

Learning and Becoming in Practice

THE INTERNATIONAL CONFERENCE OF THE LEARNING SCIENCES (ICLS) 2014

PROCEEDINGS VOLUME 1



University of Colorado Boulder
June 23-27, 2014

www.isls.org/icls2014

HOSTED BY:



University of Colorado
Boulder

SPONSORED BY:



International Society of
the Learning Sciences

Design-Based Research Process: Problems, Phases, and Applications

Matthew W. Easterday, Daniel Rees Lewis and Elizabeth M. Gerber
Northwestern University, Evanston, IL
{easterday, daniel-rees, egerber}@northwestern.edu

Abstract. Since the first descriptions of design-based research (DBR), there have been continued calls to better define DBR and increase its rigor. Here we address four uncertainties about DBR: (a) the phases of the DBR process, (b) what distinguishes DBR from other forms of research, (c) what distinguishes DBR from design, and (d) the characteristics of DBR that make it effective for answering certain types of questions. We build on existing efforts by defining DBR as an iterative process of 6 phases: focus, understand, define, conceive, build, and test, in which other scientific processes are recursively nested. By better articulating the process of DBR, this definition helps us to better craft, improve, communicate, and teach design-based research.

Introduction

Although design has existed since the beginning of human history, its rise as an educational research methodology is relatively recent. Descriptions of design-based research (DBR) in education include: Brown (1992), special issues of *Educational Researcher* (Kelly 2003), and the *International Journal of Learning Sciences* (Barab & Squire, 2004), and several edited volumes (Kelly, Lesh, & Baek, 2008; Plomp & Nieveen, 2007; Van den Akker 1999; Van den Akker, Gravemeijer, McKenney, & Nieveen, 2006). After several decades of work on DBR, some have concluded that: “as promising as the methodology is, much more effort ... is needed to propel the type of education innovation that many of us feel is required” (Anderson & Shattuck, 2012). While it is difficult to evaluate an entire research methodology (McKenney & Reeves, 2013), proponents of DBR should take these criticisms as a friendly challenge to more rigorously define DBR (Hoadley 2004).

There is general agreement that DBR should generate effective educational interventions and useful theory (Van den Akker et al., 2006, Ch. 1). We consider DBR to cover a wide range of projects, such as Margolis and Fisher’s ethnographic study of women in computer science that produced a theoretical model used to re-design a computer science department, increasing the percentage of women from 7% to 42% over 5 years (2003, p. 6). In our own work, this combination of design and research includes formative evaluations and controlled randomized experiments that result in design principles for educational games that allow us to increase both learning and interest.

DBR provides educational researchers with a process for *use-inspired basic research* (Stokes 1997; Schoenfeld 1999; Lester 2005) where researchers design and study interventions that solve practical problems in order to generate effective interventions and theory that is useful for guiding design. DBR is important because it recognizes that neither theory nor interventions alone are sufficient. The classical model of research and development, that is, basic research leading to applied research, leading to development, leading to products, does not work well (Stokes 1997). Alternatively, design, unguided by theory, is likely to be incremental and haphazard. Theory derives its purpose from application and application derives its power from theory. Our problem as DBR researchers is to devise a means of conducting DBR that reliably produces both theory and interventions.

Problems arising from the ill-definition of DBR

Unfortunately, there are many unresolved issues with DBR that arise because we lack a clear definition about what DBR is, how it is conducted, and what it produces. We describe four of these problems.

Problem 1: Uncertainty about the DBR Process

The first problem is the uncertainty about the phases of DBR--the process typically looks different depending on who conducts it. There seems to be no accepted precisely described DBR process at the level of specificity dedicated to other methodologies such as experiments or grounded theory.

Understanding the DBR process requires us to define the phases of DBR. A *phase* describes the goal of a set of *methods* within a design *process*; for example, surveys and interviews could be considered methods in a data collection phase of a research process. We need to understand the phases of design so that we can: make coherent decisions about which methods to apply and when; explain the high-level process of DBR to new researchers; effectively communicate DBR methodology in the concise form required for publication; and understand similarities and differences across different instantiations of DBR in a way that allows us to borrow

methods and improve the DBR methodology. Understanding the phases of DBR allows us to better design and to better communicate.

The integrative learning design framework (ILDF) (Bannan 2007; Bannan-Ritland 2003) is perhaps the best attempt to define the phases of DBR. However, the four phases in ILDF: exploration, enactment, local impact evaluation and broader impact evaluation, blend distinct design goals. For example, the enactment phase includes *prototyping* and the evaluation phase includes *system refinement*--both of which have a similar goal of building an intervention, but which nevertheless appear in different phases. Furthermore, the local evaluation phase and broader impact evaluation phase conflate the phase goal (of evaluating) with iteration. That is, both small-scale evaluation and large-scale evaluation serve the goal of testing, they simply occur in earlier or later iterations. Conflating phase and iteration creates a problem when we imagine an intermediate evaluation between local and broader impact that cannot be fit into the framework. Finally, it is not clear "where in this framework might randomized field trials be appropriate" (Bannan-Ritland 2003, p. 24).

Some of the most popular design processes used by practitioners like *Instructional Systems Design* (Dick, Carey, & Carey, 2008) provide a clearly articulated process and methods for designing instruction but do not attempt to define the high level phases of design or how the process might be used for research. Other popular design frameworks such as ADDIE (Analysis, Design, Development, Implementation, Evaluation) provide an umbrella term but "no real or authentic meaning" (Molenda 2003, p. 36).

Problem 2: Uncertainty about How DBR differs from Other Forms of Research

DBR is typically imagined as a form of *qualitative* research useful for building theory, that is, for addressing the problem of meaning (Kelly 2004) or used in the context of discovery (Kelly 2006, p. 177) as opposed to verifying an existing theory. While qualitative, it is distinct not just from laboratory experiments but also from ethnography and large-scale trials (Collins, Joseph, & Bielaczyc, 2004). Others argue that DBR can be productively *interleaved* with quantitative methods, for example, as a mixed methods approach crossing the field and lab (Brown 1992, p. 152-154; Kelly 2006, p. 169-171), as a point on an interleaved continuum (Hoadley 2004), or as a methodology with an agnostic stance toward quantitative and qualitative perspectives (Bannan-Ritland 2003, p. 24). Other writings describe DBR as a way to *integrate* other research methods (Collins, Joseph, & Bielaczyc, 2004, p. 39) or disciplines (Buchanan 2001) and that "methods of development research are not necessarily different from those in other research approaches" (Van den Akker 1999, p. 9). These research methods are applied in a stage appropriate manner (Bannan-Ritland 2003; Kelly 2004, 2006, p. 177). Finally, there is disagreement amongst design research theorists (outside of education) about whether design is a **science** at all, with some arguing that it is a science focused on the nature of designed artifacts (Simon 1996), others arguing that such a science is impossible because designers address problems that are not generalizable (Buchanan 1992, p. 17).

Problem 3: Uncertainty about How DBR Differs from Design, or Why Design Is Not Research

DBR proponents seek to establish DBR as a distinct and valid form of research. However, in arguing for DBR, we often ignore how DBR differs (if at all) from design as practiced in industry. Other fields, such as human-computer interaction, struggle with similar questions (e.g., Zimmerman, Forlizzi, & Evenson, 2007). Researchers claim DBR differs from design because it is: (a) research driven, that is, it addresses research questions, references literature, produces theoretical claims, and seeks to generalize beyond a specific context; and (b) involves more systematic evaluation, including formative data collection, documentation and analysis, (Bannan 2007; Edelson 2002).

Bannan (2003) points out that these are not typical attributes of practitioner methodologies like ISD (Dick, Carey & Carey, 2008). Of course, designers in industry often use qualitative methods (e.g., Beyer & Holtzblatt, 1998); develop novel, generalizable interventions described in forms such as patents or software patterns; rigorously evaluate qualitative and quantitative data through user-testing labs (Thompson 2007) and large scale experiments such as Google's A/B testing (Christian 2012). It is not clear whether there is a clear separation between design and design research or whether the distinction is artificial, or somehow peculiar to the field of education.

Problem 4: Uncertainty about What Might Make DBR Effective (If It Is)

The lack of clarity about the nature of DBR makes it difficult to justify its effectiveness as a research methodology. DBR is only useful if it allows us to reliably produce useful interventions and effective theories, "better, faster, or cheaper" than other methodologies, or to do so at least in some contexts. Without a clear description of the DBR process, we cannot make a coherent argument about the tradeoffs between DBR and other methodologies.

To increase the rigor of DBR, we need to provide a formal definition of DBR. The 4 problems arise because we do not have a clear definition about how DBR is conducted, at least not at the level of specificity provided for other methodologies. In 1992, Brown called on the field to define DBR and a decade later special issues in Educational Researcher and IJLS set out to answer that call; two decades later, we still lack a clear definition. DBR remains what organizational behavior researchers call a *low paradigm* field (or practice), where there is little technical consensus about the research questions considered important, the guiding theoretical models and, most significantly for our purposes, research methods (Pfeffer 1993). Low paradigm fields have more difficulty acquiring funding (because funders can be less certain of results), have lower journal acceptance rates (because there are greater disagreements about quality), lower collaboration and more difficulty training graduate researchers--all ultimately resulting in lower accumulation of knowledge (Pfeffer 1993; and Herrington, McKenney, Reeves, & Oliver, 2007 on DBR doctoral training). Reasonable people might disagree about the paradigmatic status of the Learning Sciences, but the calls to better define the argumentative grammar (Kelly 2004) and rigor (Hoadley 2004) of DBR suggest that we can make DBR a higher paradigmatic practice. Dede puts it bluntly: "...neither policy makers nor practitioners want what the DBR community is selling right now. We appropriately don't match the narrow conceptions of science currently in vogue at the federal level, but have much internal standard-setting to accomplish before we can put forward a defensible alternative" (2004, p.14). Twenty years on from Brown and Collins, the benefits of increased methodological consensus warrant a renewed attempt to provide a formal definition of DBR.

A Formal Definition of the Design-Based Research Process

Here we present a definition of DBR as a process that integrates design and scientific methods to allow researchers to generate useful products and effective theory for solving individual and collective problems of education. This paper focuses on describing the DBR process as part of this definition.

Design Process

The design and DBR processes consist of 6 iterative phases in which designers: *focus* the problem, *understand* the problem, *define* goals, *conceive* the outline of a solution, *build* the solution, and *test* the solution (Figure 1).

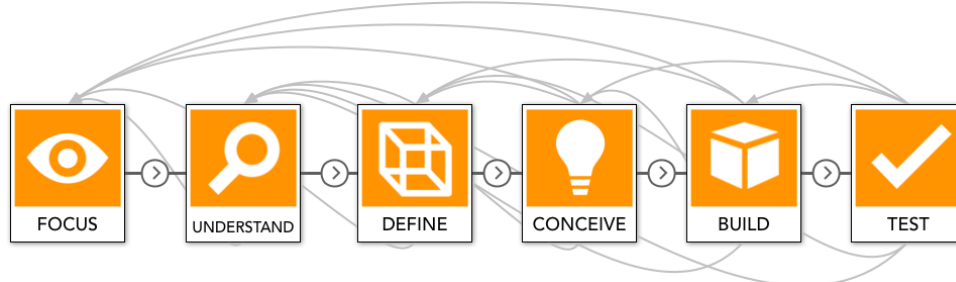


Figure 1. The design process consists of 6 iterative phases: focus, understand, define, conceive, build and test.

Focus

In the *focus* phase, designers bound the audience, topic, and scope of the project. The audience specifies whom the product serves, including learners and the other stakeholders affected, such as parents or the community. The team specifies who is designing the product and their reasons for participating. The topic specifies the general problem the product should address and how it arose. The scope specifies the constraints and the scale of the project. These issues are typically captured in a design brief.

Why: Focusing sets the direction of the project. A design is meant to achieve an intended goal and there can be no meaningful goal without some problem or opportunity to address. Focusing ensures that there is something worth designing and that the team has the expertise to succeed.

Understand

In the *understand* phase, designers study learners, domains, contexts and existing solutions. The understand phase investigates the problem through empirical methods and secondary sources, and synthesizes that knowledge into a form that can be easily used later in the process. Empirical methods include quick human-centered techniques such as observation, interviewing, surveys, data analytics, etc. Review of secondary sources focuses on: research that helps understand the problem such as models of learning and cultural contexts; analysis of current solutions to similar or related problems; and identification of design principles. The empirical data and research literature must be synthesized through methods such as identifying themes, building graphical models and creating learner personas.

Why: Typically the initial impetus for the project involves a situation in which existing solutions do not work or for which a novel solution is desirable--so designers must work to understand the nature and causes of the problem. Applicable secondary sources can be tremendously helpful in understanding the problem or

avoiding dead ends, but typically the problem arises in the first place because the root causes are unclear or because existing knowledge is insufficient to solve the problem. Furthermore, design requires detailed knowledge of user needs and context so empirical methods that can be employed quickly are almost always necessary to understand the problem.

Just as in science, discovering new features of the learning environment in the understand phase may be the core innovation of the design or theoretical contribution, such as building a better model of expertise or identifying the learning challenges in a particular domain. This includes ontological innovations, such as identifying Meta-Representational Competence (diSessa & Cobb, 2004) as a needed skill in a domain.

Define

In the define phase, designers set goals and assessments. Defining means converting an indeterminate problem, which has no solution, into a determinate problem that can be solved (Buchanan 1992). There are many ways to frame a problem. For example, suppose that the designer finds that: (a) the target learners are from immigrant communities, (b) their client wants to improve learners' performance on common core literacy and civic education standards, and (c) there are gaps in research literature about how to leverage learners' cultural resources. The problem could be defined as a question of "how might we engage students in debates about legal status?" or "how might we teach students to construct video documentaries about immigration policy?" or "how might we teach students to analyze the political values in English/Spanish-language youth media?" By completing the sentence "How might we...?" the designer selects a goal from the infinite and unknown number of goals that *could* be defined.

Why: A design focus, by definition, cannot be solved because there is no determinate (specific) goal provided--that is, there is nothing explicit to solve. It is up to the designers to define what that goal is, taking into account the goals important to the stakeholders and which can be productively solved. Only after the goal has been defined can a design be said to succeed or fail.

A novel problem definition can be the core innovation because it can lead to entirely new kinds of solutions.

Conceive

In the conceive phase, designers sketch a plan for the solution. Given a *definition* (even if implicit) the designer can plan a design intended to reach the goal. This involves imagining a solution and analyzing whether it will work. In this phase, the designer has not committed to implementing the design in a given medium, but rather creates a non-functional, symbolic or graphical representation that allows the designer to conceptually analyze the solution by determining the components of the design and how they might work together. Here, designers also develop theoretical products (diSessa & Cobb, 2004) such as design arguments (Van den Akker 1999), the underlying principles of the product, which may be of different levels of complexity (Buchanan, 2001), from communication, to artifacts, services, and systems (Penuel, Fishman, Haugan Cheng & Sabelli, 2011). The distinction between the conceive and build phase is between that of a conceptual plan constrained only by the designer's knowledge and that of a concrete prototype that is at least partially functional and constrained by a medium.

Why: Designers have a number of tools for planning, sketching, and modeling a design. These tools allow designers to test the design against their own knowledge and theory, to identify problems and improved solutions before committing to implementation in a particular medium, which can be difficult, costly, or time consuming.

Build

In the build phase, designers implement the solution. Once a design has been conceived, the designer can implement the design in a form that can be used. This implementation can be of lower or higher fidelity depending on the stage of the project and the question that the designer wants to test, which may be about a particular aspect of the educational intervention, or whether the educational intervention as conceived can achieve its goal.

Why: A design must be implemented to achieve a goal, and because a design is never completely finished, every implementation provides a prototype that can answer questions about whether the goal has been achieved.

Test

In the test phase, designers evaluate the efficacy of the solution. Iterative user-testing involves testing successive (often parallel) versions of the design at increasing levels of fidelity. Early testing of the plans produced in the conceive phase focuses on questions of relevance and consistency and then later on expected practicality, with expert reviews and walkthroughs. Later testing on prototypes constructed in the *build* phase

focus on questions of actual practicality and effectiveness using 1-1, small group, field trials and their variants (Tessmer 1993).

Testing often uses formative evaluation, which may not establish causality to the extent possible in controlled, randomized experiments, but which can quickly reject bad designs. This increases the likelihood of finding an effective design that can be verified later through summative evaluation. Some consider the boundary between formative and summative evaluation the point at which design research ends and the sciences of the artificial (Simon 1996), or in this case, rigorous evaluations testing strong causal claims of design principles, begins. We consider both valid forms of testing in DBR.

Why: Testing provides the designer with feedback about the success of the design and the validity of the theoretical propositions. It tells the designer whether the design has achieved its practical and theoretical goals.

Iteration

The design phases are not carried out in a linear sequence but rather iteratively. For example, in *building* an educational game, formative *testing* might reveal that the game is only attractive to boys, so one might return to *understand* how gender affects the likability of specific game features.

Rapid iteration is a tenet of modern human-centered design. It protects against the risks of designing interventions that are over-budget and behind schedule by quickly testing the designer's assumptions. Rather than design an entire intervention and discover only at the end that it does not work, iterative design argues for quickly building low fidelity prototypes, testing them, and re-designing--gradually evolving the intervention over time.

There is a delicate balance between planning, iteration and medium. When planning allows designers to avoid mistakes and the medium makes testing costly (e.g., building bridges), then there will be little iteration or at least a greater emphasis on lower-fidelity prototyping and modeling. However, if our ability to avoid bad designs through planning is limited and the medium makes the costs of testing low (e.g., web applications), then iteration is likely to be quick and frequent. Because education is a complex environment, our ability to predict the effect of an intervention is low. The cost of testing in education is probably relatively moderate--while the cost of implementing a lesson is low, the cost of testing may be greater depending on the type question/evaluation.

The DBR Process Includes Recursively Nested Research Processes

Scientific findings are also *products* created (or discovered) by a design process. For example, scientists may conduct an experiment in which they *focus* on a topic, *understand* the background literature, *define* a hypothesis, *conceive* of an experiment, *build* evidence by gathering and analyzing data, and finally *test* the validity of their findings, perhaps through peer-review. Qualitative research methodologies such as grounded theory follow a similar set of phases, except there the purpose is to *build theory* rather than verify a hypothesis.

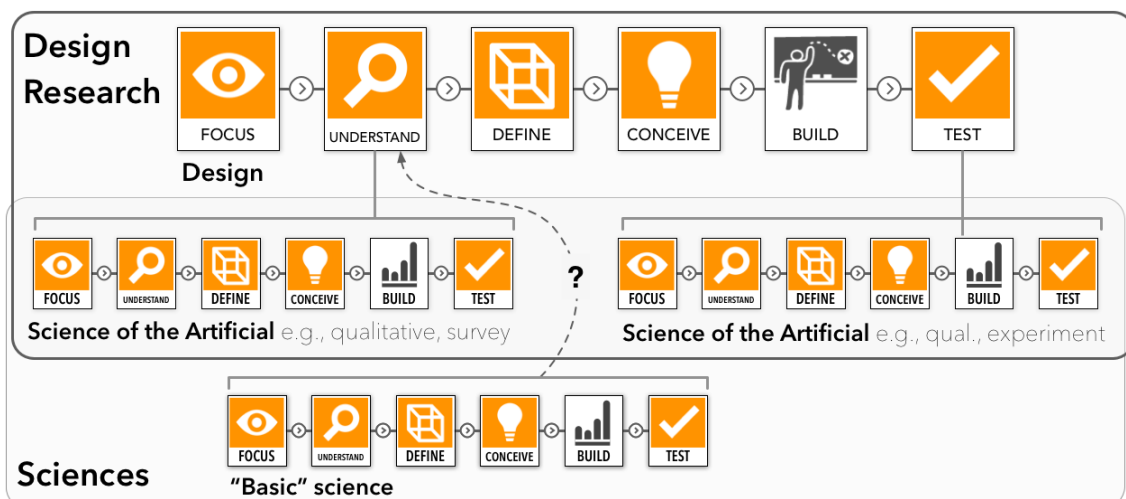


Figure 2. Scientific research methodologies (both qualitative and quantitative) follow a design process and produce products such as theories and models that can be incorporated into the design of another product such as an educational intervention.

Products that serve one purpose, such as verification of a hypothesis, can be used as components in the design of another product, such as an educational intervention (Figure 2). That means that in designing a learning environment, we might conduct other sub-design processes (such as a qualitative study or an

experiment) as part of the DBR process. For example, a DBR study of a journalism curriculum might conduct a qualitative study about learners' media practices in the understand phase, or a controlled randomized test of the curriculum in the test phase. In other words, design processes can be recursively nested within each other. This explains the shape-shifting nature of DBR--DBR looks like other forms of research because it incorporates these methodologies to do its work.

Stage Dependent Search

By understanding how design incorporates other scientific design processes, we can make a more compelling argument for why DBR can be an effective educational research methodology. Design research uses a stage-dependent search strategy (Bannan-Ritland 2003; Kelly 2004, 2006), in which designers choose different build and test methods depending on the stage of the design. In early stages of a project, such as when the problem context is poorly understood and there are few effective implementations, researchers are likely to produce unsuccessful designs, so they must choose a research and development strategy that allows them to quickly reject failures and understand the theoretical issues that must be addressed. So in the early stages of a project, researchers should focus on low-fidelity prototyping and collecting the minimal amount of data needed to quickly reject failure and identify potential successes. As researchers identify promising prototypes they can focus on theory building with qualitative methods to better understand the issues a design might address and the mechanism through which it affects learning. Once researchers have a plausible, well-grounded theory and an implementation with some evidence of success, they can conduct randomized controlled experiments to verify the efficacy of the theory and intervention. If researchers use randomized, controlled, experiments at the beginning stages of a complex design problem, they are likely to waste resources verifying a bad design. Likewise, if researchers never advance beyond theory building and radically novel designs, they are unlikely to provide strong evidence for the efficacy of an intervention or principle.

Resolving the Uncertainties

This formal definition of the DBR process resolves the uncertainties presented earlier.

Problem 1 resolution: *a clear definition of the phases of DBR*. The formal definition resolves the uncertainty about the phases of design in a way that allows us to better conduct DBR, train new researchers, improve DBR methodology, and communicate process within and outside the DBR community.

Problem 2 resolution: *DBR differs from other research in that it designs a product while using other methodologies as nested processes (sub phases) of design*. The formal definition shows how DBR differs (or rather does not differ) from other forms of research. DBR incorporates other scientific design processes into the design process for creating educational interventions in a recursive, nested manner.

Problem 3 resolution: *DBR differs from design practice in that it does not just produce an educational intervention but makes use of nested scientific processes to produce theory*. The formal definition also shows how DBR differs from "normal" design. By incorporating scientific processes, DBR produces theories connected to the literature and more rigorously tests interventions. Of course, there is no hard line separating the work of practitioners and researchers because practitioners use similar methods--the difference is one of degree and intent.

Problem 4 resolution: *DBR produces gains by deploying the appropriately nested scientific process at a given stage of development*. The formal definition shows how DBR efficiently develops theory by quickly identifying plausible interventions and constructs in early phases that are more rigorously verified in later stages.

Applying the Definition

A better understanding of the DBR process helps us to do better design research, train new researchers, improve DBR methodology, and communicate process within and outside the DBR community.

Better Design

Defining the DBR process helps us to better determine which methods to use and when. For example, when planning DBR projects, thinking about the *test* phase has prevented us from jumping to formal evaluation too early or dwelling in theory building too long. For ill-defined problems, we have used the phases to justify spending more effort applying methods from the understand phase. The phases also make clear when we have only implicitly defined the goals and design arguments for a project. DBR projects work under constraints of people, resources, and time, and the phases have allowed us to more deliberately deploy those resources.

Training New Researchers

There is a bewildering array of methods applicable in DBR projects and it is challenging for new researchers to make sense of these methods (Herrington et al., 2007). We use the DBR phases to explain how the design research process works at a high level, to help novices organize sets of research methods, and to explain the

meta-cognitive strategies we use to conduct design-based research. Just as design phases help researchers think precisely, they also serve as a tool to make design logic explicit to new DBR researchers.

Improving DBR Process

A clear definition of the DBR phases also helps us to improve the process. In struggling to consolidate learner data gathered in the *understand* phase, we have used human-centered design methods for synthesizing user data, such as personas. Or in rethinking curricula as services, we have applied *conceive* methods from service design such as journey maps, swimlanes and service blueprints. The phases allow DBR researchers to more easily borrow methods from other methodologies just as human-centered design has borrowed methods from methodologies such as ethnography. The DBR phases serve as a Rosetta Stone for translating and synthesizing design processes from other methodologies.

Likewise, we can use the phases as an analytical tool for judging design processes and potential contributions. For example, noticing that the ADDIE process does not clearly identify *focus* and *define* stages, or that the ILDF conflates phase and iteration. By identifying gaps, the design phases allow us to suggest new methods that can be applied to improve these processes. Furthermore, each phase identifies the locus of potential design research contributions when clearly defined.

Communicating Research Process

We have also used the phases to describe the choices made during a DBR project and why those were effective. In publishing research and grant applications, the phases more concisely communicate the past history or future plans of a DBR project. Unfortunately, the lack of shared vocabulary and conventional methodology creates a communication barrier, for example, in grant applications that require lengthy descriptions of planned cycles of design, iteration and testing.

Well-defined DBR phases allow us to explain the logic of DBR to other researchers. For example, quantitative psychologists may see the lack of inter-rater reliability in the early stages of a DBR project as a lack of rigor. Researchers from other disciplines will naturally judge DBR by the methodological standards of their own discipline. However, when DBR researchers explain the methodological logic of shifting from an early focus on design concepts and theory building to a later focus on verification, we've found that those outside the discipline are often sympathetic to the aims of DBR. The problem is not that researchers from other disciplines are unaware of the methodological challenges of developing new interventions and theories (which DBR was developed to address), the problem is that other researchers will only accept DBR's alternative approach to addressing these methodological challenges when DBR researchers clearly and precisely articulate the rationale behind the DBR methodology.

Conclusion

We have defined DBR as a process that integrates design and scientific methods to allow researchers to generate useful educational interventions and effective theory for solving individual and collective problems of education. This definition of the DBR process is neither "a way" nor "the way" to conduct DBR, rather, it describes the fundamental nature of all forms of DBR in order to help us better communicate and think about DBR. This definition is not just an academic exercise, but necessary to establish DBR as a high paradigm methodology, allowing us to better replicate the design process, to apply methods from other design methodologies, to better teach DBR to new design researchers, to acquire more resources, and ultimately to accumulate theory relevant to practice. By formally defining DBR, we establish its credibility as a legitimate methodology of educational research.

References

- Anderson, T., & Shattuck, J. (2012). Design-Based research A decade of progress in education research? *Educational Researcher*, 41(1), 16-25.
- Bannan, B. (2007). The integrative learning design framework: An illustrated example from the domain of instructional technology. In T. Plomp & N. Nieveen (Eds.), *An introduction to educational design research* (pp. 53-73). Netherlands Institute for Curriculum Development.
- Bannan-Ritland, B. (2003). The role of design in research: The integrative learning design framework. *Educational Researcher*, 32(1), 21-24.
- Barab, S., & Squire, K. (2004). Introduction: Design-based research: Putting a stake in the ground. *The Journal of the Learning Sciences*, 13(1), 1-14.
- Beyer, & Holtzblatt. (1998). *Contextual design: Defining customer-centered systems*. San Francisco: Morgan Kaufmann.
- Brown, A. L. (1992). Design experiments: Theoretical and methodological challenges in creating complex interventions in classroom settings. *The Journal of the Learning Sciences*, 2(2), 141-178.
- Buchanan, R. (1992). Wicked problems in design thinking. *Design Issues*, 8(2), 5-21.

- Buchanan, R. (2001). Design research and the new learning. *Design Issues*, 17(4), 3-23.
- Christian, B. (2012, April 25). The A/B test: Inside the technology that's changing the rules of business. *Wired*.
- Collins, A., Joseph, D., & Bielaczyc, K. (2004). Design research: Theoretical and methodological issues. *The Journal of the Learning Sciences*, 13(1), 15-42.
- Dede, C. (2004). If design-based research is the answer, what is the question? A commentary on Collins, Joseph, and Bielaczyc; Disessa and Cobb; and Fishman, Marx, Blumenthal, Krajcik, and Soloway in the JLS special issue on design-based research. *The Journal of the Learning Sciences*, 13(1), 105-114.
- Dick, W., Carey, L., & Carey, J. O. (2008). *The systematic design of instruction* (7 ed.). New York: Pearson.
- diSessa, A. A., & Cobb, P. (2004). Ontological innovation and the role of theory in design experiments. *The Journal of the Learning Sciences*, 13(1), 77-103.
- Edelson, D. C. (2002). Design research: What we learn when we engage in design. *The Journal of the Learning Sciences*, 11(1), 105-121.
- Herrington, J., McKenney, S., Reeves, T., & Oliver, R. (2007). Design-based research and doctoral students: Guidelines for preparing a dissertation proposal. In C. Montgomerie & J. Seale (Eds.), *Proceedings of the World Conference on Educational Multimedia, Hypermedia and Telecommunications*, 2007 (pp. 4089-4097). Chesapeake, VA: Association for the Advancement of Computing in Education.
- Hoadley, C. M. (2004). Methodological alignment in design-based research. *Educational Psychologist*, 39(4), 203-212.
- Kelly, A. (2004). Design research in education: Yes, but is it methodological? *The Journal of the Learning Sciences*, 13(1), 115-128.
- Kelly, A. E. (2003). Theme issue: The role of design in educational research. *Educational Researcher*, 32(1), 3-4.
- Kelly, A. E. (2006). Quality criteria for design research. *Educational Design Research*, 107-118.
- Kelly, A. E., Lesh, R. A., & Baek, J. Y. (2008). *Handbook of design research methods in education: Innovations in science, technology, engineering, and mathematics learning and teaching*. Abingdon, Oxon: Routledge.
- Lester, F. K. (2005). On the theoretical, conceptual, and philosophical foundations for research in mathematics education. *ZDM*, 37(6), 457-467.
- Margolis, J., & Fisher, A. (2003). *Unlocking the clubhouse: Women in computing*. Cambridge, MA: The MIT Press.
- McKenney, S., & Reeves, T. C. (2013). Systematic review of design-based research progress is a little knowledge a dangerous thing? *Educational Researcher*, 42(2), 97-100.
- Molenda, M. (2003). In search of the elusive ADDIE model. *Performance Improvement*, 42(5), 34-37.
- Penuel, W. R., Fishman, B. J., Haugan Cheng, B., & Sabelli, N. (2011). Organizing research and development at the intersection of learning, implementation, and design. *Educational Researcher*, 40(7), 331-337.
- Pfeffer, J. (1993). Barriers to the advance of organizational science: Paradigm development as a dependent variable. *The Academy of Management Review*, 18(4), 599. doi:10.2307/258592
- Plomp, T., & Nieveen, N. (2007). An introduction to educational design research. In *Proceedings of the Seminar Conducted at the East China Normal University, Shanghai (PR china), November 23-26, 2007*. Enschede: Netherlands Institute for Curriculum Development.
- Schoenfeld, A. H. (1999). Looking toward the 21st century: Challenges of educational theory and practice. *Educational Researcher*, 28(7), 4-14.
- Simon, H. A. (1996). *The sciences of the artificial* (3rd ed.). Cambridge, MA: MIT Press.
- Stokes, D. (1997). *Pasteur's quadrant: Basic science and technological innovation*. Washington, DC: Brookings Institution Press.
- Tessmer, M. (1993). *Planning and conducting formative evaluations: Improving the quality of education and training*. Abingdon, Oxon: Routledge.
- Thompson, C. (2007). Halo 3: How Microsoft labs invented a new science of play. *Wired Magazine*, 15(09).
- Van den Akker, J. (1999). Principles and methods of development research. *Design Methodology and Developmental Research in Education and Training*, 1-14.
- Van den Akker, J., Gravemeijer, K., McKenney, S., & Nieveen, N. (2006). *Educational design research*. Abingdon, Oxon: Taylor & Francis.
- Zimmerman, J., Forlizzi, J., & Evenson, S. (2007). Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 493-502).

Acknowledgments

We thank Bruce Sherin, Pryce Davis, the Delta Lab, and the anonymous reviewers for their thoughtful critiques of this paper. This work supported by the National Science Foundation Grant No. IIS-1320693 and No. IIS-1217225.