

CONTENTS

Page 1

Welcome

Welcome to the winter edition of *Intersection*, the quarterly publication of the AALHE. We cover a wide scope of topics in this issue, from libraries and HIPs to assessing creative disciplines. We are currently putting together material for a special issue on “practical assessment,” so if that topic excites you, consider writing a 700-1400 word piece to submit. The deadline for that project is March 15. If you have your own ideas to share, please see the box at the bottom of page three to learn how to contribute.

We look forward to seeing many of you at the annual AALHE conference June 6 - 8, 2016 at the Pfister Hotel in Milwaukee, Wisconsin. The call for proposals is open until January 29, 2016. You can find out more information on the conference [web site](#).

Page 2-3

Librarians and
High-Impact
Educational
Practices

For the first time, the annual conference will include a sequence of sessions designed to cover assessment basics for those new to assessment and assessment leadership, for no additional cost. These sessions will take place within the conference schedule, not as pre-conference workshops. Although the four sessions are designed to be sequential, anyone can attend any of the four that may address their professional development needs. The following provides a preliminary synopsis of these sessions:

Pages 4-6
Shaping the
Campus
Conversation,
Part Two

Assessment 101: Politics and Practices—How and Why We Got Here

Understanding the history and foundations of contemporary assessment practice in order to make sense of current debates and future challenges.

Assessment 102: Measuring Learning

Understanding data-gathering for assessment, including standardized instruments and rubric ratings. A hands-on component will consider how to assess general education outcomes.

Page 7-8
Assessing Rater
Agreement

Assessment 103: Building a Culture of Assessment

Understanding the role of faculty leadership, good practices, and the tension between assessment for improvement and assessment for accountability.

Pages 9-11
Q&A with
Joey Hoey

Assessment 104: Reporting and Accreditation Standards

Understanding reporting requirements for key external (e.g. accreditors) and internal (e.g. board of directors) constituents regarding learning outcomes and institutional effectiveness.

We are hoping to provide an easy entry into the subject of learning outcomes assessment in conjunction with the other benefits the annual meeting provides, such as networking with professionals and sampling the culinary delights in Milwaukee. Hope to see you there!

Librarians and High-Impact Educational Practices



*Rhonda Huisman and
Merinda Kaye Hensley*

As academic librarians, we are often invited into the classroom to teach information literacy skills. Information literacy refers to the critical thinking skills students need to develop in order to be successful in finding, using, evaluating, and creating new information. While information literacy instruction may expose students to a wealth of information and resources, if it is taught within a single course session and not explicitly tied to an assignment or project, the takeaway instruction can have little impact or engagement with student learning. Complicating matters further from the librarian's perspective, this type of interaction with students provides little in terms of extended learning experiences and often lacks assessment.

The librarian's role in student learning is increasing across the academy, and we are working to foster deeper relationships across disciplines with faculty, students, advisors, as well other organizations on campus that are involved with student planning and retention. This type of relationship is essential for information literacy student learning outcomes, but it can also provide a basis for building engaging experiences and active learning, which in turn can contribute to higher student success outcomes across the academy. More specifically, high-impact educational practices (Kuh, 2008) create an environment for both formal and informal learning experiences that last longer and are more integrated than the traditional course work. To align information literacy efforts with high-impact educational practices, many of our library colleagues have extended their roles beyond the classroom, including being embedded in courses and distance learning, involvement in first-year learning experiences or other introductory learning experiences, and increasing collaboration with faculty one-on-one to plan curriculum, articulating student learning outcomes, and participating in a wide variety of institutional assessment activities. With increased participation in high-impact educational practices at all types of institutions, librarians are working to leverage their expertise to provide information literacy instruction and resources that support collaborative assignments and projects, learning communities, undergraduate research, and capstone courses and projects.

Librarians have been exploring project-based learning, one kind of high-impact practice, within their own teaching. While project-based learning has been widely used in K-12, this teaching method has become more common in higher education, evolving from the previously familiar problem-based learning (mostly in STEM), in which students are presented an open-ended problem to solve within the course period, and often with limited "solutions." Project-based learning focuses on the active roles of students, with a "voice and choice" atmosphere for collaboration, group roles and norms, cross-disciplinary options for resources research, and possibilities for real-world application.

As an example of direct librarian involvement, we will describe an education course shaped entirely around the concept of project-based learning. The coursework was centered on a pedagogical example for pre- and existing teachers through a project-based learning model, and it evolved into a high-impact educational practice by engaging students in a collaborative research assignment at the undergraduate and graduate levels. The "driving question" asked students to recognize that many of our K-12 librarians and libraries are disappearing from the schools - in recognition that an essential component of project-based learning revolves around a social justice or community connection. The assignment asked students to emulate the role of a task force assigned to research the benefits (or not) of a library and librarian established in the local school. Students were directed to locate resources and research that would aid in the discovery process for a final project, which could include a standard research paper with appropriate sources or a presentation to the class on their findings. Students were also given the opportunity to present their ideas beyond a typical assignment, e.g. seeking outlets for their research such as an article in the local newspaper/blog or constructing a presentation for a school board that is considering eliminating the librarian position at their local elementary or high school. As part of this project, students engaged in a complex information literacy skill-building exercise in terms of inquiry-based collaboration and real-life applicable skills.

(Continued on page 3)

Librarians and High-Impact Educational Practices

Continued from page 2

Two documents published by the Association of College and Research Libraries (ACRL) may interest teaching faculty considering working with teaching librarians to incorporate information literacy into high-impact educational practices: the Information Literacy Competency Standards for Higher Education and the new ACRL Framework for Information Literacy for Higher Education. Both documents describe established guidelines for creating curriculum and instruction in order to structure higher-order thinking skills related to information literacy through application and assessment, which can be molded to fit both project-based learning and high-impact educational practices learning goals.

What can teaching faculty and other campus organizations expect in creating and strengthening a partnership with librarians through high-impact educational practices? First, libraries should have the space, both physical and virtual, for students and faculty to explore the collections and resources that support teaching, learning and research. Second, the services and collections should reflect the needs of all users, including faculty and administration to support the entire cycle of research including data management and publishing. Third, and perhaps most essential for student learning, librarians and teaching faculty should be collaborating to construct information literacy instruction and assessment around high-impact educational practices that align with the mission, plan, goals, and outcomes relevant to institutional accreditation. As Murray (2015) stated, “Academic libraries, with shifting focus on providing an atmosphere accommodating different academic needs, can provide an informal academic environment that may foster student engagement in HIPs.”

As the chair and vice-chair of the Association of College and Research Libraries, Student Learning and Information Literacy Committee, we welcome the sharing of ideas and discourse on collaborating and promoting the value and impact that librarians and libraries can have on the campus and in the classroom.

Further Reading:

Association of College and Research Libraries, [Information Literacy Competency Standards for Higher Education](#).

Association of College and Research Libraries, [Framework for Information Literacy for Higher Education](#).

George D. Kuh, “[High-Impact Educational Practices: What They Are, Who Has Access to Them, and Why They Matter](#)” (2008), AAC&U.

Adam Murray, “Academic Libraries and High-Impact Practices for Student Retention: Library Deans’ Perspectives.” (2015). *portal: Libraries and the Academy*, 15, no.3, 471-487.

*Rhonda Huisman is Director of University Library Services at Marian University in Indianapolis and can be reached at rhuisman@marian.edu; Merinda Kaye Hensley is Assistant Professor, Digital Scholarship Liaison and Instruction Librarian at the University of Illinois Urbana-Champaign and can be reached at mhensle1@illinois.edu.
Photo by Merinda Hensley.*

Want to Contribute?

We invite your ideas and expertise for the *Intersection* as well as [Emerging Dialogues](#), a new AALHE web site for short articles and comments. Email ideas or short pieces to intersection@aalhe.org or for Emerging Dialogues to tdowns2@delmar.edu or ddirlam@vwc.edu.

Shaping the Campus Conversation, Part Two



Ephraim Schechter

Part 1 of this article, in the [October edition](#), dealt with assumptions you may encounter about “assessment for planning.” In Part 2 we’ll look at some assumptions about “assessment for accountability,” and why assessment for planning and assessment for accountability aren’t necessarily incompatible.

One common assumption about assessment for accountability is that it should produce an overview suitable for external audiences, with broad statements and small sets of comparable measures. But external audiences, especially accreditors, often want more than these “performance indicators.” They want to see the whole “loop,” asking:

- What are we trying to do? (What are my program’s goals and outcomes? What are we trying to do for our students? For our division/school/college? For our campus?)
- How well are we doing that? How do we know? (How do we monitor progress toward those goals and outcomes? What have we found out by asking this?)
- How have we used that information to guide us? (What have we done as a result of knowing how we’re doing? This is what’s usually meant by “closing the loop.”)

And...keep asking. It’s an ongoing cycle. Do results stay stable over time? If we did something because of what we learned from our assessment, what did we want it to accomplish and how did it work? This view of accountability emphasizes improvement.

Another view is that accountability is about demonstrating that we meet minimum standards. For some, accreditation really is a bureaucratic reporting exercise to demonstrate that the program or institution functions according to those minimum standards. For others, using assessment results for planning and improvement *is* one of the minimum standards we’re supposed to meet. Articles by Alexander Astin (2014) and Milton Greenberg (2014) are examples of the improvement and minimum standards views of accreditation. Interestingly, they came out within a few weeks of each other.

Accreditation aside, another way to look at accountability that’s very popular in politics and the media these days is that it’s about demonstrating what we get for our investment. Depending on who’s talking, “we” can mean students, their parents, and/or the general public, and “investment” usually means money (whether that’s tuition and fees or public support such as student loans or state allocations). But even here, you’ll find different assumptions about what “what do we get for our investment?” really means.

It can mean “what goes on at your institution?” The assessment loop can speak to this: what are your goals for student learning and institutional performance, how well are you meeting them, and are you a self-regarding institution that uses the information in your planning? Regional accreditors all want this. So do many discipline-specific accreditors.

Or accountability can reduce to “which institution does the best job?” Recently, this often means “who has the highest graduation rate for my kind of student, and whose graduates get the best-paying jobs?” and has led to legislative/DOE focus on job-related competencies and graduates’ employment rates and initial earnings. But in discussion forums with students, parents, faculty, and employers about what college should provide, critical thinking is at least as important an outcome as job skills, and understanding the world around them is nearly as important. And employers’ survey responses consistently emphasize general abilities such as critical thinking, clear communication, integrity, and the ability to think on their feet, evaluate information, and work well with others.

We’ve looked at the assumption that accountability implies continuous improvement, the assumption that it’s about return-on-investment and what that means, and the confusion about what the public and employers really want.

(Continued on page 5)

Shaping the Campus Conversation, Part Two

Continued from page 4

Another assumption you may encounter is that accountability requires comparisons. The “who’s best?” approach certainly includes this. “What’s it like at your institution?” often does, too – it often means compared to what it’s like at other institutions. This sometimes leads us to use standardized instruments which let you compare your results against those from other institutions. But it’s important to remember that results will be useful for planning/improvement only if the instrument’s assumptions/content fit your program.

It’s also worth remembering that easily available comparisons can be misleading. Sometimes the message you give depends on how much detail you provide. For example, compare “Our seniors are better at critical thinking than 65% of seniors at other colleges” with “On average, our seniors who took the XYZ test this year were at the 66th percentile on critical thinking when compared to seniors at institutions like ours, and their scores ranged from the 35th percentile to the 84th percentile.” What different messages would the two statements send to your own audiences?

Part 1 mentioned the assumption that assessment-for-planning and assessment-for-accountability are incompatible. But there are ways to balance assessment-for-planning and assessment-for-accountability. The problem, of course, is that we want it all – rich data that’ll really be useful for our planning efforts, and quick easily-graspable summaries for accountability purposes. One way to have both is to do them separately. But that takes more work, and can lead to confusion, especially when different people work on each task. Another approach is simply to do one of them and let the other go. Most often, assessment for accountability reporting is what’s chosen, if only because the accreditors are coming – which puts us right back where this article started. In NILOA’s most recent survey of chief academic officers, most respondents report that accreditation is the primary reason for doing assessment, and fewer than in previous surveys say their campus’s planning processes incorporate assessment results. This is disappointing if you’re thinking in terms of assessment-for-planning. But picking one isn’t the only way to go.

Rich, raw information that’s useful for planning can often be collapsed into rough summaries such as the proportion of students in each program that meet your outcomes targets. And those can start at the program level and be “rolled up” into division and institutional summaries. Even when each program has its own assessment approach, so their raw data can’t be combined, the proportion meeting the program’s own targets provides a common metric.

A rough summary might include (a) the proportions of students meeting institution-wide and program targets for outcomes such as student learning, retention/progress/graduation, and employment/further education, (b) the proportion of programs closing the loop and using assessment results to inform planning, and (c) the proportion of budget/planning decisions that used assessment results. This doesn’t begin to give all the information, but it can (a) give easy to see warnings about places where things aren’t going as well as you want them to, indicating that you should dig into the details to see what’s going on, and (b) point out places where things are going particularly well and where it might be useful to see what guidance they might give for places where things aren’t going so well. It’s the dashboard approach – squeezing rich information down into easy “ok”/“not ok” flags, like the warning lights on your car’s dashboard.

This article started by talking about communication and has mentioned a variety of assumptions and potential miscommunications. The bottom line is that when you look at your statements about assessment, it helps to ask:

- Who are you talking to? Why? What do you want to say to this audience? (That is, what do you want your message to accomplish?)
- How will their assumptions lead them to interpret your message? What do they already assume? (You may need to do some preliminary assessment to find out).
- And what do the assumptions they’re starting with tell you about how to shape your message?

(Continued on page 6)

Shaping the Campus Conversation, Part Two

Continued from page 5

Further Reading on Accreditation and Assessment, from page 4

The U.S. Department of Education publishes relevant [accreditation standards](#)

A. W. Astin (Feb 2014). "[Accreditation and Autonomy](#)"

B. M. Greenberg (Jan 2014). "[It's Time for a New Definition of Accreditation](#)"

Further reading on institutional comparisons, from page 4

InsideHigherEd Feb 2014 article "[Politician-Public Divide](#)"

See "[Divided We Fail](#)" for details of these discussion forums.

Bentley University's [PreparedU Project report](#), especially the section The skills discussion: A disconnect on pp 11-12

AAC&U survey reports (pdf files):

- 2010: "[Raising the bar: Employers' views on college learning in the wake of the economic downturn](#)"
- 2013: "[It takes more than a major: Employer priorities for college learning and student success](#)"

The Chronicle of Higher Education June 2013 article "[Giving employers what they don't really want](#)"

Further reading on improvement vs accountability, from page 5

The National Institute for Learning Outcomes Assessment's 2014 survey report "[Knowing what students know and can do: The current state of student learning outcome assessment in US colleges and universities](#)"

Other Resources

L. Suskie (2009). *Assessing Student Learning: A common sense guide*. 2nd edition. Jossey-Bass Publishers, San Francisco. Chapter 14.

Association for Institutional Research "[Measuring Quality In Higher Education](#)" (2012)

Ephraim Schechter is a founding board member of AALHE and well-known consultant in assessment. He can be reached at eischech@higheredassessment.com.

Photo credit: Jacquelyn Schechter, [Mindful Photography](#).

This article is adapted from a presentation at AALHE's 2014 Annual Conference.

Assessing Rater Agreement



David Eubanks

The idea of agreement is fraught with difficulty once we get past the superficial level. An example of the type of philosophical problem that statisticians have to wrestle with, consider how we would measure agreement among observers about some phenomenon they mutually observe. If five of them agree that the cat knocked the vase over, how many agreements is that? Is it five, because there were five observers, or is it ten agreements, counting pairwise (the first observer agrees with four others, the second with three additional people, and so on for $4 + 3 + 2 + 1 = 10$)? This has consequences when we compute parameters like Cohen's Kappa or its inter-rater statistical cousins.

Perhaps worse than this conundrum is the "paradox" that emerges when we have unbalanced sets of observations. If nearly everyone agrees that your students can write well, the Kappa index can be very low, indicating not much agreement at all. This is because it is a relative, rather than absolute, measure of agreement, and these odd results point the fact that we need both types of information. For example, if retention in the *Zen Cheerleading* program has increased 100% in the last year, it's helpful to know that there are only two students in the program.

As if all this weren't complicated enough already, a recent article at phys.org entitled "[Why too much evidence can be a bad thing](#)," starts off with:

Under ancient Jewish law, if a suspect on trial was unanimously found guilty by all judges, then the suspect was acquitted. This reasoning sounds counterintuitive, but the legislators of the time had noticed that unanimous agreement often indicates the presence of systemic error in the judicial process, even if the exact nature of the error is yet to be discovered.

The authors go on to present various kinds of evidence to conclude that when judgments (read: assessments) are subject to error, we should be suspicious of too much agreement, because it may be due to systematic bias rather than meaningful agreement. This can occur in training raters to use rubrics, if we coerce them into artificial agreement. For many educational outcomes, especially those involving aesthetics, there is room for disagreement. Imagine that an assignment is to create a limerick, a verse with a strictly-defined form. A student who turns in the following would automatically be marked down by a literal-minded reviewer:

There was a young man of Milan,
Whose rhymes they would never scan;
When asked why it was,
He said, 'It's because,
I always try to cram as many words into the last line as ever I possibly can.'

I found this example in John D. Barrow's book [Impossibility: The Limits of Science and the Science of Limits](#), where he uses it as an example of the power of symmetry-breaking (read: rule-breaking in our context). We've probably all met "that professor" who gives a zero to any paper with a spelling or punctuation error. Because we can legitimately disagree about the relative importance of different aspects of student work, ratings will vary—and they should.

So how much agreement is enough? The answer depends on the context, and it helps to compare different types of ratings. Because of the rich sources of rater agreement that can be found on the internet, we have the opportunity to benchmark educational data against agreement "in the wild," such as the five-star ratings on Amazon.com or on book or movie review sites. You can find examples of such data, along with open source code and an online app on [Github](#). As an example, consider rater agreement for a set of faculty ratings of students on Creative Thinking. There were 352 matched ratings on a 0 to 4 ordinal scale, with 4 the highest rating.

(Continued on page 8)

Assessing Rater Agreement

Continued from page 7

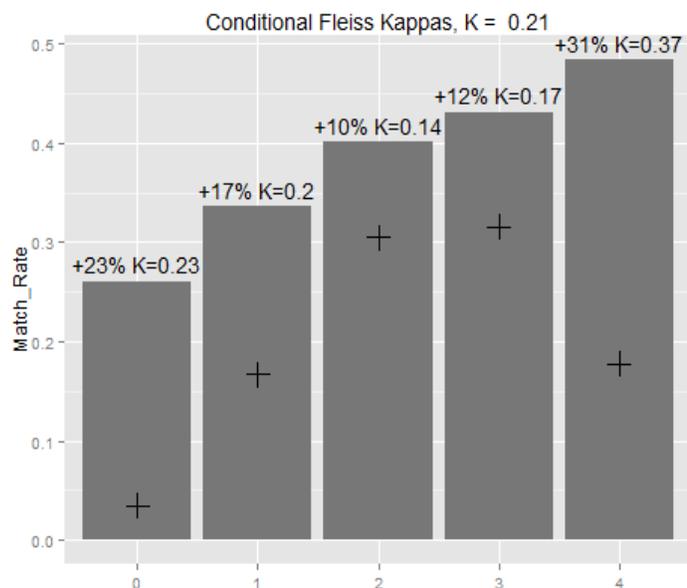


Figure 1. Rater Agreement per Rating Assignment 0-4

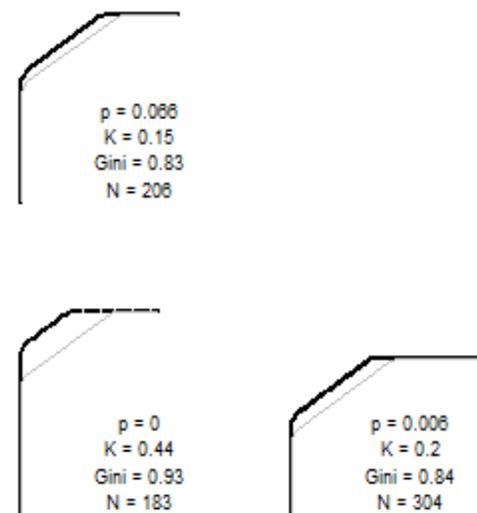


Figure 2. A Sample of Inter-Item Detail

In Figure 1 we have the whole rating scale presented, from zero (the lowest rating) to four. The bars represent the actual rater agreement for that response, and the plus sign (+) marks where we would expect random agreement rates to be if we were just pulling the ratings out of a hat. At the top of the bar the annotations tell us the actual bonus over random, and the conditional Kappa for that response. You can see that for this set of data, we get more agreement over who is a four than we do over who is a zero. Typically, agreement is worse in the middle of the range, and that is true here—it's easier to identify extreme cases.

The mini graphs in Figure 2 are a sample of the whole report, which is too large to include here. The top left and bottom right graphs compare the rating pairs 1 vs 2 and 2 vs 3, respectively. The thin line on the graph is the expected random level of agreement, and the thicker line is actual. The closer the actual line is to the random line, the less evidence for meaningful agreement, and the smaller Kappa will be. The graph for perfect agreement has a darker line in the shape of an inverted L.

The statistics in Figure 2 give a relative measure of agreement (Kappa, with one being perfect agreement) and an absolute measure (the Gini coefficient, which ranges from zero to one with one being perfect agreement), as well as sample sizes and p-values versus the null hypothesis that actual rates equal random rates. The lower left graph compares 1 vs 3, which we would expect to have better agreement because the rating levels are further separated. In fact, this is true—the 1 vs 3 distinction has a higher relative (Kappa) and absolute (Gini) rate of agreement than the adjacent ones. There is more information about how to create and interpret these graphs on the site wiki. These statistics are much higher than what we see in public ratings systems for movies or books, which might have $K = .05$ or worse for the adjacent cases, and $Gini < .6$. That comparison gives us a useful perspective on the reliability of these ratings.

What I learned from the ratings through this analysis was that while the top end of the rating scale had good agreement, the bottom end had less agreement than I expected. In a review of the rubric descriptions for the 0 and 1 categories, I found that these are not as well-defined as they could be, and that's likely the reason for the difference shown in Figure 1. That issue will be fixed the next time we use the rubric.

David Eubanks is Assistant Vice President for Assessment and Institutional Effectiveness at Furman University, and a serving board member of AALHE. He can be reached at david.eubanks@furman.edu.

Q&A with Joseph Hoey

by David Eubanks



Joseph Hoey IV, Ed.D. is Vice President, Accreditation Relations and Policy at Bridgepoint Education, and has broad experience in assessment, especially in the arts. With coauthors Jill Ferguson and David Chase, he has two books in print on that subject. Dr. Hoey can be reached at joseph.hoey@bpiedu.com.

Q: Tell us a little about your background in doing, teaching, or assessing the arts.

A: I had a very interesting and lengthy first career in classical music – specifically classical guitar performance. I studied in England for several years as well as the Hochschule der Künste in Berlin, Germany, taught guitar at the University of Reading (that’s a city, not a verb!) and several private schools in England. I completed a bachelor’s under Celin Romero at UCSD and a Master’s at Florida State under Bruce Holzman. I juried into the North Carolina Visiting Artist Program, and remained in residence at two colleges for four years, played concerts all over North Carolina, South Carolina and even the Kennedy Center in D.C. as half of the guitar duo Les Deux Amis, taught guitar at Barton College in Wilson, NC, and then decided to do a doctorate in higher education. That was in 1989.

In 1990 I took a very early course on institutional effectiveness and student learning outcomes assessment, and immediately got offered a position in a community college doing assessment. Twenty-six years later I’m still glad about that decision. So it was a natural progression to put together my knowledge of teaching applied music in the studio context and student learning outcomes assessment. Putting together an assessment office at the Savannah College of Art and Design broadened my understanding of assessment in the arts, and in 2009 we put on the first-ever conference on assessment in the arts in higher education there. We called it MUSE, for Measuring Unique Studies Effectively. That conference really opened a lot of doors and from there I started considering what kind of published works would best contribute to understanding of assessment in the arts.

Q: Can you summarize for someone not expert in the arts what are the most important ideas to keep in mind when doing outcomes assessment?

A: There are many threads to bring together, but here are a few. To start out, I need to recognize the contributions of my co-authors Jill Ferguson and David Chase. Our ongoing interactions solidified a lot of good ideas on assessment in the arts that are expressed in our books.

First, the arts are applied disciplines, so the process of becoming an artist is incredibly important – both in terms of gaining technical mastery of the craft as well as the taking the soul-searching mental journey to artistry and expression. That means that we have craft-oriented technical outcomes as well as individual, aesthetic learning outcomes to assess (or ubiquitous vs. unique outcomes, to borrow Steve Ehrmann’s nomenclature). While assessing technical, craft-oriented learning outcomes is more easily quantified at the group level, assessing unique, aesthetic outcomes in the arts is more individual in nature and in many ways similar to what we have to do in assessing doctoral education.

Second, the opportunity for students to have lots of experiences that approximate professional practice is really important - internships, design charrettes, critiques, and team projects with external sponsors for example. That aspect of pre-professional practice assessment is really important learning to assess, since it captures student learning in process. Schools like the University of Cincinnati get at this through their long-standing co-op program. Assessing those pre-professional practice experiences might include portfolios with self-reflective components; instruments that elicit information on specific learning outcomes from internship supervisors; and qualitative narrative information from external raters and industry sponsors. Increasingly, we are seeing a lot of interactive communities of practice online where critique of posted works is happening informally through social media commentary/portfolio tools, for example Behance.

(Continued on page 10)

Q&A with Joseph Hoey

Continued from page 9

Third, the design of education is hugely important in assessing the arts. That itself has several threads. In online education we use backwards design, where we start curriculum design with an explicit understanding of those outcomes that faculty, the professional field, and industry experts have identified as the most important elements of mastery – so too must we in the arts. Appropriate sequencing, ascending complexity of expectations, and scaffolding are vital to artistic growth. Master teachers in the arts have the experience-based background to understand what technical and artistic complexities to introduce and at what point in a student's development to do so. We are calling that the ability to connect the dots in arts education through seeing emergent patterns of practice. What we have found again and again is that our colleagues have a deep understanding of what the appropriate elements of mastery or learning outcomes are, and how to enable students to progress in an artistic discipline – but that in many cases they have not been asked to articulate the learning outcomes or their tacit, practice-based understanding of how to connect the dots to most effectively enable student learning in a way that makes sense for their practice. So encouraging our arts faculty to articulate the elements of mastery and understanding/articulating educational process – sequencing, scaffolding, etc. – is something that we feel has hitherto been neglected but that needs to be promoted.

Fourth is the notion that while the rest of assessment world has concerned itself primarily with Bloom's taxonomy in the cognitive domain, the fact is that for the arts we also have to concern ourselves with affective and psychomotor domains in a very deep way. Further, other learning domains like the Intrapersonal domain (dealing with capacity for self-awareness and understanding) and the conative domain (dealing with volition and the will to create) are part of the palette available to us in designing education appropriate to high achievement in the arts, and form part of the picture.

Fifth is that the arts are an area where the expression of highly informed professional judgment is fundamental. The late Elliott Eisner called this connoisseurship evaluation. It's likely to be qualitative, narrative, or oral in nature. This is related to what I said earlier about the importance of affording our students in the arts lots of pre-professional experience and opportunities for feedback through juries, critiques, competitions, and projects. For assessment, we may well have to use content analysis, video tags and other tools to understand such feedback and to relate it to the appropriate student learning outcomes.

Sixth and final – we need both individual assessment plans as well as program-level assessment plans in the arts. We can create outcomes-based plans for individual student improvement over a semester. Malcolm Knowles was talking about individual learning plans in the 1970's, and this is a present-day affirmation of that. Generalizing to the program level, we can also observe student achievement at a more aggregate program level and ask about what proportion of students moved from an elementary level to a more advanced level in technical and interpretive skills and abilities.

Q: In *Assessment in the Creative Disciplines*, it's clear that at least in the arts we can't separate the cognitive learning outcomes from personality factors. Do you have an example or two of how assessing personality traits can lead to programmatic insights?

A: Recognizing and nurturing facts of nature such as musical talent is something that needs to happen as early as possible in life and arts departments that are not doing long-term talent identification and outreach may be at risk of not attracting the talent they want.

What we call personality is reciprocally influenced by interaction and environment, and we understand it as something that can and often does change over time; I think that's one of the functions of critique. When it comes to factors such as student grit and the desire to achieve at the highest level, if a department does not select for those factors during admissions and/or makes no attempt to build or reinforce them during the course of the curriculum, I would expect to see few graduates and even fewer professional placements.

(Continued on page 11)

Q&A with Joseph Hoey

Continued from page 10

Creativity is fundamental to the arts. If we conceptualize creativity along the dimensions of the creative person, the creative process, the creative process, and the creative press or environment as did Rhodes (1961) then certainly understanding the creative environment can help lead to programmatic insights, and that's not just limited to the arts or to the students themselves. For example, if a graduate lab environment in the sciences is low-trust and low-interaction, I'd expect to see lower levels of risk-taking behaviors, less collaborative research, and thus lower overall productivity. The tone set by leadership – perhaps based on personality factors – would serve to reinforce or change such an environment.

Q: The practice of giving and receiving constructive criticism is integral to the arts. How is that taught and learned, and should we consider incorporating that practice more in the "non-art" disciplines?

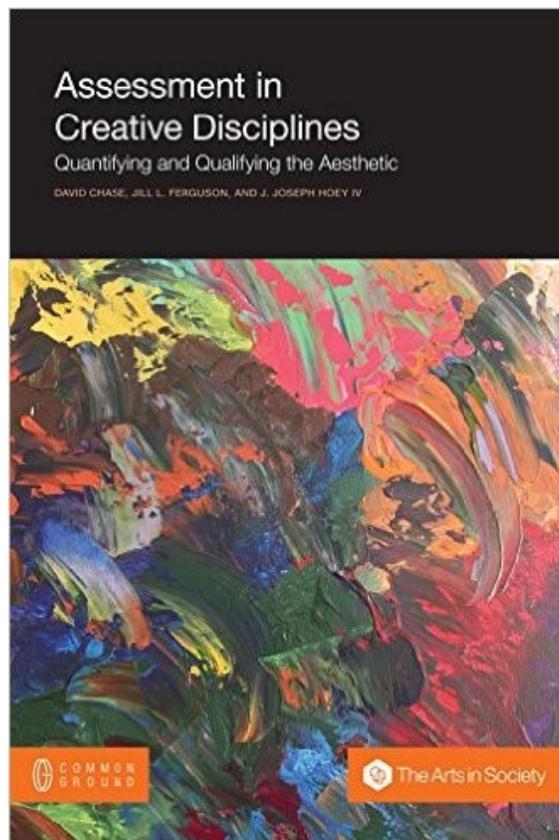
A: So for me it's important to recognize that there is art and creativity everywhere, and it's not exclusive to just a few disciplines. Yes, critique has been central to education in the arts for centuries, and students in creative disciplines are introduced to this practice of individual and collective critique very early on in their studies. Let's be clear: without a thick skin and ability to receive criticism you just don't survive in the arts – and part of the process is to receive the criticism, reflect on it, and be able to address it. The problem has been that with or without a basis in specific criteria, feedback can be suboptimal and can result in students feeling victimized or lost in an existential dilemma similar to that of a protagonist in a Kafka novel. As a student, if I don't know what to address, study or fix then I'm more apt to become discouraged and opt out. We're seeing wider acceptance of learning outcomes as a basis for critique and I'm encouraged by that. Yes, I think the practice of critique may well have a wider application in "non-art" disciplines. After all, creativity comes in many forms and guises.

Q: Are there other lessons drawn from assessing the arts that assessment leaders might find helpful in other areas?

A: Sure. The whole design of education thing I talked about earlier applies. The notion of injecting much more experiential, pre-professional education into our programs is a high-impact practice and makes a lot of sense. As we get further and further down the road to personalized learning experiences based on learner needs, the notion of a personalized assessment plan that maps to a program assessment plan makes more and more sense. Finally, creativity is not just something found in the arts. What many refer to as Big C creativity may be the provenance of the like of Da Vinci, Michelangelo and Einstein, but building a better mousetrap still involves a creative process and can be assessed.

Q: You've co-authored two books to date on assessing creativity. Where is your research taking you these days?

A: One direction we explore in both of the volumes to date is assessing learning in online arts programs. The profusion of tools, interactive fora, and public portfolio spaces on the Internet coupled with the huge advances in online visual technology is redefining how meaning and quality in the arts are understood. Understanding how to interpret and learn through digital and social media learning experiences is already huge and will only get bigger, and it's an exciting direction to explore further.



Assessment in the Creative Disciplines by Chase, Ferguson, and Hoey is a good place to learn about assessing the creative and performing arts. Cover art by Jill L. Ferguson.