



Published in final edited form as:

*J Chem Inf Model.* 2015 July 27; 55(7): 1316–1322. doi:10.1021/acs.jcim.5b00206.

## Beware of $R^2$ : simple, unambiguous assessment of the prediction accuracy of QSAR and QSPR models

D. L. J. Alexander<sup>1</sup>, A. Tropsha<sup>2</sup>, and David A. Winkler<sup>3,4,5,6,\*</sup>

<sup>1</sup>CSIRO Digital Productivity Flagship, Clayton, Australia

<sup>2</sup>UNC Eshelman School of Pharmacy, University of North Carolina at Chapel Hill, NC 27599, USA

<sup>3</sup>CSIRO Manufacturing Flagship, Clayton 3168, Australia

<sup>4</sup>Monash Institute of Pharmaceutical Sciences, Parkville 3052, Australia

<sup>5</sup>Latrobe Institute for Molecular Science, Bundoora 3046, Australia

<sup>6</sup>School of Chemical and Physical Sciences, Flinders University, Bedford Park 5042, Australia

### Abstract

The statistical metrics used to characterize the external predictivity of a model, i.e., how well it predicts the properties of an independent test set, have proliferated over the past decade. This paper clarifies some apparent confusion over the use of the coefficient of determination,  $R^2$ , as a measure of model fit and predictive power in QSAR and QSPR modelling.

$R^2$  (or  $R^2$ ) has been used in various contexts in the literature in conjunction with training and test data, for both ordinary linear regression and regression through the origin as well as with linear and nonlinear regression models. We analyze the widely adopted model fit criteria suggested by Golbraikh and Tropsha<sup>1</sup> in a strict statistical manner. Shortcomings in these criteria are identified and a clearer and simpler alternative method to characterize model predictivity is provided. The intent is not to repeat the well-documented arguments for model validation using test data, but to guide the application of  $R^2$  as a model fit statistic.

Examples are used to illustrate both correct and incorrect use of  $R^2$ . Reporting the root mean squared error or equivalent measures of dispersion, typically of more practical importance than  $R^2$ , is also encouraged and important challenges in addressing the needs of different categories of users such as computational chemists, experimental scientists, and regulatory decision support specialists are outlined.

### 1 Introduction

Although Quantitative Structure-Activity/Property Relationships (QSAR or QSPR) modelling methods have been used for more than 50 years there is still, surprisingly, considerable confusion on the best way to characterize the quality of models. These types of models are typically generated using data driven statistical or machine learning methods,

\*Corresponding author: Prof. Dave Winkler (dave.winkler@csiro.au).

The authors declare no competing financial interest.

and aim to find a quantitative relationship, often quite complex, between the molecular properties (descriptors) of a series of molecules or materials and a target property such as aqueous solubility, toxicity, drug action, cell adhesion etc<sup>2</sup>.

To make our terminology clear, we distinguish between the following (mutually exclusive) data sets:

- *Training set.* Data used to generate models. Ideally, a large training set will be available with a high degree of molecular diversity that spans a large range of the property being modelled.
- *Validation set.* Data used to estimate prediction error in order to compare models. These data are not directly used to fit the models, so give an independent measure of model predictive power. However, since models are compared using the validation set, it affects the choice of model, particularly when one model is selected from a large number of candidate models.
- *Test set.* Data used to estimate prediction error of the final chosen model. These data are used neither to fit the models nor to select between them.

All models require training data. Where sufficient data are available, it is preferable to keep some aside as a test set, in order to generate a demonstrably unbiased estimate of the magnitude of prediction error. Ideally such test sets should truly independent and drawn from a different data source to that providing the training data e.g. by synthesizing molecules predicted by a model to have specific properties and testing them. However, this type of ideal test set is rarely encountered in practice and partitioning a data set into training and test sets (once) is a practical approximation to the ideal case. Additionally, some methods use all the data in the training set (particularly when the amount of data available is insufficient to allow separate validation and test sets). The two most commonly used methods to estimate model predictivity in this case are cross-validation and bootstrapping.

Bootstrapping is a resampling method that involves taking the original training set of  $M$  data points, and sampling from it with replacement to form a new set (a bootstrap sample) that is also of size  $M$ . QSAR models are generated from the bootstrap sample. This process is repeated a large number of times. As the bootstrap sample is taken from the original data set using sampling with replacement, it is not identical with the original data set. The distribution of statistics such as  $R^2$  is then estimated by the distribution of these statistics in the bootstrap samples. Bagging (bootstrap aggregating) involves averaging these bootstrap models.

To apply cross-validation,  $N$  members of the training set of size  $M$  are removed, a model is generated from the  $M - N$  remaining data points, and the properties of the omitted molecules or materials are predicted from the model. This is done  $M/N$  times so that each member of the training set is omitted once (for leave-one-out (LOO) cross-validation, the most commonly employed implementation,  $N = 1$ ). The predictions of the properties of omitted points  $\hat{y}_{CV}$  are used to generate statistics such as  $R^2$ .

Both bootstrapping and cross-validation tend to provide overly optimistic estimates of the predictive power of the model, as the data are typically not a truly random sample of molecules. The model may fit the training set data well, but whether it would accurately or precisely predict the properties of test set molecules external to the model generation process remains unproven.<sup>1, 3</sup>

The use of an independent test set is considered the 'gold standard' for assessing the predictive power of models, and is the most stringent approach. The model may be fitted using bootstrapping or cross-validation on training set data, but its performance is measured by predictions of test set data external to the model generation process. The test set can be chosen randomly from the data set (best when the data set is large), or by some other method such as cluster analysis that chooses test set points that are representative of the data set and property range, so fall within the domain of applicability of the model (best for small data sets where the statistical fluctuations inherent in random selection are much larger).<sup>4-6</sup> When clustering is used, it can also be argued that this test set is not completely independent of the model generation process but, even so, the test set compounds and their observed property values are not used for either model development or model selection.

Regardless of how the model was fitted, a fixture of QSAR and QSPR model validation is a graph of observed versus predicted target property values derived from the training, cross-validation, and/or test sets<sup>7</sup>. A standard measure of model quality is the coefficient of determination<sup>8</sup>,  $R^2$ , but the meaning and appropriateness of this statistic depends on the context. Does the graph refer to training or test data (the model being formulated with reference only to the training data)? Is the best-fit line constrained to pass through the origin? Should the predicted or measured variable be assigned to the vertical axis? Can  $R^2$  ever be negative? Confusion or inconsistency in how  $R^2$  is derived in the various scenarios appears to be fairly common, leading to errors and possible misrepresentation of the predictive strength of a model. This has led to a recent escalation in the number of metrics used to describe the quality of prediction of the test set, adding to the confusion, especially for those relatively new to the QSAR field. For example, Golbraikh and Tropsha define  $R_o^2$  and  $R_o'^2$ , which represent the quality of test set prediction through the origin, the two values depending on whether the predicted values are plotted on the ordinate or abscissa. More recently, Roy et al.<sup>9</sup> attempted to overcome this ambiguity in  $R_o^2$  values by introducing a number of 'r<sub>m</sub>' parameters,  $r_m^2$ ,  $r_m'^2$ ,  $r_{m^2}$ ,  $r_{m^2}$ . This approach has been challenged by Shayanfar and Shayanfar,<sup>10</sup> who further introduced a different  $R_o^2$  value (see also comments by Roy et al<sup>11</sup>).

This paper details in practical terms the meaning and usage of  $R^2$  in the various contexts mentioned above, and suggests how to use it appropriately as a measure of model fit. In spirit, this paper is a sequel to the well-known study by Golbraikh and Tropsha<sup>1</sup> which identified the *inadequacy* of the leave-one-out cross-validation  $R^2$  (denoted as  $q^2$  in this case) calculated on training set data as a reliable characteristic of the model predictivity. The properties of  $q^2$  were examined, and more rigorous criteria for evaluating the predictive power of QSAR models were proposed. Herein, these previously suggested criteria are examined and improved as applied to the external predictivity of QSAR/QSPR models for independent test sets. Note that this paper does not repeat the well-documented arguments in

favour of using a test data set<sup>1, 12</sup>; instead, we focus on how to properly use  $R^2$  as a measure of how well a model can predict the properties of new data.

The following sections define  $R^2$  both in general and in the specific case of plots of observed and predicted values in QSAR or QSPR modelling. Specific comments on regression through the origin are then followed by analysis of the widely used model fit criteria of Golbraikh and Tropsha<sup>1</sup>. Particular discussion of the distinct needs of the  $R^2$  metric for optimisation purposes, as opposed to pure prediction, precedes the concluding section.

## 2 What is $R^2$ ?

$R^2$  (or  $R^2$ ) is usually defined as the square of the correlation coefficient (also called Pearson's  $r$ ) between observed and predicted values in a regression (and for this reason is the same whether observed values are regressed on predictions or the other way round). But it may not be widely known that other definitions of  $R^2$  are used, the most common of which agrees with this formula *only* if the predicted values come from an ordinary regression (or similar method) *on the observed values*. Importantly, this is not the case when applying a model to test data, or even for training data in the case of some nonlinear models<sup>13</sup>. Sections 5 and 6 demonstrate this point by examples. Care is thus required in choosing an appropriate definition of  $R^2$ .

Kvalseth<sup>13</sup> reviews no fewer than eight such definitions, recommending the following simple and informative formula:

$$R^2 = 1 - \frac{\sum (y - \hat{y})^2}{\sum (y - \bar{y})^2} \quad (1)$$

where  $y$  is the observed response variable,  $\bar{y}$  its mean and  $\hat{y}$  the corresponding predicted values. This definition applies to any prediction method, including linear regression, neural nets, etc.  $R^2$  therefore measures the size of the residuals from the model compared to the size of the residuals for a null model where all predictions are the same, i.e. the mean value  $\bar{y}$ .

The numerator in the fraction in (1) is the sum of squared residuals (SSR). The importance of this term is sometimes forgotten, but it is at the heart of the meaning of  $R^2$ : **good models give small residuals**. Squaring the residuals before summing ensures that positive and negative residuals are counted equally, rather than cancelling each other out. (Less common alternatives to  $R^2$  use the median instead of the sum<sup>13, 14</sup>, or absolute values of the residuals instead of squares.<sup>8, 13</sup>) For a good model, SSR is low. High  $R^2$  values are thus preferable; a perfect model has  $R^2 = 1$ .

Maximizing  $R^2$  for a particular dataset is equivalent to minimising SSR. Ordinary linear regression finds the best *linear* model by this criterion (hence its common name 'least squares'), but alternative methods may help when variable selection is necessary (e.g. if there are more variables than observations, so not all of them can be used), as is commonly

the case in QSAR, or if the relationship between the response and predictor variables is nonlinear, also common.

The average squared residual, MSE (mean squared error), obtained by dividing SSR by the number of observations,  $n$ , is a meaningful measure of model fit. Its square root RMSE (root mean squared error) gives a standard deviation of the residuals. (For the training set the SSR should be divided by  $n - p$ , where  $p$  is the number of parameters in the model, giving an unbiased estimate of the variance of the residuals).<sup>8</sup> Calculating  $p$  is straightforward for regression models but less so for more complicated models such as neural nets etc.<sup>3</sup>) The RMSE, or an estimate of the standard deviation of residuals from the model, should usually be reported – for example, whether a method predicts melting points with a standard deviation of 0.1 °C or 100 °C is often more relevant to potential users than various other statistics of model fit. An approximate 95 % confidence interval for predicting future data is  $\hat{y} \pm 2RMSE$ , if the model is correct and errors are normally distributed. The use of measures of dispersion such as RMSE has been strongly advocated by Burden and Winkler.<sup>6, 15</sup>

$R^2$  is a measure of how well the model fits *a particular data set*; the fraction in Equation 1 compares the variation in the residuals to that of the observations themselves. Again referring to the example of predicting melting points of compounds, achieving an RMSE of 1 °C requires a much better model if the observed data span a range of hundreds of degrees than if they cover only a few degrees. A common verbal definition of  $R^2$  is thus the proportion of variation in the response that can be explained by the model. Note though that the value of a model is generally in its overall accuracy and precision, not how successfully it explains the variation in a particular data set. RMSE, or an equivalent measure of dispersion, is often a more helpful indicator of a model's usefulness than is  $R^2$ .

Equation 1 shows that, if it were possible to augment the dataset so that the observed values have greater variation while maintaining the same model accuracy, then  $R^2$  would increase

because  $\sum (y - \bar{y})^2$  is larger. However, such a procedure would improve neither RMSE nor the practical usefulness of the model at any point in the range. Figure 1 illustrates this point: increasing the range of the data, but maintaining an identical distribution of residuals, increases  $R^2$ .

It is instructive if Equation 1 is viewed from another perspective also: the denominator is the sum of squared residuals for a *model that ignores all the predictor variables* (the 'model' that minimizes SSR in this case is simply the average response). The denominator thus acts a scaling factor, relating SSR to the overall variation in the observed data. Ordinary regression applied to a training data set can do no worse than the model  $\hat{y} = \bar{y}$ , so Equation 1 implies that  $R^2 \geq 0$  in this case. However, if the model is applied to data for which the response values  $y$  were unknown when fitting the model (as is the case in predicting observed activity values of the test set data), or if a method other than ordinary regression is used, Equation 1 can sometimes give negative values. This is often highly confusing for novices struggling with a notion that a squared value of a parameter is negative. However, in

this case the interpretation is simply that the model fit is *so poor* that the ratio in the right-hand side of the formula achieves values exceeding 1!

### 3 Assessing a model

In QSAR and QSPR studies the aim is to generate the model that gives the best predictions of properties (the dependent variable) based on other properties of molecules or materials (that is, their descriptors) in the training set. The quality of a model is assessed by a plot of the observed versus predicted dependent variable property values. This can be done for a training set, where it illustrates how well the model predicts the data used to generate it, or for a test set that contains data never used to train the model. The accuracy of prediction of the dependent variable property value for the test set data is a measure of the ability of the model to generalize to new data it has not seen before. The closer the data in such a plot lies to the line  $y=\hat{y}$  the better the model, because the predicted numerical values are very close to those measured by experiment. Including this line on the graph helps the predictive power of the model to be assessed (the graph also provides a check for outliers or trends in the data).

Equation 1 provides the formula for  $R^2$  as a measure of model fit. Note that in a plot of test data,  $y$ ,  $\hat{y}$  and  $\bar{y}$  in Equation 1 should all relate to test data, not training data. Some authors<sup>12</sup> have recommended using  $y$  and  $\hat{y}$  from the test set but  $\bar{y}$  from the training set; however using  $\bar{y}$  from the test set is not only simpler and consistent but also minimizes  $R^2$ , so is more conservative. This question does not arise in calculating RMSE, another advantage of this metric.

A regression on the observed and predicted values themselves (rather than the values for the  $y=x$  model) is sometimes reported for test set data. This, however, provides a measure not of the absolute degree of accuracy of the model, but the degree to which its predictions are *correlated* with observations. This correlation should be reported separately, as discussed in Section 6 below.

When *training set* observations are regressed on their predicted values, the fitted model is simply  $y=\hat{y}$  (by definition,  $\hat{y}$  is already the linear combination of predictors that minimizes SSR) and  $R^2$  is the same as it was for the original model; thus a separate regression is unnecessary. The same will usually hold at least approximately for other (e.g. nonlinear) prediction methods.

But this is *not* the case for test data. The regression of observed vs. predicted values in this case will have a value of  $R^2$  that is *larger* than that of the original model. The original model based on the training set data can estimate each test set observation  $y$  by a predicted value,  $\hat{y}$ ; but the linear regression of observed on predicted values maximizes  $R^2$  for a *secondary* model,  $y=a+b\hat{y}$ . The fitted model will not be  $y=\hat{y}$  in this case, since the test set is not identical with the training set; thus the secondary model will give a larger value of  $R^2$ . A large difference between the two regression  $R^2$  values would indicate a systematic error in the model, since it does not correctly predict the test data with good precision (although it may predict the trend correctly). The larger value of  $R^2$  is not a true test set estimate of

model fit, since it comes from a regression model using test set data, while the definition of test set data is that it is not used to fit a model. Instead, the test set value of  $R^2$  must be calculated directly from Equation 1.

Finally and for completeness, regressing predicted values on the actual observations has also been suggested.<sup>1</sup> Using observations in this way to predict model estimates, rather than the other way round, is counterintuitive. For training data of a regression model, Besalu et al<sup>7</sup>

show the regression line will be  $\hat{y} = \bar{y} + R^2(y - \bar{y})$  (though the reason for this is not regression to the mean as claimed). The slope in such a graph is thus (surprisingly)  $R^2$ , which may be well less than the ideal of 1, and the intercept is  $\bar{y}(1 - R^2)$ , which may be far from the ideal of zero if  $\bar{y}$  is large – even for relatively good models with  $R^2$  (for example) around 0.8. The same features will appear on graphs of test data, provided the regression model is unbiased. Again, this will usually hold at least approximately for prediction methods other than regression (provided the average predicted observation is close to the actual average observation). It is thus unnecessary to perform such a regression, and unreasonable to expect it to give a slope of 1 or zero intercept.

Some further comments on regression through the origin, recommended by Golbraikh and Tropsha<sup>1</sup>, are relevant before suggesting improvements to their criteria for a good model.

#### 4 Regression through the origin

All regression equations discussed to this point have had no constraints placed on them. But it is also possible to constrain regression equations so that the line passes through the origin. Adding this constraint invalidates the observations of Section 2; for example, the value given by Equation 1 would differ depending on whether observed values are regressed on predictions or the other way around.

Regression through the origin can also lead to negative values of Equation 1 even for the training data, since the added constraint may make the model even worse than simply estimating each observation by the sample mean. For this reason, a different formula, applicable only to regression through the origin, is used (without warning!) by various software packages<sup>10, 11</sup> including R and at least some modern versions of Microsoft Excel:

$$R^2 = 1 - \frac{\sum(y - \hat{y})^2}{\sum(y - 0)^2} \quad (2)$$

In Equation 2, SSR is again compared to the residuals from a model that ignores all predictor variables. The only such model that passes through the origin estimates each observation by the value zero. With this definition,  $R^2$  values for linear regression on training data are again non-negative. Equation 2 gives higher values of  $R^2$  than Equation 1 (much higher when the mean observation is high), so values from the different equations are not comparable.  $R^2$  values from (2) are also commonly close to 1 even when the model is



very poor, greatly reducing their utility in distinguishing good models from less good ones.<sup>13</sup>

## 5 Criteria for a good model for predictive purposes

Kubinyi<sup>16</sup> and Golbraikh and Tropsha<sup>1</sup> demonstrated clearly that  $q^2$  values on training data often bear no relation to the predictive ability of models on test data ( $q^2$  is defined as the value of Equation 1, using the training data set, where the predicted values are  $\hat{y}_{CV}$  from LOO cross-validation<sup>12</sup>). The advice to assess model fit and predictive ability using test data is extremely pertinent.

However, the measures of model fit for test data suggested in the latter paper<sup>1</sup> are not strictly statistically correct. The application of  $F$  ratios is invalid,<sup>8</sup> since it assumes both that there is only one predictor and that the mean predicted value equals the mean observation, neither of which are likely to hold; and there are other errors of interpretation. Even if the statistic was valid, it indicates only whether there is *any* relationship between predictor variables and response, not whether this fit is good.

The main criteria for validity of a model as assessed by the test set predictions suggested by Golbraikh and Tropsha<sup>1</sup>, and endorsed by Tropsha, Gramatica and Gombar<sup>12</sup>, relate to regression through the origin of observed and predicted values. These criteria are: -

1.  $q^2$  on training data  $> 0.5$
2. Regressing observed on predicted test data (or vice versa):
  - o  $R^2 > 0.6$
  - o  $R^2$  through the origin close to unconstrained  $R^2$
  - o Slope of regression through the origin close to 1

Criterion 1 is a valid condition for training data where cross validation has been applied, though as Golbraikh and Tropsha<sup>1</sup> demonstrate this gives little indication of the model's predictive power on test data.

Criterion 2, which applies to test data, ensures that no good model capturing the correlation between actual and predicted data irrespective of the absolute error will be rejected; good models have small residuals and thus the graph of observed versus predicted values will have data points clustered around the line  $y=\hat{y}$ . (Yet even a model with low  $R^2$  may still be practically useful; again, the practical usefulness of a model is better determined by its RMSE than by the value of  $R^2$ ).

As discussed in Section 3, it is more meaningful to report simply both the RMSE and the value of  $R^2$  from the  $y=\hat{y}$  model, and the squared test set correlation between observed and predicted values, as discussed in Section 6. Calculating regressions through the origin is unnecessary. Furthermore, the criteria of Golbraikh and Tropsha<sup>1</sup> identify as good some models that actually have poor fit, as the following examples demonstrate. Note also that the definitions of  $R^2$  through the origin in Golbraikh and Tropsha<sup>1</sup> are incorrect; they not only use Equation 1 rather than Equation 2 (which is merely a matter of convention), but some



equations contain typographical errors also (see Supplementary Information for details). Corrected formulas are applied in the following discussion.

Figure 2 shows two hypothetical plots of observed and predicted test data values that pass the criteria above despite representing poor model fits. In the left graph, the regression of observed on predicted data has an  $R^2$  of 0.60 and regressing predicted on observed data through the origin gives an  $R^2$  of 0.58 and slope of 1.1499, satisfying the conditions above. (Condition 1 relates to training data, not shown here.)

But, as described in Section 2, the proper measures of fit are  $RMSE = 2.1$  and  $R^2 = 0.22$ , calculated from Equation 1. The data are far from the dotted line representing  $y = \hat{y}$ ; the model fit is in fact poor.

Figure 2b shows a more extreme example; the ordinary regression and regression through the origin are identical, with  $R^2 = 1$  and a slope of 0.85. The criteria are thus again satisfied, even though the data are far from the dotted line representing  $y = \hat{y}$ . In this case  $RMSE$  is 1.3, which is very high given that the entire observed range is only 2.5. The correct  $R^2$  value is  $-1.28$ , negative in this case since the residuals are larger than the variation of the observations around their mean. This again indicates that the model fits extremely poorly, despite passing the criteria of Golbraikh and Tropsha.<sup>1</sup> (Reducing the range of the predicted values but maintaining the form of this graph, with predicted values 85% of the observed values, would decrease the value of  $R^2$  (correctly calculated) even further.)

These examples amply demonstrate a weakness in the criteria described in Golbraikh and Tropsha.<sup>1</sup> There is no value in calculating  $R^2$  for regression through the origin of observed on predicted values, or vice versa; instead the  $RMSE$  and  $R^2$  calculated from Equation 1 using  $\hat{y}$  from the original model are of much greater practical use.

Our recommended criteria for a useful model for predictive purposes are thus simply:

- High  $R^2$ , eg  $R^2 > 0.6$ , calculated for test set data using Equation 1. This ensures the model fits the data well.
- Low  $RMSE$  of test set predictions. How low  $RMSE$  must be depends on the intended use of the model. Even a model with low  $R^2$  can be practically useful if  $RMSE$  is low. A useful rule of thumb is that the test set  $RMSE$  should be less than 10% of the range of target property value.

These criteria are simpler and clearer than those proposed above and more recently by Tropsha<sup>17</sup>, due to this sharper understanding of how to calculate the statistics for the test set predictions.

## 6 Criteria for a good model for ranking molecules

The preceding sections establish improved criteria for good models *for predictive purposes*. For these purposes, we recommend plotting predicted and observed values for the test set, but calculating  $R^2$  directly via Equation 1 rather than from a line of best fit on this graph.

Such a line gives the squared test set correlation between observed and predicted values, rather than a test set measure of model fit.

However, the best fit correlation between observed and predicted values in the test set is still important in some situations. As most QSAR/QSPR practitioners are aware, the quantity and diversity of the training data are often relatively low due to the cost and difficulty in generating larger data sets, or a limitation on the number of compounds that could show appreciable activity in a biological assay. For this or other reasons, the model may not fully capture the relationship between the molecular properties (descriptors) and the dependent variable property of interest. In this rather more realistic scenario, the numerical values of the target properties of the test set data may not be predicted as accurately, but the model may still allow correct identification of *trends*. It is still of great value to know which molecules or materials are the 'best' (highest activity against a drug target, highest water solubility, lowest toxicity, lowest pathogen adhesion etc.) and which are likely to have poor properties. This allows molecules or materials to be prioritized for synthesis. In this case, the trend-capturing *correlation* between observed and predicted values for the test set may be more relevant than either RMSE or accurate numerical prediction of the property by the model.

Preferably, a regression line for the observations fitted to their predicted values should have a slope close to 1 ( $y=\hat{y}$ ) and it may pass close to the origin, as would be the case for an ideal model. However, this is not essential when the relative ranking of molecules or materials in the test set is more important than numerically accurate values for the property being modelled. The value of  $R^2$  that this secondary regression model generates is the square of the correlation coefficient between the observed and predicted values of the test set; to avoid confusion with the value of  $R^2$  for the original model, we suggest this statistic be described explicitly as the squared test set correlation between observed and predicted values. This is simple and unambiguous. This squared test set correlation need only be reported if it is significantly higher than the value of  $R^2$  calculated from Equation 1, and when relative ranking of molecules or materials suffices, rather than accurate prediction of the properties of each molecule or material.

Other measures may be even more appropriate in this case. Both observed and predicted values may be ranked from 1 to  $n$ , where  $n$  is the number of data points; the correlation between these rankings, known as Spearman's rank correlation coefficient<sup>18</sup>, assesses how well the predictions *order* the observed values, regardless of how accurate the predictions are.

$$\rho=1 - \frac{6\sum d_i^2}{n(n^2 - 1)} \quad (3)$$

where  $d_i = x_i - y_i$  is the difference in ranking  $i$  and  $n$  is the population size. If the order of the observed values exactly matches that of the predicted values after sorting, then Equation 3 will equal 1; but if some observed values are incorrectly ordered by the predictions then some of the summands in the numerator will be negative, so the index will be closer to 0.

The lowest possible value is  $-1$ , attainable only if the observed values are in the reverse order compared to the predicted values. The index is thus interpreted similarly to a correlation coefficient.

These measures are now illustrated by example. Returning to Figure 2a,  $R^2$  is only 0.22; the observed values are far from their predictions. But if relative rankings of the data are more important than accurate prediction, the squared test set correlation between observed and predicted values of 0.6 may also be mentioned; the model did at least order the test data fairly accurately, which may be all that is required. Spearman's rank correlation coefficient (0.71) is also much higher, indicating the model predictions were ordered approximately correctly even though the absolute prediction accuracy was poor.

In Figure 2b, test set observations were so far from their predicted values that  $R^2$  was negative. Yet the squared test set correlation between observed and predicted values is 1: the model ordered the data perfectly, which may suffice in some applications. Spearman's rank correlation coefficient was also 1 in this case.

A further example is provided by the melting point of ionic liquids. The data taken from Varnek et al. relate to the melting points for four chemical classes of ionic liquids.<sup>19</sup> These data were modelled by Varnek et al. as separate chemical classes but we have used molecular descriptors and a Bayesian regularized neural network to generate a model that predicts the melting points for all ionic liquids taken as a single group. The purpose is not to describe this particular model but to show how the metrics for model validation apply to a real-world problem. The QSPR model resulting from the analysis of these data is given in Figure 3. While the regression through the data gives almost identical  $R^2$  values regardless of the assignment of the observed values to the x or y axes, the regression of  $x=y$  (corresponding the numerically accurate predictions not just ranking) is very dependent on the choice of axes, as is regression through the origin.

The model tended to under-estimate low melting points and over-estimate high melting points in the test set, which is reflected in a fairly low value of  $R^2 = 0.17$  as calculated from Equation 1. This may suggest modifications to the model. However, the squared test set correlation between the observed and predicted melting points from the QSPR model is much higher: 0.47. This describes the ability of the model to correctly predict the trend in melting points rather than the correct numerical values. The Spearman rank for the predictions of the test set was 0.66. This suggests that the model may still be effective in finding which ionic liquids have the lowest or the highest melting points, even though exact numerical estimates may not be accurate.

When the purpose of a model is to find the best molecules or materials, not necessarily to predict their properties, the utility of a model is not solely dependent on quantitative accuracy.<sup>20</sup> In such cases the squared test set correlation between observed and predicted values, or the Spearman rank correlation coefficient, may be more relevant.

## 7 Conclusions

The value of Golbraikh and Tropsha's work<sup>1</sup> is in its very effective demonstration that model quality and predictivity should be assessed by test data, not only on the training data. However, the measures of model fit suggested are overly complicated and potentially misleading. Instead, researchers should simply report the  $R^2$  and RMSE or a similar statistic like the standard error of prediction for the test set, which readers are more likely to be able to interpret. The value of  $R^2$  should be calculated for the test data using Equation 1, not from a regression of observed on predicted values. However, if relative rankings of the data suffice, rather than accurate numerical prediction, then it may be relevant to report, in addition, the squared test set correlation between observed and predicted values (or an equivalent metric).

This paper elucidated some common misunderstandings surrounding the use of  $R^2$  as a measure of model fit. Much confusion could be spared, if everyone knew how to use  $R^2$ !

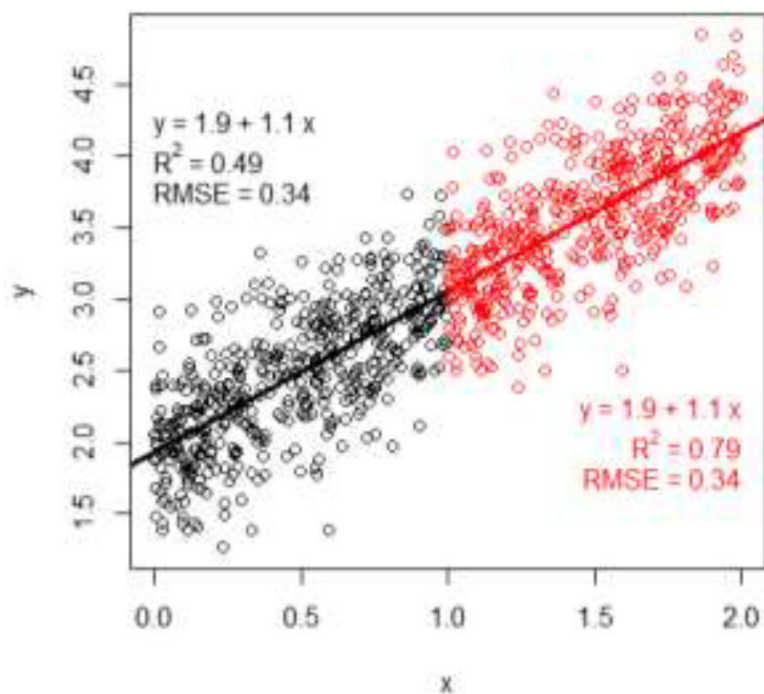
## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

## 8 References


1. Golbraikh A, Tropsha A. Beware of  $q^2$ ! *J. Mol. Graph. Model.* 2002; 20:269–276. [PubMed: 11858635]
2. Le T, Epa VC, Burden FR, Winkler DA. Quantitative Structure-Property Relationship Modeling of Diverse Materials Properties. *Chem. Rev.* 2012; 112:2889–919. [PubMed: 22251444]
3. Hastie, T.; Tibshirani, R.; Friedman, J. *The Elements of Statistical Learning*. 2 ed.. Springer; New York, USA: 2009. p. 745
4. Burden F, Winkler D. Bayesian Regularization of Neural Networks. *Meth. Mol. Biol.* 2008; 458:25–44.
5. Burden FR, Winkler DA. Robust QSAR Models Using Bayesian Regularized Neural Networks. *J. Med. Chem.* 1999; 42:3183–3187. [PubMed: 10447964]
6. Burden FR, Winkler DA. Optimal Sparse Descriptor Selection for QSAR Using Bayesian Methods. *QSAR Comb. Sci.* 2009; 28:645–653.
7. Besalu E, de Julian-Ortiz JV, Pogliani L. Trends and Plot Methods in MLR Studies. *J. Chem. Inf. Mod.* 2007; 47:751–760.
8. Seber, GAF. *Linear Regression Analysis*. John Wiley & Sons; NY: 1977. p. 465
9. Roy K, Chakraborty P, Mitra I, Ojha PK, Kar S, Das RN. Some Case Studies on Application Of “ $R(M)(2)$ ” Metrics for Judging Quality of Quantitative Structure-Activity Relationship Predictions: Emphasis on Scaling of Response Data. *J. Comp. Chem.* 2013; 34:1071–1082. [PubMed: 23299630]
10. Shayanfar A, Shayanfar S. Is Regression Through Origin Useful in External Validation of QSAR Models? *Eur. J. Pharm. Sci.* 2014; 59:31–5. [PubMed: 24721181]
11. Roy K, Kar S. The R Metrics and Regression Through Origin Approach: Reliable and Useful Validation Tools for Predictive QSAR Models (Commentary on 'Is regression through origin useful in external validation of QSAR models?'). *Eur. J. Pharm. Sci.* 2014; 62:111–114. [PubMed: 24881556]
12. Tropsha A, Gramatica P, Gombar VK. The Importance of Being Earnest: Validation is the Absolute Essential for Successful Application and Interpretation of QSPR Models. *QSAR Comb. Sci.* 2003; 22:69–77.
13. Kvalseth TO. Cautionary Note About  $R^2$ . *Am. Statistic.* 1985; 39:279–285.

14. Rousseeuw PJ. Least Median of Squares Regression. *J. Am. Stat. Assoc.* 1984; 79:871–880.
15. Burden FR, Winkler DA. An Optimal Self-Pruning Neural Network and Nonlinear Descriptor Selection in QSAR. *QSAR Comb. Sci.* 2009; 28:1092–1097.
16. Kubinyi H, Hamprecht FA, Mietzner T. Three-dimensional Quantitative Similarity-Activity Relationships (3D QSiAR) from SEAL Similarity Matrices. *J. Med. Chem.* 1998; 41:2553–2564. [PubMed: 9651159]
17. Tropsha A. Best Practices for QSAR Model Development, Validation, and Exploitation. *Mol. Inf.* 2010; 29:476–488.
18. Kendall, MG. Rank Correlation Methods. 4th ed. Griffin; London: 1970.
19. Varnek A, Kireeva N, Tetko IV, Baskin II, Solov'ev, V. P., Exhaustive QSPR Studies of a Large Diverse Set of Ionic Liquids: How Accurately Can We Predict Melting Points? *J. Chem. Inf. Model.* 2007; 47:1111–1122. [PubMed: 17381081]
20. Pearlman DA, Charifson PS. Are free Energy Calculations Useful in Practice? A Comparison with Rapid Scoring Functions for the p38 MAP Kinase Protein System. *J. Med. Chem.* 2001; 44:3417–3423. [PubMed: 11585447]



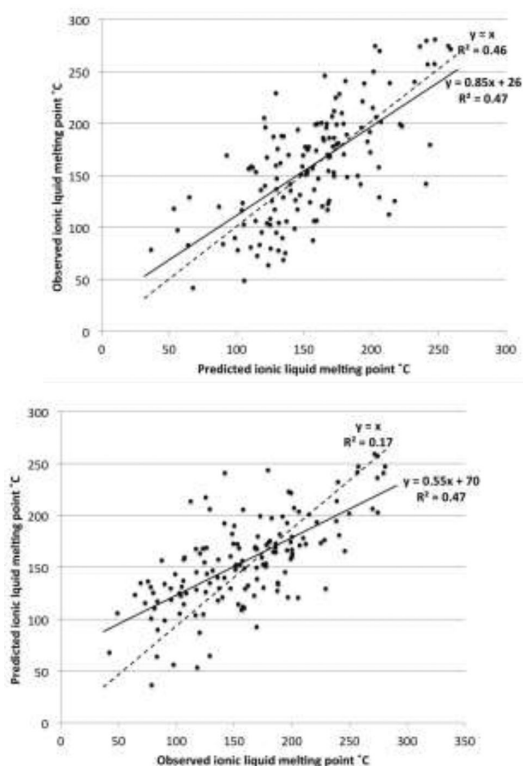
**Figure 1.**

For the black data alone,  $R^2 = 0.49$ , while for the combined data the regression line is the same but  $R^2 = 0.79$ . The RMSE is the same in each case, as the red and black residuals are identical, but increasing the range of activity values in the data increases  $R^2$ .



**Figure 2.** Illustrative plots of observed and predicted data. The dotted line shows the relationship  $y = \hat{y}$  data points for good models would lie close to this line. Solid regression lines and dashed lines for regression through the origin are also added in order to illustrate the test data conditions of Golbraikh and Tropsha,<sup>1</sup> though plotting such lines on the graphs is not advocated here.





**Figure 3.**

The test set predictions (154 compounds in test set) from a model predicting the melting point of 771 ionic liquids (ionic liquid data from Varnek et al.<sup>19</sup>) A regression line has been drawn through the data (solid line). The dotted line is for  $x=y$  and the regression values relate to fitting points to each line. The two graphs represent the plotting of the same test set data but with reversal of the assignment of the observed melting points to the axes. Note that the  $x=y$  regression gives markedly different results depending on which data are assigned to the y-axis.