

1: Regression and Correlation

In statistics, regression is a statistical process for evaluating the connections among variables. Regression equation calculation depends on the slope and y-intercept. Enter the X and Y values into this online linear regression calculator to calculate the simple regression equation line.

Contents Linear Regression Linear regression is used to predict the value of an outcome variable Y based on one or more input predictor variables X. The aim is to establish a linear relationship a mathematical formula between the predictor variable s and the response variable, so that, we can use this formula to estimate the value of the response Y, when only the predictors Xs values are known. Introduction The aim of linear regression is to model a continuous variable Y as a mathematical function of one or more X variable s , so that we can use this regression model to predict the Y when only the X is known. This mathematical equation can be generalized as follows: Collectively, they are called regression coefficients. Example Problem For this analysis, we will use the cars dataset that comes with R by default. You can access this dataset simply by typing in cars in your R console. You will find that it consists of 50 observations rows and 2 variables columns " dist and speed. Lets print out the first six observations here.. The graphical analysis and correlation study below will help with this. Graphical Analysis The aim of this exercise is to build a simple regression model that we can use to predict Distance dist by establishing a statistically significant linear relationship with Speed speed. But before jumping in to the syntax, lets try to understand these variables graphically. Typically, for each of the independent variables predictors , the following plots are drawn to visualize the following behavior: Visualize the linear relationship between the predictor and response Box plot: To spot any outlier observations in the variable. To see the distribution of the predictor variable. Ideally, a close to normal distribution a bell shaped curve , without being skewed to the left or right is preferred. Let us see how to make each one of them. Scatter Plot Scatter plots can help visualize any linear relationships between the dependent response variable and independent predictor variables. Ideally, if you are having multiple predictor variables, a scatter plot is drawn for each one of them against the response, along with the line of best as seen below. This is a good thing, because, one of the underlying assumptions in linear regression is that the relationship between the response and predictor variables is linear and additive. BoxPlot " Check for outliers Generally, any datapoint that lies outside the 1. If we observe for every instance where speed increases, the distance also increases along with it, then there is a high positive correlation between them and therefore the correlation between them will be closer to 1. The opposite is true for an inverse relationship, in which case, the correlation between the variables will be close to 0. A value closer to 0 suggests a weak relationship between the variables. A low correlation The function used for building linear models is lm. The lm function takes in two main arguments, namely: The data is typically a data. But the most common convention is to write out the formula directly in place of the argument as written below. Is this enough to actually use this model? Before using a regression model, you have to ensure that it is statistically significant. How do you ensure this? Lets begin by printing the summary statistics for linearMod. Checking for statistical significance The summary statistics above tells us a number of things. The p-Values are very important because, We can consider a linear model to be statistically significant only when both these p-Values are less than the pre-determined statistical significance level, which is ideally 0. This is visually interpreted by the significance stars at the end of the row. Null and alternate hypothesis When there is a p-value, there is a null and alternative hypothesis associated with it. In Linear Regression, the Null Hypothesis is that the coefficients associated with the variables is equal to zero. The alternate hypothesis is that the coefficients are not equal to zero. A larger t-value indicates that it is less likely that the coefficient is not equal to zero purely by chance. So, higher the t-value, the better. What this means to us? In our case, linearMod, both these p-Values are well below the 0. It is absolutely important for the model to be statistically significant before we can go ahead and use it to predict or estimate the dependent variable, otherwise, the confidence in predicted values from that model reduces and may be construed as an event of chance. How to calculate the t Statistic and p-Values? When the model co-efficients and standard error are known, the formula for calculating t Statistic and p-Value

is as follows: What R-Squared tells us is the proportion of variation in the dependent response variable that has been explained by this model. It's a better practice to look at the AIC and prediction accuracy on validation sample when deciding on the efficacy of a model. Now that's about R-Squared. What about adjusted R-Squared? As you add more X variables to your model, the R-Squared value of the new bigger model will always be greater than that of the smaller subset. This is because, since all the variables in the original model is also present, their contribution to explain the dependent variable will be present in the super-set as well, therefore, whatever new variable we add can only add if not significantly to the variation that was already explained. It is here, the adjusted R-Squared value comes to help. Adj R-Squared penalizes total value for the number of terms read predictors in your model. Therefore when comparing nested models, it is a good practice to look at adj-R-squared value over R-squared. Therefore, by moving around the numerators and denominators, the relationship between R^2 and R_{adj}^2 becomes: Both criteria depend on the maximized value of the likelihood function L for the estimated model. The AIC is defined as: The most common metrics to look at while selecting the model are:

2: Quadratic regression Calculator - High accuracy calculation

GraphPad Prism. Organize, analyze and graph and present your scientific data. MORE >

The estimated value based on the fitted regression model for the new observation at is: The prediction interval values calculated in this example are shown in the figure below as Low Prediction Interval and High Prediction Interval, respectively. The columns labeled Mean Predicted and Standard Error represent the values of and the standard error used in the calculations. Measures of Model Adequacy It is important to analyze the regression model before inferences based on the model are undertaken. The following sections present some techniques that can be used to check the appropriateness of the model for the given data. These techniques help to determine if any of the model assumptions have been violated. Coefficient of Determination R^2 The coefficient of determination is a measure of the amount of variability in the data accounted for by the regression model. As mentioned previously, the total variability of the data is measured by the total sum of squares, TSS . The amount of this variability explained by the regression model is the regression sum of squares, SSR . The coefficient of determination is the ratio of the regression sum of squares to the total sum of squares. For the yield data example, can be calculated as: It may appear that larger values of indicate a better fitting regression model. However, should be used cautiously as this is not always the case. The value of increases as more terms are added to the model, even if the new term does not contribute significantly to the model. Therefore, an increase in the value of cannot be taken as a sign to conclude that the new model is superior to the older model. Adding a new term may make the regression model worse if the error mean square, MSE , for the new model is larger than the of the older model, even though the new model will show an increased value of R^2 . In the results obtained from the DOE folio, is displayed as R^2 under the ANOVA table as shown in the figure below, which displays the complete analysis sheet for the data in the preceding table. These values measure different aspects of the adequacy of the regression model. For example, the value of S is the square root of the error mean square, MSE , and represents the "standard error of the model. The values of S , R^2 and R^2_{adj} indicate how well the model fits the observed data. Residual Analysis In the simple linear regression model the true error terms, ϵ_i , are never known. The residuals, e_i , may be thought of as the observed error terms that are similar to the true error terms. Since the true error terms, ϵ_i , are assumed to be normally distributed with a mean of zero and a variance of σ^2 , in a good model the observed error terms e_i . Thus the residuals in the simple linear regression should be normally distributed with a mean of zero and a constant variance of σ^2 . Residuals are usually plotted against the fitted values, \hat{y}_i , against the predictor variable values, x_i , and against time or run-order sequence, in addition to the normal probability plot. Plots of residuals are used to check for the following: Residuals follow the normal distribution. Residuals have a constant variance. Regression function is linear. A pattern does not exist when residuals are plotted in a time or run-order sequence. There are no outliers. Examples of residual plots are shown in the following figure. Such a plot indicates an appropriate regression model. Such a plot indicates increase in variance of residuals and the assumption of constant variance is violated here. Transformation on may be helpful in this case see Transformations. If the residuals follow the pattern of c or d, then this is an indication that the linear regression model is not adequate. Addition of higher order terms to the regression model or transformation on or may be required in such cases. A plot of residuals may also show a pattern as seen in e, indicating that the residuals increase or decrease as the run order sequence or time progresses. This may be due to factors such as operator-learning or instrument-creep and should be investigated further. Example Residual plots for the data of the preceding table are shown in the following figures. One of the following figures is the normal probability plot. It can be observed that the residuals follow the normal distribution and the assumption of normality is valid here. In one of the following figures the residuals are plotted against the fitted values, \hat{y}_i , and in one of the following figures the residuals are plotted against the run order. Both of these plots show that the 21st observation seems to be an outlier. Further investigations are needed to study the cause of this outlier. This perfect model will give us a zero error sum of squares. Thus, no error exists for the perfect model. However, if you record the response values for the same values of for a second time, in conditions maintained as strictly identical as possible to the first time,

observations from the second time will not all fall along the perfect model. The deviations in observations recorded for the second time constitute the "purely" random variation or noise. The sum of squares due to pure error abbreviated quantifies these variations.

3: How to Develop & Use a Regression Model for Sales Forecasting | Bizfluent

Linear regression is the most widely used statistical technique; it is a way to model a relationship between two sets of variables. The result is a linear regression equation that can be used to make predictions about data. Most software packages and calculators can calculate linear regression. For example: TI Excel.

The course uses the following text: John Wiley and Sons. The file follows this text very closely and readers are encouraged to consult the text for further information. Regression is used to predict the value of one variable based on the value of a different variable. Correlation is a measure of the strength of a relationship between variables. In the case of the examples used here, the data were obtained by counting the breathing rate of goldfish in a laboratory experiment. Nature of data The data for regression and correlation consist of pairs in the form x,y . The independent variable x is determined by the experimenter. This means that the experimenter has control over the variable during the experiment. In our experiment, the temperature was controlled during the experiment. The dependent variable y is the effect that is observed during the experiment. It is assumed that the values obtained for the dependent variable result from the changes in the independent variable. Regression and correlation analyses will determine the nature of this relationship, if any, and the strength of the relationship. It can be a consideration that all of the x,y pairs form a population. In some experiments, numerous observations of y are taken at each value of x . In these cases, each set of values of y taken at a particular value of x form a subpopulation of the data. Graphical representation Data are represented using a plot called a scatter plot or scatter diagram or $x-y$ plot. During analysis we try to find the equation of a line that fits the data. This is called the regression line. From algebra, we recall that points which are x,y pairs can be plotted on the Cartesian coordinate system. The slope is always the coefficient of the x term in the equation. By looking at the equation we can determine the slope and the y -intercept. Scatter diagrams can show a direct relationship between x and y . The direct relationship exists when the slope of the line b is positive. An inverse relationship exists when the slope of the line is negative. The nature of the relationship is discussed as part of correlation. We remember that m is the slope and b is the y -intercept. A regression line will have a general form.

4: Linear regression Calculator - High accuracy calculation

Simple linear regression is a way to describe a relationship between two variables through an equation of a straight line, called line of best fit, that most closely models this relationship. Linear Regression Formula.

You can use regression software to fit this model and produce all of the standard table and chart output by merely not selecting any independent variables. R-squared will be zero in this case, because the mean model does not explain any of the variance in the dependent variable: The forecasting equation of the mean model is: The sample mean has the non-obvious property that it is the value around which the mean squared deviation of the data is minimized, and the same least-squares criterion will be used later to estimate the "mean effect" of an independent variable. The error that the mean model makes for observation t is therefore the deviation of Y from its historical average value: The standard error of the model, denoted by s , is our estimate of the standard deviation of the noise in Y the variation in it that is considered unexplainable. Smaller is better, other things being equal: In the mean model, the standard error of the model is just is the sample standard deviation of Y : S denotes the sample standard deviation of X , using Excel notation. Note that the standard error of the model is not the square root of the average value of the squared errors within the historical sample of data. The accuracy of the estimated mean is measured by the standard error of the mean, whose formula in the mean model is: This is the estimated standard deviation of the error in estimating the mean. Notice that it is inversely proportional to the square root of the sample size, so it tends to go down as the sample size goes up. For example, if the sample size is increased by a factor of 4, the standard error of the mean goes down by a factor of 2, i. The accuracy of a forecast is measured by the standard error of the forecast, which for both the mean model and a regression model is the square root of the sum of squares of the standard error of the model and the standard error of the mean: This is the estimated standard deviation of the error in the forecast, which is not quite the same thing as the standard deviation of the unpredictable variations in the data which is s . It takes into account both the unpredictable variations in Y and the error in estimating the mean. In the mean model, the standard error of the mean is a constant, while in a regression model it depends on the value of the independent variable at which the forecast is computed, as explained in more detail below. The standard error of the forecast gets smaller as the sample size is increased, but only up to a point. More data yields a systematic reduction in the standard error of the mean, but it does not yield a systematic reduction in the standard error of the model. The standard error of the model will change to some extent if a larger sample is taken, due to sampling variation, but it could equally well go up or down. Confidence intervals for the mean and for the forecast are equal to the point estimate plus-or-minus the appropriate standard error multiplied by the appropriate 2-tailed critical value of the t distribution. The critical value that should be used depends on the number of degrees of freedom for error the number data points minus number of parameters estimated, which is $n-1$ for this model and the desired level of confidence. It can be computed in Excel using the T. Return to top of page. Formulas for the slope and intercept of a simple regression model: A simple regression model includes a single independent variable, denoted here by X , and its forecasting equation in real units is It differs from the mean model merely by the addition of a multiple of X_t to the forecast. The estimated coefficient b_1 is the slope of the regression line, i. The simple regression model reduces to the mean model in the special case where the estimated slope is exactly zero. The estimated slope is almost never exactly zero due to sampling variation, but if it is not significantly different from zero as measured by its t -statistic, this suggests that the mean model should be preferred on grounds of simplicity unless there are good a priori reasons for believing that a relationship exists, even if it is largely obscured by noise. Usually we do not care too much about the exact value of the intercept or whether it is significantly different from zero, unless we are really interested in what happens when X goes to "absolute zero" on whatever scale it is measured. Often X is a variable which logically can never go to zero, or even close to it, given the way it is defined. So, attention usually focuses mainly on the slope coefficient in the model, which measures the change in Y to be expected per unit of change in X as both variables move up or down relative to their historical mean values on their own natural scales of measurement. The coefficients, standard errors, and forecasts for this model are obtained as

follows. There are various formulas for it, but the one that is most intuitive is expressed in terms of the standardized values of the variables. A variable is standardized by converting it to units of standard deviations from the mean. P_X is the population standard deviation, as noted above. Sometimes the sample standard deviation is used to standardize a variable, but the population standard deviation is needed in this particular formula. The correlation coefficient is equal to the average product of the standardized values of the two variables: Also, if X and Y are perfectly positively correlated, $r = P_X P_Y$, which is equal to 1. The least-squares estimate of the slope coefficient b_1 is equal to the correlation times the ratio of the standard deviation of Y to the standard deviation of X : The ratio of standard deviations on the RHS of this equation merely serves to scale the correlation coefficient appropriately for the real units in which the variables are measured. The sample standard deviation could also be used here, because they only differ by a scale factor. The least-squares estimate of the intercept is the mean of Y minus the slope coefficient times the mean of X : This equation implies that Y must be predicted to be equal to its own average value whenever X is equal to its own average value. The sum of squared errors is divided by $n-2$ in this calculation rather than $n-1$ because an additional degree of freedom for error has been used up by estimating two parameters a slope and an intercept rather than only one the mean in fitting the model to the data. The standard error of the regression is an unbiased estimate of the standard deviation of the noise in the data, σ . Each of the two model parameters, the slope and intercept, has its own standard error, which is the estimated standard deviation of the error in estimating it. In general, the term "standard error" means "standard deviation of the error" in whatever is being estimated. P_X under the square root sign. This term reflects the additional uncertainty about the value of the intercept that exists in situations where the center of mass of the independent variable is far from zero in relative terms, in which case the intercept is determined by extrapolation far outside the data range. The standard error of the slope coefficient is given by: σ / P_X in the denominator. The terms in these equations that involve the variance or standard deviation of X merely serve to scale the units of the coefficients and standard errors in an appropriate way. This means that noise in the data whose intensity if measured by σ affects the errors in all the coefficient estimates in exactly the same way, and it also means that 4 times as much data will tend to reduce the standard errors of the all coefficients by approximately a factor of 2, assuming the data is really all generated from the same model, and a really huge amount of data will reduce them to zero. However, more data will not systematically reduce the standard error of the regression. As with the mean model, variations that were considered inherently unexplainable before are still not going to be explainable with more of the same kind of data under the same model assumptions. As the sample size gets larger, the standard error of the regression merely becomes a more accurate estimate of the standard deviation of the noise. Formulas for R-squared and standard error of the regression The fraction of the variance of Y that is "explained" by the simple regression model, r^2 . This is not supposed to be obvious. It is a "strange but true" fact that can be proved with a little bit of calculus. By taking square roots everywhere, the same equation can be rewritten in terms of standard deviations to show that the standard deviation of the errors is equal to the standard deviation of the dependent variable times the square root of $1 - r^2$: However, the sample variance and standard deviation of the errors are not unbiased estimates of the variance and standard deviation of the unexplained variations in the data, because they do not into account the fact that 2 degrees of freedom for error have been used up in the process of estimating the slope and intercept. The fraction by which the square of the standard error of the regression is less than the sample variance of Y which is the fractional reduction in unexplained variation compared to using the mean model is the "adjusted" R-squared of the model, and in a simple regression model it is given by the formula. In fact, adjusted R-squared can be used to determine the standard error of the regression from the sample standard deviation of Y in exactly the same way that R-squared can be used to determine the sample standard deviation of the errors as a fraction of the sample standard deviation of Y : You can apply this equation without even calculating the model coefficients or the actual errors! It follows from the equation above that if you fit simple regression models to the same sample of the same dependent variable Y with different choices of X as the independent variable, then adjusted R-squared necessarily goes up as the standard error of the regression goes down, and vice versa. Hence, it is equivalent to say that your goal is to minimize the standard error of the regression or to maximize adjusted R-squared through your

choice of X , other things being equal. However, as I will keep saying, the standard error of the regression is the real "bottom line" in your analysis: Adjusted R -squared can actually be negative if X has no measurable predictive value with respect to Y . If this is the case, then the mean model is clearly a better choice than the regression model. Some regression software will not even display a negative value for adjusted R -squared and will just report it to be zero in that case. Formulas for standard errors and confidence limits for means and forecasts

The standard error of the mean of Y for a given value of X is the estimated standard deviation of the error in measuring the height of the regression line at that location, given by the formula This looks like a lot like the formula for the standard error of the mean in the mean model: However, in the regression model the standard error of the mean also depends to some extent on the value of X , so the term is scaled up by a factor that is greater than 1 and is larger for values of X that are farther from its mean, because there is relatively greater uncertainty about the true height of the regression line for values of X that are farther from its historical mean value. The standard error for the forecast for Y for a given value of X is then computed in exactly the same way as it was for the mean model: In the regression model it is larger for values of X that are farther from the mean--i. The standard error of the forecast is not quite as sensitive to X in relative terms as is the standard error of the mean, because of the presence of the noise term s^2 under the square root sign. Remember that s^2 is the estimated variance of the noise in the data. In fact, s is usually much larger than $SE_{\text{mean } X}$ unless the data set is very small or X is very extreme, so usually the standard error of the forecast is not too much larger than the standard error of the regression. Finally, confidence limits for means and forecasts are calculated in the usual way, namely as the forecast plus or minus the relevant standard error times the critical t -value for the desired level of confidence and the number of degrees of freedom, where the latter is $n-2$ for a simple regression model. You can choose your own, or just report the standard error along with the point forecast. Here are a couple of additional pictures that illustrate the behavior of the standard-error-of-the-mean and the standard-error-of-the-forecast in the special case of a simple regression model. Because the standard error of the mean gets larger for extreme farther-from-the-mean values of X , the confidence intervals for the mean the height of the regression line widen noticeably at either end. The confidence intervals for predictions also get wider when X goes to extremes, but the effect is not quite as dramatic, because the standard error of the regression which is usually a bigger component of forecast error is a constant. Note that the inner set of confidence bands widens more in relative terms at the far left and far right than does the outer set of confidence bands. If the model assumptions are not correct--e. We look at various other statistics and charts that shed light on the validity of the model assumptions. The coefficients and error measures for a regression model are entirely determined by the following summary statistics: The correlation between Y and X , denoted by r_{XY} , is equal to the average product of their standardized values, i . The correlation between Y and X is positive if they tend to move in the same direction relative to their respective means and negative if they tend to move in opposite directions, and it is zero if their up-or-down movements with respect to their own means are statistically independent. The slope coefficient in a simple regression of Y on X is the correlation between Y and X multiplied by the ratio of their standard deviations: S can be used in this formula because they differ only by a multiplicative factor. In a simple regression model, the percentage of variance "explained" by the model, which is called R -squared, is the square of the correlation between Y and X . So, if you know the standard deviation of Y , and you know the correlation between Y and X , you can figure out what the standard deviation of the errors would be if you regressed Y on X . The sample standard deviation of the errors is a downward-biased estimate of the size of the true unexplained deviations in Y because it does not adjust for the additional "degree of freedom" used up by estimating the slope coefficient. An unbiased estimate of the standard deviation of the true errors is given by the standard error of the regression, denoted by s . In the special case of a simple regression model, it is: Adjusted R -squared, which is obtained by adjusting R -squared for the degrees of freedom for error in exactly the same way, is an unbiased estimate of the amount of variance explained:

5: Linear Regression With R

The formula for the best-fitting line (or regression line) is $y = mx + b$, where m is the slope of the line and b is the y-intercept. The equation itself is the same one used to find a line in algebra; but remember, in statistics the points don't lie perfectly on a line – the line is a model around which the data lie if a strong linear pattern exists.

These actions help to optimize operations and maximize profits. A regression model forecasts the value of a dependent variable -- in this case, sales -- based upon an independent variable. An Excel spreadsheet can easily handle this type of equation. Data Gathering Decide upon an independent variable. For example, suppose your company produces a product with sales that tie closely to changes in the price of oil. Your experience is that sales rise when the price of oil rises. To set up the regression, create a spreadsheet column for your annual sales over some number of previous years. Create a second column showing the percentage change in the year-over-year average price of oil in each of the sales years. To proceed, you will need the Excel Analysis ToolPak, which you can load for free by selecting "Add-ins" on the "Options" menu. Mark the range of the independent variable as the X-axis and that of the dependent variable as the Y-axis. Give a cell range for the output and mark the boxes for residuals. When you press "OK," Excel will compute the linear regression and display the results in your output range. The regression represents a straight line with a slope that best fits the data. Excel displays several statistics to help you interpret the strength of the correlation between the two variables. Video of the Day Brought to you by Techwalla Interpreting the Results The R-squared statistic indicates how well the independent variable forecasts sales. In this example, the R-squared of oil versus sales is .85. Any number above .85 indicates a strong relationship. The Y-intercept, in this example, is 15, shows the amount of product you would sell if the price of oil remained unchanged. The correlation coefficient, in this case .15, indicates that a 1 percent increase in the price of oil would drive sales up by 15 units. Using the Results The value of the linear regression depends on how well you can forecast the independent variable. For example, you might pay oil industry analysts for a private forecast that predicts a 6 percent increase in the price of oil over the next year. Multiply the correlation coefficient by 6, and add the result -- 90, -- to your Y-intercept amount of 15. The answer, 105, is the number of units you would likely sell if the price of oil rose 6 percent. You can use this prediction to prepare your production schedule for the upcoming year. You can also run the regression using different oil price movements to predict a best- and worst-case outcome. Of course, these are just predictions, and surprises are always possible. You can also run regressions with multiple independent variables, if appropriate.

6: Mathematics of simple regression

Purpose of use I was given 3 temperature measurements by a customer where X was known, and Y was given by the customer. Graphed, these measurements appeared to be linear, so this calculator allowed me to quickly compute slope/intercept values to compute the entire (useful) temperature scale from this customer's 3 measurements.

7: - The Multiple Linear Regression Model | STAT

Enter two data sets and this calculator will find the equation of the regression line and correlation coefficient. The calculator will generate a step by step explanation along with the graphic representation of the data sets and regression line.

8: Lesson 1: Simple Linear Regression | STAT

more response variables. These models are usually called multivariate regression models. In this chapter, we will introduce a new (linear algebra based) method for computing the parameter estimates of multiple regression models. This more compact method is convenient for models for which the number of unknown parameters is large.

9: Online calculator: Function approximation with regression analysis

Formula is used to calculate number of steps required to achieve desired elevation angle. Platform movement is not linear so this is perfect for adjusting the number of steps to get from one angle to another.

Bear Island reflections, Lake Winnepesaukee, New Hampshire Infrared optics and zoom lenses What was medicine like? National languages and teacher training in Africa Carolingian Cavalryman AD 768-987 (Warrior) I have not mouth and I am mean The art lover : someone who cant draw a straight line tries to defend her art-buying habit Lauren F. Winn A Database Library Textbook of orthopaedic medicine Hebrews 10:19-11:40 : the practical consequences of His superiority Letters from the darkling plain Scope of genetic engineering Experimenta circa effectum conflictus electrici in acum magneticam Regimenting the public mind : the methods of control in the propaganda system Session 1, July 4, 2005 Emerging targets in antibacterial and antifungal chemotherapy The infernal devices 2 Online resources for business Operations Gneissenau and Hammerschlag Natural Gas Conversion II Good story in english Sections and orders 11 Deterrent diplomacy The juice ladys guide to juicing for health Life is a celebration! The golden age of unstructured play, 1900-1950 Jewish participation in the leftist resistance in Nazi Germany Moths of Southern Africa Fanuc system 6m maintenance manual Maximizing Misfortune Organ Music for Lent and Easter Once on this river A different authority Certified estimating professional study guide An Urban Poets Conscious Perspective On the Kings Service Inward Glimpses of Men at Arms Christmas Catastrophe (Geronimo Stilton) Vba for modelers developing decision support systems Language in Canada Provider Report Cards