Uncertainty Representation in Visualizations of Learning Analytics for Learners: Current Approaches and Opportunities

Demmans Epp, Carrie; Bull, Susan

DOI:
10.1109/TLT.2015.2411604

Document Version
Peer reviewed version

Citation for published version (Harvard):
Demmans Epp, C & Bull, S 2015, 'Uncertainty Representation in Visualizations of Learning Analytics for Learners: Current Approaches and Opportunities', *IEEE Transactions on Learning Technologies*. https://doi.org/10.1109/TLT.2015.2411604

Uncertainty Representation in Visualizations of Learning Analytics for Learners: Current Approaches and Opportunities

Carrie Demmans Epp and Susan Bull

Abstract— Adding uncertainty information to visualizations is becoming increasingly common across domains since its addition helps ensure that informed decisions are made. This work has shown the difficulty that is inherent to representing uncertainty. Moreover, the representation of uncertainty has yet to be thoroughly explored in educational domains even though visualizations are often used in educational reporting. We analyzed 50 uncertainty-augmented visualizations from various disciplines to map out how uncertainty has been represented. We then analyzed 106 visualizations from educational reporting systems where the learner can see the visualization; these visualizations provide learners with information about several factors including their knowledge, performance, and abilities. This analysis mapped the design space that has been employed to communicate a learner's abilities, knowledge, and interests. It also revealed several opportunities for the inclusion of uncertainty information within visualizations of educational data. We describe how uncertainty information can be added to visualizations of educational data and illustrate these opportunities by augmenting several of the types of visualizations that are found in existing learning analytics reports. The definition of this design space, based on a survey of the literature, will enable the systematic exploration of how different design decisions affect learner trust, understanding, and decision making.

Index Terms—open learner models, learning dashboards, uncertainty, educational reporting, learning analytics, visual analytics

1 Introduction

This paper surveys work on visualizing uncertainty across various domains as well as the visualizations that are currently used in the educational reporting that is available to learners. This survey is the basis for the definition of a design space that will allow for the systematic study of how uncertainty information might be represented for student-users of educational reporting.

People need access to relevant information and information about the accuracy and reliability of that information to make informed decisions and understand the recommendations that an adaptive system makes. This article, therefore, aims to sensitize designers to the opportunities and challenges that surround representing model or assessment uncertainty in the visualizations that learners can see. Once sensitized to this design space, a better understanding of how learners use this information to support their decision making will be possible.

The presentation of information about the quality of the data over which people are reasoning has been gaining attention across several disciplines, including visual analytics, oceanography, meteorology, medicine, fluid flow, geography, cartography [1], and educational reporting [2]. Some of this work has focused on exploring domain-specific applications of uncertainty representation while other work has focused on the representation of uncertainty in a more discipline-independent manner. Many of the existing representations of information and its accompanying uncertainty have been visual, and while the objective of information visualization is to support user understanding, it has been found that people can struggle with understanding visual representations of uncertainty even when they have received training in their semantics, interpretation, and use [3].

Thus far, most visualizations of uncertainty in information have been explored in disciplines where users are experienced in reasoning over probabilistic or uncertain information such as in mapping [4], statistics [5]–[7], or the military and intelligence services [8], [9]. The representation of uncertainty for users who may be less comfortable or experienced with reasoning over uncertainty has only begun to be explored [2], [10]–[12].

Education is a domain where considerable work is being done in the area of visualizing student performance or knowledge in order to inform instructors, learners, the parents of learners, and other decision makers [2], [13]–[16]. The visualized information is usually based on an analysis of learner data. This analysis may involve the cognitive modelling of a learner and his/her knowledge or the statistical modelling of the learner's performance on assessments. Regardless of the approach that is used to analyze learner data, uncertainty is not typically present in the visualized data and those who are meant to interpret it are not typically trained in its use. Excluding in-

- *C. Demmans Epp is with the Department of Computer Science at the University of Toronto, 3302-10 King's College Road, Toronto, ON, CANADA M5S 3G4 E-mail:carrie@taglab.ca.*
- *S. Bull is with the School of Electronic, Electrical and Systems Engineering at the University of Birmingham, Birmingham, B15 2TT, UNITED KINGDOM E-mail: s.bull@bham.ac.uk.*

******Please provide a complete mailing address for each author, as this is the address the 10 complimentary reprints of your paper will be sent***

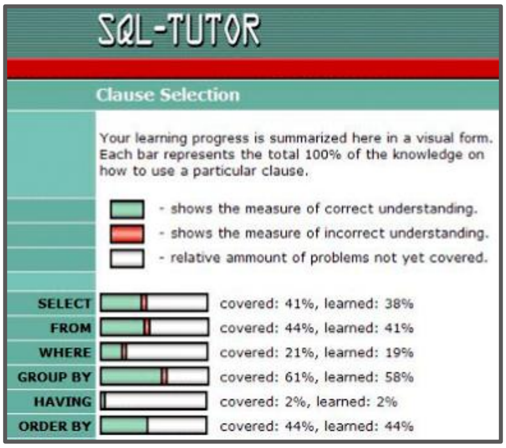

Figure 1. SQL-Tutor open learner model [21] represents student activity and knowledge level within a database programming tutor. It shows learners what they have understood and misunderstood as a proportion of their activities within the system. Source: [19]; used with permission.

formation about the reliability and consistency of data or any other form of uncertainty can negatively affect many of the tasks that visualizations of educational data aim to support. This includes the learner's self-regulation, monitoring, and decision making tasks [15], [17]–[23].

We focus this survey on a particular sub-area of information visualization within education, where learners are given access to visual representations of their knowledge, activities, abilities, assessment outcomes, or any other analytics that have been performed within their learning context. These visualizations can be thought of as graphical report cards and fall into several categories including learning dashboards [18], which are typically based on statistical models of learner performance, and open learner models [24] that require an underlying model of the learner's knowledge, abilities, beliefs, or attitudes. Both open learner models and learning dashboards report on the activities that have been performed by a learner or assessments that have been performed on a learner's activities in a technology enhanced learning environment (Figure 1, Figure 2). The data that is used to support these visualizations often comes from automated sources but could be human generated. This would be the case if the results of classroom tests, teacher assessments, peer assessments, or even self assessments were entered into the system. For the purposes of this paper, the distinction between the different types of visualizations that learners can access based on their learning activities is not relevant, what is important is that learners can see and understand reports on their activities and any inferences that have been made based on their activities.

Upon initially encountering visualizations of learning analytics, it may not be entirely obvious how they could contain uncertainty since the visualizations are largely based on objectively recorded learner activities or other educational assessments. Even if we ignore the potential for instructors or others who assess learner activities to make mistakes, be inconsistent, or be uncertain about the grade that they are assigning, we must acknowledge the uncertainty that is inherent to any modelling and analysis process because of the inevitable loss of detail or because

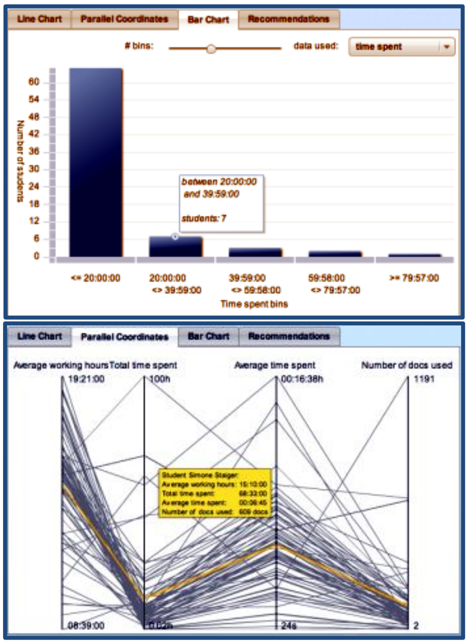

Figure 2. Student Activity Monitor (SAM). The bar chart (top) and parallel coordinate (bottom) visualizations of student activity. Source: [12]; used with permission.

of inconsistencies in the learner's behaviours or performance. It is also possible that drawing information from multiple sources, such as a teacher and logs of student activities within a technology-enhanced learning environment, may show inconsistencies because of the lens or pedagogical framework under which they operate. For example, one system may interpret high activity levels to mean increased knowledge where another system interprets those same patterns of high activity to mean a lack of knowledge because the patterns are associated with guessing or gaming the system rather than legitimate learning activities. Moreover, the visualization process can itself introduce uncertainty [25].

None of this is meant to imply that technology-enhanced learning environments and the reports that they provide are less accurate than those that are prepared by people. However, the uncertainty that is present in these environments and that results from the analysis and presentation of educational data can be accounted for and communicated to users, which could enable more informed decisions.

At present, most of the visualizations that provide information to learners about aspects of their learning do so without fully addressing the uncertainty that is part of the underlying information or the inferences that are made from this information. Rather, they tend to hide uncertainty by employing design techniques, such as the use of performance categories (e.g., low, medium, or high). This may lead to misinformed decisions and affect how learners feel about their efforts, their work products, the technology enhanced learning environment, their learning, and the visualization itself. We argue that the appropriate and thoughtful presentation of uncertainty within educational reporting can help learners to better understand their strengths and weaknesses so that they can make decisions that better support their learning goals.

To this end, we describe a design space that is based on previous work in uncertainty visualization and educational reporting for learners. To map this design space, we analyzed uncertainty representation in other domains from the perspective of the visual variables that are detailed by Bertin [26] and Gestalt [27] psychology. This analysis revealed how these variables have been manipulated to represent uncertainty and other information within 50 visualizations. We then analyzed the use of the same visual variables within 106 educational reporting visualizations that are accessible to learners. This analysis considered how these variables were used to support the high-level communication goals that the studied visualizations were attempting to convey, such as the learner's activities or mastery of a concept or his/her knowledge in contrast to that of an expert or peer. The analysis considered how the designers of the visualization handled potential uncertainty when choosing which information to present and how it was to be presented. This revealed that uncertainty information is represented in less than half of the studied visualizations and it is rarely represented directly. It is much more common for uncertainty to be communicated indirectly via the use of categories or linguistic cues.

After analyzing existing uncertainty and educational reporting visualizations, we combined the information about the use of visual variables within both sets of visualizations with information about the effectiveness of these visual variables for representing uncertainty (see 3.1 Existing Uncertainty Representations). This helped identify the variables that were available and appropriate for representing uncertainty within educational reporting. Once these opportunities were identified, we modified existing visualization approaches to illustrate how the different visual variables could be used to communicate uncertainty information in educational reporting.

The analysis of 50 visualizations that include uncertainty information and 106 educational reporting visualizations mapped the design spaces that were already occupied. This revealed the gaps that are present in these design spaces. Now that these gaps have been identified, their use for communicating information about uncertainty can be explored.

The inclusion of uncertainty information in educational reporting could allow learners to make more informed decisions that better support their desired outcomes. The gaps in this space and the examples of their exploitation may also allow the designers of these visualizations to explore this space and the relationship that the inclusion of uncertainty information has to system or report use, user trust in the system or report, and the development of learners' metacognitive skills.

The paper first discusses the elements that go into visualizing educational data. This includes elements of information visualization, a discussion of learner modelling and learning analytics, and a discussion of how educational data has been represented. Following this general background, an analysis of how uncertainty has been represented in other domains and educational reports is presented. We then present the design space that is available for representing uncertainty within visualizations of educational data for learners. This is followed by a discussion of how this design space might be explored and a current system is modified to provide an example of how this design space might be used to communicate uncertainty.

2 Visually Representing Educational Data

Very little work has been reported on how to best represent different aspects of educational data, including uncertainty, to users of educational systems. However, considerable work has been performed in the use of different types of visualizations within educational reporting and technology-enhanced learning environments. In addition to this, substantial work in the area of information visualization can be exploited to design visualizations that help learners make informed decisions about their learning or abilities; better understand their learning; or comprehend how their knowledge or abilities have changed.

We have divided the discussion of these related areas into four sections. The first explains a subset of the principles from information visualization. The second discusses how people interpret visual information, and the third presents an overview of the analysis methods that are often applied to educational data; this includes learning analytics and learner modelling. The fourth section discusses the purposes of giving learners access to educational reporting with a focus on existing approaches to visualizing educational data.

2.1 Information Visualization

To better understand how information can be communicated to learners visually, we must first understand how people perceive visual information. This leads us to explore the field of information visualization which is primarily concerned with finding ways to communicate complex information in a manner that allows the consumers of that information, in our case the learner, to more easily understand the data and make inferences based on the presented information [28]. Information visualization is further concerned with the faithful presentation of information and its associated patterns. Wainer [7] argues that effectively displayed data reminds us of the limitations of that data and prevents us from making incorrect inferences. Please see [26] and [27] for additional background in information visualization.

One of the ways that we can aid users in understanding and interpreting information is by exploiting the visual variables that psychology has shown us can be processed pre-attentively. Bertin [26] refers to these as selective. These variables reduce user cognitive load since information that is communicated through them is automatically processed without requiring focused attention. The use of the individual visual variables are detailed in the works of Gestalt psychology [27] and Bertin [26] where overlapping and sometimes different perspectives are provided. We, therefore, define the variables (see Table 1) to scaffold our later analysis (see Section 3).

Table 1. A description and example of the visual variables

***Visual Variable*: Explanation**	**Example**	***Visual Variable*: Explanation**	**Example**
Position: Changes in the x, y, or z location of an element.	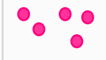	*Size*: Changes in the length, area, or repetition of elements	
Motion: This is often called common fate. Elements that share a direction, velocity, or frequency of motion tend to be perceived as belonging to the same group.		*Continuity*: Elements along a smooth and continuous path tend to be grouped together. The use of other variables, such as colour, can interfere with this visual effect.	
Hue: Changes in hue at a given value. Both Gestalt and Bertin call this colour.		*Opacity*: Changes in the amount of light that can travel through an object.	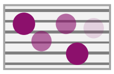
Saturation: Changes in the amount of grey in the colour space (0 = grey, 1 = full colour). This can be thought of as the purity of the colour.		*Orientation*: Changes in the alignment of elements. Rotating visual elements can result in changes in their orientation to one another.	
Value: Changes from light to dark.	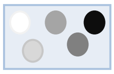	*Grain*: Variation in pattern. This is sometimes called texture.	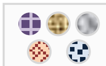
Boundary: Graphical elements tend to be grouped together when they are within a boundary or common region. A closed contour tends to be interpreted as the boundary of the object and elements within that contour as belonging to the same group.	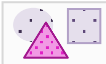	*Connectedness*: Uniform, connected regions are perceived as a single unit.	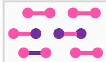
		Proximity: Graphical elements that are close together are perceived as belonging together.	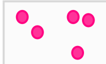
Shape: Variation in shape at a constant size. Visual elements can go from being a simple shape to have the characteristics of an icon, which limits interpretability.		*Numerosity*: Changes to the number of elements within a space tend to be grouped together. In this example, the two circles with more dots would be perceived as belonging to the same category. Our ability to distinguish between groups with different numbers of elements increases with age [29].	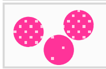
Closure: A closed contour is interpreted as a shape. If enough of the contour is shown then people perceive the shape in its entirety by filling in the missing information.			
Depth: Changes in position along the z-axis (i.e., stereoscopic depth)	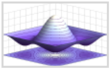	*Added Marks*: Elements with additional marks tend to be perceived as different from those without marks. Those with the same added marks also tend to be interpreted as belonging to the same class of elements.	
Arrangement: Changes in the consistency of the alignment of elements.			
Curvature: Changes in the curvature of elements. Objects with similar curvatures tend to be grouped together.	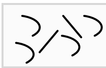	*Blur*: Changes in the clarity or fuzziness of objects. This is also known as semantic field of depth.	

While it might be possible and even tempting to use all of the visual variables, this is not recommended. Tufte advocates that all unnecessary marks be removed from a visualization [30]. However, there are cases when providing additional information can improve both the user's ability to interpret information and the memorability of the visualization's intended message [31]. There is, therefore, a balance to be found between minimizing the amount of additional information that is given to users and ensuring that the visual representation effectively communicates the intended message as well as the data's context [28]. Many metrics and heuristics for evaluating the appropriateness of visualizations have been proposed [32], [33]. A discussion of these heuristics and any accompanying metrics can be obtained from [33].

Keeping the tension between limiting unnecessary information, minimizing cognitive load by exploiting the pre-attentive processing system, and enabling system developers to reinforce their message through multiple channels resulted in the selection of a subset of the visual variables that are described by Gestalt [27] and Bertin [26] (see Table 1). The properties of each of these variables are detailed in Table 2 which describes the extent to which people can order items based on the variable without it having been assigned an order (orderable), group items based on the variable (associative), pre-attentively process the variable (selective), and quantitatively compare items based on the variable (comparable). Table 2 also details the variable's cardinality or length; this indicates the number of levels that people can distinguish and, therefore, the maximum number of levels that we can communicate using that visual variable. For example, blur has a cardinality of 4. This means that users can only effectively group items or distinguish between them without additional thought when four different levels of blur are used even though users may be capable of recognizing more levels of blur.

While these variables were chosen because they can be pre-attentively processed and any visualization can be decomposed into these variables, other considerations with respect to people's ability to interpret information need to be made.

2.2 Interpretability of Visualization Data

Manipulating the visual variables that are described in Table 1 and Table 2 provides a reasonable starting point from which to visualize data, but it is not enough to understand how these variables can be manipulated. The complex nature of human perception means that some things should be kept in mind when creating visualizations for users, especially those who may not receive training or who have lower-numeracy levels, as may be the case among learners, especially school-aged children and low performers. Moreover, carefully selecting visual representations is important since even people with high-

Table 2. Visual variable properties

	Visual Variables																				
Properties	*Added Marks*	*Arrangement*	*Blur*	*Boundary*	*Closure*	*Connectedness*	*Continuity*	*Curvature*	*Depth*	*Grain*	*Hue*	*Motion*	*Numerosity*	*Opacity*	*Orientation*	*Position*	*Proximity*	*Saturation*	*Shape*	*Size*	*Value*
Orderable	○	○	●	○	○	○	○	○	●	○	○	○	●	●	○	●	○	●	○	●	●
Associative	●	⊙	●	●	⊙	●	●	●	●	●	●	●	●	●	●	●	●	●	⊙	●	●
Selective	●	⊙	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●	⊙	⊙	●	●
Comparable (> = <)	○	○	⊙	○	○	○	○	○	○	○	○	●	●	●	○	●	○	○	○	⊙	○
Cardinality	⇩	3	4[1]	∞[2]	2[3]	2	2	⇩	∞	∞	7[4]	?	⇩[5]	?	4	∞	⇩	⇩	∞	5	7

● - True, ○ - False, ⊙ - Partially True, ⇩- Low, ∞ - Infinite, ? – Unknown

er numeracy levels can struggle with properly interpreting numerical information [2], [3], [10], [34].

While we do not address all of the challenges that people can face with respect to interpreting information, we provide some examples that are relevant to the user's ability to interpret numerical information since later discussions focus on visualizing information for learners, many of whom may have low-numeracy. One such challenge is the denominator effect, which is when people ignore the denominators in ratios [10]. This can be combated by using a common denominator, preferably of base 10, for all of the presented ratios. An alternative solution is to use icon arrays, such as the one in Figure 3, since these are known to combat the denominator effect.

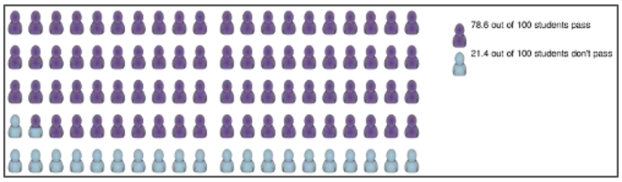

Figure 3. An icon array showing the pass rate for a course

Another common challenge faced by those trying to interpret bar graphs is the within-bar bias, where people think that points shown within the bar that are equidistant to those outside the bar are more likely to occur even though their probability of occurrence is the same (Figure 4). This type of error affects decision making and at least 27 percent of college-educated people make this mistake [3], meaning that error bar-like reporting should probably be avoided in visualizations of educational data for learners, especially in school settings or domains where students have limited mathematics and statistics training.

Designers of educational reports should also be aware that the human perceptual system gives precedence to certain visual variables (see [32] [37] [38]). Due to this and the variability with which some of the variables (e.g., hue and saturation) are interpreted [39], it is recommended that more conservative cardinalities are used for certain variables (see Table 2). This potential variability in the interpretation of different visual variables also illustrates how the visualization process can introduce uncertainty. It should be used as a cautionary note about the importance of providing additional cues within a visualization of educational data since many of the users of these visualizations may not have received training in their semantics, interpretation, or use even though it is recommended that they receive training [16].

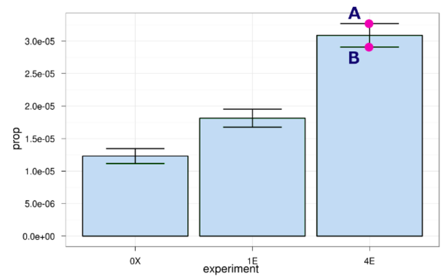

Figure 4. Within-bar bias: people who exhibit within-bar bias will interpret the value at point B as being more likely than the value at point A even though they are equally likely.

2.3 Learner Modelling and Learning Analytics

Before we can discuss the current state of the field with respect to visualizing educational data, we must discuss how this data is analyzed so that we can better understand how it has been or might be represented visually. Part of this involves understanding how data is collected within technology-enhanced learning environments. While educational reporting systems can rely on teacher-entered assessments, many technology-enhanced learning environments perform the detailed tracking of learner activities. This information could and has been combined with teacher assessments of student knowledge and its analysis is sometimes referred to as learning analytics [40], especially when it aims to describe groups of students or identify patterns in student performance. In contrast, the modelling of this data at a student level is often referred to as learner or student modelling [20], [41]–[46]. The techniques used in learning analytics are often similar to or the same as those used within learner modelling.

[1]While people can distinguish between 10 levels of blur and up to 5 levels of blur have been used effectively in at least one setting [35], it is probable that most people can only use 4 levels effectively [36].
[2]Infinity is a theoretical cardinality. There may be practical limitations on people's ability to interpret all possible values.
[3]All cardinalities of 2 are based on the definition of the variable, which is inherently binary.
[4]This is the approximate cardinality for selection, the most limited of the properties
[5]While numerosity's cardinality is unknown, it interacts with a person's cognitive development [29] and the element's contour. We, therefore, recommend using small cardinalities, especially with younger learners.

However, learning analytics usually seeks to describe learners and groups of learners to inform teaching practices, whereas learner modelling is more commonly employed to drive the personalization that occurs in adaptive learning environments. That said, learner modelling can be thought of as a special type of learning analytic. In both cases, the intended consumer of the analytic is not the learner but rather another decision maker such as a policy maker, teacher, or adaptive learning system.

Systems that employ learner models do so with the intent of using a model that represents aspects of the learner's beliefs, goals, knowledge, intentions, or cognitive state. These models can be computed using various approaches that are either data-driven, as in learning analytics and educational data mining, or based on educational and psychological theories [47].

Once a model has been computed, it can be used to adapt learning materials or activities to a learner within a technology enhanced learning environment [48]–[50]. Models can also be used to inform the learner, the teacher, or other decision makers in order to support the improvement and understanding of learning and the environments in which it occurs [14], [40], [51]. After computing a learner model or performing analytics on learner data, the results of these processes can be made available to learners at which point the veracity of the data and the inferences made from it becomes increasingly important.

We cannot assume that the data is accurate since inaccuracy could be introduced at any stage in the modelling process. Furthermore, the modelling process is inherently inaccurate because it creates an abstract representation of specific items. In spite of this limitation, models have been both useful and effective when employed to help people learn within adaptive systems [45], [52]–[56]. In many cases, the intelligent tutoring systems that employ learner models have been shown to achieve similar learning effects to those achieved through human tutoring [55].

We, therefore, have an obligation to ensure that learners can properly interpret the information that they are shown given their background preparation and abilities. This may be more difficult when communicating with learners who have different levels of knowledge and experience with respect to interpreting visual, numerical, or uncertain information [3], [5], [7], [8]. Supporting a group of students with heterogeneous abilities would not be uncommon in introductory undergraduate courses.

2.4 Visualizing Educational Data: Open Learner Models and Dashboards

Given the sometimes complex nature of the analytics that are performed on educational data, not all visualizations communicate all aspects of the underlying data or the manipulations that were performed over that data. These visualizations aim to support the monitoring of learner activities and the decision making of a range of stakeholders from the education system; all of whom have different needs and abilities. This can include policy makers, principals, parents, teachers, and students. As a result, educational reporting and the visualization of educational data can take many forms and often falls under the umbrella of learning dashboards that display everything from a student's grades on individual tasks beside a class average to learner activities and the relationships between learners in an online community [15], [24], [57]. Allowing learners to have access to these reports gives them feedback about their learning activities, performance, or knowledge. It can also help learners meet their educational goals by supporting their reflection and decision making [14].

One of the more prominent and specialized approaches to visualizing educational data for learners is called open learner modelling; these visualizations 'open' the underlying representation or model of a learner so that the learner can view all or part of the model [24], [58]. They, therefore, require the presence of an underlying learner model that has performed some form of analysis or inference on the learner's data. In the same way that learner models can be thought of as a specific type of learning analytic, open learner models can be thought of as a specific type of learning dashboard. This makes open learner models and learning dashboards a special type of information visualization that is meant to enable inferences over a representation of a person's knowledge and any inferences that have been made about that person's knowledge. Even though open learner models may not always represent what is more popularly called big data, they are intended to make the abundance of data that is available about a learner understandable. In making the information understandable, open learner models aim to achieve their primary goals of supporting the learner's metacognitive development and abilities; supporting learner reflection; providing learners with feedback about their abilities, beliefs, and knowledge; enabling the learner to plan learning activities; and supporting learner self-assessment [14], [24], [58]–[60].

Open learner models are usually meant to be used by learners themselves. However, the use of other visualizations of learner activities and abilities can be found in many types of technology enhanced learning environments, including learning management systems such as Blackboard [61] or Moodle [62]. Those that do not rely upon an underlying model of the learner are often called learning dashboards. Neither they nor open learner models typically address uncertainty in the same manner as that used in other areas of information visualization even though they may accommodate for uncertainty in the underlying model or present information about the model's accuracy [24] (see Section 3.2.2 for additional details). This may be, at least, partly due to the lack of representation of uncertainty in other forms of educational assessment and feedback. For example, students are not typically given a grade with a confidence interval, such as 85 percent plus or minus 3. The closest thing that commonly used assessment methods have to representing uncertainty is the use of grade categories (e.g., A, B, or C) that reduce the precision that assessors must use. When considering formal testing, measurement error is commonly represented in score reports using confidence intervals. Investigations into the effect that different confidenceband representation methods have on teacher and undergraduate comprehension have recently begun [2].

While most systems do not directly represent uncertainty, there have been some notable systems that either represent uncertainty [13], [63]–[65] or incorporate mechanisms that handle uncertainty. ViSMod directly represented uncertainty graphically; it showed student knowledge as a Bayesian model that included uncertainty information about the learner's concept knowledge as a probability [65]. The challenge faced when using such approaches is how to support learner understanding since their full comprehension is dependent on a fairly advanced understanding of mathematics, which renders these approaches ill-suited for use with some learner populations (e.g., primary school children or those in special education contexts).

3 Uncertainty

While some of the data analysis methods that are described in Section 2.3, account for variability or other sources of uncertainty in the data that is being analyzed, they do not explicitly discuss the concept of uncertainty, its many forms, and the implications that it has for the interpretation of the information that is presented. We must, therefore, discuss how uncertainty has been conceptualized if we are to explore its visual communication and use in educational reporting.

Uncertainty can be introduced from a variety of sources or during any level of information processing [66]. This includes the interpretation of that information [1] and its visualization [25].

Even though uncertainty can be reduced to two basic types (aleotonic and epistemic [67]), its multifaceted nature [32] has made it useful for those who work with uncertainty to detail the types of uncertainty that are often encountered in the data with which they work. These frameworks refine and sub-divide uncertainty to ensure that they account for the types of uncertainty that are prevalent in different domains. We have synthesized these frameworks into a single list and added examples from education:

- Accuracy [68]–[71]: The difference between what has been observed and reality. It includes many potential sources and types of error and could be the result of a learner slip (e.g., accidentally clicking on the incorrect multiple-choice answer) or error.
- Precision [68], [69], [71], [72]: This is the known level of accuracy of different measurement tools and includes the standard error of measurement as well as the granularity at which information can be tracked or measured. In an educational domain, a percentage grade (e.g., 87) would have a greater precision than a letter grade (e.g., A) and standardized tests have known measurement error. In a modelling context, this might be represented by the level of granularity that is provided by the statistical modelling technique or the type of constraint that a learner has violated [43]: the constraint that indicates that the rules of subtraction have been violated has a lower precision than a constraint that indicates that the rules of subtracting negative numbers have been violated.
- Completeness [69]–[72]: This is the comprehensiveness of the data which includes its coverage, selection criteria, and the availability of the desired information. This asks whether the learner has performed enough activities for us to make inferences about his/her abilities in an area. If the tutoring environment changes its intervention based on learner affect, a system that infers affective state based on logged interactions may be reasoning over incomplete information when compared to one that also uses a webcam feed of the learner's face.
- Lineage [6], [68]–[71], [73]: This is the source of the information and how the information has been manipulated. It includes any aggregations that have been performed as well as the sensitivity of any algorithms that are applied to the data. Typically, this will be measurable in educational contexts when statistical modelling techniques, such as k-means clustering, are applied since these techniques can be tested with new data and their accuracy reported. In some cases, the accuracy or sensitivity of algorithms may be unknown
- Judgement [68]–[70], [72]: This relates to the subjectivity of a data source or the inferences that are made over data as well as the trustworthiness of the data source or agent performing inferences. In an educational context, this may relate to who is performing an assessment. A student may trust an instructor's assessment over that of a peer when technology-enhanced learning environments allow peer assessments to be performed.
- Validity [68], [69], [71], [73]: The usefulness and credibility of the data for its intended purpose. Teacher comments that are tied to a rubric may seem more valid to a student than a peer's comments that are not tied to the assignment rubric.
- Currency [69]–[72]: Is the information current enough to be useful? Generally, information that is more recent is a more reliable measure of a student's abilities than old information. Some modelling systems perform a weighted average over student activities where more recent activities are given greater weight in order to account for uncertainty that might be due to currency [74]. Other systems address this by averaging performance measures and providing information about how learner performance has changed over time [48].
- Statistical Variance & Consistency [68], [69], [71], [72]: Agreement in the evidence. Learners may perform poorly on an assessment due to a distraction even though they normally perform well. Like with currency, some systems also perform different types of aggregation to account for variance in learner performance [48].

Regardless of the types of uncertainty that are being represented, people tend to reject or accept uncertain information in its entirety [71], which is why care should be taken in its representation. Well-designed visualizations may help combat people's biases towards data uncertainty and aid in their analysis [71].

To help users with decision making tasks, we should provide them with information about the data's accuracy; precision and currency information should also be given since it is important to decision making [69]. The amount of influence that uncertainty has on decision making is affected by the type of uncertainty information that is provided and its method of communication [75]. Access to uncertainty information and increased granularity in uncertainty reporting can decrease user confidence in their decisions and may only affect the decisions that are being made about items that are at the extremes (i.e., those with very high or very low uncertainty) [75]. These and other results show that the manner of presentation of uncertainty information can be exploited to encourage certain decision-making behaviours. See [10], [8], [35], [72] for information on the particular ways in which decision making can be influenced.

To date, uncertainty information has most commonly been represented numerically although it has a long tradition of being visualized [6]. Gershon [72] recommends being selective when representing uncertainty since its representation adds a dimension to the data [1], [31]. The only uncertainty that is shown should be that which is essential to the target task. Moreover, it should be presented as simply as possible while maintaining the dominance of the information that is important for a given task rather than distracting the user from his/her task [1]. In a pronunciation tutoring environment where a learner must decide which aspects of his/her pronunciation to improve, this may mean that only the items with the most extreme values (best or worst pronounced) or those with the highest accuracy are presented to learners [48]. Alternatively, those with high inaccuracy due to learner inconsistency might include uncertainty information in order to encourage the learner to increase his/her consistency in pronouncing certain characters.

To overcome this potential barrier, we would recommend the use of pre-attentive visual variables since they can be processed quickly, unlike icons[6] [69]. Others have recommended the use of what are considered to be natural cues for uncertainty: blur [4], [36], [69], [76], graduated shading or colour value [7], closure through the use of dotted or dashed lines [4], [36], and ambiguous labels when precision is lacking [7]. In contrast to what might be expected, colour hue is not recommended for communicating uncertainty because it is not pre-attentively orderable [36], but it could be used to indicate the presence of uncertainty on a binary level [1] or graduated scale when training or a legend are provided.

3.1 Existing Uncertainty Representations

To better understand how uncertainty representation might be used in educational reporting, an analysis of its use within other fields was performed. The 50 visualizations that were analyzed came from various communities including information visualization, geographic information systems, medical information systems, and statistical visualization. These visualizations were found based on a search for articles related to uncertainty and visualization. Colleagues in information visualization were also contacted to ensure that no major works had been missed. This resulted in the exploration of a community maintained reference list[7].

Table 3. A summary of the use of the visual variables for communicating uncertainty (Unc.) or other information

	Used		Unc.[8]		Unc. Only[9]	
Visual Variable	*%*	*No.*	*%*	*No.*	*%*	*No.*
Added marks	16.0	8	16.0	8	100.0	8
Arrangement	8.0	4	8.0	4	100.0	4
Blur	6.0	3	6.0	3	100.0	3
Boundary	34.0	17	8.0	4	25.0	4
Closure	12.0	6	8.0	4	66.7	4
Connectedness	14.0	7	2.0	1	12.5	1
Continuity	20.0	10	2.0	1	12.5	1
Curvature	10.0	5	2.0	1	20.0	1
Depth	10.0	5	6.0	3	60.0	3
Grain	6.0	3	6.0	3	100.0	3
Hue	54.0	27	36.0	18	66.7	18
Motion	18.0	9	10.0	5	55.6	5
Numerosity	18.0	9	16.0	8	88.9	8
Opacity	12.0	6	10.0	5	83.3	5
Orientation	24.0	12	16.0	8	72.7	8
Position	50.0	25	28.0	14	56.0	14
Proximity	16.0	8	2.0	1	12.5	1
Saturation	18.0	9	14.0	7	77.8	7
Shape	34.0	17	22.0	11	57.9	11
Size	50.0	25	36. 0	18	52.0	13
Value	18.0	9	10.0	5	50.0	5

Any visualization that represented some form of uncertainty and was not a form of educational reporting was included regardless of whether the visualization had been evaluated for its effectiveness. All of the visualizations that were analyzed had been published prior to October 2013. Papers that analyzed visual variables for their effectiveness at representing uncertainty independently of their use in a visualization were not included in the below analysis. However, they were used to inform the design space that is described in Section 4.

Each of the visualizations was analyzed for its use of the visual variables that are described in Table 1, and all representations included a visual component that was not text-based. Details of visual variable use can be seen Table 3, where the usage of visual variables is provided (Used) as are the details about the overloading of visual variables to represent uncertainty and one or more other dimensions within the data (Unc.), or the use of a variable to represent only uncertainty (Unc. Only).

[6] Icon is not meant to mean a graphic in a user interface. MacEachren et al. use the term icon to communicate the use of a symbol with a complex visual shape, such as a missile or emoticons.

[7] http://www.sci.utah.edu/~kpotter/Library/Catalogs/uncertaintyVis/

[8] Sometimes a variable was used to communicate 2 things: uncertainty and another piece of information

[9] Unc. Only indicates that the variable was only used for the purpose of communicating uncertainty information.

The most commonly used visual variables were size, hue, and position with approximately fifty percent of the visualizations using these variables. The visual variables that were most commonly used to represent uncertainty were size and hue (approximately 36 percent). Uncertainty was represented using position in 28 percent and shape in 22 percent of visualizations. The next most commonly used visual variables were saturation, orientation, numerosity, added marks, and opacity.

If we only consider the proportion of uncertainty representation to the use of visual variables then all of the uses of grain, blur, added marks, and arrangement were dedicated to uncertainty information (see Table 3). The next most used visual variables (saturation, orientation, numerosity, and opacity) represented uncertainty at least seventy percent of the time. Hue, value, shape, motion, depth, closure, position, and size were all used to represent uncertainty in at least half of the cases where these variables were used (Unc. Only). Of these highly-used variables, saturation and hue are known to be poor communicators of uncertainty; blur, position, and value have been shown to effectively communicate uncertainty; and arrangement, size, and opacity are known to be moderately effective at communicating uncertainty [69].

However, we do not yet know if these variables can be used to represent uncertainty in learning dashboards since we do not know how information has been represented in visualizations of educational data: these visual variables may already be used for communicating other information. Moreover, the varied nature of learner abilities adds another dimension to the design task since students in a single class may have heterogeneous numeracy or literacy and only some may have the background preparation that is needed to work with a particular visual representation. This is unlike the settings where most uncertainty visualization work has been done: their users are typically accustomed to working with uncertainty and trained in interpreting its various representations.

3.2 Uncertainty Representation in Educational Reporting

An analysis of 106 visualizations from the educational reports that are provided through various technology-enhanced learning environments was performed. These visualizations were found through literature searches for open learner models, learning dashboards, and educational reporting. In many cases, this also involved working backwards through the reference lists of papers and visiting authors' websites to find additional work. The selected systems have been reported in various venues including artificial intelligence in education (AIED and IJAIED); user modelling, adaptation and personalization (UMAP); intelligent tutoring systems (ITS); learning analytics (LAK); and the SIGCHI conference on human factors in computing systems (CHI). The analysis was only performed on visualizations that were reported prior to October 2013 and that were accessible to learners rather than those that were intended to be used by parents, instructors, administrators, or policy makers. The majority of visualizations were also paired with some form of adaptive tutoring system although this was not a requirement for inclusion. Any visualization of educational data to which the learner has access was analyzed since these visualizations share similar considerations even though their underlying mechanics may be different. The below-described analysis considered the message that the visualization was attempting to convey (i.e., the visualization's communication goal), its use and representation of uncertainty, and its use of visual variables (see Table 1).

3.2.1 Communication Goals

The visualizations of educational data were categorized based on their intended message (i.e., type of information that they were attempting to communicate). Visualizations were allowed to belong to multiple categories when the information being conveyed fit within the described category. Categories emerged from the data and each time that a new category was found, the previously coded visualizations were revisited. The intended message categories are defined in Table 4.

The most common communication goal was mastery with 72.9 percent of visualizations communicating learner mastery (e.g., Figure 5 and Figure 6), 38.3 percent encouraging comparison (e.g., Figure 7, Figure 8, and Figure 11), 29 percent communicating learner activities (e.g., Figure 7 and Figure 11), 20.6 percent communicating the relationship between different entities (e.g. Figure 10), and 14 percent providing social cues (e.g., Figure 7 and Figure 11). Only 1.9 percent communicated learner interest, 2.8 percent provided information about learner affect, and 0.9 percent encouraged learners to notice changes in their knowledge. The limited exploration of the representation of social cues, learner interest, learner affect, and changes in knowledge may indicate that these communication goals are not as well understood as the others. As a result, the representation of uncertainty information and the study of how learners respond to it may be premature in these types of visualizations of educational data.

Within the evaluated visualizations, bar chart style diagrams were used by 29.3 percent of visualizations, network or graph representations were used by 17.9 percent of visualizations, and textual representations were used in 28.3 percent of visualizations. Word clouds, emoticons, tree maps, spider plots, line charts, pie charts, and scatter plots were used to represent information by fewer than six percent of visualizations, and 34.9 percent used representations that fall into the Other category which includes speedometers [77], rose-like diagrams [78], animated characters [79], and magic wands [80]. Another example from the Other category is a tree that grows or dies based on changes in student knowledge [81] or a system where the visualization was paired with haptic feedback that communicated understanding [82]. Of the evaluated visualizations, 75.5 percent used only one display approach. The remaining systems combined two or more representations, with at most four display representations being used by any one visualization at a time.

Table 4. The categories of communicative intent that emerged from the studied educational visualizations.

Category	Definition
Activity	These visualizations aim to communicate what learners are doing or have done. One version of SQL-Tutor (Figure 1) combined the representation of mastery with that of activity by assigning a colour to the proportion of the activities that the learner had yet to complete. This version also used text to indicate the proportion of the learning activities that the learner had completed. Narcissus (Figure 11) communicated learner activities by using saturation and hue to show user contributions broken down by the types of activities that are commonly performed in software development teams [22].
Affect	These visualizations aim to communicate a learner's affect or motivation. My-Pet (Figure 9) changes the animated character's expressions and behaviours based on the learner's affective state [79].
Change in Knowledge	These visualizations aim to communicate how a learner's knowledge, preferences, or abilities have changed. In the case of ProTutor (Figure 8), the chart tracks the user's ability to pronounce specific Russian characters [48]. The accuracy of the learner's pronunciation of the selected characters is shown at three time points which allows the learner to see how his/her pronunciation has changed. In this case, the learner's pronunciation accuracy improved and then decreased slightly demonstrating uncertainty that is due to statistical variance and consistency (see Section 3).
Comparison	The intent of these visualizations is to encourage the comparison of two or more entities. This may mean the comparison of knowledge between two peers or the comparison of a learner's knowledge against some ideal. It could also mean the comparison of the learner's performance against a particular goal. The use of columns in Narcissus (Figure 11) enables the monitoring and comparison of individual team member's contributions against one another or against the contribution patterns that are expected given each person's role within the software development team. By placing all of the learners within one visualization (Figure 7), Comtella encourages learners to notice both their and their classmates' contributions to the learning community, and it highlights the differences in their contributions using several visual variables to communicate how learners have performed in comparison to one another [23].
Interest	These visualizations aim to communicate a learner's interest in concepts or activities. The My-Pet system (Figure 9) uses the cartoon animal to reflect the learner's observed interest in a topic back to him/her. In contrast, the Pepper system (Figure 10) represents group level interest in topics by showing which words dominate a selected discussion [21].
Mastery	These visualizations aim to communicate how much of a particular entity has been mastered by a learner (i.e., how much s/he knows or how well s/he can perform a skill). Figure 1 and Figure 5 show bar-chart style skill meters that communicate the proportion of a concept for which a learner has demonstrated knowledge or the learner's mastery of a particular competency or sub-competency.
Relationship	The intent of these visualizations is to communicate the relationship between different entities (e.g., people or concepts). The network graph used in the Next-TELL open learner model shows the relationship between the different competencies that are tracked within the underlying learner model (Figure 6). The Pepper word cloud (Figure 10) shows the relationship between discussion forums and learner interest in various topics since Pepper generates word clouds for a discussion forum or even individual posts when a user requests to see a cloud for a specified set of messages.
Social Cues	The intent of these visualizations is to create an awareness of the abilities, activities, or interests of others. Narcissus does this by showing the activities of each user in vertical columns (Figure 11), and Comtella (Figure 7) did this by showing all learners as stars within a grid that included representations of the dimensions of their contributions.

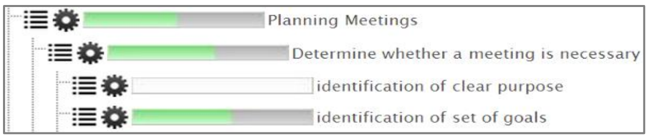

Figure 5. The Next-TELL skill meter visualization [83].

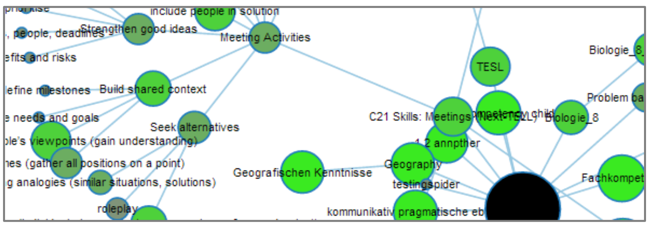

Figure 6. The Next-TELL competency network visualization. Larger, darker circles indicate higher competency levels [83].

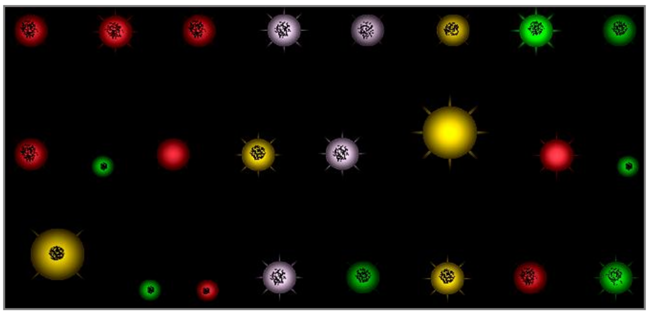

Figure 7. Comtella: each learner is represented as a star. Hue indicates learner status; star size indicates the number of contributions; star brightness represents contribution quality. A learner's offline status is communicated by adding black dots to the centre of the star. Source: [23]; used with permission.

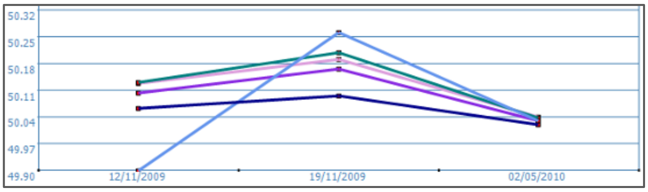

Figure 8. The pronunciation change chart from ProTutor. It shows the accuracy of a learner's pronunciation at different points in time. Source: [48]; used with permission.

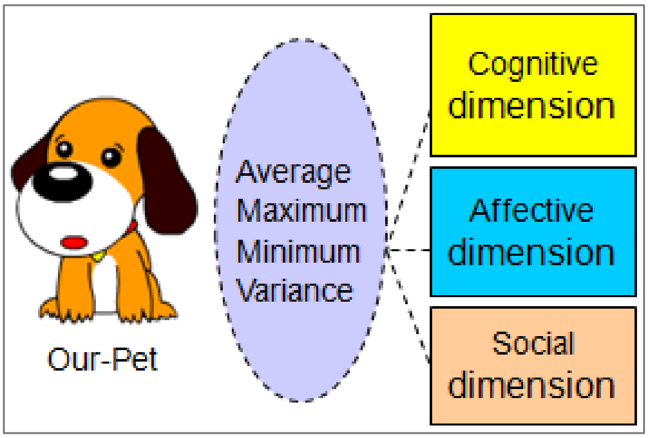

Figure 9. My-Pet uses a cartoon animal to communicate learner affective state by manipulating the facial expressions and dialogue that are used by the animal. Source: [79]; used with permission.

Figure 10. Pepper word cloud. It shows the topics that are being discussed within a specified discussion forum. [21]

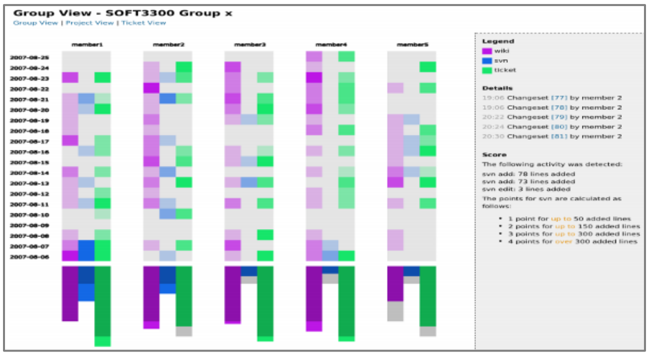

Figure 11. Narcissus shows how much of each activity type group members have performed as well as when they submitted the work product for that activity. Source: [22]; used with permission.

3.2.2 Uncertainty

Since all models and forms of aggregation contain some uncertainty, the manner in which the visualizations represented uncertainty, if it was represented, was analyzed. Of the 106 studied visualizations, 52 represented uncertainty directly (9) [13], [19], [65], [84]–[88] or indirectly (48); 5 of the 52 represented uncertainty using both indirect and direct methods [19], [65], [85], [87], [88].

Indirect representations of uncertainty involved showing the learner the evidence on which the inferences were based; text that hinted at uncertainty using hedges or other linguistic techniques (e.g., You may misunderstand concept A or "This extra contribution has been inferred from the terms visible at depth 2" [89]); or categorization of student performance into broad levels such as excellent, moderate, somewhat limited, and very limited [63]. For example, Figure 6 indirectly represents uncertainty by using performance categories that are communicated through colour; this system also allows students to see the weighted averaging of evidence that indirectly communicates uncertainty that is due to lineage, currency, or statistical variance and consistency (see Section 3 for definitions of uncertainty types). Uncertainty was directly communicated via text (e.g., "possible misconceptions" [84]), by displaying error bounds or confidence intervals [19], through the use of probabilities [65], [85], or through the use of an insufficient data category [84] (e.g., My-Pet); this is similar to the longstanding tradition of including quality statements in the legends of maps [4].

The use of categorization was most widespread: 43 percent of all visualizations and 83.6 percent of those that represented uncertainty used this technique to manage precision, lineage, or statistical variance and consistency-based uncertainty. This includes the visual representations that are shown in Figure 5, Figure 6, Figure 7, and Figure 10. All other uncertainty representations were used by less than 5 percent of the studied visualizations. However, this only tells part of the story. It fails to tell us how the visual variables were manipulated in order to communicate uncertainty.

3.2.3 Visual Variable Use

The analysis of how the studied visualizations communicated their intended messages and represented uncertainty was performed from the perspective of the visual variables that are listed in Table 1. If we consider Narcissus (Figure 11), we can see that it uses hue and saturation to communicate social cues around the types of contributions that students are making. Narcissus uses proximity and position to enable comparison and it employs the use of boundary, closure, and connectedness to communicate the activities that belong to each user. In contrast, Comtella (Figure 7) employs the boundary variable to indicate which elements belong to the visualization of student contributions to the class rather than which contributions belong to an individual student. Comtella uses position to enable comparison by making all of the student contributions visible at the same time, and it uses added marks (the black dots in the centre) to communicate social cues. Comtella communicates student activity levels using size, while hue and saturation are used to communicate social cues about the quality of a student's contribution as indicated by his/her status.

An overview of how these variables were used, with respect to the type of message that was being communicated, across all visualizations can be seen in Figure 12, and summary statistics of their use can be seen in Table 5. Figure 13 shows the variable usage as a proportion of the

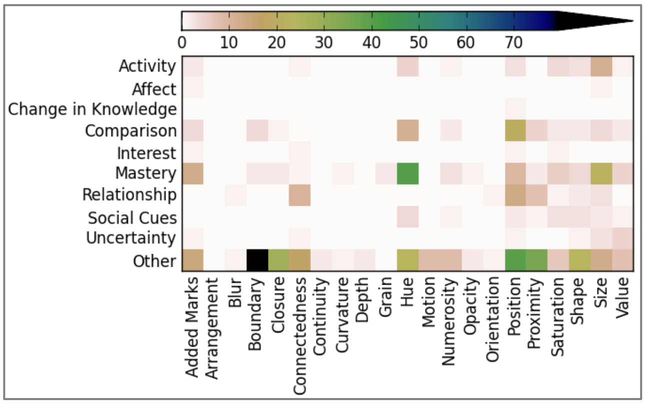

Figure 12. Visual variable use within educational reporting as a proportion of the total number of visualizations. The Other row shows the use of the variables for communicating information that does not belong to uncertainty or a particular communication goal. Dark blue indicates that a higher proportion of the visualizations use the specified variable.

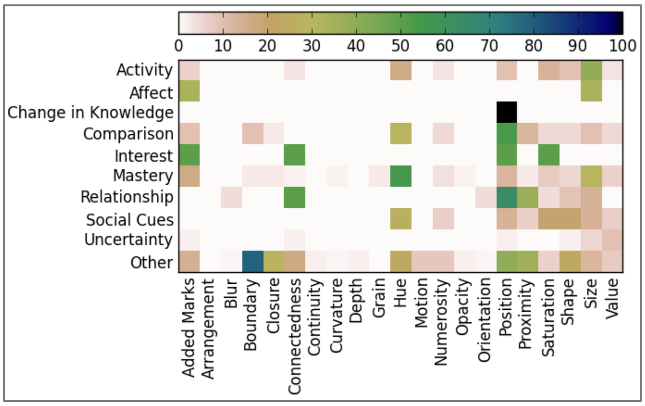

Figure 13. Visual variable use as a proportion of the number of educational reporting visualizations that communicate the message type. The Other row shows the use of the visual variables as a proportion of all visualizations, and darker colours indicate a higher frequency of use.

number of visualizations that aim to communicate a specified type of message. This allows us to see which variables tend to be used for different communication goals. For both Figure 12 and 13, dark blue indicates a higher frequency of use, greens indicate a moderate amount of use, browns indicate low usage, and white indicates no usage. The Other row indicates the use of the variable for communicating information that is not specific to uncertainty or any of the communication goals from Table 4.

The boundary visual variable is used by most visualizations (78.5 percent) even though it is not used to communicate a specific aspect of the message. Rather, it groups visual elements and thus aids in the communication of which aspects of the visualization should be considered a part of the visualization or one of its subcomponents. The heat maps shown in Figure 12 and Figure 13 also show that the mastery, activity, social cues, and comparison categories use a greater variety of variables than many of the other message types.

Table 5. A summary of the use of the visual variables within educational reporting visualizations

Visual Variable	Used %	Used No.	Uncertainty %	Uncertainty No.	Message %	Message No.
Added marks	34.9	37	2.7	1	56.7	21
Arrangement	0.0	0	0.0	0	0.0	0
Blur	1.9	2	0.0	0	50.0	1
Boundary	84.9	90	0.0	0	6.7	6
Closure	31.1	33	0.0	0	9.1	3
Connectedness	31.1	33	3.0	1	42.4	14
Continuity	1.9	2	0.0	0	0.0	0
Curvature	1.9	2	0.0	0	50.0	1
Depth	1.9	2	0.0	0	0.0	0
Grain	1.9	2	0.0	0	100.0	2
Hue	84.0	89	0.0	0	71.9	64
Motion	8.6	9	0.0	0	0.0	0
Numerosity	15.1	16	0.0	0	43.7	7
Opacity	2.8	3	0.0	0	33.3	1
Orientation	1.9	2	2.5	0	50.0	1
Position	90.6	96	1.0	1	55.2	53
Proximity	50.9	54	0.0	0	29.6	16
Saturation	22.6	24	0.0	0	70.8	17
Shape	37.7	40	0.0	1	35.0	14
Size	57.6	61	4.9	3	73.8	45
Value	20.8	22	22.7	5	40.9	9

Figure 13 reveals that position is the dominant visual variable for communicating a message that involves determining the relationship between two or more things; this includes communicating the relationship between system entities (people, concepts, or knowledge); facilitating comparison between a learner's knowledge or activities and the knowledge or activities of an expert/peer; and determining changes in knowledge. This implies that visual variables that are related to the location of an item (e.g., position or depth) should probably not be used to communicate uncertainty in a visualization that also aims to show the relationship between two or more items. However, the use of depth to represent uncertainty may be appropriate in cases where the visual analytic is attempting to communicate interest, social cues, or affect.

Like some of the visualizations from other domains (see Section 3.1), BonPatron used shape to communicate uncertainty [90]. The graph, tree, and map views from Flexi-OLM [91], the VISMod open learner model [65], and OLMlets [92] manipulated value to communicate uncertainty. Whereas, PSAT/NMSQT [13] used added marks and SIV [64] used position to communicate uncertainty. VisMod went a step beyond the manipulation of value to codify levels of certainty by also presenting the learner model as a Bayesian network that included probabilities [65]. While these systems were shown to be useful to learners, there has been little reporting on how the addition of uncertainty information affected learner decision making or understanding.

4 Opportunities for Representing Uncertainty in Educational Reporting

The above-described use of visual variables to represent uncertainty (Table 3) and educational data for learners (Table 5) reveals several opportunities for integrating uncertainty representations into existing or new visualiza-

Table 6. The appropriateness of the available visual variables for representing uncertainty

Visual Variable	Can Represent Uncertainty	Considerations
Added Marks	?	Marks could be added to information that is less certain.
Arrangement	~	Arrangement might be able to communicate uncertainty with messier arrangements communicating less certainty than more ordered ones. This might work better if changing the arrangement of hashing within a bar-chart style skill meter than it would if changing the arrangement of the elements within a star style skill meter since scattering can hinder the interpretation of a visualization in low-numeracy populations [10].
Blur	✓	Any aspect of a visualization could be blurred to communicate uncertainty. It should be noted that users disliked blur even though they could easily understand it [36]. However, they prefer the use of blur over that of value. If using blur, fuzzier images should be mapped to higher levels of uncertainty.
Closure	?	Users prefer dashed lines over blur, value, and arrangement but dashed lines are not as easily understood as blur or value [36]. It may be possible to show multiple levels of uncertainty using closure. However, it is more likely that you can only communicate the presence or absence of uncertainty. Using dashed rather than solid lines to invoke closure will require user training or the use of a legend since its use to communicate uncertainty is not automatically understood [36].
Continuity	?	Continuity would most likely only allow you to show if something is above a particular certainty threshold. Smooth curves could be used to communicate certainty and jagged lines could communicate uncertainty. This is probably most appropriate for use in graph-based visualizations.
Curvature	?	Changes in curvature could be used to communicate uncertainty but a legend will be necessary since more rounded objects could be used to communicate either certainty or uncertainty.
Depth	?	Less certain information should appear to be further away from the learner than more certain information.
Grain	?	Different textures could be used to represent different levels of uncertainty; a legend would be required.
Motion	?	Information with the same level of certainty could be given the same velocity or direction of travel. This is likely most suitable for use in visualizations that include animations or where position is not already used to communicate information other than group membership.
Opacity	~	More certain information should be more opaque. The level of transparency should increase with the level of uncertainty.
Value	✓	Value is already widely used and is naturally orderable but may require the communication of the anchors. You should, therefore, communicate whether darker values are more certain than lighter values.

✓ - Good, ~ - Okay, ? – Unknown

tions of educational data. Table 6 details the potential uses of each of the visual variables.

If we remember back to Section 3, the use of saturation, orientation, hue, and shape should be avoided when representing uncertainty because studies have shown that they do not effectively communicate uncertainty [69]. We would further recommend against the use of hue since it is already widely used for communicating other information within educational reporting (see Figure 12) and because it is neither orderable nor comparable (see Table 2). It is recommended that the use of size be avoided for communicating uncertainty since it has only been used successfully in applications where users undergo considerable training, in the order of months. Providing this level of training to student users of educational reporting is unrealistic. Should size be used, the recommended mapping is for larger objects to be more certain and smaller objects to be associated with greater uncertainty [69]. Likewise, more certain information should be presented more clearly and be less obscured than uncertain information.

The variables that are available for use fall into three groups: those that are good at communicating uncertainty (value and blur), those that are acceptable for communicating uncertainty (arrangement and opacity), and those with an unknown ability to communicate uncertainty (grain, continuity, depth, curvature, closure, added marks, and motion). These indicate potential avenues for research into how learners respond to and interpret various manipulations of the variables when attempting to communicate uncertainty information.

When choosing which variables to employ it is, obviously, important to consider their potential effectiveness at communicating uncertainty to the visualization's intended audience. Different learners will have different needs. The smiley face example from Table 7 might be appropriate for low performers in a primary school context, whereas the 3-D scatter plot might only be recommended for those who have an undergraduate education with extensive mathematics or sciences training.

It is also important to consider the communication goal of the visualization that is being created. For example, if you want to draw a learner's attention to a particular area, such as division, where s/he performs inconsistently then manipulating a variable that draws the user's attention to this aspect of the visualization would be advisable. In this case, the manipulation of the added marks or motion variables might be most effective at drawing the learner's attention to the information that is supposed to be highlighted. In a web-based environment, a blinking star could be placed beside the instructional material and practice exercises that are associated with division. This could highlight and perhaps even persuade the learner to provide the system with additional evidence of his/her abilities, which could in turn decrease the inconsistency that was observed in his/her ability to perform division. In contrast, the use of blur or opacity in this situation might obscure the intended message and result in the information being ignored.

We also recommend that a technique called brushing, which is a collection of methods for viewing multidimensional data, be used when learners can see the evidence on which the visualization is based; brushing allows the user to see the relationship between variables by identifying relevant data points across visualizations [38]. An example of brushing that communicates group membership is shown in Figure 14. Alternative techniques could be used provided they allow learners to see the source of the uncertainty beside its representation [68].

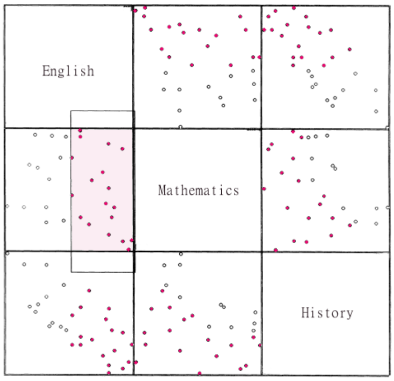

Figure 14. Brushing: the selected data points (learners) are highlighted in pink. This allows the user to see changes in learner performance based on the measures being used.

5 Adding Uncertainty Representations to Visual Analytics of Learning

Reviewing previous uncertainty representations revealed that the user's ability to understand the visualization is often overlooked [76]. This makes the careful design and evaluation of these visualizations based on the target user's ability to interpret them paramount since the interpretability of information is essential to the monitoring tasks that are supposed to be supported by open learner models and other visualizations of educational data [31]. While this paper focuses on the visualizations that are used by learners, similar approaches could be employed for other user groups such as parents, teachers, or administrators. However, the possibility exists that the design space that is available for exploration is different since the reports that target these user groups may use different visualization approaches. For example, it is unlikely that the reports that are being used by school principals contain smiley faces. It is far more likely that these reports contain error bars since it would be expected that the target population can interpret them.

Table 7 provides examples of how uncertainty information could be added to some of the types of visualizations that are commonly used in educational reporting. For the bar chart example of the skill meter, using opacity or arrangement allows the designer to communicate how consistently the learner has demonstrated evidence of a particular skill or knowledge component, which can be used as a proxy for how much the learner should trust the displayed assessment information. Moreover, using opacity to communicate something that is akin to error bars (as in example b.ii of the bar chart skill meter) could prevent within bar bias and better communicate the potential range of a learner's knowledge or abilities. The design that is proposed in b.ii has recently been externally validated for particular types of decision making tasks when they are being performed by adult members of the general public [93]. Additional guidance on the communication of error bars can be obtained from [94].

The star-based skill meter can be manipulated to communicate similar information to that shown in the bar-chart style skill meter. Opacity can be applied to icons to communicate how confident the system is in the learner's current knowledge level or its rating of the learner's affect. In 3-dimensional spaces, using depth (i.e., the z-axis) to communicate uncertainty by making certain items closer to the learner (see the 3-dimensional scatter plot example in Table 7) would help emphasize the information that has higher levels of certainty, but it could also lead to less certain items being obscured which could mean that the learner does not receive all of the information. In the case of word clouds, items can be blurred or their arrangement changed to communicate uncertainty in the underlying analytics. This uncertainty could be due to any number of things. In the case of a word cloud that is used to show a community's interest in different topics, blur could be used to communicate when small numbers of the community's members have demonstrated a deep or ongoing interest in a particular topic which has resulted in the topic being emphasized within the word cloud even though the majority of the community is not interested in that topic. This situation could easily happen in the Pepper system [21].

The decreased opacity of nodes in the graph example could be used to indicate the system's confidence in the inferences that it has made; its lack of confidence in the results of an assessment; or inconsistent evidence for a particular concept. This could signal to the learner that s/he should perform more activities within the technology enhanced learning environment. In contrast, the closure or thickness of the links between nodes could indicate the system's confidence in the relationship between the nodes. In this example, the nodes may represent learners and the links between those nodes their relationships, where a solid or thick line indicates a continued reciprocal relationship and a dashed line represents a lack of confidence in the inferred relationship. This lack of confidence could be the result of a one-way relationship or lurker-like behavior in the virtual world even though the learners may have a relationship outside of the technology enhanced learning environment.

To provide an example of how one might use the design space in a real system, we have augmented the Next-TELL open learner model because we have access to that system. The Next-TELL open learner model was also used because it allows learners to see the information on which their model is based. This provides a sufficiently complex setting in which to demonstrate the use of the identified design space.

Table 7. Examples of adding uncertainty information to the types of visual representations that are used in educational reporting.

Visualization	Uncertainty Representation	Visualization with Low Uncertainty	Visualization with High Uncertainty
Skill Meter - Bar Chart	a) map uncertainty to the arrangement of the skill meter fill b) map uncertainty to opacity for i) the entire meter when there is a lack of information ii) the area around the level of determined knowledge in a way that is similar to error bars for statistical variance and consistency	a) b.i) b.ii) 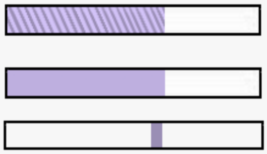	a) b.i) b.ii) 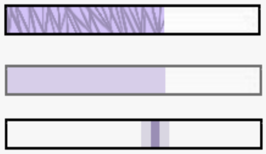
Skill Meter – Stars	map opacity to certainty		
Emoticons	map opacity to certainty		
3-Dimensional Scatter Plot	map uncertainty in node values to depth	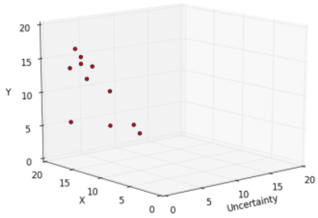	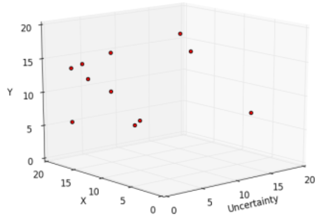
Word Cloud	a) map blur to uncertainty b) use the arrangement of words to communicate uncertainty (the messier the collection, the less certain it is) – this would only work if multiple word clouds are displayed at the same time		a)  b) 
Graph	a) map uncertainty in nodes to opacity b) map uncertainty in connections to closure c) map uncertainty in connections to line thickness where thicker lines indicate higher certainty	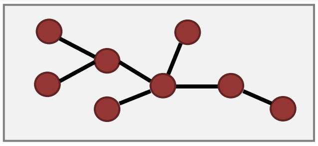	a) 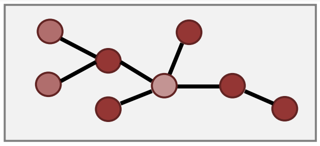 b) 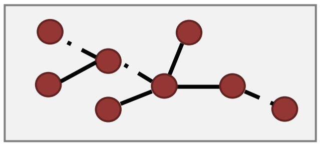 c) 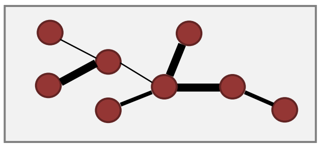

If we add uncertainty information to the Next-TELL competency network visualization by applying grain to nodes in order to communicate the presence or absence of uncertainty above an acceptable threshold, we might end up with a visualization like that shown in Figure 15 rather than the simpler visualization that is currently available (see Figure 6). This addition is possible because Next-TELL keeps a record of each assessment that is performed for a learner which allows us to determine the statistical variance and consistency in the learner's performance (i.e., a type of uncertainty – see Section 3).

Figure 16 shows what Next-TELL tracks and how it calculates its model of the learner's competency level; this evidence screen shows the results of each assessment that has been performed. The data on which this screen relies allows for the calculation of the learner's current knowledge or competency level as well as the calculation of the uncertainty associated with that competency level. In Figure 16, this uncertainty is based on the consistency of a learner's actions. Since this visualization shows both the learner's mastery of competencies and the relationship between those competencies through the connectedness variable and we have chosen to display information

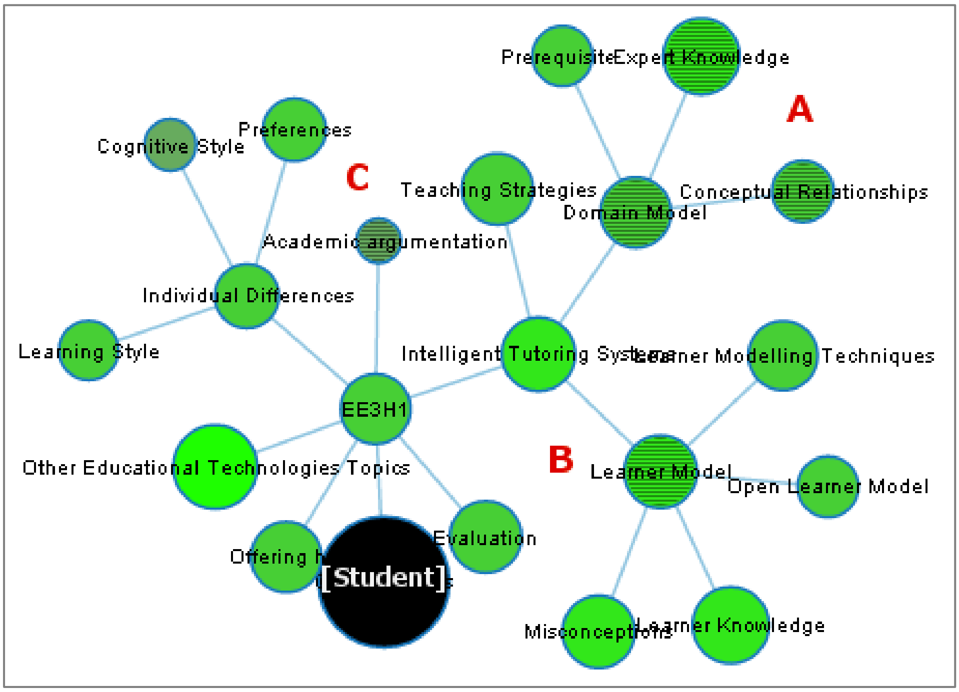

Figure 15. The Next-TELL Competency Network: it has been augmented using the grain visual variable to show the presence of uncertainty. A, B, and C show three different uncertainty propagation situations.

about the uncertainty of a student's mastery of a competency, the use of closure or curvature on the connections between competencies would be inappropriate. The existing visualization manipulated colour values to represent aspects of student mastery so hue, saturation, and value cannot be used without risking student confusion. Likewise, opacity cannot be used because it can be easily confused with saturation when background visuals are absent, as is the case in the Next-TELL competency network. We, therefore, chose to add grain since it was available for use and could be added to the visualization without overloading the semantics of the visual variables that were already being used.

Figure 15 illustrates how uncertainty at lower nodes can propagate up to parent nodes. In the case of the subgraph near A, the uncertainty that is present in the Expert Knowledge and Conceptual Relationships leaves is high enough that the level of uncertainty that is present at their parent node (i.e., Domain Model) also exceeds the acceptable threshold and is communicated by adding the grain to the Domain Model node. B shows a case where the uncertainty level at the parent node (i.e., Learner Model) exceeds that of its children (i.e., Misconceptions, Learner Knowledge, Open Learner Model, and Learner Modelling Techniques). This could be the result of the combined uncertainty of the nodes exceeding the threshold that is used for deciding when to visualize uncertainty even though it is not exceeded at the lower levels. Alternatively, it could happen because there is inconsistency in the evidence that is only associated with the Learner Model node. The last situation, C, is when the uncertainty associated with a child node (i.e., Academic Argumentation) exceeds the threshold and is visualized but the uncertainty associated with its parent node (i.e., EE3H1) does not. These situations can happen because the evidence that is used to calculate the learner's level of competency for a node includes both the evidence used to determine the child nodes' competency levels and the evidence that is only associated with the competency of the node itself. This means that the level of certainty associated with a leaf node does not necessarily result in the leaf's parent having the same level of associated uncertainty.

We also add uncertainty information to the screen where learners can inspect the modelling process (Figure 16) on which the visual representations that are seen in Figures 5, 6, and 15 are based. To ensure consistency between visualizations, grain was used to indicate uncertainty in the higher-level skill-meter visualization and hue was used to highlight the elements that contributed to the inconsistency. This use of hue emulates brushing to

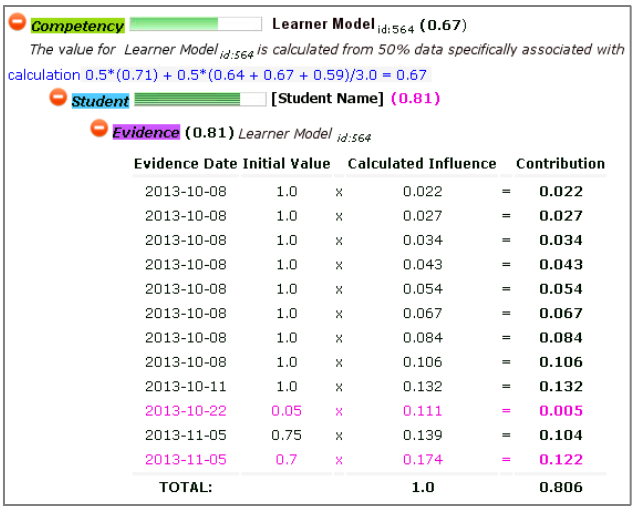

Competency Learner Model $_{id:564}$ (0.67)

The value for Learner Model $_{id:564}$ is calculated from 50% data specifically associated with

calculation 0.5*(0.71) + 0.5*(0.64 + 0.67 + 0.59)/3.0 = 0.67

Student [Student Name] (0.81)

Evidence (0.81) Learner Model $_{id:564}$

Evidence Date	Initial Value		Calculated Influence		Contribution
2013-10-08	1.0	x	0.022	=	0.022
2013-10-08	1.0	x	0.027	=	0.027
2013-10-08	1.0	x	0.034	=	0.034
2013-10-08	1.0	x	0.043	=	0.043
2013-10-08	1.0	x	0.054	=	0.054
2013-10-08	1.0	x	0.067	=	0.067
2013-10-08	1.0	x	0.084	=	0.084
2013-10-08	1.0	x	0.106	=	0.106
2013-10-11	1.0	x	0.132	=	0.132
2013-10-22	0.05	x	0.111	=	0.005
2013-11-05	0.75	x	0.139	=	0.104
2013-11-05	0.7	x	0.174	=	0.122
TOTAL:			1.0		0.806

Figure 16. Next-TELL evidence screen. Employing brushing draws attention to the pieces of evidence that are contributing to model uncertainty.

indicate the evidence (i.e., 2013-10-22 and 2013-11-05) that is associated with the uncertainty that is present in the student's skill meter. This was possible because hue is not used to communicate level of mastery in this view of the information and because hue is only being used to highlight the presence of uncertainty at a binary level.

6 Conclusion

This paper maps the use of a set of visual variables within existing visualizations of uncertain information and within visualizations of educational information. Fifty different visualizations from various disciplines were analyzed with respect to their use of 21 variables that can be manipulated when visualizing information. This analysis along with a review of the literature revealed which of the 21 variables are known to communicate uncertainty successfully, which hinder a user's understanding of uncertainty information, and which of the variables require further study in order to understand their influence on the user's ability to understand the visual representation of uncertainty information.

After describing how other disciplines have represented uncertainty, an analysis of 106 visualizations of educational data that can be seen by learners was performed. This revealed the extent to which uncertainty has been represented within educational reporting visualizations. The minimal representation of uncertainty in these visualizations was perhaps surprising given that assessment, data aggregation, the modelling of learner knowledge and abilities, and the visualization of this information all contain some element of uncertainty. The analysis describes how the same set of 21 visual variables is used with respect to the main communication goals of the 106 visualizations of educational data that were evaluated.

Unlike the visualizations of educational data, the visualizations of uncertainty from other domains were rarely used to communicate social cues; the relationship between items; or the activity, mastery, affect, or interests of a user. The uncertainty visualizations from other domains were typically used to communicate the potential changes in a variable that were due to accuracy, precision, lineage, currency, or statistical variance and consistency. The visualizations of educational data often accounted for these types of uncertainty even when they were not represented. The evaluated learning dashboards and open learner models also communicated information that was related to judgment or completeness, which was not typically observed in other domains.

The results of both analyses were combined to identify design opportunities for the inclusion of uncertainty information within the visualizations of educational data that are used for monitoring and reporting purposes. Additional considerations that are based on the current use of visual variables and their previous study in limited contexts are provided alongside these design opportunities. After describing these design opportunities, we illustrated how some might be exploited by applying them to several types or classes of visual representation that are used within existing visualizations of educational data. We further illustrated the design space by applying it to the Next-TELL open learner model. This example and the exploration of the design space demonstrate how uncertainty information could be incorporated into current visualizations and it provides a starting point for further exploration.

When exploring this design space, it is important to keep the learner's numeracy, literacy, previous training, and other abilities in mind since this affects their ability to interpret visualizations. It is also worth remembering that using some visual variables (e.g., hue, value, grain, motion, closure, arrangement, and curvature) for communicating uncertainty information can require additional support or training even though these variables hold the potential to communicate uncertainty effectively.

By exploiting this space, we can provide learners with additional information that can help them with their monitoring or decision-making tasks. Furthermore, the inclusion of uncertainty information in visualizations of educational data may influence the user's trust (either positively or negatively) of the educational report or technology enhanced learning environment. As a result, the inclusion of uncertainty information should be purposeful and its use closely monitored to ensure that it does not negatively affect learners by leading to confusion, distraction, or sub-optimal decision making.

Acknowledgment

The first (corresponding) author held a W. Garfield Weston Fellowship and a Walter C. Sumner Memorial Fellowship. We would like to thank the creators of several systems for giving us permission to use screenshots of their visualizations: Jim Hewitt, Judy Kay, Tanja Mitrovic, Eric Duval, Julita Vassileva, and Zhi-Hong Chen. Timothy Rias created the image used to represent the depth visual variable and Figure 4 builds on the work of EpochFail. Mouse by Amos Kofi Commey and Doughnut by Ema Dimitrova from the Noun Project were also used.

References

[1] K. Brodlie, R. Allendes Osorio, and A. Lopes, "A Review of Uncertainty in Data Visualization," in *Expanding the Frontiers of Visual Analytics and Visualization*, J. Dill, R. Earnshaw, D. Kasik, J. Vince, and P. C. Wong, Eds. Springer London, 2012, pp. 81–109.

[2] R. Zwick, D. Zapata-Rivera, and M. Hegarty, "Comparing Graphical and Verbal Representations of Measurement Error In Test Score Reports," *Educ. Assess.*, vol. 19, no. 2, pp. 116–138, Apr. 2014.

[3] G. E. Newman and B. J. Scholl, "Bar graphs depicting averages are perceptually misinterpreted: The within-the-bar bias," *Psychon. Bull. Rev.*, vol. 19, no. 4, pp. 601–607, May 2012.

[4] P. F. Fisher, "Visualizing Uncertainty in Soil Maps by Animation," *Cartogr. Int. J. Geogr. Inf. Geovisualization*, vol. 30, no. 2, pp. 20–27, Oct. 1993.

[5] S. Belia, F. Fidler, J. Williams, and G. Cumming, "Researchers Misunderstand Confidence Intervals and Standard Error Bars," *Psychol. Methods*, vol. 10, no. 4, pp. 389–396, 2005.

[6] M. Skeels, B. Lee, G. Smith, and G. Robertson, “Revealing Uncertainty for Information Visualization,” in *Proceedings of the Working Conference on Advanced Visual Interfaces*, New York, NY, USA, 2008, pp. 376–379.

[7] H. Wainer, “Depicting Error,” *Am. Stat.*, vol. 50, no. 2, p. 101, May 1996.

[8] S. Basapur, A. M. Bisantz, and T. Kesavadas, “The Effect of Display Modality on Decision-Making with Uncertainty,” *Proc. Hum. Factors Ergon. Soc. Annu. Meet.*, vol. 47, no. 3, pp. 558–561, Oct. 2003.

[9] A. M. Bisantz, D. Cao, M. Jenkins, P. R. Pennathur, M. Farry, E. Roth, S. S. Potter, and J. Pfautz, “Comparing Uncertainty Visualizations for a Dynamic Decision-Making Task,” *J. Cogn. Eng. Decis. Mak.*, vol. 5, no. 3, pp. 277–293, Aug. 2011.

[10] D. Spiegelhalter, M. Pearson, and I. Short, “Visualizing Uncertainty About the Future,” *Science*, vol. 333, no. 6048, pp. 1393–1400, Sep. 2011.

[11] J. Wood, P. Isenberg, T. Isenberg, J. Dykes, N. Boukhelifa, and A. Slingsby, “Sketchy Rendering for Information Visualization,” *IEEE Trans. Vis. Comput. Graph.*, vol. 18, no. 12, pp. 2749–2758, Dec. 2012.

[12] C. Demmans Epp, S. Bull, and M. D. Johnson, “Visualising Uncertainty for Open Learner Model Users,” in *CEUR Proceedings of User Modeling, Adaptation and Personalization (UMAP)*, Aalborg, Denmark, 2014, vol. 1181, pp. 9–12.

[13] J. D. Zapata-Rivera, J. S. Underwood, and M. Bauer, “Advanced Reporting Systems in Assessment Environments,” in *Workshop on Learner Modelling for Reflection, International Conference on Artificial Intelligence in Education (AIED)*, Amsterdam, Netherlands, 2005, pp. 23–31.

[14] S. Bull and J. Kay, “Chapter 15: Open Learner Models,” in *Advances in Intelligent Tutoring Systems*, R. Nkambou, J. Bourdeau, and R. Mizoguchi, Eds. Springer, 2010, pp. 301–322.

[15] E. Duval, J. Klerkx, K. Verbert, T. Nagel, S. Govaerts, G. A. Parra Chico, J. L. Santos Odriozola, and B. Vandeputte, “Learning Dashboards & Learnscapes,” in *Workshop on Educational Interfaces, Software, and Technology at the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*, Austin, TX, USA, 2012, pp. 1–5.

[16] A. F. Wise, “Designing Pedagogical Interventions to Support Student Use of Learning Analytics,” in *Proceedings of the Fourth International Conference on Learning Analytics And Knowledge*, New York, NY, USA, 2014, pp. 203–211.

[17] S. Bull and A. T. McEvoy, “An Intelligent Learning Environment with an Open Learner Model for the Desktop PC and Pocket PC,” in *International Conference on Artificial Intelligence in Education*, 2003, pp. 389–391.

[18] E. Duval, “Attention Please!: Learning Analytics for Visualization and Recommendation,” in *Proceedings of the 1st International Conference on Learning Analytics and Knowledge*, New York, NY, USA, 2011, pp. 9–17.

[19] D. Edge, E. Searle, K. Chiu, J. Zhao, and J. A. Landay, “MicroMandarin: Mobile Language Learning in Context,” in *Conference on Human Factors in Computing Systems (CHI)*, Vancouver, Canada, 2011, pp. 3169–3178.

[20] C. Hansen and G. McCalla, “Active Open Learner Modelling,” in *Workshop on Learner Modelling for Reflection (Supplemental Proceedings) at the 11th International Conference on Artificial Intelligence in Education (AIED)*, 2003, vol. 5, pp. 248–257.

[21] J. Hewitt, C. Brett, and K. MacKinnon, “A Study of Private Messaging Within an Asynchronous Discussion Environment,” in *To See the World and a Grain of Sand: Learning across Levels of Space, Time and Scale: CSCL 2013 Conference*, Sydney, Australia, 2013, vol. 2, pp. 46–49.

[22] K. Upton and J. Kay, “Narcissus: Group and Individual Models to Support Small Group Work,” in *User Modeling, Adaptation, and Personalization (UMAP)*, 2009, pp. 54–65.

[23] J. Vassileva and L. Sun, “Evolving a Social Visualization Design Aimed at Increasing Participation in a Class-Based Online Community,” *Int. J. Coop. Inf. Syst.*, vol. 17, no. 04, pp. 443–466, Dec. 2008.

[24] S. Bull and J. Kay, “Student Models that Invite the Learner in: the SMILI:) Open Learner Modelling Framework,” *Int. J. Artif. Intell. Educ. IJAIED*, vol. 17, no. 2, pp. 89–120, 2007.

[25] L. Bastin, P. F. Fisher, and J. Wood, “Visualizing uncertainty in multi-spectral remotely sensed imagery,” *Comput. Geosci.*, vol. 28, no. 3, pp. 337–350, Apr. 2002.

[26] J. Bertin, *Semiology of graphics*. Madison, WI: University of Wisconsin Press, 1983.

[27] C. Ware, *Information visualization: perception for design*. San Francisco, CA: Morgan Kaufman, 2004.

[28] C. Chen, “Information visualization,” *Wiley Interdiscip. Rev. Comput. Stat.*, vol. 2, no. 4, pp. 387–403, Jul. 2010.

[29] D. J. Greenberg and S. Z. Blue, “Visual Complexity in Infancy: Contour or Numerosity?,” *Child Dev.*, vol. 46, no. 2, p. 357, Jun. 1975.

[30] E. R. Tufte, *The visual display of quantitative information*. Cheshire, CT, USA: Graphics Press, 1999.

[31] S. Bateman, R. L. Mandryk, C. Gutwin, A. Genest, D. McDine, and C. Brooks, “Useful Junk?: The Effects of Visual Embellishment on Comprehension and Memorability of Charts,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2010, pp. 2573–2582.

[32] C. M. Wittenbrink, A. T. Pang, and S. K. Lodha, “Glyphs for visualizing uncertainty in vector fields,” *IEEE Trans. Vis. Comput. Graph.*, vol. 2, no. 3, pp. 266–279, 1996.

[33] T. D. Zuk, “Visualizing Uncertainty,” Thesis, University of Calgary, 2008.

[34] T. Gebuis and B. Reynvoet, “The Role of Visual Information in Numerosity Estimation,” *PLoS ONE*, vol. 7, no. 5, p. e37426, May 2012.

[35] A. M. Bisantz, S. S. Marsiglio, and J. Munch, “Displaying uncertainty: investigating the effects of display format and specificity,” *Hum. Factors*, vol. 47, no. 4, pp. 777–796, 2005.

[36] N. Boukhelifa, A. Bezerianos, T. Isenberg, and J.-D. Fekete, “Evaluating Sketchiness as a Visual Variable for the Depiction of Qualitative Uncertainty,” *IEEE Trans. Vis. Comput. Graph.*, vol. 18, no. 12, pp. 2769–2778, Dec. 2012.

[37] C. Collins, S. Carpendale, and G. Penn, “Visualization of Uncertainty in Lattices to Support Decision-making,” in *Proceedings of the 9th Joint Eurographics / IEEE VGTC Conference on Visualization*, Aire-la-Ville, Switzerland, Switzerland, 2007, pp. 51–58.

[38] R. A. Becker and W. S. Cleveland, “Brushing Scatterplots,” *Technometrics*, vol. 29, no. 2, p. 127, May 1987.

[39] S. Bateman, C. Gutwin, and M. Nacenta, "Seeing Things in the Clouds: The Effect of Visual Features on Tag Cloud Selections," in *Proceedings of the Nineteenth ACM Conference on Hypertext and Hypermedia*, New York, NY, USA, 2008, pp. 193–202.

[40] G. Siemens, "Learning and Academic Analytics," *Learning and Knowledge Analytics*, 05-Aug-2011. [Online]. Available: http://www.learninganalytics.net/?p=131. [Accessed: 16-Jan-2012].

[41] V. Kodaganallur, R. R. Weitz, and D. Rosenthal, "A Comparison of Model-Tracing and Constraint-Based Intelligent Tutoring Paradigms," *Int. J. Artif. Intell. Educ. IJAIED*, vol. 15, no. 2, pp. 117–144, 2005.

[42] G. McCalla, "The Ecological Approach to the Design of E-Learning Environments: Purpose-based Capture and Use of Information About Learners," *J. Interact. Media Educ.*, pp. 1–23, 2004.

[43] S. Ohlsson, "Constraint-Based Student Modelling," in *Student Modelling: The Key to Individualized Knowledge-Based Instruction*, J. Greer and G. McCalla, Eds. Springer-Verlag, 1993, pp. 165–189.

[44] J. A. Self, "Bypassing the Intractable Problem of Student Modelling," *Artif. Intell.*, vol. 41, pp. 1–26, 1990.

[45] K. VanLehn, R. Freedman, P. Jordan, C. Murray, R. Osan, M. Ringenberg, C. Rosé, K. Schulze, R. Shelby, D. Treacy, A. Weinstein, and M. Wintersgill, "Fading and Deepening: The Next Steps for Andes and Other Model-Tracing Tutors," in *Intelligent Tutoring Systems*, G. Gauthier, C. Frasson, and K. VanLehn, Eds. Springer Berlin Heidelberg, 2000, pp. 474–483.

[46] G. Weber, "Episodic Learner Modeling," *Cogn. Sci.*, vol. 20, no. 2, pp. 195–236, 1996.

[47] K. Chrysafiadi and M. Virvou, "Student modeling approaches: A literature review for the last decade," *Expert Syst. Appl.*, vol. 40, no. 11, pp. 4715–4729, Sep. 2013.

[48] C. Demmans Epp and G. I. McCalla, "ProTutor: Historic Open Learner Models for Pronunciation Tutoring," in *Artificial Intelligence in Education (AIED)*, Auckland, New Zealand, 2011, vol. 6738, pp. 441–443.

[49] W. L. Johnson, S. Marsella, and H. Vilhjálmsson, "The DARWARS Tactical Language Training System," in *Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)*, Orlando, 2004.

[50] A. Mitrovic, "SQL-Tutor: an ITS for SQL programming," *CAPIT - University of Canterbury*, 2007. [Online]. Available: http://www.cosc.canterbury.ac.nz/tanja.mitrovic/sql-tutor.html. [Accessed: 03-Mar-2013].

[51] R. S. J. D. Baker and K. Yacef, "The State of Educational Data Mining in 2009: a Review and Future Visions," *J. Educ. Data Min.*, vol. 1, no. 1, pp. 3–17, 2009.

[52] D. Barrow, A. Mitrovic, S. Ohlsson, and M. Grimley, "Assessing the Impact of Positive Feedback in Constraint-Based Tutors," in *Intelligent Tutoring Systems (ITS)*, Montreal, Canada, 2008, pp. 250–259.

[53] M. Mathews and A. Mitrovic, "Investigating the effectiveness of problem templates on learning in intelligent tutoring systems," *BSc Honors Rep. Univ. Canterb. Dep. Comput. Sci. Softw. Eng. Christch. N. Z.*, 2006.

[54] A. Mitrovic, K. R. Koedinger, and B. Martin, "A Comparative Analysis of Cognitive Tutoring and Constraint-Based Modelling," in *9th International Conference on User Modeling (UM)*, Johnstown, PA, USA, 2003, pp. 313–322.

[55] K. VanLehn, "The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems," *Educ. Psychol.*, vol. 46, no. 4, pp. 197–221, 2011.

[56] A. Corbett, A. Wagner, S. Lesgold, H. Ulrich, and S. Stevens, "The impact on learning of generating vs. selecting descriptions in analyzing algebra example solutions," in *Proceedings of the 7th international conference on Learning sciences*, 2006, pp. 99–105.

[57] S. Govaerts, K. Verbert, J. Klerkx, and E. Duval, "Visualizing Activities for Self-reflection and Awareness," in *Advances in Web-Based Learning – ICWL 2010*, Berlin, Heidelberg, 2010, vol. 6483, pp. 91–100.

[58] A. Paiva, J. A. Self, and R. Hartley, "Externalising Learner Models," in *World Conference on Artificial Intelligence in Education (AIED)*, Washington, D.C., 1995, pp. 509–516.

[59] J. Kay, "A Scrutable User Modelling Shell for User-Adapted Interaction," PhD, University of Sydney, 1999.

[60] A. Mitrovic and B. Martin, "Evaluating the Effect of Open Student Models on Self-Assessment," *Int. J. Artif. Intell. Educ. IJAIED*, vol. 17, no. 2, pp. 121–144, Jan. 2007.

[61] *Blackboard*. Washington, D.C.: Blackboard Inc.

[62] *Moodle*. Perth, Australia: Moodle Trust.

[63] N. Ahmad and S. Bull, "Do Students Trust Their Open Learner Models?," in *Adaptive Hypermedia and Adaptive Web-Based Systems*, 2008, pp. 255–258.

[64] J. B. Uther, "On the Visualisation of Large User Models in Web Based Systems," PhD, University of Sydney, Sydney Australia, 2001.

[65] J. D. Zapata-Rivera and J. Greer, "Interacting with inspectable bayesian student models," *Int. J. Artif. Intell. Educ.*, vol. 14, no. 2, pp. 127–163, 2004.

[66] A. Paramythis, S. Weibelzahl, and J. Masthoff, "Layered Evaluation of Interactive Adaptive Systems: Framework and Formative Methods," *User Model. User-Adapt. Interact.*, vol. 20, no. 5, pp. 383–453, Nov. 2010.

[67] K. Potter, P. Rosen, and C. R. Johnson, "From Quantification to Visualization: A Taxonomy of Uncertainty Visualization Approaches," in *Uncertainty Quantification in Scientific Computing*, A. M. Dienstfrey and R. F. Boisvert, Eds. Springer Berlin Heidelberg, 2012, pp. 226–249.

[68] C. Correa, Y.-H. Chan, and K.-L. Ma, "A framework for uncertainty-aware visual analytics," in *IEEE Symposium on Visual Analytics Science and Technology, 2009. VAST 2009*, 2009, pp. 51–58.

[69] A. M. MacEachren, R. E. Roth, J. O'Brien, B. Li, D. Swingley, and M. Gahegan, "Visual Semiotics & Uncertainty Visualization: An Empirical Study," *IEEE Trans. Vis. Comput. Graph.*, vol. 18, no. 12, pp. 2496–2505, 2012.

[70] K. Silius, T. Miilumaki, J. Huhtamaki, T. Tebest, J. Merilainen, and S. Pohjolainen, "Students' Motivations for Social Media Enhanced Studying and Learning," *Knowl. Manag. E-Learn. Int. J. KMEL*, vol. 2, no. 1, pp. 51–67, Feb. 2010.

[71] J. Thomson, E. Hetzler, A. MacEachren, M. Gahegan, and M. Pavel, "A typology for visualizing uncertainty," 2005, pp. 146–157.

[72] N. Gershon, “Visualization of an Imperfect World,” *IEEE Comput Graph Appl*, vol. 18, no. 4, pp. 43–45, Jul. 1998.

[73] S. K. Lodha, A. Pang, R. E. Sheehan, and C. M. Wittenbrink, “UFLOW: visualizing uncertainty in fluid flow,” in *Visualization ’96. Proceedings.*, 1996, pp. 249–254.

[74] S. Bull, M. D. Johnson, C. Demmans Epp, D. Masci, M. Alotaibi, and S. Girard, “Formative Assessment and Meaningful Learning Analytics – an Independent Open Learner Model Solution,” in *IEEE International Conference on Advanced Learning Technologies (ICALT)*, Athens, Greece, 2014, pp. 327–329.

[75] S. Deitrick and R. Edsall, “The Influence of Uncertainty Visualization on Decision Making: An Empirical Evaluation,” in *Progress in Spatial Data Handling*, D. A. Riedl, P. W. Kainz, and P. G. A. Elmes, Eds. Springer Berlin Heidelberg, 2006, pp. 719–738.

[76] C. R. Johnson and A. R. Sanderson, “A Next Step: Visualizing Errors and Uncertainty,” *IEEE Computer Graphics and Applications*, pp. 6–10, 2003.

[77] C. Munteanu, J. Lumsden, H. Fournier, R. Leung, D. D’Amours, D. McDonald, and J. Maitland, “ALEX: Mobile Language Assistant for Low-Literacy Adults,” in *Proc. International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI)*, Lisbon, Portugal, 2010, pp. 427–430.

[78] I.-H. Hsiao, F. Bakalov, P. Brusilovsky, and B. König-Ries, “Open Social Student Modeling: Visualizing Student Models with Parallel Introspective Views,” in *User Modeling, Adaption and Personalization (UMAP)*, J. A. Konstan, R. Conejo, J. L. Marzo, and N. Oliver, Eds. Springer Berlin Heidelberg, 2011, pp. 171–182.

[79] Z.-H. Chen, C.-Y. Chou, Y.-C. Deng, and T.-W. Chan, “Active Open Learner Models As Animal Companions: Motivating Children to Learn Through Interacting with My-Pet and Our-Pet,” *Int J Artif Intell Ed*, vol. 17, no. 2, pp. 145–167, Apr. 2007.

[80] S. Bull, M. Mangat, A. Mabbott, A. S. Abu Issa, and J. Marsh, “Reactions to Inspectable Learner Models: Seven Year Olds to University Students,” in *Proceedings of Workshop on Learner Modelling for Reflection, International Conference on Artificial Intelligence in Education 2005*, Amsterdam, Netherlands, 2005, pp. 1–10.

[81] S. J. H. Lee and S. Bull, “An Open Learner Model to Help Parents Help their Children,” *Technol. Instr. Cogn. Learn.*, vol. 6, no. 1, pp. 29–51, 2008.

[82] T. Lloyd and S. Bull, “A haptic learner model,” *Int. J. Contin. Eng. Educ. Life Long Learn.*, vol. 16, no. 1, pp. 137–149, Jan. 2006.

[83] S. Bull, M. D. Johnson, M. Alotaibi, W. Byrne, and G. Cierniak, “Visualising Multiple Data Sources in an Independent Open Learner Model,” in *Artificial Intelligence in Education*, 2013, pp. 199–208.

[84] N. Ahmad and S. Bull, “Learner Trust in Learner Model Externalisations,” in *Artificial Intelligence in Education (AIED)*, 2009, pp. 617–619.

[85] N. Van Labeke, P. Brna, and R. Morales, “Opening up the interpretation process in an open learner model,” *Int. J. Artif. Intell. Educ.*, vol. 17, no. 3, pp. 305–338, 2007.

[86] J. Kay and B. Kummerfeld, “Creating Personalized Systems That People Can Scrutinize and Control: Drivers, Principles and Experience,” *ACM Trans. Interact. Intell. Syst.*, vol. 2, no. 4, pp. 24:1–24:42, Jan. 2013.

[87] S. Bull and H. Pain, “Did I say what I think I said, and do you agree with me?’: Inspecting and questioning the student model,” in *World Conference on Artificial Intelligence in Education (AIED)*, Washington, DC, 1995, pp. 501–508.

[88] S. Mohanarajah, R. Kemp, and E. Kemp, “Opening a Fuzzy Learner Model,” in *Workshop on Learner Modelling for Reflection at the International Conference on Artificial Intelligence in Education (AIED)*, 2005, pp. 62–71.

[89] J. Kay and A. Lum, “Exploiting Readily Available Web Data for Scrutable Student Models,” in *Artificial Intelligence in Education (AIED)*, Amsterdam, Netherlands, 2005, pp. 338–345.

[90] J. Burston, “BonPatron: An Online Spelling, Grammar, and Expression Checker,” *Comput. Assist. Lang. Instr. Consort. CALICO J.*, vol. 25, no. 2, pp. 337–347, 2008.

[91] A. Mabbott and S. Bull, “Student Preferences for Editing, Persuading, and Negotiating the Open Learner Model,” in *Intelligent Tutoring Systems*, 2006, pp. 481–490.

[92] S. Bull and M. Britland, “Group Interaction Prompted by a Simple Assessed Open Learner Model that can be Optionally Released,” in *PING Workshop at User Modeling (UM)*, 2007.

[93] M. Correll and M. Gleicher, “Error Bars Considered Harmful: Exploring Alternate Encodings for Mean and Error,” *IEEE Trans. Vis. Comput. Graph.*, vol. 20, no. 12, pp. 2142–2151, 2014.

[94] G. Cumming, F. Fidler, and D. L. Vaux, “Error bars in experimental biology,” *J. Cell Biol.*, vol. 177, no. 1, pp. 7–11, Apr. 2007.

Carrie DEMMANS EPP has a BSc with honors in Computer Science and Russian (2007) and an MSc in computer science with a focus on open learner models within computer assisted language learning (2010). Her PhD work focuses on the use of mobile devices for supporting language learning.

Susan BULL received her PhD in Artificial Intelligence in Education from the University of Edinburgh in 1997. She is now based at the University of Birmingham, where her research focuses on open learner models, visual and learning analytics, and adaptive systems. She has authored over 100 academic papers in these and related areas.