but it is actually a powerful formalism, especially when dealing with multiple inner products on the same spaces (notice that T^* depends not only on T but also on $\langle\cdot,\cdot\rangle_1$ and $\langle\cdot,\cdot\rangle_2$).

It is convenient sometimes to think of T^* as a map from $\mathbb{V}_2^*$ to $\mathbb{V}_1^*$; this matches the corresponding situation when it is generalized to Banach spaces. In our setting this distinction might be considered moot, since all spaces considered are naturally isomorphic to their duals, but we will continue using the star notation; the diagrams, and all the matrix operations obtained from them, work equally well if the stars are dropped from the spaces. Then we have the following:

$$\begin{array}{ccc} \mathbb{R}^{p*} & \xrightarrow{\mathbf{X}} & \mathbb{R}^n \\ \wr\| & & \wr\| \\ \mathbb{R}^p & \xleftarrow[\mathbf{X}^*(=\mathbf{X}^T)]{} & \mathbb{R}^{n*} \end{array}$$

Since we are considering the standard inner products on $\mathbb{R}^p$ and $\mathbb{R}^n$ (and their dual spaces), $\mathbf{X}^*$ corresponds just to the transpose $\mathbf{X}^T$ of $\mathbf{X}$. Thus, $\mathbf{X}^*\mathbf{X} = \mathbf{X}^T\mathbf{X}$, $\mathbf{X}\mathbf{X}^* = \mathbf{X}\mathbf{X}^T$, these two symmetric matrices have the same eigenvalues (except possibly for zeros to account for the difference between p and n), and the two sets of eigenvectors can be used to form the singular value decomposition (SVD) of $\mathbf{X}$ [Golub and Van (1996)].

2.2. *The general duality diagram.* Consider now the situation where the inner products (i.e., the geometries) on $\mathbb{R}^p$ and $\mathbb{R}^n$ are not standard. That is, assume that

there are symmetric, positive definite matrices $\mathbf{Q}_{p\times p}$ and $\mathbf{D}_{n\times n}$ such that the inner products

$$\langle \mathbf{u}, \mathbf{v}\rangle_{\mathbf{Q}} := \mathbf{u}^T\mathbf{Q}\mathbf{v} \qquad \forall \mathbf{u}, \mathbf{v} \in \mathbb{R}^p$$

and

$$\langle \mathbf{w}, \mathbf{z}\rangle_{\mathbf{D}} := \mathbf{w}^T\mathbf{D}\mathbf{z} \qquad \forall \mathbf{w}, \mathbf{z} \in \mathbb{R}^n$$

somehow make more sense for a particular data analysis than the standard inner products. A typical example is when $\mathbf{D}$ is a diagonal matrix of (positive) weights, one for each individual, down-weighting individuals that are known to have been measured with a larger error; another example is when $\mathbf{Q}$ is the diagonal matrix containing the reciprocals of the sample variances for the columns of $\mathbf{X}$, which corresponds to standardizing the variables (assuming they are already centered); a related example is when $\mathbf{Q}$ is the inverse of the sample variance–covariance matrix obtained from $\mathbf{X}$, in which case the new geometry corresponds to the *Mahalanobis distance*. Often, we want to consider the case in which $\mathbf{Q} = \mathbf{L}\mathbf{L}^T$ and we are interested in a set $\{U_1, \ldots, U_p\}$ of transformed variables obtained from $\{V_1, \ldots, V_p\}$ by multiplication by $\mathbf{L}$, leading to a transformed data matrix $\mathbf{Y} = \mathbf{X}\mathbf{L}$.

Different multivariate procedures can be obtained by appropriately choosing $\mathbf{Q}$ and $\mathbf{D}$; see Section 3 for some examples.

Instead of $\mathbf{X}$ and its adjoint, consider now the transformation $\mathbf{X}\mathbf{Q}: \mathbb{R}^p \to \mathbb{R}^n$. That is, a vector $\mathbf{v}$ of coefficients is first transformed into $\mathbf{L}^T\mathbf{v}$, which is in the scale of the transformed data matrix $\mathbf{Y} = \mathbf{X}\mathbf{L}$, and then used to create a linear combination of the variables $U_1, \ldots, U_p$. Then, for all $\mathbf{u} \in \mathbb{R}^p$, $\mathbf{z} \in \mathbb{R}^n$,

$$\langle \mathbf{X}\mathbf{Q}\mathbf{u}, \mathbf{z}\rangle_{\mathbf{D}} = (\mathbf{X}\mathbf{Q}\mathbf{u})^T\mathbf{D}\mathbf{z} = \mathbf{u}^T\mathbf{Q}\mathbf{X}^T\mathbf{D}\mathbf{z} = \langle \mathbf{u}, \mathbf{X}^T\mathbf{D}\mathbf{z}\rangle_{\mathbf{Q}},$$

so $(\mathbf{X}\mathbf{Q})^* = \mathbf{X}^T\mathbf{D}$. Then, $(\mathbf{X}\mathbf{Q})^*\mathbf{X}\mathbf{Q} = \mathbf{X}^T\mathbf{D}\mathbf{X}\mathbf{Q}$ and $\mathbf{X}\mathbf{Q}(\mathbf{X}\mathbf{Q})^* = \mathbf{X}\mathbf{Q}\mathbf{X}^T\mathbf{D}$ are self-adjoint operators on $(\mathbb{R}^p, \langle\cdot,\cdot\rangle_{\mathbf{Q}})$ and $(\mathbb{R}^n, \langle\cdot,\cdot\rangle_{\mathbf{D}})$, respectively, but they are not necessarily symmetric matrices. Nevertheless, they have real eigenvalues (which match, except for zeros to account for the difference between p and n), because they are similar to symmetric matrices by way of positive definite matrices; for example,

$$\mathbf{Q}^{1/2}(\mathbf{X}^T\mathbf{D}\mathbf{X}\mathbf{Q})\mathbf{Q}^{-1/2} = \mathbf{Q}^{1/2}\mathbf{X}^T\mathbf{D}\mathbf{X}\mathbf{Q}^{1/2},$$

where $\mathbf{Q}^a$ is obtained by replacing each eigenvalue λ of $\mathbf{Q}$ with λ^a.

The eigenvectors of $\mathbf{X}^T\mathbf{D}\mathbf{X}\mathbf{Q}$ and $\mathbf{X}\mathbf{Q}\mathbf{X}^T\mathbf{D}$ are also real; however, they need not be orthogonal. Nevertheless, the eigenvectors of $\mathbf{X}^T\mathbf{D}\mathbf{X}\mathbf{Q}$ can be taken to be orthogonal with respect to $\langle\cdot,\cdot\rangle_{\mathbf{Q}}$, and those of $\mathbf{X}\mathbf{Q}\mathbf{X}^T\mathbf{D}$ to be orthogonal with respect to $\langle\cdot,\cdot\rangle_{\mathbf{D}}$. (This can be interpreted as leading to a generalized version of the SVD.)

In diagram form, we have the following:

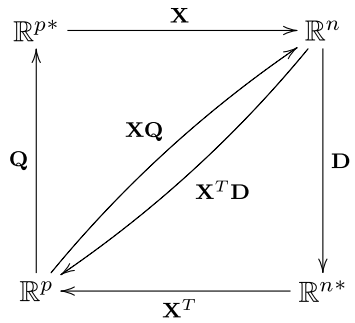

This way, the triplet of matrices $(\mathbf{X}, \mathbf{Q}, \mathbf{D})$ defines a multivariate data analysis setup, in which the main strategy is the computation of the eigendecompositions of the matrices $\mathbf{X}^T\mathbf{DXQ}$ and $\mathbf{XQX}^T\mathbf{D}$. It is also customary to denote $\mathbf{V} = \mathbf{X}^T\mathbf{DX}$ and $\mathbf{W} = \mathbf{XQX}^T$, so that the two operators of interest become $\mathbf{VQ}$ and $\mathbf{WD}$.

The eigendecomposition can be computed for the smaller of the two matrices (which is usually much smaller than the other), and, if needed, the eigenvectors for the other one can be easily obtained: for example, if $\mathbf{X}^T\mathbf{DXQv} = \lambda\mathbf{v}$, then $\mathbf{w} := \mathbf{XQv}$ satisfies

$$\mathbf{XQX}^T\mathbf{Dw} = \mathbf{XQ}(\mathbf{X}^T\mathbf{DXQv} = \mathbf{XQ}(\lambda\mathbf{v})) = \lambda\mathbf{w}.$$

Furthermore, orthogonality is also preserved among eigenvectors: if $\mathbf{v}_1, \mathbf{v}_2$ are $\mathbf{Q}$-orthogonal eigenvectors for $\mathbf{X}^T\mathbf{DXQ}$, then

$$\langle \mathbf{XQv}_1, \mathbf{XQv}_2 \rangle_{\mathbf{D}} = \mathbf{v}_1^T\mathbf{QX}^T\mathbf{DXQv}_2 = \mathbf{v}_1^T\mathbf{Q}(\lambda\mathbf{v}_2) = \lambda\langle \mathbf{v}_1, \mathbf{v}_2 \rangle_{\mathbf{Q}} = 0.$$

Thus, whole eigendecompositions are easily transferred.

2.3. *Connections with kernel methods.* The operator $\mathbf{XQX}^T\mathbf{D}$ can be seen as a precursor of the more general kernel matrices used in today's kernel PCA type methods [Schölkopf, Smola and Muller (1998)].

In the kernel approach to data analysis one assumes that the data is provided as a $n \times n$ matrix $\mathbf{K}$ containing proximity scores for each pair of individuals; these scores might have been computed from measured variables (as in the case of the matrix $\mathbf{XQX}^T\mathbf{D}$, the *linear kernel*; nonlinear functions of the variables offer a great variety of other possibilities), or by directly comparing the individuals; see Schölkopf, Tsuda and Vert (2004) for a review of the theory and examples of applications in computational biology.

Kernel matrices are similarity matrices, that is, the proximity score for two individuals is high when they are similar and low (even negative) when they are dissimilar. (Distance matrices, on the other hand, have higher values for dissimilar pairs of individuals.) What is the meaning of $\mathbf{K}$ as an operator on $\mathbb{R}^n$? A useful

interpretation is to consider it as a smoothing operator acting on real-valued functions f defined on the set of individuals: the value of $\mathbf{K}f$ at the ith individual is a weighted sum of all the values of f, with higher weights for those individuals more similar to the ith; in other words, f is averaged locally (up to a multiplicative constant). This offers one explanation of why the eigendecomposition of $\mathbf{K}$ is useful, since repeated application of $\mathbf{K}$ to $f \in \mathbb{R}^n$ (which should produce a very smooth function) converges toward an eigenvector of the leading eigenvalue, and "very smooth" functions of the data points can be used as coordinates.

Purdom (2011) explores the similarities between the duality diagram and kernel approaches in Appendix B, for the case of kernel Canonical Correspondence Analysis.

3. Examples of well-known methods as particular cases of the diagram. Here we briefly describe how some well-known multivariate methods can be expressed as particular cases of the duality diagram, by appropriately choosing $\mathbf{Q}$ and $\mathbf{D}$. We will assume that $\mathbf{X}$ is centered by columns (i.e., the mean has been subtracted for each variable).

3.1. *Principal components analysis* (*PCA*). PCA seeks to find linear combinations of the variables that explain most of the variability in the data; see Mardia, Kent and Bibby (1979), for example, for more details.

Take $\mathbf{Q} = \mathbf{I}_p$, and $\mathbf{D} = \frac{1}{n}\mathbf{I}_n$. This corresponds to PCA in the original scales; it is equivalent to a straightforward SVD on $\mathbf{X}$ (except for the factor $1/n$).

If one standardizes the variables, as it is often appropriate to eliminate unit scale effects, then $\mathbf{Q}$ is taken to be the diagonal matrix containing the reciprocals of the sample variances of the columns of $\mathbf{X}$ (so $\mathbf{L}$ contains the reciprocals of the standard deviations). While the ith eigenvector $\mathbf{v}_i$ of the ($\mathbf{D}$-weighted) sample covariance matrix $\mathbf{Y}^T\mathbf{D}\mathbf{Y}$ provides the loadings of the variables $U_1, \ldots, U_p$ for the ith principal component (so that the actual components have to be obtained by $\mathbf{p}_i = \mathbf{Y}\mathbf{v}_i/\sqrt{\lambda_i}$), $\mathbf{p}_i$ can be obtained directly as an eigenvector for $\mathbf{X}\mathbf{Q}\mathbf{X}^T\mathbf{D}$: indeed,

$$\begin{aligned}\mathbf{X}\mathbf{Q}\mathbf{X}^T\mathbf{D}\mathbf{p}_i &= \mathbf{X}\mathbf{Q}\mathbf{X}^T\mathbf{D}\mathbf{Y}\mathbf{v}_i/\sqrt{\lambda_i} = \mathbf{X}\mathbf{L}\mathbf{L}^T\mathbf{X}^T\mathbf{D}\mathbf{Y}\mathbf{v}_i/\sqrt{\lambda_i} \\ &= \mathbf{Y}\mathbf{Y}^T\mathbf{D}\mathbf{Y}\mathbf{v}_i/\sqrt{\lambda_i} = \mathbf{Y}\lambda_i\mathbf{v}_i/\sqrt{\lambda_i} = \lambda_i\mathbf{p}_i.\end{aligned}$$

Computing the principal components $\mathbf{p}_i$ does not require the explicit decomposition of $\mathbf{Q}$ as $\mathbf{L}\mathbf{L}^T$.

3.2. *Correspondence analysis* (*CA*). A total of m observations are classified according to two categorical variables, one with n categories or levels, and the other with p, producing a $n \times p$ matrix $\mathbf{N}$ of counts for each combination of levels (a *contingency table*). One wants to study how the counts differ from the expected counts under the assumption of independence between the two variables. To cast CA as a duality diagram, we first define the frequency matrix $\mathbf{F} = \mathbf{N}/m$ and

the marginal frequency vectors $\mathbf{c} = \mathbf{F}^T \mathbf{1}_{n\times 1}$ and $\mathbf{r} = \mathbf{F}\mathbf{1}_{p\times 1}$; then, the expected counts (conditionally on the marginals) are given by $n\mathbf{r}\mathbf{c}^T$. Using the matrices $\mathbf{D}_r = \operatorname{diag}(\mathbf{r})$ and $\mathbf{D}_c = \operatorname{diag}(\mathbf{c})$, we can standardize F by

$$\mathbf{X} := \mathbf{D}_r^{-1}(\mathbf{F} - \mathbf{r}\mathbf{c}^T)\mathbf{D}_c^{-1} = \mathbf{D}_r^{-1}\mathbf{F}\mathbf{D}_c^{-1} - \mathbf{1}_{n\times p}.$$

The matrix $\mathbf{X}$ seems like a reasonable choice to study by eigendecomposition. However, all rows and columns have been reduced to the same importance, while, heuristically, categories with larger marginal counts should provide more accurate information on the distribution of the other variable, and thus should be given greater weight. This can be achieved by defining the triplet $(\mathbf{X}, \mathbf{D}_c, \mathbf{D}_r)$. Notice that actually $\mathbf{X}$ is centered by rows and by columns with respect to the inner products given by $\mathbf{D}_r$ and $\mathbf{D}_c$. This approach matches the traditional definition of CA. Purdom (2011) shows how the information about relationships between the rows of the contingency table can be incorporated into the duality diagram in the special case where there are binary trees that connects the rows of the abundance matrix.

3.3. *Variance*, *inertia*, *co-inertia*. The study of variability of one continuous variable is done through the use of the variance; this notion is generalized in several different directions to accommodate the complexities of dealing with multiple tables, graphs, etc., through the concept of inertia. As in physics, we define inertia as a weighted sum of squared distances of the weighted points. For each of the diagrams studied above, the inertia designates the trace of the operator WD, and we have $\text{Inertia}_{\text{total}} = \operatorname{tr}(WD) = \operatorname{tr}(VQ)$. As pointed out in Purdom (2011), in the case of CA, the inertia is proportional to the χ^2 statistic, whereas in ordinary PCA it is just the total variance of all the variables. In discriminant analysis, the inertia is decomposed into between-groups and within-group components; these are also used in the BCA analysis [Thioulouse (2011), Dray and Jombart (2011)].

The weighted distances between columns have another interpretation in ecology and Purdom (2011) shows how they can be associated to different measures of diversity.

The decomposition of total inertia can be seen as a generalization to MANOVA which is the special case of a variance decomposition. Purdom (2011) uses this effectively to show how to decompose the total diversity across all locations into the average diversity of individual locations and plus the average of pairwise dissimilarities of locations.

Dray and Jombart (2011) use similar decompositions to show what part of the inertia can be assigned to spatially local variation in their BCA approach to multivariate spatial data. They also show how the graphical relationships between rows can be encoded in a special metric $\mathbf{D}$ built from the weighted connectivity matrix. (In their paper, they call these weights W.)

4. One more level of complexity: Comparing diagrams. Interesting results can be obtained by combining two or more triplets. The usual assumption is that two (or more) sets of variables are measured on the same set of n individuals; thus, the matrix $\mathbf{D}$ is assumed to be common, but each set of variables has its own version of $\mathbf{Q}$, of the appropriate size.

For example, one of the triplets might contain data from variables measured on each of the individuals, while the other might encode known relationships between the individuals.

4.1. *The RV coefficient.* A key element in the comparison of the operators arising from two duality diagrams is the RV coefficient. It can be considered as a generalization of the squared correlation coefficient by using the Froebenius matrix product.

Given two symmetric matrices $\mathbf{A}, \mathbf{B}$ of the same size, we define $\mathrm{COVV}(\mathbf{A}, \mathbf{B}) = \mathrm{tr}(\mathbf{AB})$, and

$$\mathrm{RV}(\mathbf{A}, \mathbf{B}) = \frac{\mathrm{tr}(\mathbf{AB})}{\sqrt{\mathrm{tr}(\mathbf{AA})\,\mathrm{tr}(\mathbf{BB})}},$$

whenever $\mathbf{A}, \mathbf{B} \neq \mathbf{0}$. Many nice properties of these definitions arise from the fact that $\mathrm{tr}(\mathbf{AB})$ defines an inner product on the vector space of symmetric matrices of a given size. This can be adapted to the general setting of multiple duality diagrams: having $\mathbf{D}$ fixed, call $S(\mathbf{D})$ the vector space of $\mathbf{D}$-symmetric matrices, that is, matrices satisfying $\mathbf{DA} = \mathbf{A}^T\mathbf{D}$ (equivalently, $\mathbf{A}$ is self-adjoint with respect to $\langle \cdot, \cdot \rangle_{\mathbf{D}}$). Then $\mathrm{tr}(\mathbf{AB})$ defines an inner product on $S(\mathbf{D})$.

When comparing two duality diagrams $(\mathbf{X_1}, \mathbf{Q}_1, \mathbf{D})$, $(\mathbf{X_2}, \mathbf{Q}_2, \mathbf{D})$, then numbers p_1, p_2 of variables might be different, yielding matrices $\mathbf{V}_1\mathbf{Q}_1, \mathbf{V}_2\mathbf{Q}_2$ of different size; however, we will be comparing the matrices (operators) $\mathbf{W}_1\mathbf{D}$ and $\mathbf{W}_2\mathbf{D}$, which are of the same size and $\mathbf{D}$-symmetric. We define the RV coefficient of the two diagrams as $\mathrm{RV}(\mathbf{W}_1\mathbf{D}, \mathbf{W}_2\mathbf{D})$.

Some immediate properties of the RV coefficient are as follows: its values are always in $[0, 1]$; it equals 1 only when $\mathbf{W}_1 = \alpha\mathbf{W}_2 \neq \mathbf{0}$, for some nonzero scalar α; and it equals 0 only when $\mathbf{X}_1^T\mathbf{D}\mathbf{X}_2 = \mathbf{0}$ (provided $\mathbf{Q}_1, \mathbf{Q}_2$ are nonsingular). The proofs are not too hard; more details can be found in Escoufier (2006).

The RV coefficient between diagrams (or triplets) can be used for justifying the use of eigenvalues and eigenvectors in this setting. For example, performing PCA based on $(\mathbf{X}, \mathbf{Q}, \mathbf{D})$ and selecting the q top components is equivalent to finding a matrix $\mathbf{Z}_{n\times q}$ such that the RV coefficient between $(\mathbf{X}, \mathbf{Q}, \mathbf{D})$ and $(\mathbf{Z}, \mathbf{I}_q, \mathbf{D})$ is maximized.

4.1.1. *PCA with respect to instrumental variables.* When one data set $\mathbf{Y}$ has the special status of a response that we would like to predict or explain from the other data set $\mathbf{X}$ of explanatory variables, we can generalize ordinary regression to a multivariate response through the same diagram framework. This is called *PCA*

with respect to instrumental variables, abbreviated PCA-IV (also known as redundancy analysis, RDA), first described by Rao (1964). In terms of the comparison of duality diagrams and RV coefficients, this problem can be rephrased as that of finding the metric $\mathbf{M}$ to associate to $\mathbf{X}$ so that $(\mathbf{X}, \mathbf{M}, \mathbf{D})$ is *as close as possible* to $(\mathbf{Y}, \mathbf{Q}, \mathbf{D})$ in the RV sense. That is, we want to maximize $RV(\mathbf{XMX}^T\mathbf{D}, \mathbf{YQY}^T\mathbf{D})$. We abbreviate the cross-products by writing

$$\mathbf{X}^T\mathbf{DX} = \mathbf{S}_{xx}, \qquad \mathbf{Y}^T\mathbf{DY} = \mathbf{S}_{yy}, \qquad \mathbf{X}^T\mathbf{DY} = \mathbf{S}_{xy}$$

and

$$\mathbf{R} = \mathbf{S}_{xx}^{-1}\mathbf{S}_{xy}\mathbf{QS}_{yx}\mathbf{S}_{xx}^{-1}.$$

Then for any $\mathbf{R}$

$$\|\mathbf{YQY}^T\mathbf{D} - \mathbf{XMX}^T\mathbf{D}\|^2 = \|\mathbf{YQY}^T\mathbf{D} - \mathbf{XRX}^T\mathbf{D}\|^2 + \|\mathbf{XRX}^T\mathbf{D} - \mathbf{XMX}^T\mathbf{D}\|^2.$$

The first term on the right-hand side does not depend on $\mathbf{M}$, and the second term will be zero for the choice $\mathbf{M} = \mathbf{R}$.

If we add the extra constraint that we only allow ourselves a rank q approximation, with $q < \min\{\operatorname{rank}(\mathbf{X}), \operatorname{rank}(\mathbf{Y})\}$, the optimal choice of a positive definite matrix $\mathbf{M}$ is to take $\mathbf{M} = \mathbf{RBB}^T\mathbf{R}$ where the columns of $\mathbf{B}$ are the eigenvectors of $\mathbf{X}^T\mathbf{DXR}$ with

$$\mathbf{B} = \left(\frac{1}{\sqrt{\lambda_1}}\boldsymbol{\beta}_1, \ldots, \frac{1}{\sqrt{\lambda_q}}\boldsymbol{\beta}_q\right)$$
$$\text{such that } \begin{cases} \mathbf{X}^T\mathbf{DXR}\boldsymbol{\beta}_k = \lambda_k\boldsymbol{\beta}_k, \\ \boldsymbol{\beta}_k^T\mathbf{R}\boldsymbol{\beta}_k = \lambda_k, \\ \lambda_1 > \lambda_2 > \cdots > \lambda_q. \end{cases} \qquad k = 1, \ldots, q,$$

The PCA with regards to instrumental variables of rank q is equivalent to the PCA of rank q of the triple $(\mathbf{X}, \mathbf{R}, \mathbf{D})$ where

$$\mathbf{R} = \mathbf{S}_{xx}^{-1}\mathbf{S}_{xy}\mathbf{QS}_{yx}\mathbf{S}_{xx}^{-1}.$$

4.2. *Comparing more than two diagrams.* Consider k diagrams $(\mathbf{X_1}, \mathbf{Q}_1, \mathbf{D})$, $\ldots$, $(\mathbf{X_k}, \mathbf{Q}_k, \mathbf{D})$. This could correspond, for example, to k different studies on the same subjects, using different variables, or the same variables measured on the same units at different points in time (time course study); a review of this problem in the setting of community ecology is found in Thioulouse (2011). It is often important to summarize the relationships between the diagrams in a compact and intelligible way. The RV coefficients in fact allow us to consider this as performing a PCA of the PCAs. We compute the multivariate correlation coefficients between tables and use those as the matrix to be diagonalized, similarly to what happens in ordinary PCA.

The values of the pairwise computations of the COVV and RV coefficients are arranged into $k \times k$ symmetric matrices $\mathbf{C}$ and $\mathbf{R}$, respectively, and the eigendecomposition of these matrices can lead to useful low-dimensional representations, just as in the case of PCA using the covariance or correlation matrices, respectively. In this case, a 2- or 3-dimensional plot can be created in which each point represents one of the studies (diagrams).

Furthermore, since $\mathbf{C}$ and $\mathbf{R}$ have nonnegative entries, the eigenvector $\mathbf{u}_1$ corresponding to the largest eigenvalue can be taken to have only nonnegative entries, adding up to 1. Then, defining $\mathbf{W} = \sum_{i=1}^{k} u_{i1}\mathbf{W}_i$, the operator $\mathbf{WD}$ can be taken as a compromise or summary of all the diagrams, and one can study how far, in the RV sense, different studies are from the compromise.

These steps are part of the so-called STATIS procedure [Escoufier (1980)]. One can think of these data sets as a data cube, with three indices; then a similar procedure can be used to compare two or more such cubes.

5. Conclusions. The duality diagram is a useful formalism that allows one to easily compare many classical multivariate methods, revealing what they have in common and where they differ. But, furthermore, it has become a valuable tool for dealing with two problems that have become very common: (1) combining and amalgamating data which, although collected from different sources and using different methods, shed light on different aspects of the same phenomenon; and (2) taking advantage of complex, nontraditional data types, like tree and network information. These two problems are closely related, as the data to be amalgamated are often of complex type.

The overhead in effort to understand the abstract definitions in the duality diagram approach to data analysis is amply offset by the clearer picture that is gained and by the wealth of applications that become available. In this article we have tried to reduce that overhead by laying out arguments that show that those definitions are actually quite natural. The three articles [Dray and Jombart (2011), Thioulouse (2011), Purdom (2011)] in this group are excellent examples of the power of this approach, but are only a small sample from a large and growing body of work.

Recently, Shinkareva et al. (2008) have used the RV coefficient and STATIS approaches to explore fMRI brain activation in conjunction with stimulations such as images of tools.

A series of papers [Culhane et al. (2002), Culhane, Perrière and Higgins (2003), Fagan, Culhane and Higgins (2007)] have applied BCA and Co-Inertia analyses to the problem of integrating multiple sources of data from heterogeneous gene expression and proteomic studies.

Baty et al. (2006) used the PCAIV method to identify special genes in microarray data and Baty et al. (2008) used bootstrap and permutation type tests for evaluating the stability of the gene identifications produced.

Most of the methods presented in these papers have been coded into functions for the statistical computation environment R [Ihaka and Gentleman (1996)], many

available in the library `ade4`, for which exemplary presentations have been published [see Chessel, Dufour and Thioulouse (2004), Dray, Dufour and Chessel (2007), Dray and Dufour (2007)]. In the case of Thioulouse (2011), you can even run in an interactive way through all the commands generating each and every plot through the reproducible website at http://pbil.univ-lyon1.fr/SAOASOPET/.

REFERENCES

BATY, F., FACOMPRÉ, M., WIEGAND, J., SCHWAGER, J. and BRUTSCHE, M. (2006). Analysis with respect to instrumental variables for the exploration of microarray data structures. *BMC Bioinformatics* **7** 422.

BATY, F., JAEGER, D., PREISWERK, F., SCHUMACHER, M. and BRUTSCHE, M. (2008). Stability of gene contributions and identification of outliers in multivariate analysis of microarray data. *BMC Bioinformatics* **9** 289.

BENZÉCRI, J.-P. (1973). *L'analyse des données: Leçons sur l'analyse factorielle et la reconnaissance des formes, et travaux du Laboratoire de statistique de l'Université de Paris VI*. Dunod, Paris.

CAILLIEZ, F. and PAGES, J. P. (1976). *Introduction à l'analyse des données.* SMASH, Paris.

CHESSEL, D., DUFOUR, A. and THIOULOUSE, J. (2004). The ade4 package, I: One-table methods. *R News* **4** 5–10.

CULHANE, A., PERRIERE, G., CONSIDINE, E., COTTER, T. and HIGGINS, D. (2002). Between-group analysis of microarray data. *Bioinformatics* **18** 1600.

CULHANE, A., PERRIÈRE, G. and HIGGINS, D. (2003). Cross-platform comparison and visualisation of gene expression data using co-inertia analysis. *BMC Bioinformatics* **4** 59.

DRAY, S. and DUFOUR, A. (2007). The ade4 package: Implementing the duality diagram for ecologists. *J. Statist. Softw.* **22** 6.

DRAY, S., DUFOUR, A. and CHESSEL, D. (2007). The ade4 package—II: Two-table and k-table methods. *R News* **7**(2) 47–52.

DRAY, S. and JOMBART, T. (2011). Revisiting Guerry's data: Introducing spatial constraints in multivariate analysis. *Ann. Appl. Statist.* **5** 2278–2299.

ESCOUFIER, Y. (1980). L'analyse conjointe de plusieurs matrices de données. In *Biométrie et Temps* (E. Jolivet, ed.) 59–76. Societe Francaise de Biométrie, Paris.

ESCOUFIER, Y. (2006). Operator related to a data matrix: A survey. In *COMPSTAT 2006—Proceedings in Computational Statistics* 285–297. Physica, Heidelberg. MR2330545

FAGAN, A., CULHANE, A. and HIGGINS, D. (2007). A multivariate analysis approach to the integration of proteomic and gene expression data. *Proteomics* **7** 2162–2171.

GIFI, A. (1990). *Nonlinear Multivariate Analysis*. Wiley, Chichester. MR1076188

GOLUB, G. H. and VAN LOAN, C. F. (1996). *Matrix Computations*, 3rd ed. Johns Hopkins Univ. Press, Baltimore, MD. MR1417720

HOLMES, S. (2006). Multivariate data analysis: The French way. In *Probability and Statistics: Essays in Honor of David A. Freedman* (D. Nolan and T. Speed, eds.) 219–233. IMS, Beachwood, OH. MR2459953

IHAKA, R. and GENTLEMAN, R. (1996). R: A language for data analysis and graphics. *J. Comput. Graph. Statist.* **5** 299–314.

MARDIA, K. V., KENT, J. T. and BIBBY, J. M. (1979). *Multivariate Analysis*. Academic Press, London. MR0560319

PURDOM, E. (2011). Analysis of a data matrix and a graph: Metagenomic data and the phylogenetic tree. *Ann. Appl. Statist.* **5** 2326–2358.

RAO, C. R. (1964). The use and interpretation of principal component analysis in applied research. *Sankhyā A* **26** 329–359. MR0184375

SCHÖLKOPF, B., SMOLA, A. and MULLER, K.-R. (1998). Nonlinear component analysis as a kernel eigenvalue problem. *Neural Comput.* **10** 1299–1319.

SCHÖLKOPF, B., TSUDA, K. and VERT, J.-P. (2004). *Kernel Methods in Computational Biology*. MIT Press, Cambridge, MA.

SHINKAREVA, S., MASON, R., MALAVE, V., WANG, W., MITCHELL, T. and JUST, M. (2008). Using fMRI brain activation to identify cognitive states associated with perception of tools and dwellings. *PLoS One* **3** e1394.

THIOULOUSE, J. (2011). Simultaneous analysis of a sequence of paired ecological tables: A comparison of several methods. *Ann. Appl. Statist.* **5** 2300–2325.

DEPARTMENT OF STATISTICS
STANFORD UNIVERSITY
SEQUOIA HALL
STANFORD, CALIFORNIA 94305
USA
E-MAIL: odlc@stanford.edu
susan@stat.stanford.edu
URL: http://www-stat.stanford.edu/~susan/