

Reasons to Worry: The creation of curriculum materials is a design task

I start from the premise that the creation of materials is a design task. This leads me to worry that the whole game of having a fixed set of criteria for curriculum materials is extremely problematic. Here's the reasons for this worry.

1. There are always multiple solutions to a design task. For example, we might have the idea that a house can be built by putting wood together in certain ways. But we wouldn't want to use this as a set of criteria for evaluating every possible house design. Someone might reasonably have a design for a house that uses bricks instead of wood. Similarly it's extremely plausible that someone could design a curriculum that did not alert teachers to commonly held student ideas, or that did not explicitly present the unit purposes to teachers.

And the situation is even worse than (1) applies, for a few reasons:

2. A given set of design parameters that works in one situation might not work in another.
3. Any design criteria that we might specify are only going to take on precise meanings in a particular circumstance. Even a criterion such as "alert teachers to commonly held student ideas" will mean very different things in different circumstances.
4. Design involves tradeoffs. Even if we were able to agree on a single set of design criteria, we would have to keep in mind that design always involves making tradeoffs; in any particular instance, we'll have to decide how to weight the various criteria.

Resolving the worries: Develop heuristics rather than fixed design criteria

I believe that the way to resolve these worries, and to still make progress, is to think of criteria such as "alert teachers to commonly held student ideas" as design *heuristics*. This move entails all of the implications that are generally entailed by something being heuristic:

1. Heuristics don't always apply. You might even have a set of heuristics that appear contradictory.
2. When heuristics apply, they only give partial guidance.
3. This is the big one. Heuristics are not right are wrong. *Heuristics are judged based on the extent to which they are useful for some population who we would like to make use of the heuristics.*

Accepting these entailments has a number of implications that I'll follow throughout the rest of these notes. Here I will just mention one, for illustration: The sort of quasi-experimental studies proposed by Nancy Songer in her short presentation just don't make sense. She proposed to do experiments to figure out, for example, what kinds of conceptual sequencing are best. In we believe my heuristic perspective, there are no right answers to questions of this sort, and experiments of this sort might not make sense.

Two kinds of functions for design heuristics

There are two kinds of functions for these design criteria, *evaluative* and *generative* functions. If we accept that the criteria are heuristic, then there are significant implications for how we should think about each of these functions.

1. *Evaluative functions.* One thing that we might do with our design criteria is use them to evaluate a set of curriculum materials that have been designed. If we accept that the creation of a curriculum is a design task, and that criteria are heuristic, there is a strong implication: We must have a procedure for evaluating curriculum materials that is sufficiently *flexible*. For example, because design involves tradeoffs, we must not necessarily think negatively of a curriculum only because it does not seem to meet some of our design criteria. Another example: We must be careful not to have an evaluative procedure that presumes to anticipate every possible design solution for curriculum materials.
2. *Generative functions.* Another thing we might do with a set of design criteria is use them to guide the design of a set of curriculum materials. If we really accept that criteria are heuristic, then the criteria won't be either right or wrong; they will be better or worse heuristics to the extent that they have power for the individuals that make use of them. To find out whether the heuristics have power for some set of individuals, we need to do some empirical work. Here's a simple form that this empirical work might take. First, we decide which population of individuals we want to use our criteria as generative design heuristics. Then we give them some candidate sets of design heuristics, and we look at what sort of curricula they design.

Recommendation: What the center should and should not strive to produce

I think that everyone should agree that the primary output of the center should not be an evaluation of specific curricula.

I also think that the primary output of the center should not be specific design principles/criteria/heuristics.

Instead, the center should be engaged in an exploration of the kinds of forms that design heuristics can take. We should seek to explore, empirically, the affordances of specific varieties of design heuristics.

We should strive to explore a large part of the space of kinds of heuristic guidance

The type of heuristic guidance embodied in the AAAS analysis procedure, and in Danny Edelson LfU analysis procedure (which was really concocted for this meeting), live in just one small corner of the space of kinds of design heuristics. There is a huge range of kinds of heuristic guidance that we might invent, and whose affordances we might explore.

Here's a simple example: Okhee Lee, when looking at the chemistry unit, looked at it with a lens that had a parallel set of criteria in mind, criteria having to do with linguistic and cultural diversity.

Jim Minstrell and David Hammer seem to design curricula with a very different kind of design heuristics in mind. Both of them seem to, under some circumstances, think of themselves as setting up instruction to explore some "space" of ideas. (Jim used the metaphor of going on a field trip.) And this space of ideas is determined both by the nature of the domain, and by the ideas that students bring with them. I have advocated a similar view in a book chapter that I wrote (and that I'll attach with these notes).

About standards, learning performances, goals, etc.

Something I want to be clear about is that any descriptions of benchmarks, learning goals, learning performances, etc. must themselves be thought of as serving heuristic functions for the evaluation or generation of designs. Thus we should be seeking to examine, empirically, how well they serve these heuristic functions.

Jennifer Schwarz, as part of her dissertation, is designing a way to make content descriptions more heuristically powerful by first specifying a particular kind of taxonomy of *types* of content targets.

The AAAS standards have a very, very particular form, and they are only one candidate among many possibilities of heuristics that might serve similar functions. As I'm sure they will agree, they embody very specific assumptions about grainsize, language, etc.

One example: The AAAS benchmarks are represented in the form of short paragraphs. This has implications for what sort of benchmarks we end up constructing, and the nature of the heuristic guidance provided.

Another example: The AAAS benchmarks, and the Michigan learning performances, carve up the content in a very particular way. In particular, they carve up the world of content by beginning with the sort of expert parceling of a domain that might be found in a textbook. This has both strengths and weaknesses that we must be aware of.

This, I believe, is an absolutely fundamental point: The descriptions of content embodied by standards are incredibly far removed from a cognitive account of understanding. Suppose that we were somehow able to build moderately complete cognitive accounts of the mental

representations involved in the understanding of some sub-portion of a discipline. It is almost certainly the case that, if we start putting these mental representations into categories, those categories will not line up in any simple way with standards.

I'm not saying that we can't proceed without having a complete account of the mental representations that constitute understanding in some sub-domain. (Some of us have spent years and years trying to build a cognitive account of understanding of some tiny sub-portions of domains such as Newtonian mechanics.) But there are very serious implications of the observation that standards are not part of a cognitive account:

1. We must realize that descriptions of the sort embodied in the standards cannot provide reliable guidance for the determination of sequences of prerequisites. For that we need something that strives to be closer to a cognitive account.
2. A content map like the AAAS atlas does not map the space of "ideas," where "idea" is understood to be a cognitive entity. There are thus some particular things that the map can't tell us about where instruction might go.

For illustration, here's a metaphor that might replace the idea of progressing through content space with some ideas serving as prerequisites to others:

- Instruction as hopping from lily pad to lily pad. Each lily pad is a partial understanding. It might not even be strictly correct. Each lily pad is not a prerequisite to some later understanding. Rather, a lily pad is a place that you can jump to from your current lily pad, and that puts you in a position to jump to some other lily pad that you want to get to.

Looking where the light is brightest (this is Brian's language)

I have some more specific complaints concerning the AAAS analysis procedure. These are complaints not about the whole idea of the analysis procedure, but rather some complaints about the specific criteria and how they are weighted.

The first complaint has to do with the grading down of curriculum materials because they have statements or representations that might be misleading to students. As far as I know, there is not much evidence that these sorts of misleading representations are really much of a problem for students. Take, as an example, the assertion in the literature that perspective drawings of the solar system lead students to think that the Earth's orbit is much more elliptical than it actually is. That is, as far as I know, not much more than a "Just So" story. It makes a nice story, but it might not turn out to be true. (Maybe there is research I don't know of.)

And there are lots of reasons to think that stories of this sort are not particularly plausible. First, we *do* know that children are quite sophisticated about representations very early in life. (See, for example, Judy DeLoache's studies of the understanding of representations by very young children. According to DeLoache, much of the development in children's understanding of representations happens between the ages of 2 and 3.)

Second, an individual statement or representation is a tiny pinprick of experience working against a large background of experience of the student, and a large background of existing knowledge.

This leads to a question: Even if these misleading representations aren't that big a problem, shouldn't we still strive to get them to be as un-misleading as possible? My answer is "yes," but with a number of caveats.

First, if we are going to apply the *misleading representations* heuristic, we should weight it less than other heuristics. This relates to a point that Brian Reiser has made periodically (at least to me). He says that we have to be careful about "looking where the light is brightest." Representations are a very visible feature of printed curriculum materials. But this does not necessarily tell us much about how important they are in the chain that leads to student learning.

(Another way to say this is that we must be careful not to give up too much in the service of reliability. Reliability is really an extremely weak constraint. If we just coded every other line in materials with a letter "A," and the other lines with the letter "B," we would have perfect reliability.)

Second, by their nature, representations do not veridically capture all aspects of the things they are intended to represent. So, it's quite tricky to know when a representation will be misleading and when it won't. For example, suppose the reader for a chemistry unit has a picture of some atoms forming a molecule. Should we worry that the students will think that real atoms are the same size as the atoms as they are represented in the diagram? This might be a worry in some cases. But, my guess is that this will usually not be a worry for middle school students

The point is judging when a representation will be misleading is a very complex judgement, and knowing that some particular misconceptions have been reported in the literature is not nearly enough. We *do* want to be trying to make these judgments. But we need to be careful about setting ourselves up as the ones who will make them.

Final comment: We've learned a lot, but it's a tiny, tiny percentage of what we need to know

There's a tendency, particularly in the AAAS rationale, to speak as if there are some established facts about how science learning happens, and what works pedagogically. I just don't think that this is the case. There's lots of incredibly interesting and informative research, but there's nothing like consensus.

For example, I don't believe that the book by Driver can be used as an authoritative source concerning student misconceptions about specific domains. The book has a long list of citations that are roughly relevant to individual domains. But it doesn't contain a critical account of those citations. I happen to think that just about everything there is hugely problematic, even though it might be a productive contribution to ongoing research.

To take just one important critique: Some of us believe that much of what is called “student misconceptions” are sometimes answers that students generate, in the moment, to very specific interview prompts.

Another critique: It doesn’t make any sense to say that students have the misconception that “dissolving is a chemical reaction.” How can they have this misconception, if they don’t really have the categories of “dissolving” and “chemical reaction.”

Some of us think that it’s rarely helpful to think of student ideas as interfering, or as individual nuggets that must be “addressed.”

I don’t mean to say that my school of thinking is necessarily correct. But to speak of these things as if they are relatively well-established facts in the literature is just not right.