

Bridging learning research and teaching practice for the public good: The learning engineer

Candace Thille, Stanford University



About this Research

Higher education's dual missions of research and teaching ideally position the sector to make rapid progress in discovering and implementing the most effective processes for teaching and learning. Historically, though, a gap has existed between learning research and teaching practice in higher education.

To help campus leaders understand how that gap can be bridged, the TIAA Institute invited this work by Candace Thille, who outlines a new academic role—the *learning engineer*. Learning engineers, in collaboration with researchers and practitioners, can design learning environments and data systems that yield predictive and explanatory models of student learning that support course improvement, instructor insight, and student feedback. Further, they can support the selection of useful knowledge modeling approaches for specific students, contexts, and learning goals. In short, learning engineers can facilitate rapid progress in the basic science of human learning.

Citation for this paper: Thille, Candace M. (2016). *Bridging Learning Research and Teaching Practice for the Public Good: The Learning Engineer*. New York, NY: TIAA Institute.

About the TIAA Institute

The TIAA Institute helps advance the ways individuals and institutions plan for financial security and organizational effectiveness. The Institute conducts in-depth research, provides access to a network of thought leaders, and enables those it serves to anticipate trends, plan future strategies and maximize opportunities for success. To learn more, visit www.tiaainstitute.org.

Executive Summary

The pressure on institutions of higher education to address the challenges of increasing access, reducing cost, and educating a greater number and diversity of students is intense. Higher education's dual missions of research and teaching ideally position the sector to address these challenges rapidly by discovering and enacting the most effective processes for teaching and learning. However, resource constraints and many of the traditional structures and processes in higher education impede a functional, bi-directional relationship between research and practice. Further, students most in need of a robust personalized academic support system often are enrolled at institutions with the most resource constraints.

Personalized and adaptive educational technologies have great potential for good, but there is also potential for harm if careful, rigorous thought is not devoted to understanding the learning process, specifying the outcomes of interest, designing how and from whom the data are collected, and choosing how data are modeled and represented. A new academic role, the *learning engineer*, is needed to bridge the chasm between learning research and teaching practice in higher education. Learning engineers, in collaboration with researchers and practitioners, will design learning environments and data systems that provide student and instructor feedback, support continuous improvement learning design and facilitate rapid progress in the science of human learning.

Any opinions expressed herein are those of the author, and do not necessarily represent the views of TIAA, the TIAA Institute or any other organization with which the author is affiliated.

Key Takeaways

- Calls for accountability on the part of higher education give rise to more fundamental questions about how people learn and how one knows when learning is happening.
- While learning research has produced results that could be used to enhance education, the results of that research often have not translated into successful changes in teaching practice or student learning.
- The dual missions of research and teaching ideally position colleges and universities to make rapid progress in discovering and enacting the most effective processes for achieving human learning, but traditional structures and processes often impede a functional, bi-directional relationship between research and practice.
- In all sectors, advances in machine learning, data science, crowdsourcing and computation are enabling a much larger part of human processes and decision-making to be done by machines, which are rapidly becoming a core part of the teaching process in higher education.
- The design of educational technologies, and the modeling and interpretation of data for pedagogical decision making, are active areas of research. Without transparency and peer review, such efforts are alchemy, not science.
- A new academic role, the *learning engineer*, is needed to bridge the historic chasm between learning research and teaching practice in higher education. Learning engineers, in collaboration with researchers and practitioners, will design learning environments and data systems that yield predictive and explanatory models of student learning that support course improvement, instructor insight, student feedback, and the basic science of human learning.

American higher education today faces multiple challenges that are at once immense and vexing. Global demand far exceeds capacity. The achievement gap between rich and poor is widening. Institutions of higher education are seeing increasing variability in the student population's background knowledge, relevant skills, and future goals. Tuition and fees are outstripping inflation at an alarming rate. Completion rates among populations most in need are unacceptably low, especially if, as a society, our intention is to graduate students who have developed the knowledge and skills needed not only to secure employment but also to be engaged citizens.

Public colleges and universities in the United States are experiencing the pressure to serve not just more students, but a greater variety of students. In the face of shrinking state budgets, they are being asked to increase attainment rates and reduce the cost of instruction. They are pinched for money, squeezed for space, and find themselves under unprecedented pressure to see to it that the growing number of students who matriculate on their campuses each year complete their degrees. Public and private institutions are coming under increased pressure from policymakers to increase admission and improve graduation rates in order to maintain access to the \$180 billion a year that the federal government invests in student loans, grants and tax benefits¹.

These pressures have given rise to discussions about the fundamental purpose of higher education: Is the principal goal of higher education's teaching mission to prepare graduates to participate effectively in a diverse democracy, or to prepare graduates to participate effectively in the workforce? Or as some argue, is it "to hone the skills of analysis and critical inquiry, while helping students develop the habits of mind and capacities of creative problem solving and forming independent judgments about complex questions"—that is, the skills and habits of mind that are critical for both goals (Pasquerella, In Press). Debates for and against a narrow vocational focus versus a broad liberal approach to higher education have been waged since the founding of our nation. As Roth (2014) documents, "tensions between the lofty and practical ideals for higher education" have been expressed in the writings of Thomas Jefferson and Benjamin Franklin, as well as in the great debates between W.E.B. Du Bois and Booker T. Washington.

Irrespective of the merits on each side of the debate, the question remains: Are the processes of higher education effective in achieving their purpose? This seemingly simple question about the effectiveness of instructional practices leads to more fundamental questions about how people learn and how one knows when learning is happening.



Questions about the effectiveness of instructional practices lead to more fundamental questions about how people learn and how one knows when learning is happening.

1. See, for example, proposed legislation that would tie access to federal aid to admission and graduation rates <https://www.washingtonpost.com/news/grade-point/wp/2016/09/22/ipartisan-senate-bill-takes-a-carrot-and-stick-approach-to-boosting-college-access-graduation-rates/>

Learning is an active process of meaning construction in which the learner maps new information onto prior knowledge.



Studies have shown the critical role of prior knowledge and the cognitive mechanisms of chunking and knowledge reorganization.

The Science of Learning

As it happens, creating new knowledge and answering fundamental questions is the primary aim of higher education's research mission. Creating new knowledge and answering fundamental questions about human learning is the focus of the science of learning—an interdisciplinary field comprising cognitive science, neuroscience, education, psychology, sociology, economics and computer science. Learning researchers historically have followed a systematic and empirical approach to achieving the goal of understanding, predicting and explaining human learning. Much of what is known about learning comes from an accumulation of evidence from multiple studies in laboratories and classrooms.

Results from the science of learning can help resolve controversies about common education practices that often are based on ideology and opinion. The science of learning has demonstrated that learning is an active process of meaning construction in which the learner maps new information onto prior knowledge. Simply being exposed repeatedly to material, such as re-reading text or replaying recorded lectures, is not a successful strategy for producing learning. Lecturing is a very common instructional practice that has come under fire for being anti-constructivist and ineffective. Research by Schwartz and Bransford (1998) showed that lecturing can be effective when the students have sufficient prior knowledge to allow them to construct meaning from the lecture. The researchers found that students learned from lectures after engaging with challenging problems that highlighted the precise issues the lecture addressed. Note that while the research demonstrated that the common instructional practice of lecturing can be quite effective when paired with problem solving, the sequence of problem solving and lecture is the reverse of the sequence that is commonly used both in traditional practice and in the new “flipped classroom” models.

While some learning research can help identify the conditions under which common instructional strategies will likely be most effective, other research produces counter-intuitive results that require either going against one's intuitions or deviating from widely accepted instructional practices. For example, research by Richland, Zur, and Holyoak (2007) and by Schwartz, Chase, Oppenzo, and Chin (2011) showed that providing students with multiple analogous examples and having them induce the underlying structure is much more effective than the common instructional practice of giving students a single example and explanation of the underlying principle.

Most learning research uses external behavioral measures to make inferences about changes in the learner knowledge state. Using these inferential methods, many studies have shown the critical role of prior knowledge and the cognitive mechanisms of chunking and knowledge reorganization. As students develop expertise, researchers can observe how they group smaller pieces of information into larger chunks and reorganize their knowledge

into categories aligned to solution-relevant principles in a domain. These learning processes help explain why some instructional strategies are more effective for some students than others, and can provide guidance about how to differentiate instruction to meet the needs of each learner. Novices who have limited prior knowledge learn better from a step-by-step demonstration of how to perform a task or how to solve a problem that helps them build meaningful knowledge structures, whereas more advanced students who have sufficient prior knowledge benefit more from open-ended problems.

Neuroscience is revealing that learning can be observed through changes in the brain. Learning relies on the ability to flexibly integrate information across specialized regions of the brain. Recent research by Shine et al. (2016) demonstrated that integration across separate neural regions enabled faster and more accurate performance on cognitive tasks, confirming a direct link between cognitive performance and the dynamic reorganization of the network structure of the brain.

Moving beyond examining learning processes from the perspective of an individual's cognitive processes or brain structure, social science research has shown the critical roles that environmental and social factors play in learning. Many studies have demonstrated how cues in the learning context can either facilitate or threaten one's sense of belonging and social identity. Stereotype threat, one such phenomenon that has been widely studied for over two decades, is "the concrete real-time threat of being judged and treated poorly in settings where a negative stereotype about one's group applies" (Steele, 2002, pg. 385). Stereotype threat has been shown to contribute to systematic underperformance of individuals in a variety of contexts². The negative consequences go well beyond poor performance on a task and include hindering learning (Taylor & Walton, 2011); undermining cognitive capacity (Schmader & Johns, 2003); reducing self-regulatory abilities (Baumeister, DeWall, Ciarocco, & Twenge, 2005); depressing motivation (e.g., Steele et al., 2002; Steele, 1997); self handicapping (Stone, 2002; Keller, 2002); distancing oneself from the stereotyped group (Cohen & Garcia, 2005; Pronin, Steele, & Ross, 2004); and redirecting a learner's aspirations or career paths (Steele, James, & Barnett, 2002; Davies, Spencer, Quinn, & Gerhardstein, 2002). Much of the same research has revealed multiple strategies through which performance deficits and other negative consequences due to stereotype threat can be reduced or eliminated.

A great deal of learning research has produced results that could be used to enhance education; however, the results of that research often have not translated into successful changes in teaching practice or student learning.

Novices with limited prior knowledge learn better from a step-by-step demonstration of how to perform a task or solve a problem that helps them build meaningful knowledge structures.

More advanced students who have sufficient prior knowledge benefit more from open-ended problems.

2. Systematic underperformance on tasks from stereotype threat have been shown on groups and tasks as diverse as African American students (Cohen, Garcia, Purdie-Vaughns, Apfel, & Brzustoski, 2009; Steele & Aronson, 1995; Steele, 1997), first-generation college students (e.g., Harackiewicz et al., 2014; Stephens, Fryberg, Markus, Johnson, & Covarrubias, 2012), and women in historically male-dominated engineering programs (Miyake et al., 2010; Spencer, Steele, & Quinn, 1999; Walton, Logel, Peach, Spencer, & Zanna, 2015). white athletes (Stone, Lynch, Sjomering, & Darley, 1999), women in negotiation (Kray, Galinsky, & Thompson, 2002), gay men in childcare (Bosson, Haymovitz, & Pinel, 2004), and women in driving (Yeung & von Hippel, 2008).

The Chasm Between Learning Research and Teaching Practice

One explanation for this lack of impact is that the burden for changing the practice of teaching has fallen largely on the shoulders of individual faculty members, who are being asked to incorporate research results into their classroom practice. Research results are not easily accessible to practitioners, as research published in refereed journals typically does not provide clear, conclusive answers on most issues of practice. Dense, jargon-laden academic publications are likely to be ignored by faculty outside of the psychology or education domains. Moreover, faculty members teaching in their disciplines rarely have time to conduct thorough searches for learning research results, or to synthesize relevant research findings to address specific learning challenges.

Over the past decade, numerous books, reports and articles directed at practitioners and students have attempted to address this barrier by translating results from learning research into usable guiding principles for teaching and learning (see Schwartz, Tsang, & Blair, 2016; Ambrose, Bridges, DiPietro, Lovett, & Norman, 2012; Benassi, Overson, & Hakala, 2014; Clark & Mayer, 2008; Dunlosky, Rawson, Marsh, Nathan, & Willingham, 2013; Mayer, 2011). While these guiding principles are a good place to start in supporting students and faculty to engage in generally better teaching and learning practices, instructional practices that have shown evidence of effectiveness in experimental settings, or in specific contexts, often have not shown effectiveness in new or different contexts.

Learning is complex, and the failure of research results to demonstrate effectiveness in new contexts can be because the research results have been oversimplified in their translation to practice, or because the theory or model is not sufficiently robust to account for the complexity of the new context. In either case, the traditional linear technology transfer model, which assumes a non-problematic relationship between the research base and teaching practice, has not been optimal.



The dual missions of research and teaching ideally position higher education institutions to make rapid progress in discovering and enacting the most effective processes for achieving human learning.

It would be logical to conclude that the dual missions of research and teaching would ideally position institutions of higher education to make rapid progress in discovering and enacting the most effective processes for achieving human learning. After all, the researchers (neuroscientists, cognitive scientists, and social scientists), the practitioners (the faculty teaching in the sciences, humanities, and professional schools), and those who would benefit from the progress in research and practice (the students), are geographically and temporally colocated. Some examples exist of researchers and practitioners taking advantage of this geographic and temporal collocation to simultaneously address the challenge of education practice and make progress on fundamental questions of human learning³.

3. The collaborative work of chemistry faculty and learning scientists in the development of activities to teach chemical equilibrium in the Open Learning Initiative introductory chemistry course resulted in both an equilibrium module that significantly improved learning outcomes, particularly for the lowest-performing students, as well as a contribution to the body of learning theory suggesting how prior knowledge and the conceptual content of diagrams influences multimedia learning.

However, such examples are not common, many of the current structures and processes in higher education such as funding models, competition among institutions, and processes for tenure and promotion, are barriers to taking advantage of the obvious benefit of this geographic and temporal colocation. Longer-term structural and policy changes will take time and the path to making those changes is not simple.

The Learning Engineer

One step that institutions of higher education can take that would simultaneously address the short-term challenge of improving teaching and learning practice and lay the groundwork for the longer term path to changing the relationship between learning research and teaching practice is to create a new academic role, the learning engineer.

The Learning Engineer: An Interdisciplinary Role

Learning engineers will have deep understanding across disciplines that comprise the complex tasks of designing, supporting, and continuously improving instructional methods and technologies. They will be knowledgeable in those aspects of neuro, cognitive, education, and social science that are material to understanding human learning in diverse contexts. Learning engineers will be competent in using the methods and tools of assessment, cognitive science, computer science, and data science; and practiced in the art of unpacking expertise to create effective educational technologies, data models and analytic systems. Further, they will be able to develop fluency in the discourse of the subject domains in which they are working and be competent at communication with, and facilitation of, interdisciplinary teams. In addition to translating results from research into practice, learning engineers will assure that research is use-inspired and informed by practice. They will help converge the divergent thinking and language of those on both sides of the research and practice chasm. Through this inherently collaborative role, the learning engineer will be the catalyst for a much-needed change in the interplay of learning research and teaching practice in higher education.

While the suggestion to bring engineers into education to increase efficiency was made over a century ago (Munroe, 1912), the role of the learning engineer was introduced by the economist and computer scientist Herbert Simon fifty years ago. In a speech to the presidents forum of the American Council on Education, Simon (1967) suggested creating the role of a learning engineer to mitigate the problem of a faculty largely untrained in the

profession of teaching. He described the primary responsibility of the learning engineer as “working in collaboration with members of the faculty whose interest they can excite, they design and redesign learning experiences in particular disciplines” (p. 77) In addition, their role was “to bring the campus into contact with the lively and significant current activity in cognitive psychology, and with developments related to learning machines and computer-aided instruction” (p.77).

The role of the learning engineer that Simon and others have imagined is grounded in the traditional linear technology transfer model of research to practice. The role is unidirectional and translational—to help faculty apply results from learning research to foster improvement in their teaching practice. However, it is just as important that questions explored in learning research be informed by teaching practice. The role of the learning engineer is no longer unidirectional, but rather a role that facilitates changing the suboptimal linear technology transfer model to a model in which research and practice form a virtuous cycle of continuous improvement.

Advances in information technology—along with the explosion in development of educational technologies that mediate the teaching and learning process—make the need for learning engineers even more pressing. In 2016, the role of the learning engineer is not simply to assist in designing more effective classroom practices, but also to build the educational technology and back-end data systems that support instructional practice, student learning, and learning research.

Several advances in computer science are revolutionizing other fields and are fueling the interest in using technology not only to provide greater access to education but also to transform instruction. Foremost among them is the maturing of machine learning—the study and construction of algorithms that can learn from and make predictions on data. Computer scientists have made progress in designing learning algorithms, as well as scaling existing algorithms, to work with extremely large data sets and build models based on data. A branch of machine learning, reinforcement learning, is a framework that shifts the focus of machine learning from simple pattern recognition to experience-driven sequential decision making. Recurrent neural networks, also called “deep learning,” belong to a class of dynamic models that connect artificial neurons over time. Adaptive, artificial neural networks are trained using a method called backpropagation. The use of recurrent neural networks has rapidly advanced progress on several time series tasks, such as speech recognition and image captioning (LeCun, Bengio, & Hinton, 2015). Learning researchers and computer scientists are currently debating the value of deep learning for tracing human knowledge development (Piech et al., 2015; Khajah, Lindsey, & Mozer, 2016). Learning engineers can design the analytic research systems that improve knowledge modeling, and can support selection of useful knowledge modeling approaches for specific students, contexts, and learning goals.

Learning engineers can design analytic research systems that improve knowledge modeling and support selection of useful approaches for specific students, contexts, and learning goals.

Increases in the performance of information processing algorithms have been accompanied by significant progress in hardware technology, cloud computing resources, widespread web-based data gathering, and crowdsourcing methods that devise innovative ways to harness distributed human intelligence. Educational technology provides the opportunity to capture data during the learning process to feed the machine learning models.

Networked online learning environments already collect massive amounts of student interaction data; however, the insights into student learning that can be gleaned from those data are limited by the type of interaction that is observable and by the meaning associated with the data generated by the interaction. In other words, current machine learning algorithms, modeling on the data generated by student interactions with current systems, are powerful for making predictions but not for generating explanations. This limits their capacity to support transparent pedagogical decision making.

The learning engineer would design learning environments and data systems that generate explanatory models of a student's learning and support course improvement, instructor insight, student feedback, and the basic science of human learning.

The data collected from technology-mediated learning environments can provide a detailed record of the students' learning process, making that process amenable to scientific study. In designing activities in technology-mediated environments, the learning engineer would design a range of tasks—grounded in current theories of human learning—that structure performances, automatically collecting enough pieces of evidence that can be identified and aggregated to provide a reasonably coherent picture of the learners' knowledge state and of the learning process. The data generated by such student interactions provide granular detail. Aggregating meaningful fine-grained evidence is easier than trying to break down coarse-grained evidence post-hoc into smaller pieces. The learning engineer must be skilled in assessment, and take a systematic approach to designing, collecting, analyzing and interpreting information to foster students' learning and development.

Adaptive and Personalized Educational Technologies

Much of the excitement in using educational technology to transform education is the promise that adaptive systems, with limited human intervention, can personalize instruction for large numbers of students. Adaptive systems collect data from the student interactions with the technology and then model those data to make predictions about the student knowledge state. Historically, the cognitive models that drive these systems have been built by mapping how experts in the discipline under study think and perform. Learning engineers will need to be adept at extracting and modeling the implicit and declarative aspects of how experts think, by using traditional methods of cognitive task analysis and new methods of using data to build and refine cognitive models.

Educational technology to transform education holds the promise that adaptive systems, with limited human intervention, can personalize instruction for large numbers of students.



Large data sets generated by students in thousands of contexts, combined with new machine learning algorithms, provide an unprecedented opportunity to discover new patterns predictive of student success.

Adaptive systems make recommendations about what a student should do next based on the prediction generated by the model. Hence, the data that are collected from the student interactions in these systems are modeled and used for pedagogical decision making, either so the system can make autonomous decisions (e.g. decide what learning task to give the student next) or to give information to the instructor to support their decision making (e.g. learning dashboards that present a visual representation of a student's predicted competence on specified learning outcomes).

The large data sets generated by use of educational technologies by thousands of students in thousands of contexts, combined with new machine learning algorithms, provide an unprecedented opportunity to discover new patterns that are predictive of student success. Given that we know that instructional methods can be differentially effective for different groups and individuals in different contexts, it would be problematic if the models are trained on unrepresentative populations in narrow contexts. It would also be problematic if the models are trained on data resulting from unconscious biases (e.g. gender, class and/or racial bias) that influenced the outcome decisions. Pattern recognition and prediction tools have the potential to provide new kinds of transparency about data and inferences. Such tools may be used by learning engineers to detect, remove or reduce human bias. If such patterns are not critically evaluated before they are used to inform the design of learning environments, we run the risk of building tools to reinforce existing norms that reproduce inequality.

Currently, adaptive educational systems are being designed and built mostly outside of the academy and sold into the academy as tools and products to facilitate innovation in teaching and learning. This work, however, encompasses numerous areas of research in the science and engineering of learning, including: how to design the technology to support teaching and learning; which data to collect; the factors to include in a predictive model; how to weight those factors; which modeling approaches and algorithms to use; how new patterns that are revealed in the data should be interpreted and used; the ideal divisions of tasks between humans and machines based on their differing capabilities and costs; the boundaries for when to use predictions for autonomous decision-making or for supporting human decision making; and what information to represent, and how to represent it, in support of human decision making. In an academic context, all of these areas of research must be transparent and subject to peer review and challenge—or the process is better described as alchemy, not science.

The Opportunity and The Risk: Big Data and Educational Data Mining

Educational data mining (EDM), knowledge modeling, and the teaching and learning decision-support systems that can be built from data and models hold tremendous potential to improve instruction and student learning outcomes.

EDM detects statistical relationships in a dataset. The accumulated set of discovered relationships—or models—encompass what a learner knows, a learner's affective state, and a learner's behaviors and motivations. EDM techniques have enabled detection of a wide range of constructs. When patterns in the data can be discerned and interpreted, the processes of classifying learning activities, classifying learners, predicting learning outcomes, and recommending appropriate actions can be automated.

The systems and algorithms used to model the data are not neutral. Any system built using data will reflect the biases and decisions made when collecting that data, as well as the behaviors and judgements of the groups and individuals from whom the data are collected. Within education, there is evidence that diverse populations must be used to develop detectors because not all models work for all students even when overall model quality is high (cf. Ocumpaugh, Baker, Gowda, Heffernan, & Heffernan, 2014). There are multiple examples from other sectors showing the negative impact of systems built on biased or unrepresentative data.

Higher education has the opportunity to use research in human learning, large data sets, data mining, data modeling, and the design of reporting systems to detect and counterbalance unconscious implicit biases. For example, mining of large data sets in one study already has revealed that significant gendered performance differences are ubiquitous in large introductory STEM lecture courses. This has led to hypotheses that evaluation methods used in STEM lecture courses interact with stereotype threat to create gendered performance differences (Koester, Grom, & McKay, 2016). If models are not transparent and critically evaluated before they are built into predictions systems, data mining and the resulting decision support systems will simply reproduce existing patterns, inherit the prejudice of prior decision-makers, and further entrench biases in the education system.

It may be possible to bridge the longstanding gap between learning research and teaching practice, creating educational technologies that improve learning and support progress in the fundamental sciences of human learning.

Conclusion

Learning engineers can serve as lynchpins in effectively shifting the relationship between learning research and teaching practice away from the current suboptimal, linear technology transfer model. In collaboration with learning researchers and practitioners, learning engineers, will be able to bridge the longstanding gap between learning research and teaching practice in higher education, and create educational technologies that improve learning and simultaneously support progress in the fundamental sciences of human learning.

Some (Staton, 2013; Craig, 2015) have argued that information technology is a disruptive force that will make education more efficient by forcing the “unbundling” of the multiple complex services of the university—including separating the teaching and research missions. The same technology, with the support of the learning engineer, can instead be used to leverage the strengths of higher education’s dual mission. The future is clear: technology will be a core part of the teaching, learning and research processes of higher education. Yet a basic tenet of any successful business strategy is that one does not outsource its core business process. If research and teaching institutions continue to outsource educational technology design, data collection and data modeling, they not only run the risk of violating that basic tenet, but also jeopardize the opportunity to transform higher education to support all students.

Decisions made today and in the near term that address how educational technology and learning analytic systems are designed and implemented are likely to have long-lasting influences on the nature and directions of such developments. Thus it is critically important for social scientists, learning scientists, learning engineers, and policymakers to balance the imperative to innovate with mechanisms to ensure that the economic and social benefits of that innovation are broadly shared across society and, likewise, that they contribute to fulfilling the multifaceted mission of higher education.

About the Author

Candace Thille is the founding director of the Open Learning Initiative (OLI) at Carnegie Mellon University and at Stanford University. She is a senior research fellow in the Office of the Vice Provost for Teaching and Learning and an assistant professor of Education at Stanford University. Thille serves on the board of directors of the AAC&U, as a fellow of the International Society for Design and Development in Education, on the Assessment 2020 Task Force of the American Board of Internal Medicine, and on the advisory council for the NSF Education and Human Resources Directorate. She served on the working group of the President's Council of Advisors on Science and Technology (PCAST) that produced the 2012 *Engage to Excel* report, and on working groups of the U.S. Department of Education's Office of Educational Technology, co-authoring the 2010 and 2015 *National Education Technology Plans*. Thille recently was named one of The Sixteen Most Innovative People in Higher Education by the *Washington Monthly*.



Candace Thille

References

- Ambrose, S. A., Bridges, M.W., DiPietro, M., Lovett, M.C., & Norman, M.K. (2010). *How learning works: Seven research-based principles for smart teaching*. San Francisco, CA: Jossey-Bass.
- Baumeister, R. F., DeWall, C. N., Ciarocco, N. J., & Twenge, J. M. (2005). Social exclusion impairs self-regulation. *Journal of Personality and Social Psychology*, 88(4), 589–604. doi:10.1037/0022-3514.88.4.589.
- Benassi, V., Overson, C.E., & Hakala, C., (Eds.) (2014). *Applying science of learning in education: Infusing psychological science into the curriculum*. Retrieved from <http://teachpsych.org/ebooks/asle2014/index.php>.
- Clark, R. C., & Mayer, R. E. (2008). *e-Learning and the science of instruction, 3rd edition*. San Francisco: John Wiley & Sons, Inc.
- Cohen, G. L., Garcia, J., Purdie-Vaughns, V., Apfel, N., & Brzustoski, P. (2009). Recursive processes in self-affirmation: intervening to close the minority achievement gap. *Science*, 324(5925), 400–403.
- Craig, R. (2015). *College disrupted: The great unbundling of higher education*. New York: St. Martin's Press.
- Davies, P. G., Spencer, S. J., Quinn, D. M., & Gerhardstein, R. (2002). Consuming images: How television commercials that elicit stereotype threat can restrain women academically and professionally. *Personality and Social Psychology Bulletin*, 28,1615-1628.
- Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., & Willingham, D. T. (2013). *Improving students' learning with effective learning techniques: Promising directions from cognitive an educational psychology*. *Psychological Science in the Public Interest*, 14, 4-58.
- Harackiewicz, J. M., Canning, E. A., Tibbetts, Y., Giffen, C. J., Blair, S. S., Rouse, D. I., & Hyde, J. S. (2014). Closing the social class achievement gap for first-generation students in undergraduate biology. *Journal of Educational Psychology*, 106(2), 375–389. doi:10.1037/a0034679.
- Keller, J. (2002). Blatant stereotype threat and women's math performance: Self-handicapping as a strategic means to cope with obtrusive negative performance expectations. *Sex Roles*, 47, 193–198.
- Khajah, M., Lindsey, R. V., & Mozer, M. C. (2016). How deep is knowledge tracing? In T. Barnes, M. Chi, & M. Feng (Eds.), *Proceedings of the Ninth International Conference on Educational Data Mining* (pp. 94-101). Educational Data Mining Society Press.
- Koester, B. P. Grom, G., McKay, T. A. (2016). Patterns of gendered performance difference in introductory STEM courses. Physical review physics education research. arXiv:1608.07565 [physics.ed-ph].

- Kray, L. J., Galinsky, A. D., & Thompson, L. (2002). Reversing the gender gap in negotiations: An exploration of stereotype regeneration. *Organizational Behavior and Human Decision Processes*, 87, 386-409.
- LeCun, Y. Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521:436–444.
- Mayer R.E. (2011) *Applying the science of learning*. Boston, MA: Pearson.
- Miyake, A., Kost-Smith, L. E., Finkelstein, N. D., Pollock, S. J., Cohen, G. L., & Ito, T. A. (2010). Reducing the gender achievement gap in college science: a classroom study of values affirmation. *Science*. 330(6008), 1234–1237.
- Munroe, J. P (1912). *New Demands in Education*. New York: Macmillan. 20-21.
- Ocupaugh, 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.
- Pasquerella, L. (In Press). Rewriting the dominant narrative: How liberal education can advance racial healing and transformation. *Liberal Education*. Fall 2016 102(4).
- Piech, C., Bassen, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L. J., & Sohl-Dickstein, J. (2015). Deep knowledge tracing. In C. Cortes, N. D. Lawrence, D.D. Lee, M. Sugiyama, & R. Garnett, R. (Eds.) *Advances in Neural Information Processing Systems* 28, 505–513.
- Pronin, E., Steele, C., & Ross, L. (2004). Identity bifurcation in response to stereotype threat: Women and mathematics. *Journal of Experimental Social Psychology*, 40, 152-168.
- Richland, L. E., Zur, O., & Holyoak, K. (2007). Cognitive supports for analogies in the mathematics classroom. *Science* 316(5828), 1128-1129.
- Roth, M. S. (2014). *Beyond the university: Why liberal education matters*. New Haven: Yale University Press.
- Schmader, T., & Johns, M. (2003). Converging evidence that stereotype threat reduces working memory capacity. *Journal of Personality and Social Psychology*, 85(3), 440–452.
- Schwartz, D. L. & Bransford, J. D. (1998). A time for telling. *Cognition and Instruction*, 16(4), 475-522
- Schwartz, D. L., Chase, C. C., Oppezzo, M. A. & Chin, D. B. (2011). Practicing versus intervening with contrasting cases: The effects of telling first on learning and transfer. *Journal of Educational Psychology*, 103(4), 759-775.
- Schwartz, D. L., Tsang, J.M., Blair, K.P (2016). *The ABCs of how we learn: 26 scientifically proven approaches, how they work and when to use them*. New York: W. W. Norton & Company, Inc.
- Shine, J.M., Bissett, P.G., Bell, P.T., Koyejo, O. , Balsters, J. H., Gorgolewski, K. J., Moodie, C.A., & Poldrack, R.A. (In Press). The dynamics of functional brain networks: Integrated network states during cognitive task performance. *Nueron*. DOI: <http://dx.doi.org/10.1016/j.neuron.2016.09.018>.

- Simon, H. A. (1967). The job of a college president. *Educational Record* 48, no. 1: 68–78. The American Council on Education, Washington, D.C.
- Spencer, S. J., Steele, C. M., & Quinn, D. M. (1999). Stereotype Threat and Women's Math Performance. *Journal of Experimental Social Psychology*, 35(1), 4–28.
- Stanton, M. (2013). Unbundling higher education: Taking apart the components of the college experience. In A.P. Kelly & C. Carey (Eds.) *Stretching the higher education dollar: How innovation can improve access, equity, and affordability*. Harvard Education Press. Cambridge, MA.
- Steele, C. M. (1997). A threat in the air. How stereotypes shape intellectual identity and performance. *The American Psychologist*, 52(6), 613–629.
- Steele, C. M., & Aronson, J. (1995). Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology*, 69(5), 797–811.
- Steele, J., James, J. B., & Barnett, R. (2002). Learning in a man's world: Examining the perceptions of undergraduate women in male-dominated academic areas. *Psychology of Women Quarterly*, 26, 46-50.
- Stephens, N. M., Fryberg, S. A., Markus, H. R., Johnson, C. S., & Covarrubias, R. (2012). Unseen disadvantage: How American universities' focus on independence undermines the academic performance of first-generation college students. *Journal of Personality and Social Psychology*, 102(6), 1178–1197.
- Stone, J., Lynch, C. I., Sjomeling, M., & Darley, J. M. (1999). Stereotype threat effects on black and white athletic performance. *Journal of Personality and Social Psychology*, 77, 1213-1227.
- Taylor, V. J., & Walton, G. M. (2011). Stereotype threat undermines academic learning. *Personality and Social Psychology Bulletin*, 37(8), 1055–1067.
- Walton, G. M., Logel, C., Peach, J. M., Spencer, S. J., & Zanna, M. P. (2015). Two brief interventions to mitigate a “chilly climate” transform women's experience, relationships, and achievement in engineering. *Journal of Educational Psychology*, 107(2), 468–485.
- Yeung, N. C. J., & von Hippel, C. (2008). Stereotype threat increases the likelihood that female drivers in a simulator run over jaywalkers. *Accident Analysis & Prevention*, 40, 667-674.



TIAA-CREF Individual & Institutional Services, LLC, Teachers Personal Investors Services, Inc., and Nuveen Securities, LLC, Members FINRA and SIPC, distribute securities products.

TIAA Institute is a division of Teachers Insurance and Annuity Association of America (TIAA), New York, NY.

©2016 Teachers Insurance and Annuity Association of America-College Retirement Equities Fund (TIAA-CREF), 730 Third Avenue, New York, NY 10017