Tech Know Learn (2014) 19:205–220
DOI 10.1007/s10758-014-9223-7

INTEGRATIVE REVIEW

Educational Data Mining and Learning Analytics: Applications to Constructionist Research

Matthew Berland · Ryan S. Baker · Paulo Blikstein

Published online: 3 May 2014

Abstract Constructionism can be a powerful framework for teachingcomplex content to novices. At the core of constructionism is the suggestion that by enabling learners to build creative artifacts that require complex content to function, those learners will have opportunities to learn this content in contextualized, personally meaningful ways. In this paper, we investigate the relevance of a set of approaches broadly called "educational data mining" or "learning analytics" (henceforth, EDM) to help provide a basis for quantitative research on constructionist learning which does not abandon the richness seen as essential by many researchers in that paradigm. We suggest that EDM may have the potential to support research that is meaningful and useful both to researchers working actively in the constructionist tradition but also to wider communities. Finally, we explore potential collaborations between researchers in the EDM and constructionist traditions; such collaborations have the potential to enhance the ability of constructionist researchers to make rich inferences about learning and learners, while providing EDM researchers with many interesting new research questions and challenges.

Keywords Constructionism · Educational data mining · Learning analytics · Design of learning environments · Project-based learning

M. Berland (✉)
Department of Curriculum and Instruction, University of Wisconsin–Madison, Madison, WI, USA
e-mail: mberland@wisc.edu

R. S. Baker
Teacher's College, Columbia University, New York, NY, USA
e-mail: baker2@exchange.tc.columbia.edu

P. Blikstein
Graduate School of Education, Stanford University, Stanford, CA, USA
e-mail: paulob@stanford.edu

In recent years, project-based, student-centered approaches to education have gained prominence, due in part to an increased demand for higher-level skills in the job market (Levy and Murnane 2005), positive research findings on the effectiveness of such approaches (Baker et al. 2008), and a broader acceptance in public policy circles, as shown, for example, by the Next Generation Science Standards (NGSS Lead States 2013). While several approaches for this type of learning exist, Constructionism is one of the most popular and well-developed ones (Papert 1980). In this paper, we investigate the relevance of a set of approaches called "educational data mining" or "learning analytics" (henceforth abbreviated as 'EDM') (Baker and Yacef 2009; Romero and Ventura 2010; Baker and Siemens in press) to help provide a basis for quantitative research on constructionist learning which does not abandon the richness seen as essential by many researchers in that paradigm. As such, EDM may have the potential to support research that is meaningful and useful both to researchers working actively in the constructionist tradition and to the wider community of learning scientists and policymakers. EDM, broadly, is a set of methods that apply data mining and machine learning techniques such as prediction, classification, and discovery of latent structural regularities to rich, voluminous, and idiosyncratic educational data, potentially similar to those data generated by many constructionist learning environments which allows students to explore and build their own artifacts, computer programs, and media pieces. As such, we identify four axes in which EDM methods may be helpful for constructionist research:

1. EDM methods do not require constructionists to abandon deep qualitative analysis for simplistic summative or confirmatory quantitative analysis;
2. EDM methods can generate different and complementary new analyses to support qualitative research;
3. By enabling precise formative assessments of complex constructs, EDM methods can support an increase in methodological rigor and replicability;
4. EDM can be used to present comprehensible and actionable data to learners and teachers in situ.

In order to investigate those axes, we start by describing our perspective on compatibilities and incompatibilities between constructionism and EDM.

At the core of constructionism is the suggestion that by enabling learners to *build* creative artifacts that require complex content to function, those learners will have opportunities to learn that complex content in connected, meaningful ways. Constructionist projects often emphasize making those artifacts (and often data) public, socially relevant, and personally meaningful to learners, and encourage working in social spaces such that learners engage each other to accelerate the learning process. diSessa and Cobb (2004) argue that constructionism serves a *framework for action*, as it describes its own praxis (i.e., how it matches theory to practice). The learning theory supporting constructionism is classically constructivist, combining concepts from Piaget and Vygotsky (Fosnot 2005). As constructionism matures as a constructivist framework for action and expands in scale, constructionist projects are becoming both more complex (Reynolds and Caperton 2011), more scalable (Resnick et al. 2009), and more affordable for schools following significant development in low cost "construction" technologies such as robotics and 3D printers. As such, there have been increasing opportunities to learn more about how students learn in constructionist contexts, advancing the science of learning. These discoveries will have the potential to improve the quality of all constructivist learning experiences. For example, Wilensky and Reisman (2006) have shown how constructionist modeling and simulation can make science learning more accessible, Resnick (1998) has shown how constructionism can reframe programming as art at scale, Buechley and Eisenberg (2008) have used e-textiles to engage female students in robotics, Eisenberg (2011) and Blikstein (2013a, b, 2014) use constructionist digital fabrication to successfully teach programming, engineering, and electronics in a novel, integrated way. The findings of these research and design projects have the potential to be useful to a wide external community of teachers, researchers, practitioners, and other stakeholders. However, connecting findings from the constructionist tradition to the goals of policymakers can be challenging, due to the historical differences in methodology and values between these communities. The resources needed to study such interventions at scale are considerable, given the need to carefully document, code, and analyze each student's work processes and artifacts. The designs of constructionist research often result in findings that do not map to what researchers, outside interests, and policymakers are expecting, in contrast to conventional controlled studies, which are designed to (more conclusively) answer a limited set of sharply targeted research questions. Due to the lack of a common ground to discuss benefits and scalability of constructionist and project-based designs, these designs have been too frequently sidelined to niche institutions such as private schools, museums, or atypical public schools.

To understand what role EDM methods can play in constructionist research, we must frame what we mean by constructionist research more precisely. We follow Papert and Harel (1991) in their situating of constructionism, but they do not constrain the term to one formal definition. The definition is further complicated by the fact that constructionism has many overlaps with other research and design traditions, such as constructivism and socio-constructivism themselves, as well as project-based pedagogies and inquiry-based designs. However, we believe that it is possible to define the subset of constructionism amenable to EDM, a focus we adopt in this article for brevity. In this paper, we focus on the constructionist literature dealing with students learning to construct understandings by constructing (physical or virtual) artifacts, where the students' learning environments are designed and constrained such that building artifacts in/with that environment is designed to help students construct their own understandings. In other words, we are focusing on creative work done in computational environments designed to foster creative and transformational learning, such as NetLogo (Wilensky 1999), Scratch (Resnick et al. 2009), or LEGO Mindstorms.

This sub-category of constructionism can and does generate considerable formative and summative data. It also has the benefit of having a history of success in the classroom. From Papert's seminal (1972) work through today, constructionist learning has been shown to promote the development of deep understanding of relatively complex content, with many examples ranging from mathematics (Harel 1990; Wilensky 1996) to history (Zahn et al. 2010).

However, constructionist learning environments, ideas, and findings have yet to reach the majority of classrooms and have had incomplete influence in the broader education research community. There are several potential reasons for this. One of them may be a lack of demonstration that findings are generalizable across populations and across specific content. Another reason is that constructionist activities are seen to be time-consuming for teachers (Warschauer and Matuchniak 2010), though, in practice, it has been shown that supporting understanding through project-based work could actually save time (Fosnot 2005) and enable classroom dynamics that may streamline class preparation (e.g., peer teaching or peer feedback). A last reason is that constructionists almost universally value more deep understanding of scientific principles than facts or procedural skills even in contexts (e.g., many classrooms) in which memorization of facts and procedural skills is the target to be evaluated (Abelson and diSessa 1986; Papert and Harel 1991). Therefore, much of what is learned in constructionist environments does not directly translate to test scores or other established metrics.

Constructionist research can be useful and convincing to audiences that do not yet take full advantage of the scientific findings of this community, but it requires careful consideration of framing and evidence to reach them. Educational data mining methods pose the potential to both enhance constructionist research, and to support constructionist researchers in communicating their findings in a fashion that other researchers consider valid. Blikstein (2011, p. 110) made the argument that “one of the difficulties is that current assessment instruments are based on products […], and not on processes, due to the intrinsic difficulties in capturing detailed process data for large numbers of students. […] However, new data collection, sensing, and data mining technologies […] are enabling researchers to have an unprecedented insight into the minute-by-minute development of several activities.”

By enabling scalable and precise assessments of more complex constructs than can be typically assessed through traditional assessment instruments (such as multiple-choice tests), EDM methods support an increase in methodological rigor and replicability, while maintaining much (though not all) of the richness of qualitative methods. EDM methods do not require constructionists to abandon qualitative and meaningful evaluation for simplistic multiple-choice tests; instead, EDM can add some of the benefits of quantitative work to rich qualitative understanding. Furthermore, EDM has the possibility to generate new understandings of how students learn in constructionist learning environments and how to adapt our environments to those new understandings.

Importantly, EDM provides a powerful set of methods that can be used to present actionable data to learners and teachers, by which we can give learners the tools to help themselves and use their own data.

Though this paper, we will examine that potential in terms of current work in EDM and constructionism, potential research overlaps, and open questions generated by bringing them together.

1 Grading and Assessment

The limitations of traditional tests and assessments are well-known (Baker et al. 2010), but those tests remain standard in most schooling, due to the ease of administration and the perceived need for assessment of student success and teacher quality.

Regarding alternative forms of assessments for constructionist learning, Papert (1980) suggested detailed peer critiques (or *crits*) in an art class or actual use of a student’s tool in an authentic setting can provide meaningful feedback. This is undoubtedly true, but the feedback received in these formats are not very precise and well-defined, and take much longer than other forms of automated feedback (e.g., feedback of a compiler about bugs in the code). There is no reason why broader assessments such as *crits* cannot live alongside more fine-grained assessments such as compiler feedback or the types of process assessments that EDM can generate.

However, EDM can support continual and real-time assessment on student process and progress, in which the amount of formative feedback is radically increased. This allows for faster progress overall (Black and Wiliam 1998; Shute 2008), more opportunity for teacher insight into students’ learning (Roschelle et al. 2005), and can provide a more constructive basis for continual assessment. This is important, as teachers frequently feel challenged in using constructionist tools in public school settings as districts frequently mandate a minimum number of grades per week. This may then unnecessarily impede teachers' incorporation of constructionist practices as they may find it very difficult to grade a large-scale project 2–3 times per week as an artifact, unless the design process is broken down into artificially small subcomponents. Anecdotally, when instructing practicing teachers in constructionist practices, the first author has heard complaints from teachers that they are required to give at least two grades per week per assignment, even in projects spanning weeks or months; the teachers found such assessments to be difficult for projects that required exploration and creativity. Unfortunately, these rules are often a reality in contemporary classrooms, and they can hinder good project-based learning and teaching (Blumenfeld et al. 2000). Fortunately, educational data mining can serve to support teachers in supporting such learning, which owing to professed reasons of practicality is often found only in more affluent schools (Warschauer and Matuchniak 2010), by providing access to more data to support students' progress monitoring and teachers' continual assessment of progress. This is by no means a concrete solution to the problem of overly aggressive assessment, but it may provide the teachers concrete resources to argue against or (at least) nominally comply with the policy.

2 What Educational Data Mining Can Bring to the Table

Some of these goals for increasing the rigor of constructionist research and providing more valid assessment may be achieved by integrating methods from the emerging discipline of educational data mining and learning analytics (EDM). EDM has become a useful method for research in other educational paradigms, with the potential to offer both richness and rigor. EDM has been defined as "an emerging discipline, concerned with developing methods for exploring the unique types of data that come from educational settings, and using those methods to better understand students, and the settings which they learn in" (IEDMS 2009).

EDM typically consists of research to take educational data and apply data mining techniques such as prediction (including classification), discovery of latent structure (such as clustering and q-matrix discovery), relationship mining (such as association rule mining and sequential pattern mining), and discovery with models to understand learning and learner individual differences and choices better (see Baker and Yacef 2009; Romero and Ventura 2010; Baker and Siemens in press, for reviews of these methods in education). Prediction modeling algorithms automatically search through a space of candidate models to find the model which best infers a single predicted variable from some combination of other variables. These models are developed on some set of data, typically validated for their ability to make accurate predictions for new students, but ideally also for new content (cf. Baker et al. 2008)—and new populations of students (cf. Ocumpaugh et al. 2014). As such, developing a prediction model depends on knowing what the predicted variable is for a small set of data; a model is then created for this small set of data, and validated so that it can be applied at greater scale. For instance, one may collect data on whether 140 students demonstrated a scientific inquiry strategy while learning, develop a prediction model to infer whether the inquiry behavior occurred, validate it on sub-sets of the 140 students that were not included when creating the prediction model, and then use the model to make predictions about new students (e.g. Sao Pedro et al. 2010, 2012). As such, prediction models can be used to analyze the development of a student strategy or behavior in a fine-grained fashion, over longitudinal data or many students, in an unobtrusive and non-disruptive way. This allows much (though not all) of the richness of qualitative analysis, while being much more feasible to conduct at scale than qualitative analysis is. As such, it may prove useful for constructionist research, but relatively little work has been done in creating predictive models of creative constructionist learning environments. To date, it has been largely used to model student strategies (Amershi and Conati 2009; Sao Pedro et al. 2010, 2012), student behaviors associated with disengagement (Baker et al. 2008), student emotions (Dragon et al. 2008; D'Mello et al. 2010; Worsley and Blikstein 2011), longer-term student learning (Baker et al. 2011), and participation in future learning (e.g. dropout) (Arnold 2010).

Other EDM methods accomplish different goals, but have the same virtue of enabling analysis of student behavior and learning at scale but in a richer fashion than traditional quantitative methods. For example, cluster analysis finds the structure that emerges naturally from data, allowing researchers to search for patterns in student behavior that commonly occur in data, but which did not initially occur to the researcher. Relationship mining methods (such as sequential pattern mining) find sequences of learner behavior that manifest over time and are seen repeatedly or in many students. In all cases, once a model or finding obtained via data mining is validated to generalize across students and/or contexts, it can be applied at scale and used in discovery with models analyses that leverage models at scale to infer the relationship between (for instance) student behaviors and learning outcomes, or student strategies and evidence on student engagement.

While EDM research has been conducted on a range of different types of educational data, a large proportion of EDM research has involved more restrictive (or, at least, less creative) online learning environments. Early research in EDM often involved very structured learning environments, such as intelligent tutoring systems (cf. Baker et al. 2004; Beck and Woolf 2000; Merceron and Yacef 2004). Data from these structured learning environments was a useful place to start research in EDM, as the structure of the learning environment makes it easier to infer structure in the data. For example, these environments privilege clearly defined 'skills' that map onto student responses, each of which will be clearly and a priori identified as correct or incorrect. That focus makes it easier to accomplish acceptable-quality inference of those defined skills, a task which can be a significant challenge in other types of learning environments. For this reason, data from structured learning environments remains a considerable part of the research literature in EDM.

However, in recent years, EDM research has increasingly involved open-ended online learning environments. In the first issue of the Journal of Educational Data Mining, Amershi and Conati (2009) published an analysis of the strategic behaviors employed by successful and unsuccessful learners in a fully exploratory online learning environment, using cluster analysis to discover patterns in student behavior. In their environment, students explore the workings of a range of common search and other AI algorithms. Amershi and Conati discovered that 'less successful' learners are less likely to pause and self-explain during execution of an algorithm, and after completing algorithm execution. Less successful learners were also less likely to break down domain spaces into sub-spaces. It remains an open question whether this pattern would apply to, say, novices learning the Scratch programming language, and whether design modifications could help those novices better create more substantive artifacts.

In another example of research in a more open-ended online learning environment, Sao Pedro et al. (2010, 2012) analyzed student experimentation behaviors in a physical science simulation environment, as mentioned above. Through a combination of human annotation of log files and the use of prediction modeling to develop automated detectors that could replicate the judgments being made by the human coders, they were able to identify whether students were demonstrating skill in designing sequences of experiments, and infer latent experimentation skill in those students. A third example can be found in work by Lynch et al. (2008) to classify the structure of students' argumentation strategies. They used decision trees—a type of prediction modeling—to infer which attributes of students' argumentation processes in an online legal reasoning system where students argue about U.S. Supreme Court cases were predictive of students' eventual scores on a legal reasoning test.

These specific environments were not constructionist. However, the move towards conducting EDM in more open-ended online learning environments, and the growth in understanding how to discover and exploit the structure in data from these environments, creates enabling conditions for extending these methods to constructionist learning.

The process of extending EDM methods to constructionist data is not and will not be trivial; every new type of learning environment has required a learning process for EDM researchers. Typically, that learning process has involved a collaborative dialogue between experts in EDM and experts in the specific learning domain and online learning environment being studied. However, the successes in applying EDM methods to new domains and online learning environments gives hope that the process of extending EDM to constructionism will be quite tractable. That is not to say that EDM can or should tackle all research questions. Pure qualitative methods remain the standard for the exploration of possibilities, and pure quantitative methods remain the standard for confirmatory studies and larger scale hypothesis testing. However, EDM can provide a third way to reap many of the benefits from both more traditional qualitative and quantitative analyses.

The move towards identifying student meta-cognition, and self-regulatory skill within structured learning environments, using EDM, is of potential value to researchers in the constructionist paradigm, where issues of learners learning to actively participate in and drive their own learning and complex performance is of strong interest. For instance, Jeong et al. (2010) have identified patterns of students' transitions between problem-solving, self-assessment, and backtracking to reconsider previously learned material that distinguish between more successful and less successful learners.

EDM enables rigorous, replicable, and precise description of learner behavior, as well as analysis of how those behaviors interact with other constructs of interest. Learner behavior can be tracked in how it grows and changes over time. This approach plays a key role in Jeong et al.'s (2010) research into students' patterns of self-regulation over time. EDM methods have even been used to predict students' preparation for future learning of new and different materials from other paradigms (cf. Baker et al. 2011), providing a tool for linking analyses of learning within constructivist learning environments to a student's learning progression. Generally, EDM methods allow for linking assessments of various aspects of student learning and learning processes to a range of other constructs; they also support the linking of various aspects of student process and learning to each other. These types of research fall squarely into the type of EDM research referred to as *discovery with models*, where EDM models of various constructs are studied in relationship to one another and to assessments of other constructs.

At the same time, by using solely data already being collected by the learning environment, EDM enables ecologically valid research in that no interventions or interruptions to authentic student process are necessary to collect the data required to conduct EDM research. Implementing the data collection into the learning environment is often reasonably trivial, and many such learning environments already log (and subsequently ignore) much of the data for debugging purposes.

As such, EDM can be used to evaluate student methods, processes, and roles, helping us understand the strategies that learners develop as they participate in constructionist learning activities. EDM can be useful for studying processes of construction and development as well as the problem-solving and exploration domains in which it has been most used. In particular, EDM methods and related learning analytics methods have been used to study programming and the development of programming skills, including experts' and novices' patterns in program construction, compilation and debugging (Berland et al. 2013; Blikstein 2009, 2011), modeling programmers' trajectories within an assignment using Hidden Markov Models (Piech et al. 2012), inference of what a student is trying to program (Vee et al. 2006), and prediction of whether the student is at risk of failing to acquire programming skill (Dyke 2011; Tabanao et al. 2011).

3 Automated Feedback from EDM

As well as supporting understanding of learning in a range of learning interactions, educational data mining and learning analytics methods can support the provision of automated feedback to learners. Perhaps the highest-profile example of this is the Purdue SIGNALS project, which uses automated algorithms to assess the probability of student failure, and then informs instructors and students when a student is at risk of failing (Arnold 2010). Automated model based feedback of this type has been used for a variety of applications: to encourage students to engage in more effective help-seeking strategies (Roll et al. 2011), to provide feedback to students on how to solve problems more effectively (Stamper et al. 2011), to respond constructively when students game the system (Baker et al. 2006; Walonoski and Heffernan 2006), and to scaffold students' emotions (D'Mello et al. 2010).

For constructionist learning environments, it is of particular importance that learners be given process feedback to help them learn and build what they want to build. As well as supporting analysis of learner data, EDM is also well suited to giving feedback in situ, through providing a basis for understanding where and when students need support. Real-time feedback has long been a hallmark of constructionism. In particular, Logo has been successful (in part) because learners can see their understanding instantiated concretely in the artifacts they created, and EDM can simply provide more methods and tools for better feedback. Furthermore, as constructionist projects often feature public or shared artifacts, EDM can provide information to support students in helping each other in real-time, an emerging area of value for constructionist learning environments. However, the solution space of most constructionist projects is very wide, providing a challenge to giving targeted feedback. In building a robot or writing a computer program, for example, there are infinite possible paths for students to take. It is unfeasible to try to predict all possible mistakes that students could possible make, and have predetermined feedback for each of those cases—thus, the application of EDM techniques to constructionist learning will require building on methods designed for large solution spaces (e.g., Stamper et al. 2011). Some recent work on this field, however, provides indication real-time feedback might be possible even in more open-ended tasks. For example, Piech et al. (2012) captured tens of thousands of code snapshots from college students creating a computer program in the Karel language. By using a variety of techniques from machine learning, they were able to build a state machine and identify "sink" states from which students would only exit with great difficulty—this information could be used to provide instruction and feedback for students just before they are about to enter such problematic "sink" states.

4 Steps Towards Using EDM in Constructionist Learning Environments

Using EDM to study and improve constructionist learning environments will involve challenges, including bringing together two research communities without a strong history of collaboration with each other, and with different conceptions of what learning is and how it can be measured. However, it is our opinion that this is both feasible and desirable. In this section, we suggest some directions for how EDM research could be incorporated into a constructionist paradigm.

Before EDM methods can be applied to constructionist learning environments, the data from those learning environments must be placed into a form for which EDM methods can be effectively applied. One important challenge for this interdisciplinary field will be to create standard data formats to allow researchers to generate sharable data and replicable experiments. Several data formats are amenable to EDM analysis, and data formats are typically inter-convertible. For instance, Berland et al. (2013) created a database for their programming learning environment, IPRO, in which they catalog every edit made by any student in the environment. In that case, they record all changes to every primitive, as well as compiles, tests, rearrangements of code, and even simple aesthetic changes. That generates a massive store of data, from which large numbers of data features can be distilled for later mining. Furthermore, each data-point fully describes a discrete point in time for each student, allowing the data to be analyzed both at one point in time and over time, post hoc. In short, our experience suggests that collecting as many discrete data points at an exceptionally small granularity makes EDM much more tractable.

The work on *multimodal learning analytics* (Blikstein 2013a, b; Worsley and Blikstein 2011, 2012, 2013, in press; Worsley 2012) was one of the early attempts to apply EDM techniques for constructionist learning, merging machine learning techniques and multimodal data collection using data from a variety of synchronized sources: skin conductivity sensors, video, audio, gesture tracking, and eye-tracking. For example, they use video and gesture tracking to study students building simple physical structures with everyday materials. Students were previously classified based on their perceived level of knowledge in the domain of engineering design. A coding scheme was developed and agreed upon by a team of research assistants. Both video and gesture data were captured, and in analysis, many different approaches were attempted. The analyses ranged from a simple count of the number and duration of the codes to a cluster analysis of the temporal action sequences. The final algorithm was able to attain 70 % accuracy in classifying students' previously determined level of expertise. Schneider and Blikstein (2014) used gesture tracking within a learning activity in science in which students used tangibles interfaces, and were able to predict students' performance in a post-test only by examining their gesture data. Blikstein and Worsley also explored text mining, since a variety of features can be extracted from text or transcripts with prosodic, linguistic, semantic, or sentiment analysis. In one study, undergraduate students were invited to solve a series of design challenges during a think-aloud interview session. The data was analyzed using different machine learning techniques in order to predict the expertise level of the subjects. The data revealed counter-intuitive aspects of expertise in engineering, for example, certainty words were more significant than content words for the prediction of expertise. Indeed, Sherin (under review) employed text mining, topic mining, and clustering methods to identify how conceptual elements are activated in a set of semi-clinical (open-ended) interviews.

In some EDM research on programming, semantic actions have been construed as compile attempts (cf. Tabanao et al. 2011; Blikstein 2011; Piech et al. 2012), whereas in other research, semantic actions have been construed as the use of specific operators (cf.

Berland et al. 2013; Corbett and Anderson 1995). From simple features, more complex features can be distilled and abstracted. Rather than listing these features here, we encourage readers to look at some of the cited research to see examples of features used for specific domains and research questions. Typically, the process of engineering relevant features with construct validity is one of the largest challenges in the entire EDM process. While this process is often invisible to the reader of the resultant paper, studying features used in past models can be invaluable for developing a "feature engineering intuition"—a sense for which features will provide meaningful evidence for a set of research questions. It is worth noting that it is typically not desirable to simply develop thousands of very similar features and select between them automatically; doing so typically results in models that are over-fit (Marzban and Stumpf 1998; Mitchell 1997), working well on a specific data set but not generalizing to new data sets. There are automated algorithms which attempt to select good features which are not overly correlated to one another (cf. Yu and Liu 2004); these methods are a useful part of any data miner's toolbox, but are no substitute for conducting thoughtful feature engineering in the first place. Instead (or in addition), it is desirable to attempt to develop a set of a few dozen relatively different features with some construct validity, and select among these. An example of the benefits of selecting features with construct validity in mind can be found in Sao Pedro et al. (2012), where features selected based on construct validity as well as fit led to better performance at detecting student scientific inquiry skill within a new data set.

Once the data set is in a workable format, many approaches can be used to analyze the data. A full suite of EDM methods are discussed by Baker and Yacef (2009) and Romero and Ventura (2010). Repeating this discussion is outside the scope of the current paper. However, one key step for many (but not all) approaches is defining or discovering semantically meaningful constructs in data. One example of this is identifying internal or intermediate states of student constructions. By identifying these constructs, we can then visualize and investigate the pathways to powerful ideas and what types of behaviors and artifacts specifically make up those pathways. Constructionist research is often predicated on the suggestion that what students actually make matters and that their constructions are important, and that investigating the relationships between and intermediate states of those constructions deepens that commitment.

There are broadly two approaches in data mining to labeling data with semantically meaningful constructs: more bottom-up "unsupervised" approaches such as clustering, and more top-down "supervised" prediction approaches such as classification and regression (note that regression in the EDM sense is distinct from regression approaches used in traditional statistics; the mathematical underpinnings are similar, but the way the models are chosen, used, and validated is quite different).

Clustering can be conducted with completely unlabeled data, allowing bottom-up discovery of common patterns within the data.[1] These patterns can then be studied by a human analyst and correlated with other constructs to understand their meaning, as in Amershi and Conati (2009). Unsupervised clustering may be chosen based on a theoretical commitment to let student data lead the way and to "listen" to their actual process rather than impose artificial educational constructs. A problem with unsupervised clustering is that it can be difficult to make sense of that process—it can generate analyses that require considerable work to interpret, compared to prediction models. In some cases, however, this level of analyses can be more helpful than supervised ones. It allows for extremely quick feedback for researchers and teachers about the space of the students' constructions.

[1] For more information, consult Witten et al. (2011).

Using clustering can also help refine feature selection and aid in better prediction later. There is usually value in better understanding and mapping raw data; unsupervised clustering can often be thought of as a somewhat arbitrarily divided "viewable map" of the data.

Prediction models, by contrast, require a certain degree of human-labeled data. Prediction methods discover models that can infer the human labels, so that the model can then be used to label typically much more extensive unlabeled data. Human labels can be generated through hand-labeling log files (cf. Baker et al. 2006), field observations (cf. Baker et al. 2004; Dragon et al. 2008; Walonoski and Heffernan 2006; Worsley and Blikstein 2013), through the use of external tests (cf. Baker et al. 2011; Lynch et al. 2008; Muehlenbrock 2005), through other attributes of the data set such as future correctness (cf. Corbett and Anderson 1995), or through other sources such as teacher evaluations. Once a modeling algorithm is selected, models can be generated to predict the human labels. A modeling approach can be validated for reliability through multi-level cross-validation, where the model is repeatedly tested on new data at multiple levels (such as new data from the same students, new data from new students, new data from new content, and so on). It is worth noting that algorithms which are less prone to over-fitting—such as linear, logistic, and step regression, J48 decision trees, and K*, have historically been more successful in educational applications than more complex algorithms such as neural networks and support vector machines with complex kernel functions.[2]

Once semantically meaningful categories have been defined or discovered, they can be analyzed further through approaches that can infer rich relationships, such as association rule mining, sequential pattern mining, and the wide range of potential discovery with models approaches.

It is worth noting that framing student constructions with partially supervised or supervised methods—in terms of both existing understandings of how people learn, and giving feedback based on the results of those methods, does not imply that students must be undesirably constrained to following one of a small set of approaches. Supervised methods can generate more human understandable and more broadly applicable models of students' processes. As we begin to know what we are looking for, we can use EDM to understand it better. At this point, this can take the shape of formalizing big ideas in terms of what students are doing. One potential area of application, for example, would be in the study of recursion, repeatedly described by Papert (1980, 2000) as a *powerful idea*. It may be possible to study learners' developing understanding of recursion using EDM in this fashion, through labeling data in terms of recursion strategies, building EDM models, applying those models to more data, and conducting discovery with models to analyze how recursion grows, find where students have problems, find when and how students make breakthroughs, and finally, study where students use recursion in their final artifacts.

5 Open Questions in Constructionism and EDM

Much of constructionist research is exploratory rather than confirmatory. This makes it an excellent fit for EDM when compared to traditional statistical methodologies. In particular, constructionism is itself resistant to grading, ranking, and classifying children as bad/underachievers, because inherent in the idea of constructionism is that all students can engage in deep learning if the environment, tools, and facilitation are well-designed (Papert

[2] For detailed descriptions of these algorithms and their application, consult Witten et al. (2011).

1980). Several subsets of constructionist research each contain research questions potentially amenable to EDM.

How does constructionism manifest itself at a micro-genetic level and how do different constructionist experiences engender different micro-behaviors? Most constructionist projects require students to construct something—to learn to make something with electronics, programming, art and crafts supplies, or other materials. In the past, much of the possible feedback that students received from the artifact they were building was either coarse-grained or not aligned to constructionist pedagogical goals. For example, if a student makes a syntax error in Logo, the compiler would "complain" about the error, but it is remarkably difficult to make these messages contextually relevant for a creative project. One of the historical problems in providing better feedback was that the data that teachers, facilitators, and students could access was too simplistic. However, modern toolkits or integrated development environments ("IDEs") where constructionist projects occur (such as cloud-based apps or version-controlled saving) can capture process data at a very fine level. Indeed, it is no longer onerous to keep every single action, change, or keystroke that students might input. That process-based data is an incredibly rich source of data about learning. For instance, Berland et al. (2013) and Blikstein (2009, 2011) have mined those data to better understand how students learn to program. Beyond this, additional data sources can provide further leverage; Stevens et al. (2008) argue that a lot of that the most important procedural data is lost when over-relying on logs. Video of classrooms or informal spaces produce voluminous data that is easy to cross-reference with log-data, as do field observation methods (cf. Baker et al. 2004). There is a massive store of data to be mined, and an almost inexhaustible number of research questions that can be answered using a combination of video-analysis and log analysis. For instance, these methods may enable us to answer the question: to what degree do different types of explanations of a particular programming concept affect how students use that concept in their own programs?

How can micro-genetic analysis of conceptual change in constructionist work benefit from more complex models and analyses of behavior? There exist many unanswered questions about how and why students come to understand complex content. diSessa's (1993) knowledge-in-pieces framework has provided a backbone for how many constructionists model conceptual change. However, this work is difficult, and it requires careful, laborious, and close analysis of transcript data. Sherin (2012) has found that EDM techniques can help streamline the work of that type of analysis. By categorizing snippets of text based on language regularities and vocabulary, it becomes easier to understand changes in students' cognition. Duncan and Berland (2012) also suggest that by combining EDM with careful qualitative analysis, it is possible to strengthen qualitative arguments that use discourse analysis by exploring regularities or similarities.

What relationships exist between constructionist play, deep understanding, and complex inter-related behaviors across many students? For instance, NetLogo (Wilensky 1999) allows students to explore not only virtual simulations of scientific phenomena, but to construct new simulations or modify existing ones to better understand them, and even do so collaboratively. Kafai and Peppler (2011) use interactive social games to allow students take part in a constructive virtual world. Game logs have proven amenable to EDM methods in the past (e.g. Andersen et al. 2010; Conati and Maclaren 2005; Liu et al. 2011), but constructionist games can offer uniquely rich data, by enabling students to build something novel within them. There exist numerous possibilities for novel research using EDM to understand or visualize how play and social interaction in online communities can frame or change understanding or behavior.

The introduction of student-centered, project-based learning is a century-old challenge for educators worldwide. We envision that the integration of constructionist pedagogical approaches with EDM will pave a way for a wider adoption of student-centered approaches since this new interdisciplinary subfield could make assessment more feasible in large scale, enable the building of smarter technologies for real-time feedback, streamline and optimize the process of giving feedback to students, and offer researchers deeper insight into the learning processes in constructionist learning environments.

Acknowledgments Berland would like to thank the Complex Play Lab for help with this work, Don Davis for editorial help, and National Science Foundation Awards #SMA-1338508 and #EEC-1331655. Baker would like to thank support from the Bill and Melinda Gates Foundation, Award #OPP1048577, and from the National Science Foundation through the Pittsburgh Science of Learning Center, Award #SBE-0836012. Blikstein would like to thank the National Science Foundation through the CAREER Award #1055130, the AT&T Foundation, and the Lemann Foundation.

References

Abelson, H., & diSessa, A. (1986). *Turtle geometry: The computer as a medium for exploring mathematics*. Cambridge, MA: The MIT Press.

Amershi, S., & Conati, C. (2009). Combining unsupervised and supervised classification to build user models for exploratory learning environments. *Journal of Educational Data Mining, 1*(1), 18–71.

Andersen, E., Liu, Y.-E., Apter, E., Boucher-Genesse, F., & Popovic, Z. (2010). Gameplay analysis through state projection. In *Proceedings of the fifth international conference on the foundations of digital games* (pp. 1–8). Monterey, California: ACM.

Arnold, K. E. (2010). Signals: Applying academic analytics. *Educause Quarterly, 33*(1), n1.

Baker, E. L., Barton, P. E., Darling-Hammond, L., Haertel, E., Ladd, H. F., Linn, R. L., et al. (2010). *Problems with the use of student test scores to evaluate teachers*. Washington, DC: Economic Policy Institute.

Baker, R.S., Corbett, A.T., Koedinger, K. R. (2004). Detecting Student Misuse of Intelligent Tutoring Systems. In *Proceedings of the 7th international conference on intelligent tutoring systems* (pp. 531–540).

Baker, R., Corbett, A., Koedinger, K., Evenson, S., Roll, I., Wagner, A., et al. (2006). Adapting to when students game an intelligent tutoring system. In M. Ikeda, K. Ashley, & T.-W. Chan (Eds.), *Intelligent tutoring systems*, Lecture Notes in Computer Science (Vol. 4053, pp. 392–401). Berlin/Heidelberg: Springer. Retrieved from http://www.springerlink.com.libweb.lib.utsa.edu/content/t3103564632g7n41/abstract/.

Baker, R. S. J. D., Corbett, A. T., Roll, I., & Koedinger, K. R. (2008). Developing a generalizable detector of when students game the system. *User Modeling and User-Adapted Interaction, 18*(3), 287–314.

Baker, R., Gowda, S., & Corbett, A. (2011). Towards predicting future transfer of learning. In G. Biswas, S. Bull, J. Kay, & A. Mitrovic (Eds.), *Artificial intelligence in education*, Lecture Notes in Computer Science (Vol. 6738, pp. 23–30). Berlin: Springer. Retrieved from http://www.springerlink.com.libweb.lib.utsa.edu/content/41325h8k111q0734/abstract/.

Baker, R., & Siemens, G. (in press). Educational data mining and learning analytics. To appear in Sawyer, K. (Ed.), *Cambridge Handbook of the Learning Sciences: 2nd Edition.* Cambridge, UK: Cambridge University Press.

Baker, R., & Yacef, K. (2009). The state of educational data mining in 2009: A review and future visions. *Journal of Educational Data Mining, 1*(1), 3–17.

Beck, J. E., & Woolf, B. P. (2000). High-level student modeling with machine learning. In G. Gauthier, C. Frasson, & K. VanLehn (Eds.), *Intelligent tutoring systems*, Lecture Notes in Computer Science (pp. 584–593). Berlin, Heidelberg: Springer. Retrieved from http://link.springer.com.libweb.lib.utsa.edu/chapter/10.1007/3-540-45108-0_62.

Berland, M., Martin, T. et al. (2013). Using learning analytics to understand the learning pathways of novice programmers. *Journal of the Learning Sciences, 22*(4), 564–599.

Black, P., & Wiliam, D. (1998). *Inside the black box: Raising standards through classroom assessment*. London: King's College London.

Blikstein, P. (2009). *An atom is known by the company it keeps: Content, representation and pedagogy within the epistemic revolution of the complexity sciences*. PhD. dissertation, Northwestern University, Evanston, IL.

Blikstein, P. (2011). Using learning analytics to assess students' behavior in open-ended programming tasks. In *Proceedings of the I learning analytics knowledge conference (LAK 2011)*, Banff, Canada.

Blikstein, P. (2013a). Digital fabrication and 'making' in education: The democratization of invention. In J. Walter-Herrmann & C. Büching (Eds.), *FabLabs: Of machines, makers and inventors*. Bielefeld: Transcript Publishers.

Blikstein, P. (2013b). Multimodal Learning Analytics. In: *Proceedings of the III learning analytics knowledge conference (LAK 2013)*, Leuven, Belgium.

Blikstein, P. (2014). Bifocal modeling: Promoting authentic scientific inquiry through exploring and comparing real and ideal systems linked in real-time. In A. Nijholt (Ed.), *Playful user interfaces* (pp. 317–352). Singapore: Springer.

Piech, C., Sahami, M., koller, D., Cooper, S., & Blikstein, P. (2012). Modeling how students learn to program. In *Proceedings of the 43rd ACM technical symposium on Computer Science Education* (pp. 153–160). Raleigh, North Carolina, USA: ACM.

Blumenfeld, P., Fishman, B. J., Krajcik, J., Marx, R. W., & Soloway, E. (2000). Creating usable innovations in systemic reform: Scaling up technology-embedded project-based science in urban schools. *Educational Psychologist, 35*(3), 149–164.

Buechley, L., & Eisenberg, M. (2008). The lilypad arduino: Toward wearable engineering for everyone. *Pervasive Computing, IEEE, 7*(2), 12–15.

Conati, C., & Maclaren, H. (2005). Data-driven refinement of a probabilistic model of user affect. In L. Ardissono, P. Brna, & A. Mitrovic (Eds.), *User modeling 2005*, Lecture Notes in Computer Science (pp. 40–49). Berlin: Springer. Retrieved from http://link.springer.com.libweb.lib.utsa.edu/chapter/10.1007/11527886_7.

Corbett, A. T., & Anderson, J. R. (1995). Knowledge decomposition and subgoal reification in the ACT programming tutor. In *Proceedings of the 7th world conference on artificial intelligence in education*.

D'Mello, S. K., Lehman, B., & Person, N. (2010). Monitoring affect states during effortful problem solving activities. *International Journal of Artificial Intelligence in Education, 20*(4), 361–389. doi:10.3233/JAI-2010-012.

diSessa, A. A. (1993). Toward an epistemology of physics. *Cognition and Instruction, 10*(2–3), 105–225. doi:10.1080/07370008.1985.9649008.

diSessa, A. A., & Cobb, P. (2004). Ontological innovation and the role of theory in design experiments. *Journal of the Learning Sciences, 13*(1), 77–103. doi:10.1207/s15327809jls1301_4.

Dragon, T., Arroyo, I., Woolf, B. P., Burleson, W., Kaliouby, R. el, & Eydgahi, H. (2008). Viewing student affect and learning through classroom observation and physical sensors. In B. P. Woolf, E. Aïmeur, R. Nkambou, & S. Lajoie (Eds.), *Intelligent tutoring systems*, Lecture Notes in Computer Science (pp. 29–39). Berlin: Springer. Retrieved from http://link.springer.com.libweb.lib.utsa.edu/chapter/10.1007/978-3-540-69132-7_8.

Duncan, S., & Berland, M. (2012). Triangulating Learning in Board Games: Computational thinking at multiple scales of analysis. In *Proceedings of Games, Learning, & Society 8.0* (pp. 90–95). Madison, WI, USA.

Dyke, G. (2011). Which aspects of novice programmers' usage of an IDE predict learning outcomes. In *Proceedings of the 42nd ACM technical symposium on computer science education* (pp. 505–510). Dallas, TX: ACM.

Eisenberg, M. (2011). Educational fabrication, in and out of the classroom. *Society for Information Technology & Teacher Education International Conference, 2011*(1), 884–891.

Fosnot, C. T. (2005). *Constructivism: Theory, perspectives, and practice*. New York: Teachers College Press.

Harel, I. (1990). Children as software designers: A constructionist approach for learning mathematics. *Journal of Mathematical Behavior, 9*(1), 3–93.

IEDMS. (2009). *International Educational Data Mining Society*. Retrieved April 22, 2013, from http://www.educationaldatamining.org/.

Jeong, H., Biswas, G., Johnson, J., & Howard, L. (2010). Analysis of productive learning behaviors in a structured inquiry cycle using hidden markov models. *Manuscript submitted for publication*.

Kafai, Y. B., & Peppler, K. A. (2011). Youth, technology, and DIY developing participatory competencies in creative media production. *Review of Research in Education, 35*(1), 89–119.

Levy, F., & Murnane, R. J. (2005). *The new division of labor: How computers are creating the next job market*. Princeton University Press.

Liu, Y.-E., Andersen, E., Snider, R., Cooper, S., & Popović, Z. (2011). Feature-based projections for effective playtrace analysis. In *Proceedings of the 6th international conference on foundations of digital games* (pp. 69–76). Bordeaux, France: ACM.

Lynch, C., Ashley, K., Pinkwart, N., & Aleven, V. (2008). *Argument graph classification with Genetic Programming and C4.5* (pp. 137–146). Presented at the educational data mining 2008: 1st international conference on educational data mining, Proceedings, Montreal, Quebec, Canada.

Marzban, C., & Stumpf, G. J. (1998). A neural network for damaging wind prediction. *Weather and Forecasting, 13*(1), 151–163. doi:10.1175/1520-0434(1998)013<0151:ANNFDW>2.0.CO;2.

Merceron, A., & Yacef, K. (2004). Mining student data captured from a web-based tutoring tool: Initial exploration and results. *Journal of Interactive Learning Research, 15*(4), 319–346.

Mitchell, T. M. (1997). *Machine learning*. New York: McGraw-Hill.

Muehlenbrock, M. (2005). Automatic action analysis in an interactive learning environment. In *Proceedings of the workshop on usage analysis in learning systems at the 12th international conference on artificial intelligence in education AIED 2005* (pp. 73–80).

NGSS Lead States. (2013). *Next generation science standards: For states, by states*. Washington, DC: The National Academies Press.

Ocumpaugh, J., Baker, R., Gowda, S., Heffernan, N., & Heffernan, C. (2014). Population validity for educational data Mining models: A case study in affect detection. *British Journal of Educational Technology, 45*(3), 487–501.

Papert, S. (1972). Teaching children to be mathematicians versus teaching about mathematics. *International Journal of Mathematical Education in Science and Technology*. Retrieved from http://www.informaworld.com/index/746865236.pdf.

Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. NYC: Basic Books.

Papert, S. (2000). What's the big idea? Toward a pedagogy of idea power. *IBM Systems Journal*. Retrieved from http://llk.media.mit.edu/courses/readings/Papert-Big-Idea.pdf.

Papert, S., & Harel, I. (1991). Situating constructionism. *Constructionism*. Retrieved from http://namodemello.com.br/pdf/tendencias/situatingconstrutivism.pdf.

Piech, C., Sahami, M., Koller, D., Cooper, S., & Blikstein, P. (2012). Modeling how students learn to program. In *Proceedings of the 43rd ACM technical symposium on computer science education*, Raleigh, North Carolina, USA.

Resnick, M. (1998). Technologies for lifelong kindergarten. *Educational Technology Research and Development, 46*(4), 43–55.

Resnick, M., Maloney, J., Monroy-Hernandez, A., Rusk, N., Eastmond, E., Brennan, K., et al. (2009). Scratch: Programming for all. *Communications of the ACM, 52*(11), 60–67.

Reynolds, R., & Caperton, I. H. (2011). Contrasts in student engagement, meaning-making, dislikes, and challenges in a discovery-based program of game design learning. *Educational Technology Research and Development, 59*(2), 267–289.

Roll, I., Aleven, V., McLaren, B. M., & Koedinger, K. R. (2011). Improving students' help-seeking skills using metacognitive feedback in an intelligent tutoring system. *Special Section I: Solving information-based problems: Evaluating sources and information Special Section II: Stretching the limits in help-seeking research: Theoretical, methodological, and technological advances, 21*(2), 267–280. doi:10.1016/j.learninstruc.2010.07.004.

Romero, C., & Ventura, S. (2010). Educational data mining: a review of the state of the art. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on, 40*(6), 601–618.

Roschelle, J., Penuel, W. R., Yarnall, L., Shechtman, N., & Tatar, D. (2005). Handheld tools that "Informate" assessment of student learning in Science: A requirements analysis. *Journal of Computer Assisted learning, 21*(3), 190–203. doi:10.1111/j.1365-2729.2005.00127.x.

Sao Pedro, M. A. S., Baker, R. S. J. D., & Gobert, J. D. (2012). Improving construct validity yields better models of systematic inquiry, even with less information. In J. Masthoff, B. Mobasher, M. C. Desmarais, & R. Nkambou (Eds.), *User modeling, adaptation, and personalization*, Lecture Notes in Computer Science (pp. 249–260). Berlin, Heidelberg: Springer. Retrieved from http://link.springer.com.libweb.lib.utsa.edu/chapter/10.1007/978-3-642-31454-4_21.

Sao Pedro, M. A. S., Gobert, J. D., & Raziuddin, J. J. (2010). Comparing pedagogical approaches for the acquisition and long-term robustness of the control of variables strategy. In *Proceedings of the 9th international conference of the learning sciences* (Vol. 1, pp. 1024–1031). Chicago, IL: International Society of the Learning Sciences.

Schneider, B. & Blikstein, P. (2014). Unraveling students' interaction around a tangible interface using multimodal learning analytics. In *Proceedings of the 7th international conference on educational data mining*. London, UK.

Sherin, B. (2012). Using computational methods to discover student science conceptions in interview data. In *Proceedings of the 2nd international conference on learning analytics and knowledge* (pp. 188–197). ACM.

Sherin, B. (under review). A computational study of commonsense science: An exploration in the automated analysis of clinical interview data.

Shute, V. J. (2008). Focus on formative feedback. *Review of Educational Research, 78*(1), 153–189. doi:10.3102/0034654307313795.

Stamper, J., Eagle, M., Barnes, T., & Croy, M. (2011). Experimental evaluation of automatic hint generation for a logic tutor. In G. Biswas, S. Bull, J. Kay, & A. Mitrovic (Eds.), *Artificial intelligence in education*, Lecture Notes in Computer Science (Vol. 6738, pp. 345–352). Berlin/Heidelberg: Springer. Retrieved from http://www.springerlink.com.libweb.lib.utsa.edu/content/f0w1t2642k4t6128/abstract/.

Stevens, R., Satwicz, T., & McCarthy, L. (2008). In-game, in-room, in-world: Reconnecting video game play to the rest of kids' lives. *The ecology of games: Connecting youth, games, and learning, 9*, 41–66.

Tabanao, E. S., Rodrigo, M. M. T., & Jadud, M. C. (2011). Predicting at-risk novice Java programmers through the analysis of online protocols. In *Proceedings of the seventh international workshop on computing education research* (pp. 85–92). Providence, Rhode Island, USA: ACM.

Vee, M. H. N. C., Meyer, B., & Mannock, K. L. (2006). Understanding novice errors and error paths in Object-oriented programming through log analysis. In *Proceedings of workshop on educational data mining at the 8th international conference on intelligent tutoring systems (ITS 2006)* (pp. 13–20).

Walonoski, J.A., Heffernan, N.T. (2006): Prevention of off-task gaming behavior in intelligent tutoring systems. In *Proceedings of the 8th International Conference on Intelligent Tutoring Systems* (pp. 722–724). Berlin: Springer.

Warschauer, M., & Matuchniak, T. (2010). New technology and digital worlds: Analyzing evidence of the equity in access, use and outcomes. *Review of Research in Education, 34*(1), 179–225.

Wilensky, U. (1996). Making sense of probability through paradox and programming: A case study of connected mathematics. In Y. Kafai & M. Resnick (Eds.), *Constructionism in practice* (pp. 269–296). Mahwah, NJ: Lawrence Erlbaum Associates.

Wilensky, U. (1999). *NetLogo*. Retrieved from http://ccl.northwestern.edu/netlogo.

Wilensky, U., & Reisman, K. (2006). Thinking like a wolf, a sheep, or a firefly: Learning biology through constructing and testing computational theories—An embodied modeling approach. *Cognition and Instruction, 24*(2), 171–209.

Witten, I. H., Frank, E., & Hall, M. A. (2011). *Data mining: Practical machine learning tools and techniques*. Morgan Kaufmann.

Worsley, M. (2012). Multimodal learning analytics: Enabling the future of learning through multimodal data analysis and interfaces. In *Proceedings of the international conference on multimodal interfaces*, Santa Monica, CA.

Worsley, M., & Blikstein, P. (2011). What's an Expert? Using learning analytics to identify emergent markers of expertise through automated speech, sentiment and sketch analysis. In *Proceedings for the 4th annual conference on educational data mining*.

Worsley, M., & Blikstein, P. (2012). An Eye For Detail: Techniques for using eye tracker data to explore learning in computer-mediated environments. In *Proceedings of the international conference of the learning sciences* (pp. 561–562). Sydney, Australia.

Worsley, M., & Blikstein, P. (2013). Towards the Development of Multimodal Action Based Assessment. In *Proceedings of the III learning analytics knowledge conference (LAK 2013)*, Leuven, Belgium.

Worsley, M., & Blikstein, P. (in press). Analyzing engineering design through the lens of learning analytics. *Journal of Learning Analytics*.

Yu, L., & Liu, H. (2004). Efficient feature selection via analysis of relevance and redundancy. *The Journal of Machine Learning Research, 5*, 1205–1224.

Zahn, C., Krauskopf, K., Hesse, F. W., & Pea, R. (2010). Digital video tools in the classroom: Empirical studies on constructivist learning with audio-visual media in the domain of history. In *Proceedings of the 9th international conference of the learning sciences* (Vol. 1, pp. 620–627). Chicago, Illinois: International Society of the Learning Sciences.