

Chapter 1

Introduction

Johann Ari Larusson and Brandon White

Supporting teaching and learning with technology is becoming as commonplace as chalk in today's educational institutions. However, simply making technology available or requiring students to use it does not necessarily guarantee success. How does one effectively explore online learning communities so as to get an accurate description of the complex interactions taking place? What methods for analysis are available? What does a method of analysis even look like? What is the unit of analysis? How can an institution effectively organize its data? How does the information collected enrich students' learning experiences? How can we positively impact the teachers' pedagogical practices? How does one even design for successful implementations of educational technology that report back data rich enough to affect subsequent implementations? These questions are the ones required to better inform an educational agenda not only for "teaching with technology," but simply for teaching in the first place.

These processes, whatever their form, are inherently complex. The teaching and learning themselves might be taking place in a classroom, but are all nevertheless unfolding in an intangible time and space—inside a "black box," so to speak—producing enormous volumes of data where the vision of what data to collect, how to collect it, and how to explore it is not necessarily clear. In recent years, learning analytics (LA) has emerged as a field that seeks to provide answers to questions such as the ones highlighted above. Learning analytics can be summarized as the collection, analysis, and application of data accumulated to assess the behavior of educational communities. Whether it be through the use of statistical techniques and predictive

J.A. Larusson (✉)

Center for Digital Data, Analytics and Adaptive Learning at Pearson,
Boston, Massachusetts, USA
e-mail: johann.larusson@pearson.com

B. White

Department of English, University of California, Berkeley, Berkeley, California, USA
e-mail: brandonw@berkeley.edu

modeling, interactive visualizations, or taxonomies and frameworks, the ultimate goal is to optimize both student and faculty performance, to refine pedagogical strategies, to streamline institutional costs, to determine students' engagement with the course material, to highlight potentially struggling students (and to alter pedagogy accordingly) to fine-tune grading systems using real-time analysis, and to allow instructors to judge their own educational efficacy. In every case, learning analytics gives all stakeholders insight into what is taking place from Day 1 to Day X of a given class irrespective of the type of activity taking place. In short, learning analytics is broadly defined as the effort to improve teaching and learning through the targeted analysis of student demographic and performance data (Elias 2011; Fritz 2010). The contents of the "black box," in other words, become that much more visible, with their various markers sampled, collected, evaluated, and replayed in a legible form.

Learning analytics encompasses a range of cutting-edge educational technologies, methods, models, techniques, algorithms, and best practices that provide all members of an institution's community with a window into what actually takes place over the trajectory of a student's learning. Involvement in LA technologies and pedagogies allows educators and scholars to engage in a contemporary and innovative approach to an educational issue that is already an integral part of higher education.

In many ways, the field of learning analytics should be considered new. The field itself has come into being largely thanks to the proliferation of digital data produced by educational institutions' increasing tendency to produce, submit, and assess academic work in electronic form (Greer and Heaney 2004; Hirst 2011). While the first formal conference on LA, held in 2011, is evidence of its growing relevance in educational circles on an international scale, the fact that such a conference had not existed previously is sign enough of LA's relative infancy.

Learning analytics ideally attempts to leverage data to provide insight into the activities taking place within the classroom. What metrics are derived can then be fed back into pedagogy or applied with consequences even well outside the classroom itself. Several higher education institutions in particular have begun applying learning analytics to evaluate crucial aspects of the learning process and pedagogical practice, alongside institutional aims like student retention and cost reduction (Siemens and Long 2011). Holistic descriptions of several of these practices can be found in (Siemens and Long 2011) and (Ferguson 2012). A recent U.S. Department of Education brief held that learning analytics prioritizes the "*human tailoring* of responses, such as through adapting instruction content, intervening with at-risk students, and providing feedback" (Bienkowski et al. 2012, p. 13). This approach "does not emphasize reducing learning into components but instead seeks to understand entire systems and to support human decision making" (ibid). Yet for all the budding interest in LA, its earliest implementations have evolved from older models and methods, from raw data mining (cf. Baker and Yacef 2009) and learning community studies (cf. Dawson 2010) to the broader field of academic analytics (Goldstein and Katz 2005; Campbell et al. 2006). As institutions and educators increasingly begin to install learning analytics systems, or learning analytics enabled systems, they often tend to employ frameworks inherited from several of these other fields. Even nominal attempts to directly improve learning and teaching practice tend to digest institutional systems data with limited understanding of how that data

could or should inform pedagogy. Although these other inquiries remain vital and valuable fields, the purpose of this volume is to help situate LA's unique priorities, unique intended benefits, and unique ranges of personnel capable of putting that technology into practice. Growing the field of learning analytics requires making sure that it remains distinct from what came before and that its purpose remains rigorously clear.

Up until recent years, research and practice in this area has been hampered by a lack of definition, with work in the field dispersed throughout a number of journals and conferences, making it more difficult for experts to share results, get a real sense of what is new and innovative, or to identify the best practices, strategies, or tools to use. The birth of learning analytics as a field of study in its own right, through a now annual conference, the recently established Society for Learning Analytics Research, and with workshops and symposiums being organized around the world, has now made it possible to consolidate research once taking place along its periphery under one umbrella.

Learning analytics is uniquely positioned as a field with the potential to guide the efforts of any of a number of institutional actors or stakeholders, from students to instructors, IT professionals to educational administrators. While the inputs of learning analytics derive primarily from the classroom, any one of these stakeholders may well be charged with evaluating the results, putting changes into action, and weighing the impact that results. This book attempts to provide the first comprehensive reference book for LA, with the aim of helping scholars, researchers, developers, IT professionals, chief technology and information officers, university administrators, or anyone and everyone interested in advancing the field of learning analytics by showcasing the latest results, strategies, guidelines, methods, models, and tools. Collecting all of this information in one volume will allow scholars and researchers to take stock of ongoing efforts in the field, helping to illuminate what areas remain to be explored, and thus pushing the field yet further forward.

The purpose of this volume is, simply put, to provide an entry point into the field for any one of these actors depending upon their unique institutional interests. As a field with a broad appeal, simply navigating the extant literature of learning analytics, let alone attempting to put any of those principles into practice, can prove daunting. The chapters that follow each attempt to consolidate much of the available literature while putting forth best practice guidelines or model case studies that might prove of interest to particular types of readers. As such, this book is organized not around common problems or the mounting complexity of its efforts, but rather around the kinds of communities that each chapter attempts to address. It is the hope of the editors that this approach will allow different kinds of readers an opportunity to easily identify those chapters most likely to offer immediate insights. In the remainder of this introduction, each chapter's possible contributions to the field are thus suggested alongside discussion of its possible appeal to different classes of readers. Rather than simply summarizing what follows, the reader can think of this chapter as a map to the different ways in which the book itself might be read.

These myriad types of engagement are what the complexity of the field of learning analytics requires. What is learning analytics? The answer is not simply our own, as editors, but the one that the book itself, through each of our various

contributors, comes to suggest. These inquiries *themselves are* learning analytics. By probing the field's theoretical investments and by classifying its possible components, by exploring its history relative to educational data mining (EDM) and by helping to place it within a broader institutional context, by applying case studies to educators and students and academic advisors alike, these chapters take stock of all the many stakeholders that learning analytics might attempt to benefit, and thus most comprehensively demonstrate its power and its promise.

1.1 Preparing for Learning Analytics

The first section of the book, "Preparing for Learning Analytics," looks to clarify the stakes of learning analytics by supplying suggestions for the field's domain, potential, and possible points of emphasis. In each of the three chapters, these suggestions take the form of guidelines for what the development of a learning analytics application might require. Whereas later chapters begin with established technologies or established pilot programs at host universities, these chapters take none of that for granted, investigating instead the very foundations of learning analytics practice.

An initial entry point for nearly any reader can be found in Chap. 2, Abelardo Pardo's "Designing Learning Analytics Experiences." Pardo synthesizes the results and proposals of dozens of research findings to suggest five phases of design and execution through which an LA intervention might pass. These phases form a flexible framework that might be applied to any LA endeavor, providing readers with a sense of the kinds of decisions, dependencies, and trade-offs that are encountered in taking an analysis tool from conceptualization to subsequent enhancement.

The first stage, "capture," corresponds to the earliest collection of student data. The second stage, "report," delivers that data to a specifically defined set of stakeholders. The third stage, "prediction," deploys any of a number of techniques to provide non-intuitive answers to frequently encountered educational questions, such as the likelihood of an individual student failing a course or failing to graduate altogether. The "act" stage that follows offers the possibility of issuing automated solutions or implementing manual ones that have the potential, ideally, to reverse the most dire consequences of the earlier prediction. In the final stage, "refinement," the efficacy of the resulting actions is assessed anew so that the long-term viability of the analysis can itself be modified as need be.

Each stage in this process is presented not just as a single phase, with a pre-defined beginning and end, but as intimately bound up with choices that might have been made in earlier stages. Taking for instance only the "report" phase, Pardo singles out LA systems aimed at different classes of stakeholders. A process designed to deliver data directly to instructors requires a different set of investments than one generating metrics for IT professionals. Rather, therefore, than prescribing specific guidelines for what every stage ought to entail, Pardo instead offers a series of

questions at each stage that might inform how an LA implementation might be successfully designed and executed.

In a way, Pardo's chapter can be considered as a heuristic for much of the work of this volume as a whole. All of the subsequent chapters offer some specific engagement with one or more of the steps that Pardo outlines, and each is inevitably a product of choices that must have been made in any particular stage. Certain chapters, like that of Ryan Baker and Paul Inventado (see Chap. 4), are concerned principally with a specific phase: in this case, the kinds of calculations that are typically embraced during the "prediction" stage by learning analytics and EDM communities, respectively. Two case study chapters, as offered by Andrew Krumm et al. (see Chap. 6) and Brandon White and Johann Ari Larusson (see Chap. 8) provide considerations of two separate phases of analysis, carried out over a number of years. It might be suggested that these chapters have specific implications for what Pardo calls "refinement." Every attempted "prediction," however, is inevitably dependent on the stages that must have come before, every "refinement" only as good as the various acts that have been executed along the way.

Readers of this volume might thus do well to begin with Chap. 2, and to think of its many insights when taking up any of the chapters that follow. Such schematization of the assumptions underlying any particular technology will only lead to better questions, more pointed questions, and thus more opportunities for further refinement. Getting more LA systems to this final stage can be considered one of this volume's explicit goals. As Pardo mentions, because of LA's only relatively recent adoption and expansion, few LA systems have graduated to the "refinement" stage of analysis. There simply is not much longitudinal data on how a system might go through several iterations. It is our hope as editors that the next generation of LA research will see yet further instances of Pardo's framework being brought full circle.

Chapter 3, "Harnessing the Currents of the Digital Ocean" by John T. Behrens and Kristen E. DiCerbo, extends the discussion of the previous chapter to address the abundance of electronic information that now characterizes many educational efforts. Behrens and DiCerbo contrast this "digital ocean" with what they call the "digital desert," the pre-digital environment of the late twentieth century where data was rare, expensive to obtain, and as such was only amenable to limited, if any, analytic applications. While the potential consequences of such a shift can be seen as an impetus for the instantiation of learning analytics as a field in the first place, Behrens and DiCerbo suggest that the prevalence of digital data requires perhaps a more fundamental reshaping that the field might yet need to undergo. In their account, the simple technological limitations of the "digital desert" were themselves responsible for the kinds of activities, like multiple choice quizzes, that were developed for analysis. These activities tended—and still often tend—to be presented in a fixed form, with static questions matched to fixed answers to measure the correct response. As we stand on the shore of the "digital ocean," however, these same standards for collection, evaluation, and dissemination needn't remain a constraint.

Behrens and DiCerbo argue that the potential of learning analytics lies in its ability to reconceptualize the educational space, allowing us to think of user activity as

an ever-modulating stream of inputs from which certain attributes can be observed over time rather than requiring any moment-by-moment measure of correctness. The consequences for this shift in worldview would thus alter not only the ways in which data is collected, but also the very way in which it is understood. If the current understanding of summative evaluations (like final exams) interspersed with formative exercises (like homework, quizzes, midterms, or papers) can be likened to an autopsy conducted following a series of routine checkups, a naturalistically embedded assessment by way of learning analytics might instead be compared to a heart monitor that regularly and automatically generates feedback on the conditions at hand. Rather than seeing the analytic interface as something that delivers content to students, this understanding would instead see learning analytics as something allowing students themselves to create, explore, and reinforce the conditions of their own learning. As in Chap. 2, this discussion would situate learning analytics as an embedded system of continued research and refinement.

This chapter takes as its explicit endeavor an attempt to help readers rethink some of the underlying assumptions regarding how data and data analysis might be structured in a computationally complex space. Readers looking to expand their sense of what learning analytics might attempt would benefit from using this chapter as a primer on the possibilities while also helping to root the field's ambit in a broader history of educational theory. The chapter itself concludes with a wealth of suggestions for future research, not strictly in the *applications* of learning analytics, but on the types of thinking and training that learning analytics might require of the researchers themselves. These suggestions are invaluable as a basis from which to evaluate the core assumptions of learning analytics, and in so doing further suggest the kinds of practices, principles, and applications that await on the horizon.

Despite the only recent consolidation of learning analytics as a specific field of inquiry, LA approaches and LA methods have not developed in a vacuum. Chapter 4, Ryan Baker and Paul Inventado's "Educational Data Mining and Learning Analytics," examines learning analytics relative to the abutting discipline of EDM. Baker and Inventado provide a historical contextualization for EDM's growth, both as a research community and as a specific field of scientific inquiry.

Their particular focus, however, lies in the identification of several key methods that EDM has traditionally deployed that are perhaps more foreign to LA research, although they needn't necessarily be. These methods fall into the broad categories of "prediction models," "structure discovery," "relationship mining," and "discovery with models." For each category, Baker and Inventado not only identify relevant applications of the method, but discuss how each method has historically been a part of EDM research, and to what extent it remains so to this day.

By exploring each category more closely, Baker and Inventado are able to provide broad contextualizations of what EDM-type analysis might attempt. Their chapter could, in this way, be considered as a complement to Abelardo Pardo's chapter on learning analytics design (see Chap. 2). Yet the two projects—and, consequently, two learning communities—needn't be entirely at cross-purposes. The diachronic exploration of EDM's evolution that Baker and Inventado present is useful on the one hand for illuminating the areas where learning analytics researchers

have pursued EDMs through different means. It is useful in turn, on the other hand, for suggesting areas of inquiry that LA has heretofore left mostly untapped.

Readers looking to answer a specific subset of research questions might do well to consult this chapter as a kind of guide to what other strategies might continue to augment LA research. Taking a cue from John T. Behrens, one of the other authors in this volume (see Chap. 3), Baker and Inventado note that learning analytics and EDM essentially have their names reversed: that while learning analytics tends to focus on educational outcomes, EDM is more often than not concerned with the immediate products of learning. What this chapter potentially suggests is that a long-term implementation of LA and EDM methods in concert would ultimately wind up informing one another intimately, with improved learning coming to ensure continually optimal educational outcomes. At the point in the near future when both fields have readily demonstrated enough success that they can be readily installed, run, and refined—at the point, in other words, where research becomes practice—the difference between the two becomes virtually indistinguishable.

1.2 Learning Analytics for Communities

The second section of this volume explores learning analytics that speak to the specific interests of learning communities beyond the immediate teacher–student relationship. These chapters ask what it means to conceive of learning analytics at a large scale, either by discussing the implications of learning analytics for institutions as a whole, or by empowering a different level of stakeholders to leverage analytic insights.

One of the common concerns of several of this volume’s authors lies in the granularity—or specificity—of the reporting data that an LA system might produce (see, for instance, Chap. 2). Data must be specific enough that its insights are made intelligible, but general enough that the end user isn’t overwhelmed by abundant detail. Many of the chapters in this volume, such as Chap. 8 or Chap. 6, describe technologies meant to be put in the hands of on-the-ground users, be they instructors or students advisors. In a more fully integrated LA landscape, however, one can easily imagine any number of classrooms interventions taking place side by side. As soon as decisions about LA use need to be made beyond the individual classroom, a different series of questions immediately need to be considered. Chapter 5, “Learning Analytics at an Institutional Level,” by Matthew D. Pistilli, James E. Willis, III, and John P. Campbell, describes the way in which an institutional actor, such as an administrator, a technology officer, or a system administrator, might go about the process of implementing and overseeing an LA architecture.

Building off of Tinto’s theory of student departure, Astin’s theory of student involvement, and Chickering and Gamson’s principles for good practice in undergraduate education, the authors suggest a framework for where an institutional attempt at LA might even be committed. The standard that the authors put forth is ultimately a measure of a student’s place in his or her educational environment.

Learning communities inevitably extend well outside the classroom, and even factors as casual as a student's frequency of contact with an instructor or involvement in the extracurricular games taking place nearby can stimulate a student's investment in his or her educational institution, increasing the likelihood that he or she will remain enrolled, excel in classes, and work towards a degree. Institutions themselves are thus ideally positioned to leverage observations of these interrelated interactions through analytics. Such a model of analytical practice takes stock of a diverse array of factors, gathered from a variety of different interactions, and uses the data from these interactions to suggest altered approaches that might improve a student's comfort, confidence, and capability in his or her educational setting.

Several consequences emerge from this analysis, the first of which is suggested even by the use cases sketched above. Pistilli et al. foremost suggest a renovation in the ways in which institutions even come to think of analytics in the first place, urging institutions to take stock of ambient data based on existing interactions between students, faculty, and supporting staff rather than going out and creating data sets from the ground up. The second suggestion informs the way in which such data might ultimately be used. Policies governing the privacy of information collected and disseminated in such a context are alone an important consideration for any such implementation, especially considering the varying standards for how confidential information might be handled at different universities even within the same country, city, or state. But what the authors ultimately put forward is a means for an institution to consider the interests of every other stakeholder concerned. It is not only, for instance, that faculty need to be sensitive to how they deploy analytics in their interactions with students, but that they also need to remain cognizant of the ways in which students could actually be *discouraged* by the result, leading to a negative feedback loop which is far from any LA implementation's intended purpose.

Readers of this volume with a particular interest in institutional efficacy would do well to consult this chapter early as a baseline look at what the commitments of an institution in an LA context are or could be. What this chapter suggests is that the most stable sense of analytics' place in an institution's daily life can only be understood holistically, as an aggregate consideration of the benefits accrued to every one of its individual actors.

Chapter 6, Andrew E. Krumm, R. Joseph Waddington, Stephanie D. Teasley, and Steven Lonn's "A Learning Management System-Based Early Warning System for Academic Advising in Undergraduate Engineering" reports on an ongoing case study working hand-in-hand with stakeholders to develop a system capable of informing academic advisors of students in need of additional support. It shouldn't escape notice that this study, while directly dependent on student data, is the only chapter of this volume that doesn't use the individual course instructor as the primary instigator of interventions. This configuration of stakeholders thus suggests one immediate application of the kind of discussion found in Chap. 5. The relevant stakeholders here, and the ones who the authors approached with considerations for the second phase of their study, are the academic advisors who more often than not function as gatekeepers between instructional and institutional requirements.

The authors thus come to extend the conception of what a pedagogical intervention, properly carried out, might be.

As an early warning system, the authors' project was designed to alert academic advisors to whether students were likely to require further encouragement based on a number of data markers culled from a learning management system: graded activities, frequency of log-ins, and relative contextualization of a student's performance based on that of his or her own peers. This method thus combined real-time data with longitudinal tracking, creating a kind of self-correcting system. Since the intended recipients of this information were not involved in the day-to-day work of instruction, what the study ultimately attempted to improve upon was its own necessary granularity as determined by the frequency of the reports that advisors would receive.

Two possible classes of readers might be most interested in this volume's case studies (Chaps. 6–8), but in this chapter in particular: those readers looking even for a specific sense of what an LA application might entail in the first place, and those readers who, having surveyed the other contents of the book, now want to test how other broader principles might be put into practice. Readers of Chap. 6 would especially benefit from considering the chapter alongside Chap. 2. The early warning system described and refined in Chap. 6 remains the only instance of what Pardo calls “refinement” carried to its utmost, with the results of live analysis coming to actively reinfect the production cycle of a subsequent iteration of the tool. As such, the study discussed here is of value not only for how it might model what a successful refinement entails, but for what it might suggest about the future of learning analytics, when certain systems have become an established enough part of educational practice that they can routinely produce results but also be routinely improved.

1.3 Learning Analytics for Teachers and Learners

The final section of this book details attempts to use analytics to explore the environment most familiar to academic practice: the classroom. The reliance of learning analytics on student data has been a consistent theme throughout this volume. These chapters describe ways in which that data might be deployed, allowing instructors to use analytics to reshape or refine pedagogy.

Chapter 7, by Christopher Brooks, Jim Greer, and Carl Gutwin, “The Data-Assisted Approach to Building Intelligent Technology Enhanced Learning Environments” serves as a possible bridge between more theoretically oriented material and the case studies that follow. As the authors describe, intelligent tutoring systems have often been deployed as a means of scaling educational materials to better suit student performance. Such systems, however, are possibly unwieldy, requiring not just one expert in a given discipline or domain—the individual that we would ordinarily think of as a course instructor—but a separate pedagogical expert to weigh the multiple possible responses to a learner's mistake, and yet another series of experts to build tiers of content suitable to any number of learners. Such a tutoring system, in other words, becomes magnitudes more labor intensive than the

simple intensive instruction that the intelligent tutoring system might have been designed to ease. The authors thus propose a “data-assisted approach” to intelligent tutoring systems that acknowledges the instructor’s place in conducting classes and manually performing pedagogical interventions.

Although the authors roughly categorize classes of implementations that might be pursued using a data-assisted approach, the heart of the chapter lies in an enumeration of specific motivating scenarios that describe three different educational contexts, three different technologies, and three different types of data for which a data-assisted intelligent technology would be of immediate use. These scenarios can most productively be considered as case studies of three different applications in their own right. In each instance, the authors not only showcase the application’s functionality, but provide a supplementary consideration for how each application can be deployed at a different level of granularity, with a different specific focus, or with a different degree of buy-in from users.

Readers of almost any specific interest would do well to consult this chapter as a kind of test case for the possible range of even a single type of LA engagement. Those readers new to learning analytics may want to consult this chapter in concert with Chap. 2 as a means of measuring LA’s possible prerogatives against how those prerogatives are developed in practice. One of the chapter’s many merits is the way in which it suggests that the many choices confronting a new LA implementation needn’t be binary choices at all, but decisions that the *right* implementation can pursue in parallel.

Chapter 8, “Identifying Points for Pedagogical Intervention Based on Student Writing: Two Case Studies for the ‘Point of Originality,’” by Brandon White and Johann Ari Larusson, provides a second case study chapter. This chapter showcases a computational method and tool called the “Point of Originality,” which measures a student’s ability to put key course concepts into his or her own words as a course progresses. With the mounting trend in higher education towards larger and larger gateway courses, especially in the early phases of a student’s academic career, the Point of Originality is proposed as a way to let instructors quickly diagnose which students are likely to be struggling, and to use that information to conduct specific pedagogical interventions. As in Chap. 6, the use case for the tool is not an uncommon one: the only required inputs are the kinds of regular, iterative writing activities (be it a blog, discussion board, or other written assignment) that many instructors already use. Once an instructor has input a series of course concepts—either a few terms that might be likely to come up on an approaching midterm, or else a string of every key term found on the course syllabus—a custom algorithm calculates how proximately related every word in every student’s writing sample is to those terms.

The chapter follows two different proof of concept case studies, the second conducted as a larger-scale elaboration of the first. The earlier case study uses actual course data to imagine a not unconventional scenario: an instructor is preparing to distribute an assignment midway through a semester, and wants to know whether students will likely be equipped to answer the prompt. Using the terms of the prompt as the initial “query term” concepts, the Point of Originality tool weighs the degree to which students have been able to put those same concepts into their own words at

any point during the semester, and provides a graphical and numerical representation of which students are likely to succeed and which are likely to struggle. The second case study adapts these findings for an entirely different course environment, one with many more students and with a highly technical scientific subject matter. The movement between these two case studies suggests the application of a principle suggested, in different ways, by each of the chapters in this volume (and perhaps most pointedly so in Chap. 3)—that the integrity of analytics efforts come from their ability to be universalized, used as well in one course context as in another. In an effort to streamline the possible applicability of the Point of Originality yet further, this second case study even removes the requirement that an instructor limit his or her query terms to those related to a specific assignment: rather, this case study simply made use of more than a 100 terms relevant to the full course syllabus, offering instructors insight into how students were approaching the course material in the most general terms. Both case studies demonstrated a strong degree of correlation between the metrics produced by the tool and a student's eventual performance in the class. The suggestion is that use of the Point of Originality tool would have singled out the struggling students well in advance, and would have allowed an instructor to take action as need be, either by working with those students or even by looking at the results to determine *which* of the concepts in circulation met with the most difficulty.

As in Chap. 6, readers might well be interested in this chapter as a way of determining how the principles outlined elsewhere in this volume can be applied to the design of an actual learning analytics platform, or for simply determining what a learning analytics system looks like to begin with. The chapter itself provides a slightly different shift in inquiry from several of the other chapters by examining a different *type* of data set than used elsewhere. As such, the chapter might usefully be considered alongside like attempts at analysis, not as a competing method, but as a tool that might be used in concert with a handful of others.

It is the possible harmony between several of these methods that forms one of the final suggestions of this volume. As learning analytics matures as a field, there will come a point where diagnostic methods are liable to overlap, where one analytics tool might be used by instructors while another relays a different set of information to institutional actors. The several contributions of this volume are all ultimately cross-compatible. Learning analytics depends on particular methods, particular metrics, particular tools, but the only learning analytics *solution* might well be a holistic solution, one that speaks equally to the experiences of learners, educators, and administrators.

References

- 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.
- Bienkowski, M., Feng, M., & Means, B. (2012). *Enhancing teaching and learning through educational data mining and learning analytics: An issue brief*. Washington, DC: U.S. Department of Education Office of Educational Technology.

- Campbell, J. P., Finnegan, C., & Collins, B. (2006). Academic analytics: Using the CMS as an early warning system. *WebCT impact conference*, Chicago, IL.
- Dawson, S. (2010). Online forum discussion interactions as an indicator of student community. *Australasian Journal of Educational Technology*, 22(4), 495–510.
- Elias, T. (2011). *Learning analytics: Definitions, processes and potential*. Unpublished Internal Whitepaper of Athabasca University, Canada. Retrieved from https://landing.athabascau.ca/mod/file/download.php?file_guid=43713
- Ferguson, R. (2012). Learning analytics: Development, drivers, and challenges. *International Journal of Technology Enhanced Learning*, 4(5–6), 304–317.
- Fritz, J. (2010). Classroom walls that talk: Using online course activity data of successful students to raise self-awareness of underperforming peers. *The Internet and Higher Education*, 49, 89–97.
- Goldstein, P. J., & Katz, R. N. (2005). *Academic Analytics: The Uses of Management Information and Technology in Higher Education*. Boulder: EDUCAUSE.
- Greer, L., & Heaney, P. J. (2004). Real-time analysis of student comprehension: An assessment of electronic student response technology in an introductory earth science course. *Journal of Geoscience Education*, 52, 345–351.
- Hirst, T. (2011). Pragmatic analytics: Insight, representation, and structure. *Keynote presented at the First International Conference Learning Analytics and Knowledge*, Banff, AB.
- Siemens, G., & Long, P. (2011). Penetrating the fog: Analytics in learning and education. *Educause Review*, 46(5), 31–40.