

The TwoKey Plot for Multiple Association Rules Control

Antony Unwin², Heike Hofmann¹, and Klaus Bernt²

¹ AT&T Research, Florham Park NJ 09732, USA
heike@att.research.com

² Department of Computer-Oriented Statistics and Data Analysis University of
Augsburg, 86135 Augsburg, Germany
{antony.unwin,klaus.bernt}@math.uni-augsburg.de

Abstract. Association rules have been popular in theory, though it is unclear how much success they have had in practice. Very many association rules are found in any application by any approach and they require effective pruning and filtering. There has been much research in this area recently, but less with the goal of providing a global overview and summary of all rules, which may then be used to explore the rules and to evaluate their worth. The unusual feature of association rules is that those with the highest objective values for the two key criteria (support and confidence) are not usually those with the most subjective interest (because we know the obvious results already). The TwoKey plot is a way of displaying all discovered association rules at once, while also providing the means to review and manage them. It is a powerful tool in order to get a first overview of the distribution of confidence and support. Features such as separate groups of rules or outliers are detected immediately. By exploiting various ancestor relationship structures among the rules, we can use the TwoKey Plot also as a visual assessment tool, closely related to pruning methods – e.g. those proposed by Bing Liu (1999). The concept will be illustrated using the interactive software MARC (Multiple Association Rules Control).

1 Introduction

Analysis by association rules is one of a number of techniques in the field of Data Mining, the study of very large data sets. Simple methods like Association Rules are popular because they can be applied to such large data sets, while some more traditional statistical methods do not scale up. It is typical of Data Mining analyses that very many results are produced, not only through the application of a single method but also because several different techniques may be applied. Tools for organising and managing large numbers of results are very necessary and some of the ideas in this paper will apply to results from other Data Mining approaches as well.

Association rules were introduced formally by Agrawal et al (1993) and have been discussed widely since (see Bruzzese & Davino (2001) or Hofmann & Wilhelm (2001) for recent references). They are a simple approach to analysing the

association between large numbers of categorical variables and part of their attraction is that the method is easy to explain and that individual results are readily understandable. What is not so easy to deal with is that for any realistic data set there are two main problems: firstly, that far too many association rules tend to be reported and most software packages just display them all in a fairly incomprehensible long list (see Table 1 for the first lines from a typical example); secondly, that the obvious way round the first problem “to just report the ‘best’ rules discovered” doesn’t work. Association rules that are best on the two key criteria of support and confidence are likely to be trivial or well known already. The most interesting rules will tend to be ones, which have good, but not outstanding, values for support and confidence. Unfortunately, there will usually be a large number of these. There has been research on pruning and filtering sets of rules, but while these methods look promising they still leave many rules to be evaluated. Background information has to be brought into play and it makes sense to try to concentrate on discussing a small group of rules. Some of these will have high criteria values, but be of no practical application. What remains will ideally be of value to the data set owners.

Table 1. Sample output from association rules software. (Only the first four rules found are shown, there are many more to follow!)

```

age>33 & weeks worked in year>8 → income=50000+
  [Cov=0.307 (91870); Sup=0.051 (15308); Conf= 0.167; Lift=2.69]
age>33 & wage/hour≤0 & weeks worked in year>8 → income=50000+
  [Cov=0.277 (82808); Sup=0.049 (14634); Conf=0.177; Lift=2.85]
age>33 & hispanic origin=All other & weeks worked in year>8 → income=50000+
  [Cov=0.276 (82712); Sup=0.049 (14585); Conf=0.176; Lift=2.84]
age>33 & wage/hour≤0 & hispanic origin=All other & weeks worked in year>8
  → income=50000+
  [Cov=0.248 (74276); Sup=0.047 (13931); Conf= 0.188; Lift=3.02]

```

2 Analysis by Association Rules

It is assumed that the data set comprises 0/1 variables reflecting the absence or presence of attributes in transactions. For non-binary data we will assume an appropriate recoding. An association rule [Agrawal et al., 1993] is defined as an implication of the form $X \rightarrow Y$, where X and Y are mutually exclusive itemsets. An association rule $X \rightarrow Y$ holds with confidence $c = c(X, Y)$, if $c\%$ of transactions in the data set that contain X also contain Y . $X \rightarrow Y$ has support $s = s(X, Y)$ in the data set, if $s\%$ of transactions contain X and Y . Confidence is equivalent to the conditional probability of Y given X . Support is equivalent to the joint probability of X and Y . In this paper we shall restrict ourselves to rules with single item conclusions (i.e. the itemset Y on the Right Hand Side always contains just one item) for expository purposes. Simple rules

with just one assumption (when itemset X on the Left Hand Side has only one member) will be referred to as level two rules, because two items are involved. By extension, a rule with $(m - 1)$ assumptions on the LHS will be a level m rule.

All Association Rules software will produce some kind of listing of rules (as in Table 1) and most packages will also provide summaries. Few appear to offer substantial organisation and management capabilities. To emphasise how important this is, it is worth noting that 5807 rules of between 2 and 8 levels were found for the 34 variable data set considered in this paper, even though support had to be more than 60% and confidence higher than 80%. The MARC software produces an initial summary window giving the parameter values chosen, the number of rules found (in total and by level) and a listing of the individual variables with their frequency of occurrence and the numbers of rules in which they are involved as either assumption or conclusion. There is also a set of controls to enable flexible selections of variables for detailed analysis and the capability to sort on any of the numbers produced. Sorting is a much-underrated facility in statistical software. That the rules should be sorted on confidence and support is obvious, but sorting on variables is also informative for grouping rules with the same variables as assumptions or conclusions.

3 Data Set Description

Association rules are supposed to be ideal for analysing market basket data, that is customer shopping transactions. Which products are bought together is an important question for store management. These data sets have large numbers of products (just think of how many different products your local supermarket stocks) and almost all are bought relatively rarely. Either the method is successful, but the results have not been released for reasons of confidentiality, or there has been less practical application of this kind than theoretical development.

Like most other researchers in this field we have therefore not analysed a real market basket data set (there are several analyses of artificial data sets in the literature) but have experimented with applying the method of association rules to another kind of binary data. The data set contains information on birdsbreeding habits by geographic region [Buckland et al., 1990]. 395 regions in North East Scotland were classified as breeding or non-breeding grounds for each of 34 different kinds of bird. There is interest in which groups of birds breed in the same regions. Do some birds share the same breeding grounds? Does the breeding presence of certain birds imply that a further species breeds there too? While this data set is not large and does not have a large number of variables, it is a real data set and it is both interesting to see what association rules might suggest and also as an example to illustrate how large numbers of rule results can be displayed and analysed.

One unexpected result of the analysis was the realisation that not only is effective post-filtering essential to cut down the number of rules discovered to a manageable group, but that pre-filtering is necessary as well. The high support threshold chosen immediately excluded all birds, which bred in less than 60% of

the regions from being in any of the rules, whether as assumption or conclusion. In many ways it is the rarer birds, which are more interesting. Setting the support threshold much lower and then selecting the rules in which the rarer birds were conclusions would always be possible, but huge numbers of rules of no interest would be produced. It is more practical to both lower the threshold and to have the option of excluding the more frequent birds from appearing as conclusions in any rules.

4 What Is a TwoKey Plot?

A TwoKey plot shows the two keys of confidence and support for all discovered association rules. Its lower limits are determined by the threshold criteria specified by the user. In Fig. 1 these were a minimum support of 60% and a min-

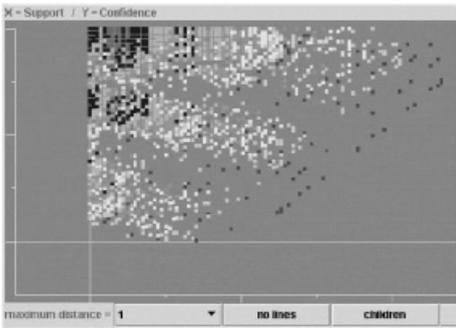


Fig. 1. A scatterplot of all 5807 rules found. No density estimation has been used. In the colour version the points are coloured by the level of the rule and this is partially seen in the shading

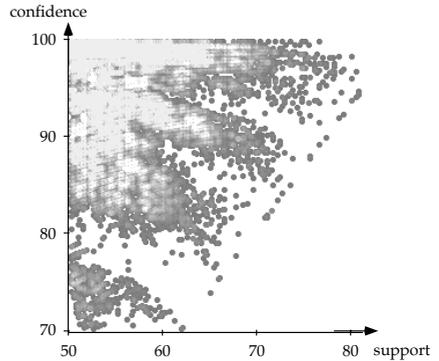


Fig. 2. MANET plot of confidence against support with pointsize = 2 and maximum brightness at pixels with 12 or more points

imum confidence of 80%. The maximum values are, of course, potentially 100% and this may be reached on the Y axis (confidence). It is more than a standard scatterplot because it has extensive interactive features, but it is more than an interactive scatterplot, because it has tools specifically designed for association rules. A TwoKey plot includes basic interactive features (see Unwin [1999]) such as querying (to find out which rules are represented), zooming (to study a subset of rules in more detail), selecting (to highlight rules of certain types) and linking (to garner information from other displays). There are also density estimation tools and line connections to display relationships between rules.

Several features can be identified in Fig. 1. The darker points on the right are all 2-level rules. They form straight lines because the LHS assumption is the same in all cases. The reason for this is simple. For an association rule $X \rightarrow Y$

the fraction of confidence and support is constant in X , i.e. for each LHS X , there is a straight line, on which all data points $X \rightarrow Y$ appear with increasing confidences (starting from the origin).

In Fig. 1 there are three blocks of rules, which have the appearance of clouds moving to the left leaving exhaust trails behind them. Rules involving a particular variable as conclusion will sometimes be in just one of these blocks. This arises when there is at least one good 2 level rule (i.e. good for that conclusion) and adding further assumptions lowers support but adds little to confidence.

The colour coding of rule levels is very informative in an initial overview. Low-level rules tend to be to the right (high support) and lower (less confidence), but individual rules of different types then stand out, especially higher-level rules with more support. To get an impression of how many rules are located in different areas of the plot, a density estimation version of the display is needed (without therefore the colour coding of levels). This cannot yet be carried out in MARC, but is available in the general interactive graphics software, MANET, written by Heike Hofmann [Hofmann, 2000b]. Figure 2 shows the same plot as Fig. 1 but in MANET and with a fast (though crude) density estimate. There are two parameters controlling the density estimation, the size r of individual points and an upper limit L on the density represented. In Fig. 1 $r = 2$ and $L = 12$. Each pixel at which 12 or more points overlap is coloured bright white and lower densities are shaded in a proportional fashion. (At many locations in the top left hand corner of the plot there are substantially more than 12 rules, as is readily found by direct querying of the locations). Although the method is simple, it is fast enough to be interactive so that changing r and L reveals different aspects of the density of rules. This is particularly useful when examining different parts of the plot. Parameter values suitable for high density areas leave lower density areas looking quite uniform and vice versa.

TwoKey plots can be used as standard interactive displays, where they may be linked to other graphic displays of the data: a barchart of level (for the same reason as the colour coding above), barcharts of the individual variables (to show in which groups of rules variables are involved) or mosaic plots [Hofmann, 2000a] of combinations of the variables (to select rules with certain combinations or, more usefully, to show which combinations arise in a selected group of rules). Linking is a powerful two-way tool. The information it delivers depends on the diagram where you make the selection and which diagrams you link to. Graphical interaction provides an easily understandable interface to the results for smaller numbers of variables. For working with larger numbers, the control window in MARC is more suitable, where many variables can be selected simultaneously. It could be argued that detailed interactive approaches are most relevant when studying subsets of results after extensive filtering, but both kinds of selection facilities should be provided. Using subjective filtering (e.g. just display the rules with a particular conclusion or display all rules involving any of a subset of variables) is an effective exploratory tool, but remains subjective. Selection rules based on objective filtering criteria must still be added, but it is not yet

clear which. There are several promising alternatives and they need to be put into a common framework. We envisage a three-part process:

1. Draw the TwoKey plot for all rules to gain a global overview.
2. Apply one or more “objective” filtering rules to reduce the number of rules under consideration.
3. Use subjective selection tools to uncover the most interesting rules amongst the remainder. This would not be a rigid three-stage process, but a continuing exploration to identify rules and subgroups of rules that are worth following up.

In an interactive environment of linked graphics (Becker et al 1987, Unwin et al 1996, Theus 1996, Wilhelm 2000), further information can be brought into consideration using selection & highlighting.

In this data set, spatial information about the items is available. We can extract two rules and have a closer look at them. The rules Mallard → Pheasant (54.4, 81.4) and Blue Tit, Curlew, Great Finch, Oystercatcher, Wren → Pheasant (50.1, 97.1) are shown on the corresponding map (see Fig. 3).

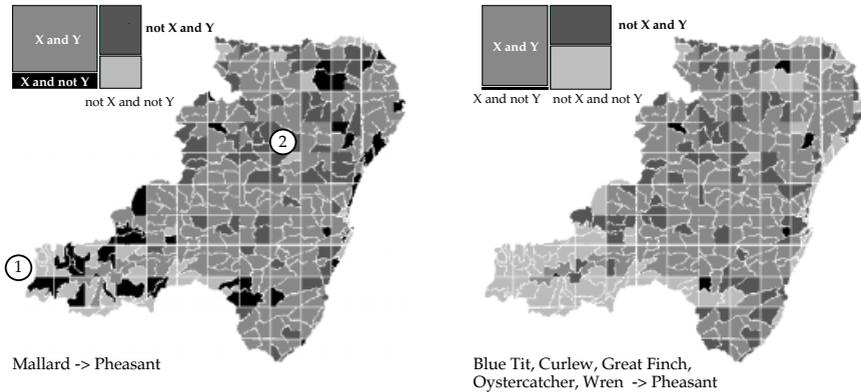


Fig. 3. Maps of the bird watching area. Gray shading corresponds to different combination of presence/absence of left and right hand side of rules

Gray shading is used to encode different combinations of left and right hand side of two rules. Areas, in which both birds of left and right hand side of the rules have been observed are coloured middle gray, areas, in which only birds from the right hand side but not of the left hand side have been observed are coloured black. The darkish gray represents all those areas where birds of the right hand side are present, but not all birds of the left hand side have been observed.

In this example you can see, that the simple rule Mallard → Pheasant shows some spatial relationship: several areas in the south-eastern part of the map are falsely assigned to Pheasant (marked by the circled 1), whereas the rule misses

out some areas in the centre (marked by 2). This problem is overcome by the more specific rule shown on the right of Fig. 3. The areas in the south-eastern part are fitted nicely by it and the rule detects most of the areas in the centre, too.

Using interactive graphics opens the way to find structural behaviour in association rules, e.g. separate spatial clusters or subgroups among the target population.

5 Special Tools in TwoKey Plots for Analysing Association Rules

5.1 Children and Descendants

Given a typical simple two level rule that $X \rightarrow Y$, it is clearly worth investigating what happens when more conditions are added to the left hand side. $X \rightarrow Y$ could be an obvious relationship and we might be interested in seeing how it might be extended. In MARC's TwoKey plot you can select any rules (not just two level rules) and have lines drawn from each rule to its immediate children. A child is defined as having one more condition on the left hand side. Figure 4

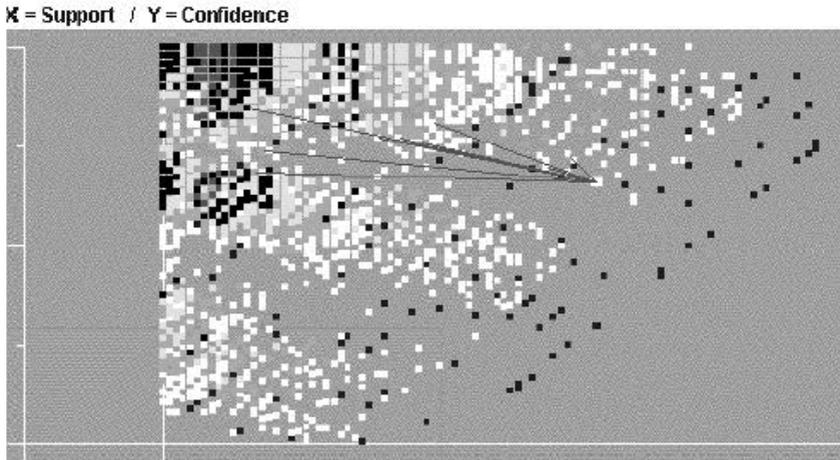


Fig. 4. The children of a 3-level rule have been linked by lines to the rule

shows a plot of this kind with one three level rule selected. Naturally, only rule children are shown which have sufficient confidence and support values to appear in the plot. It is immediately clear which additional group of rules it might be worth looking at if we start from the selected rule. All rule children must lie to the left (they cannot have higher support) and only rules with higher confidence will be worth considering. There are likely to be two kinds: a few rules which have

somewhat higher confidence, but only a little less support and some rules which have a lot less support, but rather more confidence. It very much depends on the client's goals and the particular rules, which group might be more important. We will have a closer look at this situation in Sect. 6. In Fig. 4 we can see that only 3 children are worth considering and they have very different reductions in support. Note that the display would also make clear if there were a number of closely related alternatives or if one rule child was substantially better than the others. It is possible in MARC to display descendants of 2 or more generations, but the display rapidly becomes unclear. It is a topic of current research to see how this might be improved.

5.2 Parents and Ancestors

Just as it can be informative to study rule descendants, it is helpful to be able to examine rule predecessors. For a single m level rule there can be at most $(m - 1)$ parents. Lines may of course be drawn in MARC from selected rules to their parents. Parents must lie to the right (higher support). A rule with a parent with higher confidence will not be a good rule. There is an argument for filtering out all such rules, but, initially at any rate, it seems appropriate to display all rules that pass the user-specified thresholds.

5.3 Neighbours

Individual association rules are rather blinkered. They tell you nothing about closely related rules. Relatives cover one form of closeness, but the term neighbour is used in MARC to cover all others.

A rule $R2: X' \rightarrow Y'$ is a neighbour of distance d of another rule $R1: X \rightarrow Y$ if one of the following holds (where a step is an addition or a deletion):

- (a) the RHS is the same ($Y' = Y$) and the LHS's are d steps apart;
- (b) the RHS is different and the LHSs are $d - 2$ steps apart.

With this definition the only immediate neighbours ($d = 1$) are parents and children. Rules with only the RHS different are neighbours of distance 2.

Linking neighbours of a selected rule by lines in MARC helps in various ways. For instance, for distance $d = 2$, the rules on a straight line are those with the same assumptions but different conclusions. The other rules of the same colour have the same conclusion but one different assumption. The remaining rules of different colours have the same conclusion but either two fewer assumptions (grandparents) or two more (grandchildren).

MARC produces tables of results for the selected subgroup of rules to provide a textual overview to complement the graphic display. Further developments will include the incorporation of statistical tests in these tables to identify rules, which are "significantly" different from the originally selected rules, and this may also be colour-coded on the lines in the graphic display. Significance would only be used as a rough guideline. The interdependence of the tests and the multiple testing involved precludes anything more, but it is valuable to attempt to assess relationships in a formal statistical way.

6 Relationships among Rules

Selecting rules from the set of results is a crucial issue, with the help of interactive selection methods and linking we have a lot more choices for criteria of acceptance, all of which may be sensible in the background of a specific application.

Of particular interest is, how we may compare rules of the form $X_1 \rightarrow Y$ and $X_2 \rightarrow Y$, where the first rule has higher confidence, but less support than the second rule. Is there a way to tell, which of these rules is “better” than the other?

Between each two of the rules in fig. 5 that fulfill an ancestor relationship a line is drawn.

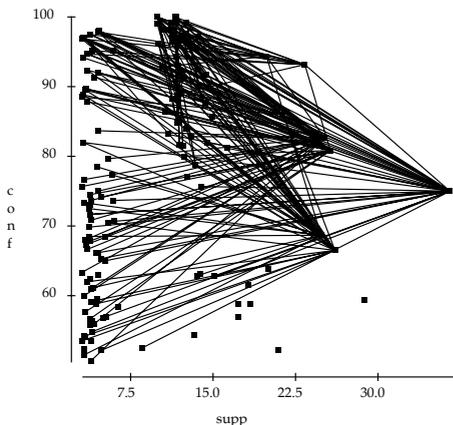


Fig. 5. Lines between each ancestor/successor pair of rules

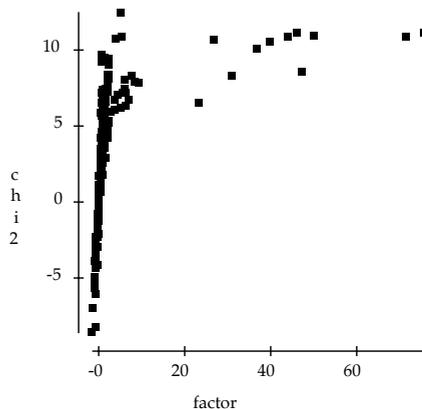


Fig. 6. Slope sl between two related rules vs. their test-statistic T

For these lines, the angle between a rule and its successor is of interest: it is clear, that an angle $\leq 0^\circ$ connects a rule to a successor with less confidence - something, which marks the successor rule as a total failure. Lines with a very steep angle, on the other hand, show those successors, which have only little less support than their ancestor but more confidence.

Can the angles of these lines therefore be exploited for a visual pruning mechanism? For this, we have a look at the (negative) slope sl between $X \rightarrow Z(s_1, c_1)$ and its successor $X \cap Y \rightarrow Z(s_2, c_2)$:

$$sl = \tan(\alpha) = \frac{c_2 - c_1}{s_1 - s_2}$$

This expression for the slope between two related rules recalls another test-statistic, which is used widely for pruning rules based on their ancestor relation-

ship structure [Bing Liu et al., 1999]: Let T be the test statistic of $P(Z | X, Y)$ vs $P(Z | X, \neg Y)$:

$$T = \sqrt{nP(X)} \frac{P(Y \cap Z | X) - P(Y | X)P(Z | X)}{\sqrt{P(Z | X)P(\neg Z | X) \cdot P(Y | X)P(\neg Y | X)}}$$

T is approximately normally distributed (it's the square root of a χ^2 test of conditional independence of Y and Z given X). Prominent features in Fig. 6 are the strong linear relationship between the lower values of the statistics and the cloud of points with high slope. Within this cloud of points, T and sl do take very different values. The two rules mined from the SAS Assoc Data, herring & baguette \rightarrow heineken and chicken & coke & sardines \rightarrow heineken, have similar values in T , but different slope values:

rule	ancestor	T	sl
herring&baguette \rightarrow heineken	baguette \rightarrow heineken	10.72	4.13
chicken&coke&sardines \rightarrow heineken	chicken&coke \rightarrow heineken	10.82	71.41

In order to explain the differences in T and slope, we rewrite the slope as:

$$sl = [P(Y \cap Z | X) - P(Y | X)P(Z | X)] \cdot \frac{1}{P(Y | X)P(X \cap Z \cap \neg Y)}$$

The first term of this product gives a measure for the conditional independence of Y and Z given X and is well known from the test statistic T . The second term is more tricky: $P(Y | X)$ will be large, if $P(X \cap \neg Y)$ is small; this measures the loss in the support of $X \rightarrow Z$ by adding the item Y . $P(X \cap Z \cap \neg Y)$ is linear in the number of result hits that are thrown away by adding Y .

Some properties of the slope:

- according to the value of the slope that rule $X \cap Y \rightarrow Z$ is chosen, which cuts off small slices from the support while gaining maximum confidence.
- It ignores the values of the ancestor rule $X \rightarrow Z$.
- It's not additive:

$$sl_3 \leq sl_1 + sl_2,$$

where sl_1 is the slope between rules $X \rightarrow Z$ and $\rightarrow Z$, sl_2 is the slope between rules $X \cap Y \rightarrow Z$ and $X \rightarrow Z$ and sl_3 is the slope between rules $X \cap Y \rightarrow Z$ and $\rightarrow Z$,

For each ancestor rule, the results of sl and T show a strong linear relationship (cf. Fig. 7).

Conclusion: for single ancestor rules the results from the test statistic and the slope are approximately linear dependent. The angles between the ancestor rule and its successors therefore may be used as a measure for the strength of a successor rule. This leads to the following criterion: A rule $X \cap Y_1 \rightarrow Z$ is dominated by a rule $X \cap Y_2 \rightarrow Z$,

$$\begin{aligned} &\iff sl(X \cap Y_1 \rightarrow Z) < sl(X \cap Y_2 \rightarrow Z) \\ &\iff \frac{P(Z | X, Y_2) - P(Z | X)}{P(Z \cap X) - P(Z \cap X \cap Y_2)} > \frac{P(Z | X, Y_1) - P(Z | X)}{P(Z \cap X) - P(Z \cap X \cap Y_1)}. \end{aligned}$$

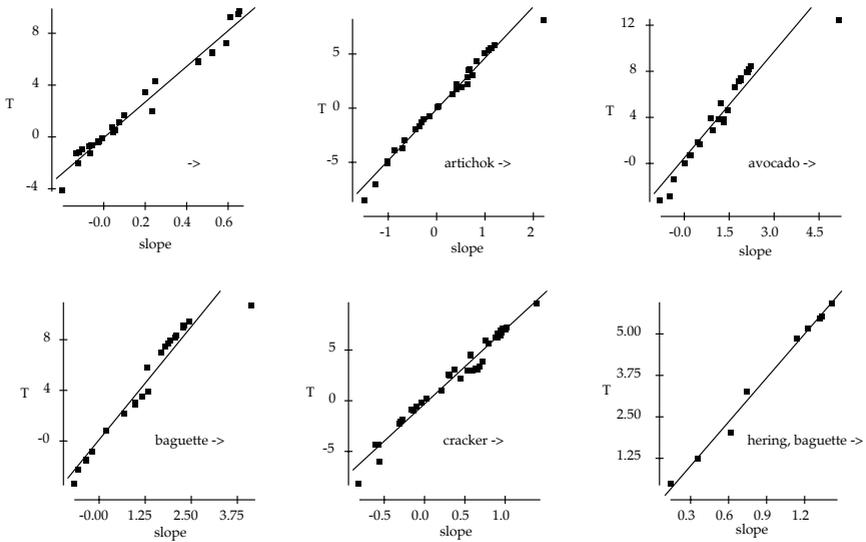


Fig. 7. Linear dependencies between sl and T for fixed ancestors

7 Summary

The concept of TwoKey Plots is a very simple but very effective one. It has been introduced as a means of providing a global overview of large amounts of association rules discovered in an analysis and of providing the tools to explore and evaluate them. The range and flexibility of the resulting graphic displays make this a very attractive method for working with association rules.

Data Mining methods can produce many different outputs and these need to be organised, managed and filtered. Pruning or filtering methods are an important first step, but they leave many rules still to be assessed. Exploiting the ancestor relationship structure of rules gives a visual approach to filter out potentially interesting rules. Here, we have seen that a visual analysis of the standard ancestor-successor relationship coincides with a χ^2 -test of independence as proposed by Bing et al (1999).

The effective application of any method of analysis lies in the combination of objective results and subjective knowledge. The TwoKey plot is a step in this direction for association rules. Its particular strength lies in interactivity, so that users can incorporate both their background knowledge and their foreground goals in evaluating the rules that have been discovered.

Software. The MARC (Multiple Association Rules Control) software used for the work in this paper is being developed by the authors at Augsburg. Klaus Bernt is responsible for the system design and programming. MARC is written in Java and it is planned to make the software available to other researchers in the near future. (Check our website www1.math.uni-augsburg.de for further

details.) Like all Augsburg software, MARC is named after a painter close to the Impressionists. Franz Marc was born in Munich, near to Augsburg, and was an influential member of the Blaue Reiter group.

References

- [Agrawal et al., 1993] Agrawal R., Imielinski T. & Swami A. (1993) *Mining Associations between Sets of Items in Massive Databases*. Proc. of the ACM-SIGMOD 1993 Int'l Conference on Management of Data, Washington D.C., pp. 207 - 216.
- [Becker et al., 1987] Becker R.A., Cleveland W.S. & Wilks A. (1987) Dynamic Graphics For Data Analysis. In: Cleveland & McGill, 1988, pp.1-50.
- [Bing Liu et al., 1999] Bing Liu, Wynne Hsu & Yiming Ma (1999) Pruning and Summarizing the Discovered Association Rules. In Proc. of the 5th ACM SIGKDD Conference: KDD99, pp.125-134, 1999.
- [Buckland et al., 1990] Buckland, S.T., Bell, M.V. and Picozzi, N. (eds) (1990) *The Birds of North-East Scotland*. North-East Scotland Bird Club, Aberdeen. 473pp.
- [Bruzzeze & Davino, 2001] Bruzzeze, D., Davino, C. (2001) Statistical Pruning of Discovered Association Rules. Computational Statistics (to be published)
- [Cleveland & McGill, 1988] Cleveland W. & McGill R. (1988) *Dynamic Graphics for Statistics*. Pacific Grove, CA: Wadsworth & Brooks, Inc.
- [Hofmann, 2000a] Hofmann H. (2000) Exploring categorical data: interactive mosaic plots. *Metrika*, 51(1), 11-26.
- [Hofmann, 2000b] Hofmann H. (2000) MANET www1.math.uni-augsburg.de/Manet/. In Augsburg: Rosuda.
- [Hofmann et al., 2000] Hofmann H., Siebes A., Wilhelm A. (2000) Visualizing Association Rules with Interactive Mosaic Plots, In Proc. of the 6th Int'l Conference of SigKDD, pp. 227-235.
- [Hofmann & Wilhelm, 2001] Hofmann, H., Wilhelm, A. (2001) Visual Comparison of Association Rules. Computational Statistics (to be published)
- [Mannila H., 1997] Mannila, H (1997) Methods and Problems in Data Mining. In F. Afrati Kolaitis, P. (Ed.), *International Conference on Database Theory*, . Delphi: Springer.
- [Piatetsky-Shapiro & Masand, 1999] Piatetsky-Shapiro, G. & Masand, B. (1999) Estimating Campaign Benefits and Modeling Lift. In 5th ACM SIGKDD, (pp. 185-193).
- [Theus, 1996] Theus M. (1996) *Theorie und Anwendung Interaktiver Statistischer Graphik*. Augsburg: mathematisch-naturwissenschaftliche Schriften, 14, Wißner Verlag, Augsburg, Germany.
- [Unwin, 1999] Unwin, A. R. (1999) Requirements for interactive graphics software for exploratory data analysis. *Computational Statistics*, 14, 7-22.
- [Unwin et al., 1996] Unwin A. R., Hawkins G. Hofmann H. & Siegl B. (1996) Interactive Graphics for Data Sets with Missing Values - MANET. *Journal of Computational and Graphical Statistics*, 5, 2, pp. 113 - 122.
- [Wilhelm, 2000] Wilhelm A. (2000) *Interactive Statistical Graphics: The Paradigm of Linked Views*. Habilitationsschrift, Universität Augsburg.