Revisiting interestingness of strong symmetric association rules in educational data

Agathe Merceron[1], Kalina Yacef[2]

[1]University of Applied Sciences TFH Berlin, Media and Computer Science Department,
Luxemburgerstr. 10,13353 Berlin, Germany
merceron@tfh-berlin.de
[2]School of Information Technologies, University of Sydney
NSW 2006, Australia
kalina@it.usyd.edu.au

Abstract. Association rules are very useful in Educational Data Mining since they extract associations between educational items and present the results in an intuitive form to the teachers. Furthermore, they require less extensive expertise in Data Mining than other methods. We have extracted association rules with data from the Logic-ITA, a web-based learning environment to practice logic formal proofs. We were interested in detecting associations of mistakes. The rules we found were symmetrical, such as X→Y and Y→X, both with a strong support and a strong confidence. Furthermore, P(X) and P(Y) are both significantly higher than P(X,Y). Such figures lead to the fact that several interestingness measures such as lift, correlation or conviction rate X and Y as independent. Does it mean that these rules are not interesting? We argue in this paper that this is not necessarily the case. We investigated other relevance measures such as Chi square, cosine and contrasting rules and found that the results were leaning towards a positive correlation between X and Y. We also argue pragmatically with our experience of using these association rules to change parts of the course and of the positive impact of these changes on students' marks. We conclude with some thoughts about the appropriateness of relevance measures for Educational data.

Keywords: Association rules, Interestingness measures.

1 Introduction

Association rules are very useful in Educational Data Mining since they extract associations between educational items and present the results in an intuitive form to the teachers. In [1], association rules are used to find mistakes often made together while students solve exercises in propositional logic. [2] and [3] used association rules, combined with other methods, to personalise students' recommendation while browsing the web. [4] used them to find various associations of student's behavior in their Web-based educational system LON-CAPA. [5] used fuzzy rules in a personalized e-learning material recommender system to discover associations between students' requirements and learning materials. [6] combined them with genetic programming to discover relations between knowledge levels, times and scores that help the teacher modify the course's original structure and content.

Compared with other Data Mining techniques, association rules require less extensive expertise. One reason for that is that there is mainly one algorithm to extract association rules from data. The selection of items and transactions within the data remains intuitive. In comparison with a classification task for example, there are many classifiers that, with the same set of data, can give different results. The data preparation and most importantly the definition of concepts specific to a particular algorithm (such as the concept of distance between elements) can be complex and it is often not easy to understand which the right choice is and why it works or not. [4] is a good example of a complex application of classification in Educational Data Mining.

However association rules also have their pitfall, in particular with regard to the extraction of interesting rules. This is a common concern for which a range of measures exist, depending on the context [7, 8]. We explore in this paper a few measures in the context of our data. We extracted association rules from the data stored by the Logic-ITA, an intelligent tutoring system for formal proof in propositional logic [9]. Our aim was to know whether there were mistakes that often occurred together while students are training. The results gave symmetric strong associations between 3 mistakes. Strong means that all associations had a strong support and a strong confidence. Symmetric means that X→Y and Y→X were both associations extracted. Puzzlingly, other measures of interestingness such as lift, correlation, conviction or Chi-square indicated poor or no correlation. Only cosine was systematically high, implying a high correlation between the mistakes. In this paper, we investigate why these measures, except cosine, do poorly on our data and show that our data have a quite special shape. Further, Chi-square on larger datasets and contrasting rules introduced in [10] give an interesting perspective to our rules. Last but not least, we did not dismiss the rules found as 'uninteresting', on the contrary. We used them to review parts of the course. After the changes, there was no significant change in the associations found in subsequent mining, but students' marks in the final exam have steadily increased [9, 11].

2 Association rules obtained with the Logic-ITA

We have captured 4 years of data from the Logic-ITA [9], a tool to practice logic formal proofs. We have, among other analysis, extracted association rules about the mistakes made by our students in order to support our teaching. Before we describe the elements of this data, let us first present the basic concepts that we use about the association rules.

2.1 What can association rules do?

Association rules come from basket analysis [12] and capture information such as if customers buy beer, they also buy diapers, written as beer→diapers. Two measures accompany an association rule: support and confidence. We introduce these concepts now.

Let I = $\{I_1, I_2, ..., I_m\}$ be a set of m items and T = $\{t_1, t_2, ..., t_n\}$ be a set of n transactions, with each t_i being a subset of I.

An *association rule* is a rule of the form X→Y, where X and Y are disjoint subsets of I having a support and a confidence above a minimum threshold.

Support: *sup*(X→Y) = |{t_i such that t_i contains both X and Y}| / n. In other words, the support of a rule X→Y is the proportion of transactions that contain both X and Y. This is also called P(X, Y), the probability that a transaction contains both X and Y. Support is symmetric: *sup*(X→Y) = *sup*(Y→X).

Confidence: *conf*(X→Y) = |{t_i such that t_i contains both X and Y}| / |{ t_i containing X}|. In other words, the confidence of a rule X→Y is the proportion of transactions that contain both X and Y among those that contain X. An equivalent definition is : *conf*(X→Y) = P(X, Y) / P(X), with P(X) = |{ t_i containing X}| / n. Confidence is not symmetric. Usually *conf*(X→Y) is different from *conf*(Y→X).

Support makes sure that only items occurring often enough in the data will be taken into account to establish the association rules. Confidence is the proportion of transactions containing both X and Y among all transactions containing X. If X occurs a lot naturally, then almost any subset Y could be associated with it. In that case P(X) will be high and, as a consequence, *conf*(X→Y) will be lower.

Symmetric association rule: We call a rule X→Y a *symmetric* association rule if *sup*(X→Y) is above a given minimum threshold and both *conf*(X→Y) and *conf*(Y→X) are above a given minimum threshold. This is the kind of association rules we obtained with the Logic-ITA.

2.2 Data from Logic-ITA

The Logic-ITA was used at Sydney University from 2001 to 2004 in a course formerly taught by the authors. Over the four years, around 860 students attended the course and used the tool. An exercise consists of a set of formulas (called premises) and another formula (called the conclusion). The aim is to prove that the conclusion can validly be derived from the premises. For this, the student has to construct new formulas, step by step, using logic rules and formulas previously established in the proof, until the conclusion is derived. There is no unique solution and any valid path is acceptable. Steps are checked on the fly and, if incorrect, an error message and possibly a tip are displayed.

All steps, whether correct or not, are stored for each user and each attempted exercise. In case of incorrect steps, the error message is also stored. A very interesting task was to analyse these mistakes and try and detect associations within them. This is why we used association rules. We defined the set of items I as the set of possible mistakes or error messages. We defined a transaction as the set of mistakes made by one student on one exercise. Therefore we obtain as many transactions as exercises attempted with the Logic-ITA during the semester, which is about 2000.

2.3 Association rules obtained with Logic-ITA

We used association rules to find mistakes often occurring together while solving exercises. The purpose of looking for these associations was for the teacher to ponder and, may be, to review the course material or emphasize subtleties while explaining concepts to students. Thus, it made sense to have a support that is not too low. The strongest rules for 2004 are shown in Table 1. The first association rule says that if students make mistake *Rule can be applied, but deduction incorrect* while solving an exercise, then they also made the mistake *Wrong number of line references given* while solving the same exercise. As we can see in the small subset of 3 pairs of rules shown in this table, the rules are symmetric and display comparable support and confidence. Findings were quite similar across the years (2001 to 2004).

Table 1. Some association rules for Year 2004.

M11 ==> M12 [sup: 77%, conf: 89%] M12 ==> M11 [sup: 77%, conf: 87%] M11 ==> M10 [sup: 74%, conf: 86%] M10 ==> M11 [sup: 78%, conf: 93%] M12 ==> M10 [sup: 78%, conf: 89%] M10 ==> M12 [sup: 74%, conf: 88%]	M10: Premise set incorrect M11: Rule can be applied, but deduction incorrect M12: Wrong number of line reference given

3 Measuring interestingness

Once rules are extracted, the next step consists in picking out meaningful rules and discarding others. We will first present some available measures and then compare them on a series of datasets.

3.1 Some measures of interestingness

It is a fact that strong association rules are not necessarily interesting [7]. Several measures, beside confidence, have been proposed to better measure the correlation between X and Y. Here we consider the following measures: lift, correlation, conviction, Chi-square testing and cosine.

lift(X→Y) = *conf*(X → Y) / P(Y). An equivalent definition is: P(X, Y) / P(X)P(Y). Lift is a symmetric measure. A lift well above 1 indicates a strong correlation between X and Y. A lift around 1 says that P(X, Y) = P(X)P(Y). In terms of probability, this means that the occurrence of X and the occurrence of Y in the same transaction are independent events, hence X and Y not correlated.

Correlation(X→Y) = P(X, Y)-P(X)P(Y) / sqrt(P(X)P(Y)(1-P(X))(1-P(Y))). Correlation is a symmetric measure. A correlation around 0 indicates that X and Y are not correlated, a negative figure indicates that X and Y are negatively correlated and a positive figure that they a positively correlated. Note that the denominator of the division is positive and smaller than 1. Thus the absolute value |*cor*(X→Y)| is greater than |P(X, Y)-P(X)P(Y)|. In other words, if the lift is around 1, correlation can still be significantly different from 0.

Conviction(X→Y) = (1 − P(Y)) / (1 − *conf*(X→Y)). Conviction is not a symmetric measure. A conviction around 1 says that X and Y are independent, while conviction is infinite as *conf*(X→Y) is tending to 1. Note that if P(Y) is high, 1 − P(Y) is small. In that case, even if *conf*(X, Y) is strong, *conviction*(X→Y) may be small.

To perform the Chi-square test, a table of expected frequencies is first calculated using P(X) and P(Y) from the contingency table. The expected frequency for (X and Y) is given by the product P(X)P(Y). Performing a grand total over observed frequencies versus expected frequencies gives a number which we denote by Chi. Consider the contingency table shown in Table 2. P(X) = P(Y) = 550/2000. Therefore the expected frequency (Xe and Ye) is 550 x 550 / 2000 = 151.25 as shown in Table 3. We calculate the other frequencies similarly. The grand total for Chi is therefore:
Chi = $(500-151.25)^2$ / 151.25 + $(50-398.75)^2$ / 398.75 + $(50-398.75)^2$ / 398.75 + $(1400-1051.25)^2$ / 1051.25 = 1529.87.

Table 2. A contingency table.

	X	not X	Total
Y	500	50	550
not Y	50	1400	1450
Total	550	1450	2000

Table 3. Expected frequencies for low support and strong confidence.

	Xe	not Xe	Total
Ye	151.25	398.75	550
not Ye	398.75	1051.25	1450
Total	550	1450	2000

The obtained number Chi is compared with a cut-off value read from a Chi-square table. For the probability value of 0.05 with one degree of freedom, the cut-off value is 3.84. If Chi is greater than 3.84, X and Y are regarded as correlated with a 95% confidence level. Otherwise they are regarded as non-correlated also with a 95% confidence level. Therefore in our example, X and Y are highly correlated.

Cosine(X→Y) = P(X, Y) / sqrt(P(X)P(Y)), where sqrt(P(X)P(Y)) means the square root of the product P(X)P(Y). An equivalent definition is: Cosine(X→Y) = |{t_i such that t_i contains both X and Y}| / sqrt (|{t_i containing X}| |{t_i containing Y}|). Cosine is a number between 0 and 1. This is due to the fact that both P(X, Y) ≤ P(X) and P(X, Y) ≤ P(Y). A value close to 1 indicates a good correlation between X and Y. Contrasting with the previous measures, the total number of transactions n is not taken into account by the cosine measure. Only the number of transactions containing both X and Y, the number of transactions containing X and the number of transactions containing Y are used to calculate the cosine measure.

3.2 Comparing these measures

Measures for interestingness as given in the previous section differ not only in their definition but also in their result. They do not rate the same sets the same way. In [7], Tan et al. have done some extensive work in exploring those measures and how well they capture the dependencies between variables across various datasets. They considered 10 sets and 19 interestingness measures and, for each measure, gave a ranking for the 10 sets. Out of these 10 sets, the first 3 sets (for convenience let us call them E1, E2 and E3 as they did in their article) bear most similarities with the data we have obtained from Logic-ITA because they lead to strong symmetric rules. However there is still a substantial difference between these 3 sets and our sets from the Logic-ITA. In [7]'s datasets E1, E2 and E3, the values for P(X, Y), P(X) and P(Y) are very similar, meaning that X and Y do not occur often one without the other. In contrast, in the sets from the Logic-ITA, P(X) and P(Y) are significantly bigger than P(X, Y). As we will see this fact has consequences both for correlation and conviction.

Since the datasets from [7] did not include the case of our datasets, we also explored the interestingness measures under different variant of the datasets. In the following we take various examples of contingency tables giving symmetric association rules for a minimum confidence threshold of 80% and we look at the various interestingness results that we get. The set S3 and S4 are the ones that match best our data from the Logic-ITA. To complete the picture, we included symmetric rules with a relatively low support of 25%, though we are interested in strong rules with a minimum support of 60%. This table is to be interpreted as follows. 2000 exercises have been attempted by about 230 students. (X, Y) gives the number of exercises in which both mistakes X and Y were made, (X, not Y) the number of exercises in which the mistake X was made but not the mistake Y, and so on. For the set S3 for example, 1340 attempted solutions contain both mistake X and mistake Y, 270 contain mistake X but not mistake Y, 330 contain mistake Y but not mistake X and 60 attempted solutions contain neither mistake X nor mistake Y. The last 3 lines, S7 to S9, are the same as S2 to S4 with a multiplying factor of 10.

Table 4. Contingency tables giving symmetric rules with strong confidence

	X,Y	X, not Y	not X, Y	not X, not Y.
S1	500	50	50	1400
S2	1340	300	300	60
S3	1340	270	330	60
S4	1340	200	400	60
S5	1340	0	0	660
S6	2000	0	0	0
S7	13400	3000	3000	600
S8	13400	2700	3300	600
S9	13400	2000	4000	600

For each of these datasets, we calculated the various measures of interestingness we exposed earlier. Results are shown in Table 5. Expected frequencies are calculated assuming the independence of X and Y. Note that expected frequencies coincides with observed frequencies for S6, though Chi square cannot be calculated. We have put in bold the results that indicate a positive dependency between X and Y. We also highlighted the lines for S3 and S4, representing our data from the Logic-ITA and, in a lighter shade, S8 and S9, which have the same characteristics but with a multiplying factor of 10.

Table 5. Measures for all contingency tables.

	sup	confXY confYX	lift	Corr	convXY convYX	Chi	cos
S1	0.67	0.90	**3.31**	**0.87**	**7.98** **7.98**	**1522.88**	**0.91**
S2	0.67	0.82 0.82	1.00	-0.02	0.98 0.98	0.53	**0.82**
S3	0.67	0.83 0.82	1.00	-0.01	0.98 0.99	0.44	**0.82**
S4	0.67	0.87 0.77	1.00	0	1.00 1.00	0,00	**0.82**
S5	0.67	1.00 1.00	1.49	**1**	- -	**2000**	**1**
S6	1.00	1.00 1.00	1.00	-	- -	-	**1**
S7	0.67	0.82 0.82	1.00	-0.02	0.98 0.98	**5.29**	**0.82**
S8	0.67	0.83 0.80	1.00	-0.01	0.98 0.99	**4.37**	**0.82**
S9	0.67	0.87 0.77	1.00	0	1.00 1.00	0.01	**0.82**

We now discuss the results. First, let us consider the lift. One notices that, when the number X and Y increase in Table 4 and consequently P(X) and P(Y) increase, mechanically the lift decreases. As an illustration of this phenomenon, let us consider that a person is characterized by things she does everyday. Suppose X is 'seeing the Eiffel tower' and Y is 'taking the subway'. If association rules are mined considering the Parisians, then the lift of X→Y is likely to be low because a high proportion of Parisians both see the Eiffel tower everyday and take the subway everyday. However if association rules are mined taking the whole French population, the lift is likely to be high because only 20% of the French are Parisians, hence both P(X) and P(Y) cannot be greater then 0.20. The ranking for the lift given in (Tan and al.) is rather poor for their sets E1, E2 and E3, the closest matches with our data. They give strong symmetric association rules and both P(X) and P(Y) are high.

Let us now consider the correlation. Note that P(X) and P(Y) are positive numbers smaller than 1, hence their product is smaller than P(X) and P(Y). If P(X, Y) is significantly smaller than P(X) and P(Y), the difference between the product P(X)P(Y) and P(X, Y) is very small, and, as a result, correlation is around 0. This is exactly what happens with our data, and this fact leads to a strong difference with [7]'s E1, E2 and E3 sets, where the correlation was highly ranked: except for S1 and S5, our correlation results are around 0 for our sets with strong association rules.

Another feature of our data is that 1-P(X), 1-P(Y) and 1-conf(X→Y) are similar, hence conviction values remain around 1.

It is well known (see S7 to S9) that Chi-square is not invariant under the row-column scaling property, as opposed to all the other measures which yielded the same results as for S2 to S4. Chi-square rate X and Y as independent for S2 and S3, but rate them as dependent in S7 and S8. As the numbers increases, the Chi-square finds increasing dependency between the variables. This leads us to explore the calculation of Chi-square on a larger population, cumulating 4 years of data.

Finally cosine is the only measure that always rate X and Y as correlated. This is due to the fact that cosine calculation is independent of n, the size of the population, and considers only the number of transactions where both X and Y occur, as well as the number of transactions where X occur and Y occur.

3.3 Cumulating Data over 3 years and Chi-square.

We have mined association rules for four consecutive years and obtained stable results: the same symmetric rules with a support bigger than 60% came up. What would happen if we merge the data of these 4 years and mine the association rules on the merged data? Roughly, we would obtain contingency tables similar to S3 and S4 but with bigger figures: each figure is multiplied by 4. Because proportions do not change, such a table gives the same association rules, with same support, lift, correlation, conviction and cosine for S3 and S4. The difference is that the Chi-square increases. As illustrated with S7, S8 and S9 Chi-square is not invariant under the row-column scaling property. Due to a change in the curriculum, we have not been able to mine association rules over more years. However one can make the following projection: with a similar trend over a few more years, one would obtain set similar to S8 and S9. Chi-square would rate X and Y as correlated when X and Y are symmetric enough as for S3 and S8.

3.4 Contrast rules

In [10], contrast rules have been put forward to discover interesting rules that do not have necessarily a strong support. One aspect of contrast rules is to define a neighborhood to which the base rule is compared. We overtake this idea and consider the neighborhood {not X → Y, X → not Y, not X → not Y} assuming that X → Y is a symmetric rule with strong support and strong confidence. Taking the set S3, we get:

sup(not X→Y) = 0.17. *sup*(X→not Y) = 0.17. *sup*(not X→not Y) = 0.03
conf (not X→Y) = 0.85 *conf* (X→not Y) = 0.17. *conf* (not X→not Y)=0.15

These rules give complementary information allowing to better judge on the dependency of X and Y. They tell us that from the attempted solutions not containing mistake X, 85% of them contain mistake Y, while from the attempted solutions containing mistake X only 15% do not contain mistake Y. Furthermore, only 3% of the attempted solutions contain neither mistake X nor mistake Y. The neighborhood. {not Y → X, Y → not X, not Y → not X} behaves similarly.

3.5 Pedagogical use of the rules

We have shown in earlier papers how the patterns extracted were used for improving teaching [9, 11, 13]. Note that since our goal was to improve the course as much as possible, our experiment did not test the sole impact of using the association rules but the impact of all other patterns found in the data. After we first extracted association rules from 2002 and 2001 data, we used these rules to redesign the course and provide more adaptive teaching. One finding was that mistakes related to the structure of the formal proof (as opposed to, for instance, the use and applicability of a logic rule) were associated together. This led us to realise that the very concept of formal proofs was causing problems and that some concepts such as the difference between the two types of logical rules, the deduction rules and the equivalence rules, might not be clear enough. In 2003, that portion of the course was redesigned to take this problem into account and the role of each part of the proof was emphasized. After the end of the semester, mining for mistakes associations was conducted again. Surprisingly, results did not change much (a slight decrease in support and confidence levels in 2003 followed by a slight increase in 2004). However, marks in the final exam questions related to formal proofs continued increasing. We concluded that making mistakes, especially while using a training tool, is simply part of the learning process and this interpretation was supported by the fact that the number of completed exercises per student increased in 2003 and 2004 [9].

4 Conclusion

In this paper we investigated the interestingness of the association rules found in the data from the Logic-ITA, an intelligent tutoring system for propositional logic. We used this data mining technique to look for mistakes often made together while solving an exercise, and found strong rules associating three specific mistakes.

Taking an inquisitive look at our data, it turns out that they have quite a special shape. Firstly, they give strong symmetric association rules. Strong means that both support and confidence are high. Symmetric means that both $X \rightarrow Y$ and $Y \rightarrow X$ are rules. Secondly, $P(X)$ and $P(Y)$, the proportion of exercises where mistake X was made and the proportion of exercises where mistake Y was made respectively, is significantly higher than $P(X, Y)$, the proportion of exercises where both mistakes were made. A consequence is that many interestingness measures such as lift, correlation, conviction or even Chi-square to a certain extent rate X and Y as non-correlated. However cosine, which is independent of the proportions, rate X and Y as positively correlated. Further we observe that mining associations on data cumulated over several years could lead to a positive correlation with the Chi-square test. Finally contrast rules give interesting complementary information: rules not containing any mistake or making only one mistake are very weak. So, while a number of measures may have led us to discard our association rules, other measures indicate the opposite. Additionally, the use of these rules to change parts of our course seemed to contribute to better learning.

This really indicates that the notion of interestingness is very sensitive to the context. Since Education data often has relatively small number of instances, measures based on statistical correlation may not be relevant for this domain. Our experience tends to say so. We think that it is highly dependent on the way the rules will be used. In an educational context, is it really important to be certain of the probabilistic dependency of, say, mistakes? When the rule X→Y is found, the pragmatically-oriented teacher will first look at the support: in our case, it showed that over 60% of the exercises contained at least three different mistakes. This is a good reason to ponder. The analysis of whether these 3 mistakes are statistically correlated is in fact not necessarily relevant to the remedial actions the teacher will take and may even be better judged by the teacher. As a future work we would like to investigate how subjective interestingness measures would work on our data.

References

1. Merceron, A., Yacef, K., *Mining Student Data Captured from a Web-Based Tutoring Tool: Initial Exploration and Results.* Journal of Interactive Learning Research (JILR), 2004. **15**(4): p. 319-346.
2. Wang, F., *On using Data Mining for browsing log analysis in learning environments*, in *Data Mining in E-Learning. Series: Advances in Management Information,*, Romero, C., Ventura, S., Editors. 2006, WITpress. p. 57-75.
3. Wang, F.-H., Shao, H.-M., *Effective personalized recommendation based on time-framed navigation clustering and association mining.* Expert Systems with Applications, 2004. **27**(3): p. 365-377.
4. Minaei-Bidgoli, B., Kashy, D.A., Kortemeyer, G., Punch, W.F. *Predicting student performance: an application of data mining methods with the educational web-based system LON-CAPA*. *ASEE/IEEE Frontiers in Education Conference*. 2003. Boulder, CO.
5. Lu, J. *Personalized e-learning material recommender system*. *International conference on information technology for application (ICITA'04)*. 2004. China, 374–379.
6. Romero, C., Ventura, S., de Castro, C., Hall, W., Ng, M.H., *Using Genetic Algorithms for Data Mining in Web-based Educational Hypermedia Systems*, in *Adaptive Systems for Web-based Education*. 2002: Malaga, Spain.
7. Tan, P.N., Kumar, V., Srivastava, J. *Selecting the Right Interestingness Measure for Association Patterns*. *8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2001. San francisco, USA, 67-76.
8. Brijs, T., Vanhoof, K., Wets, G., *Defining interestingness for association rules.* International journal of information theories and applications, 2003. **10**(4): p. 370-376.
9. Yacef, K., *The Logic-ITA in the classroom: a medium scale experiment.* International Journal on Artificial Intelligence in Education, 2005. **15**: p. 41-60.
10. Minaei-Bidgoli, B., T., P-N., Punch, W.F. *Mining Interesting Contrast Rules for a Web-based Educational System*. *International Conference on Machine Learning Applications (ICMLA 2004)*. 2004. Louisville, KY, USA,
11. Merceron, A., Yacef, K. *A Web-based Tutoring Tool with Mining Facilities to Improve Learning and Teaching*. *11th International Conference on Artificial Intelligence in Education.* 2003. Sydney: IOS Press, 201-208.
12. Agrawal, R., Srikant, R. *Fast Algorithms for Mining Association Rules*. *VLDB*. 1994. Santiago, Chile,
13. Merceron, A., Yacef, K. *Educational Data Mining: a Case Study*. *Artificial Intelligence in Education (AIED2005)*. 2005. Amsterdam, The Netherlands: IOS Press, 467-474.