Visualizing Uncertainty in the Prediction of Academic Risk

Xavier Ochoa
Escuela Superior Politécnica del Litoral
Vía Perimetral, Km. 30.5
Guayaquil, Ecuador
xavier@cti.espol.edu.ec

ABSTRACT
This work proposes a generic visual representation to help relevant decision-makers to effectively address the inherent uncertainty present in the prediction of academic risk based on historical data. The three main sources of uncertainty in this type of prediction are visualized: the model predictive power, the data consistency and the case completeness of the historic dataset. To demonstrate the proposed visualization technique, it is instantiated in a real-world scenario where the risk to fail at least one course in an academic semester is predicted and presented in a student-counseling system. This work also proposes how this visualization technique can be evaluated and applied to other Visual Learning Analytics tools.

Categories and Subject Descriptors
K.3.1 [**Computing Milieux**]: Computers and Education-Computer Uses in Education

Keywords
Visual Learning Analytics, Uncertainty Visualization, Academic Risk

1. INTRODUCTION
The main goal of the Learning Analytics field is to provide relevant information to the actors of the learning process (students, instructors and administrators) to help them take better learning-related decisions. A considerable amount of research effort [6] has been invested in find ways to analyze the large amount of traces that are a by-product of the learning process to convert it into that relevant information. An equal important, but lesser researched, area of Learning Analytics explores the best ways in which that relevant information is presented to the final user to maximize its usefulness for decision-making. This second area is often called "Visual Learning Analytics" given that it is very related to the field of Visual Analytics, that focuses on "analytical reasoning facilitated by interactive visual interfaces" [14]. Visual Analytics differentiates from simple data visualization because its purpose is not only presenting the information resulting from a predefined analysis process, but empowering the decision-maker to control the analysis process and interact with the multiple dimensions that the resulting information could have to gain a deep understanding of the implications that those results have in the decision at hand.

Currently, there are very few early examples of Visual Learning Analytics, in contrast to simple visualization of Learning Analytics results: Lemo [7] is a system that use interactive visualization to help instructors understand the activity logs of LMSs. The end-user is capable of exploring the dataset through selecting and filtering the desired information in a variety of visualization options. Gomez et al. [8] also create a system to explore in deeper detail the academic and non-academic data stored in the LMS system through the use of interactive visualizations.

One virtually unexplored avenue of Visual Learning Analytics is how to make explicit the uncertainty that is inherent in any analysis process in a way in which is meaningful for the decision-maker. Moreover, if possible, the decision-maker should also be able to manipulate the analysis process to adjust the uncertainty to a level where he or she finds appropriate. This kind of techniques to present and manage the uncertainty are common in more mature fields such as meteorology (e.g. hurricane path prediction uncertainty [13]), medicine (e.g. uncertainty in the effect of medical interventions [10]) and economy (e.g. uncertainty in the prediction of future growth [13]). There exists, however, some examples of the visualization of uncertainty in Open Learner Models [5] that could be consider a precursor in the field of Visual Learning Analytics.

This work will focus on how Visual Learning Analytics techniques could be used to visualize and control the inherent uncertainty in the prediction of academic risk. The organization of this paper is as follows: Section 2 explores how academic risk is usually obtained and which are the main sources of uncertainty in this type of analysis. Section 3 discusses how the prediction value, together with the main uncertainty values should be visualized. Section 4 presents a case-study where the visualization techniques are instantiated to help counselors give advice about the risk to fail a semester to individual students. Finally, the paper finishes with conclusions about the work and guides for further work to evaluate the technique and how to adapt it to other Visual Learning Analytics tools.

2. PREDICTING ACADEMIC RISK

In the context of this work, the term "academic risk" is defined as the probability of a student to reach an unfavorable outcome in their studies. This unfavorable outcome could be as benign as the failure to submit a homework or as costly as dropping-out of a program. As very little can be done once the unfavorable outcome has been already reached, especially for the more costly forms (e.g. failing a course or dropping-out), there is a strong incentive to being able to estimate the academic risk of the student, or what is equivalent, predict the probability that the student will, without intervention, reach the unfavorable outcome. Due to its importance, predicting different forms of academic risk has been one of the oldest forms of Learning Analytics [11].

There are several current examples of systems that seek to estimate different kinds of academic risks: Signals [1] is arguably the poster-boy of learning analytics systems to predict academic risk. Using historical and current information about the behavior of a student in a course, it is able to predict the probability that the student has of fail the course. Another, more simple approach is taken by StepUp! [12] that just compares the activity of a student with the activity of their peers and assigns a ranking value that could be seen as a fuzzy academic risk predictor. Finally, there are several modern drop-out risk predictors from which the work of Dekker et al. [4] could be considered a good representative. This system uses a classification tree trained over historical data in order to obtain rules to assess the risk of a student to dropping-out from a university program.

All of the mentioned systems used data collected from previous or current students to create a prediction model. This model could be built with statistical or data-mining methods. Once the model has been built, it is fed with the information from the student target of the prediction and an estimation of the academic risk is produced. This estimation is normally presented to the instructor, counselor or the student through some form of visualization technique.

In all of the steps of the above-mentioned process there are inherent uncertainties that are propagated and contribute to the uncertainty that is present in the estimated value of academic risk. The following subsection discusses the nature of these sources of uncertainty and their relative importance for the prediction.

2.1 Uncertainty Sources

To facilitate the analysis of the different sources of inherent uncertainty in the prediction of academic risk, they are classified in two group according to their origin: predictive model limitations and dataset limitations. The following subsections sub-classify these two groups into more concise and measurable uncertainty values.

2.1.1 Predictive Model Limitations

Perhaps the most obvious source of uncertainty introduced in any type of prediction is the one introduced by the imperfections of the predictive model. In general, predictive models are built to take as input a group of predictor variables and to produce a predicted value. Given that models are only an approximation and simplification of reality, it is expected that the predicted values differ, in different degrees, from the real values. A whole area of Statistics is devoted to measure the predictive power of different types of models. The best example of the measure of the predictive power is the R-squared statistic used to score regression models. This measurement establishes what percentage of the variance in the real values of the predicted quantity are explained by the model.Different models usually have different predictive power depending on the predictor variables used, the type of algorithm and the amount and quality of data used to build them. It is a common practice to evaluate different competing models and select the one with the best predictive power according to an appropriate scoring function.

2.1.2 Dataset Limitations

Given that most academic risk predictors are built based on historical or current data, the characteristics of the data and its limitations play a major role in the overall uncertainty of the predicted value of that risk. The work of Thomson et al. [15] established a detailed typology for the limitations of data that affect certainty in predictive models: accuracy, precision, completeness, consistency, lineage, currency, credibility, subjectivity and interrelatedness. All these types of limitations are usually defined at the dataset level and their effect in uncertainty is usually propagated into the final predictive power of the model that was built with that dataset.

Given the nature of academic datasets, the most important of these dimensions are consistency and subjectivity. Historical academic data, for example final grades of students, is generally accurate (there is a significant cost of registering a grade wrongly), precise (it has enough resolution to separate passing and failing students), complete (all students should have grades or at least a pass/fail at the end of a course), current (the grades are producing during the course or at least very close to the ending of the course) and credible (the academic institutions will have serious problems if their academic records are not credible). Also, academic records have no major problems with lineage (the grades are rarely processed after the instructor records them) and the records do not suffer from interrelatedness (instructors do not copy the grades from one student to another or among them). However, consistency of academic data could introduce uncertainty in the prediction of academic risk. As academic programs evolve, they also change: the courses offered could change, the grading rules could become more strict or more relaxed, different instructors will imprint their own characteristic in the courses, among other changes. Depending on the nature and magnitude of the changes, the academic records of a current student and one that studied ten years ago could not be comparable or, more dangerously for prediction models, could provide a false sense of similarity when in reality the values in those records are not measuring the same students characteristics. Another possible limitation of historical academic data is its subjectivity. Grades, scores and student evaluations are commonly assigned according to the criteria of the instructor. Even during the same course, students that did a similar level of work could receive different grades. While the effect of consistency errors in the overall prediction uncertainty could be limited by only considering comparable years of the academic program in the dataset, the uncertainty produced by the subjectivity could not be reduced if it is already present in the data.

Due to the fact that most academic risk predictors compare current students to previous similar students that were in a similar context, another type of data limitation plays a role in the overall uncertainty of the prediction: case completeness. For example, predictive model A estimates the academic risk of failing a course based on number of other courses taken at the same time and the GPA of the student; predictive model B estimates the academic risk of failing a course based on the number of courses taken at the same time, the GPA of the student, the fact that the student has an external job, if the student is married, the number of children the student has, the distance from his house to the university and the number of courses taken before the current one. Both models estimate the academic risk of failing the course as the percentage of similar students that have failed the course in the past. A hypothetical prediction power analysis shows that model B is less uncertain that model A. However, this prediction power is calculated for the general population, for some students model A could be less uncertain than model B. Lets suppose that student A is taking 3 other courses, has a GPA of 3.5, has an external job, is married, has 5 children, lives 100 km from the university and has taken just one course before the current one. Lets suppose too that this is a very unusual combination of values for the students of this specific course. If the model A is applied, only the number of other courses that the student is currently taking (3) and his or her GPA (3.5) are considered. These two values, by themselves, are not unusual, so it is probable that there will be several previous students that could be considered similar. The prediction of academic risk for the hypothetical student will be drawn from a large pool of previous experiences. If the model B is applied, due to the unusual values of the rest of variables, the model could only find one other student close enough to be considered similar in the dataset. In this situation, the prediction of academic risk for the student will be 100%, if the previous student failed the course or 0% if he or she passed. While, in general, model B has more predictive power than model A, for this particular student the approximate estimation of model A will be much more less uncertain than the one provided by model B, due to the lack of similar cases in the dataset. The prediction for "outlier" students, that is, students that have few similar students in the dataset, is less certain than the prediction for "mainstream" students, that has a large collection of similiar cases. Simple models have less similarity dimensions, and the number of possible cases is lower than in complex models with larger dimensions sets. The variety and quantity of cases in the dataset, that is the case completeness of the dataset, introduce a uncertainty factor that varies from student to student and depends on the complexity of the model.

3. VISUALIZING UNCERTAINTY

As mentioned in the introduction, the visualization of uncertainty is already an established feature in more mature fields. In Visual Learning Analytics, however there are still no thoroughly evaluated techniques. The most recommended path in this case will be to adapt uncertainty visualization techniques that are common and proved useful [3] in other fields to represent the predicted value, together with the different uncertainty produced by the sources described in the previous section: the model predictive power, the data consistency and the case completeness. The goal of the visualization of those values is to present the most information about the prediction in an interpretable and useful way. The following subsection proposes various techniques for each one of these elements in detail.

3.1 Predicted Risk Value

The value of the academic risk of a student, being just a scalar that can be expressed as an easily interpretable numeric value between 0 and 1 (as probability) or from 0% to 100% (as relative frequency) can be presented using a large variety of visualization techniques such as textual, progress arc, gauge or bullet graphs. Figure 1 shows an example of this type of visualizations. Attached to the visualization of the value, all of these types of visualization present the decision-maker with a pre-defined guide to assess the level of risk described depending on the magnitude of the value. In the case of textual and arch representations, the color of the text or the arch (e.g. green, yellow and red) or an additional iconic representation (e.g. traffic light) could be used to provide an indication of the severity of the risk. In the case of the gauge and bullet graphs, different ranges can be color-coded to also provide this information. Some previous implementations of visualization of academic risk, such as Signals [1], use only an iconic representation (the traffic light approach) to represent the predicted value. Representing only the range in which the value is, instead of the actual value is used to account for the uncertainty of the prediction. However, in most cases, those ranges are crisp, meaning that a single unit change in the predictive value can cause the color to change, defeating the purpose of presenting only ranges in the first place. For example, a student with a risk of 0.50 will be coded with green, while a student with a risk of 0.51 will be coded yellow. With just the iconic representation, there is no way for the decision-maker to establish if the students is closer to green or to red. Moreover, the span of the ranges (what values are considered to be green, yellow or red) is often also unknown to the decision-maker. Using only the iconic representation is discouraged given that this work present other ways to deal with the inherent uncertainty in the prediction.

3.2 Model Predictive Power

Similarly to the predicted risk value, the model predictive power is also an scalar magnitude. Contrary to risk probability, the meaning of the output of the different model-scoring techniques (such as R-squared, BIC, AIC, Brier score, etc.) are far from being easy to interpret by non-statisticians. To effectively communicate the predictive power of the model, or what is the same, the level of uncertainty that a given model will introduce in the prediction, the expert analyst in charge of the academic risk prediction should define a set of iconic representations (e.g. traffic lights, happy-sad faces, plus signs, etc.) to correspond with different values of predictive power. Given that usually there are no model with "bad" power (otherwise it will not be used in the analysis), it is recommended that a plus signs textual representation ("+"

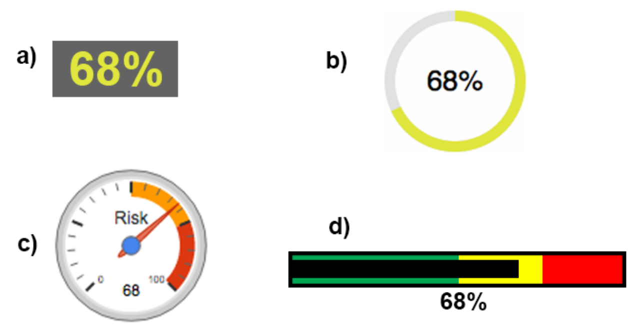

Figure 1: Predicted value visualization: a) textual representation, b) progress arc graph, c) gauge graph and d) bullet graph

for lower scoring models, "++" for medium scoring models and "+++" for the best scoring models) is used to represent different levels of power. The words "Good", "Very Good" and "Excellent" could be complement or replace this visualization. An example of this visualization could be seen in Figure 2.

It is important to note that this visualization is only necessary when the decision-maker can select between different models or the system chooses the model based on the available data. If the predictive risk is using a single model, the value of presenting this extra information is diminished.

Figure 2: Model predictive power visualization

3.3 Data Consistency

The representation of uncertainty introduce by the data inconsistency is challenging given that there is no way to precisely measure it. In the case of academic datasets, the consistency is related to the changes in different aspects of the study program or course over time. It is expected that the closer in time the historic data is, the greater the level of consistency and the lower the level of uncertainty. If there exists a record of major changes in the academic program (course changes, evaluation policies changes, etc) or the courses (syllabus change, pre-requisites changes, instructor change, etc), they can be plotted in a timeline that span over the whole data range of the historical data. In this way, instructors and counselors that are familiar with the history of the program or course could recognize the changes and adjust their perception of the uncertainty introduced in the prediction, while students or users not familiar with the history of the program or course could just count the number of changes to form their own estimation of the uncertainty in the prediction, although less precise than the ones with previous knowledge. An example of this type of visualization can be seen in Figure 3

3.4 Case Completeness

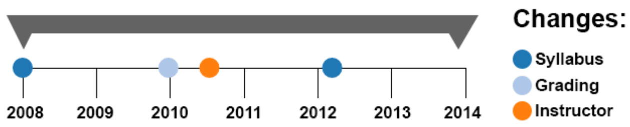

Figure 3: Data consistency visualization for a course historical data

In most predictive models is easy to obtain a measure of how many "similar" elements are considered at the moment of obtaining the predictive value for a given element. In the case of academic data, the case completeness could be measured as the number of records that are directly used to calculate the academic risk of a given student. This number could go from 0 to the total number of records in the dataset. A low value is an indication of a high uncertainty in the predicted value. Higher values, usually larger than 30, are enough to discount the number of cases as a source of uncertainty. The recommended visualization technique for this value is an iconic representation with icons that represent alert states at different number of different cases pre-defined by the expert behind the analysis (e.g. a red stop sign for values between 0 and 5, a yellow exclamation mark for values between 5 and 30 and a green check for values higher than 30). Together with the icon, a textual representation of the number of cases could be included to improve understandability (e.g. This prediction is based only on 3 previous cases). Figure 4 presents an example of this visualization.

Figure 4: Case completeness visualization based on iconic representation

3.5 Interaction

The visualization described in the previous sub-section could help the decision-maker to better understand the inherent uncertainty of the risk value prediction. However, if the decision-maker is not confortable with the uncertainty of the prediction the only course of action is to discard the prediction. As mentioned in Section 2, the uncertainty of the prediction depends on several factors such as the model used, the length of historical data used and the number of similar cases used by the model to generate the prediction for a given student. The trade-off between these parameters is decided by the expert in charge of the prediction. Usually the model selected will be the one with greatest predictive power and the range of historical data will be selected to maximize this number. This selection is bound to be sub-optimal for some students, specially those with special cases. The use of interactive visualization transfer the control of the analysis parameters to the decision-maker. He or she could adjust them in order to reach the lowest level of uncertainty possible for a given student and the domain knowledge that the decision-maker has about the academic program or course.

Very simple interactive controls could be added to the visualization in order to control the main parameters affecting uncertainty factors. Each time a new value is selected on those controls, the uncertainty visualizations should be updated enabling the exploration of the uncertainty space by the decision-maker. To control the uncertainty resulting from the predictive power of the model, the decision-maker could be presented with a set of widgets where the model algorithm or parameters could be selected. To control the uncertainty resulting from the lack of consistency in the historical records, the timeline where this information is presented could be complemented with a selection bar to select subsets of the whole time period. The uncertainty produced by the lack of similar cases could not be affected directly, but it will change its response to the changes in the model used and the selected time period.

4. CASE-STUDY: RISK TO FAIL

To illustrate the ideas presented in the previous sections, they will be applied to a real-world academic risk prediction application. This application is part of a larger counseling system used regularly by professors and students at a mid-size university in Ecuador. The goal of this application is to determine the academic risk of failing at least in the next semester based on the planned course selection and study load. To produce this prediction the application uses a variety of models that cluster the student and the planned semester with similar students and semesters in the historical dataset. The models calculate the risk based on the previous frequency of similar students in similar semesters that failed at least one course. The counselor could interact with the visual analysis by selecting the courses that the student will take the next semester, the type of clustering that is applied to select similar students and semesters and the time period used to obtain similar cases. The counselor is presented with a prediction of the probability of the student failing the course and the visualization of the uncertainty produced by the model, the data consistency and case completeness. The counselor use the information received to recommend the student to take more or less study load in the coming semester.

4.1 Dataset

The dataset used for this application was built based on a Computer Science program at the target university. All the courses taken by CS students each semester and the grades obtained in those courses were stored since the first semester of 1978 to the second semester 2013. The courses that have changed name were grouped together according to the transition rules during those changes. A total of 30.929 semesters were taken by 2.480 different students.

4.2 Predictions Models

A multi-level clustering approach was used to build different models to find similar students and calculate the academic risk value. Two main variables controlled the generation of the different models: the student similarity and the semester similarity. The students were clustered at three levels: No clustering at all (all the students were considered similar), clustering based on GPA values (five clusters based on range) and clustering based on similarity of grades in the different courses (the Fuzzy C-means (FCM) algorithm [2] was used to create 10 clusters). The semesters were clustered at five levels (all using Fuzzy C-means): Level 1, based on the total load of the courses calculated from their difficulty [9]; Level 2, based on the typology of courses; Level 4, based on the grades that the students obtain in the courses [9]; Level 4, based on the knowledge area of the courses; Level 5, based on the actual name of the courses. The intersection of the level student and semester clustering defines a predictive model. For example, a model is created by finding similar students based on their GPA taking similar semesters based on the difficulty of courses taken (Level 2). The predictive power of the models was obtained computing the Brier score [16] of the forecast made for the last semester (2013-2) with the models built from the data from all the previous semesters.

4.3 Visualizing the Prediction

Figure 5 presents the interactive visualization created for the case-study academic risk prediction application. All the elements discussed in Section 3 are present. The predicted value is presented using a bullet graph with a 0%-100% scale, a yellow interval between 50% and 75% and a red interval between 75% and 100%. The model prediction power is shown with an iconic representation of one, two or three plus signs, together with a textual description. The data consistency is represented with an interactive timeline indicating the major events that changed the Computer Science program during the analyzed period. The case completeness of the dataset for the target student is presented using an iconic representation of group of different amounts of people related to a color (one individual in red to indicate a large amount of uncertainty, few people in yellow to represent middle values and a green crowd to represent low values. Finally, selection boxes are presented to the decision-maker to define the levels of clustering (for students and semesters) that determine the model that will be used for the prediction. All of these visualizations and controls are implemented with easy-to-use D3 Javascript visualization library [1].

5. CONCLUSIONS AND FURTHER WORK

Visualizing the uncertainty in the prediction of academic risk, specially in an interactive way, has the potential to improve the usefulness of this type of systems. Even simple techniques are able to present to the decision-maker with the information needed to assess the uncertainty of the prediction for different selections of model and historical training data. With an interactive visualization the decision-maker, with their domain-expertise knowledge, becomes a co-designer of the analytic process, instead of a simple user of the results of the analysis. Implementing this visualization in real-world scenarios is simple given that the sources of uncertainty are well understood and could be measured or estimated.

The main task to be completed in this research is the real-world evaluation of the visualization to establish the answers to two main questions: 1) Is the visualization contributing to the understanding of the inherent uncertainty of the prediction of academic risk? and 2) Is the knowledge about the uncertainty helping the decision-maker to make better

[1]D3.JS visualization library - http://d3js.org

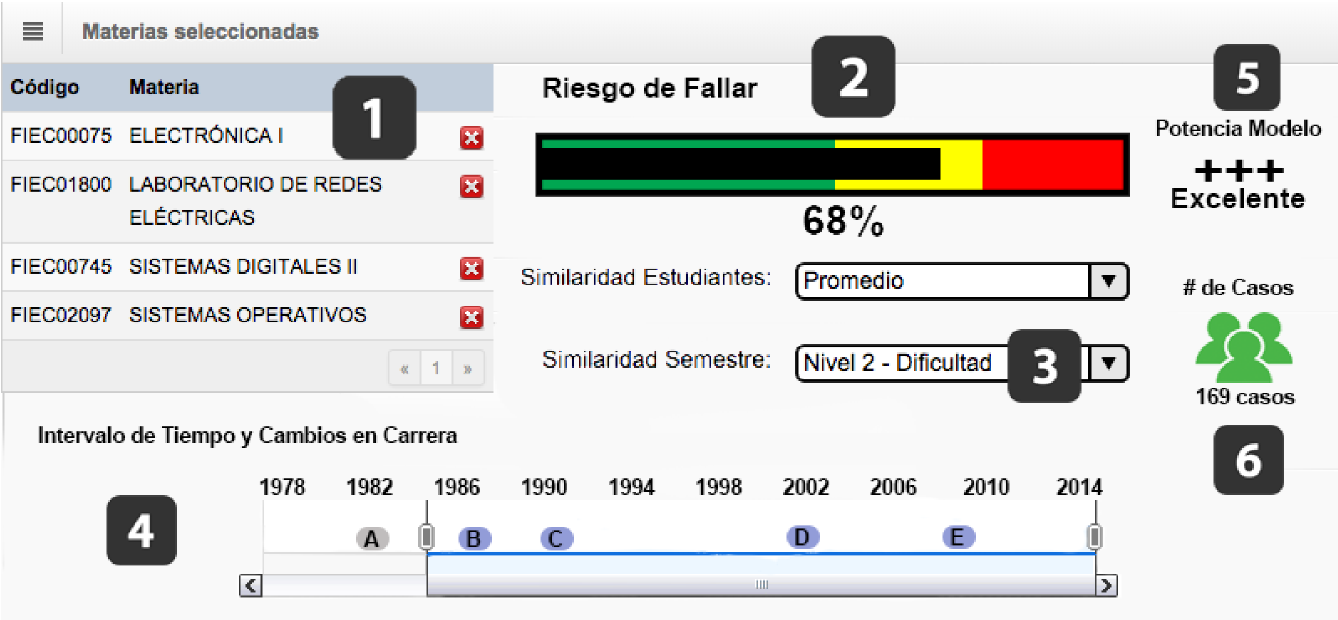

Figure 5: Example of visualization integrated in the counseling system: 1) Course selector, 2) Predicted academic risk value visualization, 3) Model selector, 4) Time period selector and consistency visualization, 5) Model predictive power visualization and 6) Case completeness visualization

decisions or to provide better advice? To answer these questions, the tool presented in the case study will be used in two experimental groups of counselors. One group will see the prediction and the uncertainty visualization. The second group will see only the prediction visualization. A third control group will continue to use the counseling system without the academic risk predictor application. The average failure rate for each counselor will be recorded at the end of the semester and compared with the failure rate between experimental and control group and also with the failure rate from previous semesters. Surveys will be conducted just after the counseling sessions in order to establish the level of understanding of the uncertainty in the prediction.

Finally, the ideas presented in this paper could be adapted to other types of Visual Learning Analytics tools, especially those focused on prediction and forecasting. The methodology followed in this paper could be a general framework for these adaptations: 1) exploring the main sources of uncertainty in the analysis, 2) establishing methods to measure or estimate the uncertainty contribution of those sources, 3) using existing visualization techniques to present the uncertainty values in a way that will be easy to interpret by the end-user, 4) provide control to the end-user through interactive visualizations to change the parameters to the models and to select the desired data and 5) evaluate the impact of the visualization. Visualizing the uncertainty is a way to empower the user of Visual Learning Analytics tools, stressing that automatic analysis could support, but not replace, human judgment.

6. ACKNOWLEDGMENTS

The author wants to acknowledge the contribution of Secretaría Nacional de Educación Superior, Ciencia y Tecnología (SENESCYT) in Ecuador and the Fonds Wetenschappelijk Onderzoek - Vlaanderen (FWO) in Belgium through the funding of the “Managing Uncertainty in Visual Analytics” project, to which this work belongs.

7. REFERENCES

[1] K. E. Arnold and M. D. Pistilli. Course signals at purdue: Using learning analytics to increase student success. In *Proceedings of the 2nd International Conference on Learning Analytics and Knowledge*, pages 267–270. ACM, 2012.

[2] J. C. Bezdek. *Pattern recognition with fuzzy objective function algorithms*. Kluwer Academic Publishers, 1981.

[3] S. Deitrick and R. Edsall. *The influence of uncertainty visualization on decision making: An empirical evaluation*. Springer, 2006.

[4] G. W. Dekker, M. Pechenizkiy, and J. M. Vleeshouwers. Predicting students drop out: A case study. In *International Conference on Educational Data Mining (EDM)*. ERIC, 2009.

[5] C. Demmans-Epp, S. Bull, and M. Johnson. Visualising uncertainty for open learner model users. In *CEUR Proceedings associated with UMAP 2014*, 2014.

[6] R. Ferguson. Learning analytics: drivers, developments and challenges. *International Journal of Technology Enhanced Learning*, 4(5):304–317, 2012.

[7] A. Fortenbacher, L. Beuster, M. Elkina, L. Kappe, A. Merceron, A. Pursian, S. Schwarzrock, and B. Wenzlaff. Lemo: A learning analytics application focussing on user path analysis and interactive visualization. In *Intelligent Data Acquisition and Advanced Computing Systems (IDAACS), 2013 IEEE 7th International Conference on*, pages 748 – 753,

2013.

[8] D. Gomez, C. Suarez, R. Theron, and F. Garcia. *Advances in Learning Processes*, chapter Visual Analytics to Support E-learning. InTech, 2010.

[9] G. Méndez, X. Ochoa, and K. Chiluiza. Techniques for data-driven curriculum analysis. In *Proceedings of the Fourth International Conference on Learning Analytics And Knowledge*, LAK '14, pages 148–157, New York, NY, USA, 2014. ACM.

[10] M. C. Politi, P. K. Han, and N. F. Col. Communicating the uncertainty of harms and benefits of medical interventions. *Medical Decision Making*, 27(5):681–695, 2007.

[11] C. Rampell. Colleges mine data to predict dropouts. *The chronicle of higher education*, 54(38):A1, 2008.

[12] J. L. Santos, K. Verbert, S. Govaerts, and E. Duval. Addressing learner issues with stepup!: an evaluation. In *Proceedings of the Third International Conference on Learning Analytics and Knowledge*, pages 14–22. ACM, 2013.

[13] D. Spiegelhalter, M. Pearson, and I. Short. Visualizing uncertainty about the future. *Science*, 333(6048):1393–1400, 2011.

[14] J. Thomas and P. C. Wong. Visual analytics. *IEEE Computer Graphics and Applications*, 24(5):0020–21, 2004.

[15] J. Thomson, E. Hetzler, A. MacEachren, M. Gahegan, and M. Pavel. A typology for visualizing uncertainty. In *Electronic Imaging 2005*, pages 146–157. International Society for Optics and Photonics, 2005.

[16] D. S. Wilks. *Statistical methods in the atmospheric sciences*, volume 100. Academic press, 2011.