Chapter 5

A Friendly Introduction to “Knowledge in Pieces”: Modeling Types of Knowledge and Their Roles in Learning

Andrea A. diSessa

Abstract Knowledge in Pieces (KiP) is an epistemological perspective that has had significant success in explaining learning phenomena in science education, notably the phenomenon of students’ prior conceptions and their roles in emerging competence. KiP is much less used in mathematics. However, I conjecture that the reasons for relative disuse mostly concern historical differences in traditions rather than in-principle distinctions in the ways mathematics and science are learned. This article aims to explain KiP in a relatively non-technical way to mathematics educators. I explain the general principles and distinguishing characteristics of KiP, I use a range of examples, including from mathematics, to show how KiP works in practice and what one might expect to gain from using it. My hope is to encourage and help guide a greater use of KiP in mathematics education.

Keywords Knowledge in pieces • Conceptual change • Complex systems

5.1 Introduction

5.1.1 Overview

Knowledge in Pieces (KiP) names a broad theoretical and empirical framework aimed at understanding knowledge and learning. It sits within the field of “conceptual change” (Vosniadou 2013), which studies learning that is especially difficult. While KiP began in physics education—in particular to provide a deeper

Some parts of this chapter are based on text in diSessa (2017), “Knowledge in Pieces: An evolving framework for understanding knowing and learning,” in T. G. Amin and O. Levrini (Eds.), *Converging perspectives on conceptual change: Mapping an emerging paradigm in the learning sciences*, Routledge. Used by kind permission of Taylor and Francis/Routledge.

A. A. diSessa (✉)
University of California, Berkeley, USA
e-mail: disessa@berkeley.edu

understanding of the phenomenon of “prior conceptions” (misleadingly labeled as “misconceptions”; Smith et al. 1993)—it has since engaged other areas, such as mathematics, chemistry, ecology, computer science, and even views of race and racism (Philip 2011).

I aim to produce a relatively non-technical introduction to KiP that can be understood by those who are not experts in the field of conceptual change, KiP’s “home discipline.” I emphasize breadth and “big ideas” over depth, while still pointing in the direction of KiP’s distinctive fine structure and technical precision. A longer but still general introduction to KiP for those who want to pursue these ideas more deeply is diSessa et al. (2016).

Before beginning discussion in earnest, I would like to make two points about my strategy of exposition. First, the initial examples I give will be from physics, KiP’s “home turf.” I beg the (mathematical) reader’s indulgence in doing so, but it allows me to select some of the best and most accessible examples of KiP analyses, where its core features are transparent, and where some competitive advantages over contrasting points of view are easiest to see. These examples are at the high-school level, so I do not expect them to be too much of a conceptual challenge. Mathematical examples will follow in Sects. 5.3 and 5.4. Second, with respect to mathematical examples, there are, of course, perspectives in the mathematics literature that bear on the same topics. While I will mention some of these (see comments and references in Sects. 5.3 and 5.4), careful comparative analysis is too complex for the scope of this paper. Readers who already know the relevant perspectives from mathematics education, of course, should be prepared to elaborate their own comparisons and conclusions.¹

KiP is essentially epistemological: It aims to develop a modern theory of knowledge and learning capable of comprehending both short-term phenomena—learning in bits and pieces (hence the name, Knowledge in Pieces)—and long-term phenomena, such as conceptual change, “theory change,” and so on. It aims to build a solid two-way bridge between, on the one hand, theory, and, on the other hand, data concerning learning and intellectual performance. “Two-way” implicates that (a) the theory is strongly constrained by and built out of observation, but also that (b) the theory can “project” directly onto what learners actually do as they think and learn, giving general meaning to their actions. KiP is, thus, a reaction against theories that are *a priori*, very high level, and consequentially are difficult to apply to the messiness of real-world learning.

KiP shares important features with two major progenitors. The first is Piagetian and neo-Piagetian developmental psychology, epitomized in mathematics education by Les Steffe, Ernst von Glasersfeld, Robbie Case, and many others.

¹While I provide specific hints later for more detailed comparisons on a per-topic basis, probably the most effective single hint I can provide for reader-developed comparisons is to consider (a) whether work on the same topic identifies intuitive pre-cursor ideas in detail (few do), including their productive as well as problematic nature, and (b) whether data analysis includes extensive examination and explanation of students’ in-process thinking, in addition to long-term comparisons. The presentation of distinctive KiP themes, just below, elaborates these points.

The core unifying feature of KiP with this work is constructivism, the focus on how long-term change emerges from existing mental structure. The second progenitor is cognitive modeling, such as in the work of John Anderson (e.g., his work on intelligent tutoring of geometry), or Kurt vanLehn (e.g., his work on students’ “buggy” arithmetic strategies). The relevant common feature with KiP in the case of cognitive modeling is accountability to real-time data. A key distinctive feature of KiP, however, is its attempt to combine *both* long-term and short-term perspectives on learning. Piagetian psychology, in my view, was never very good at articulating what the details of students’ real-time thinking have to do with long-term changes. In complementary manner, I judge that cognitive modeling has not done well comprehending difficult changes that may take years to accomplish.

I now introduce a set of interlocking themes that characterize KiP as a framework. These will be elaborated in the context of examples of learning phenomena to illustrate their meaning in concrete cases and their importance.

Complex systems approach—KiP views knowledge, in general, as a complex system of many *types* of knowledge elements, often having many particular *exemplars* of each type. Two contrasting types of knowledge are illustrated in the next main section.

Learning is viewed as a transformation of one complex system into another, perhaps with many common elements across the change, but with different organization. For example, students’ *intuitive knowledge* (see the definition directly below) is fluid and often unstable, but *mature concepts* must achieve more stability through a broader and more “crystalline” organization, even if many of the same elements remain in the system. The pre-instructional “conceptual ecology” of students must usually be understood with great particularity—essentially “intuition by intuition”—in order to comprehend learning; general properties go only so far. A number of such particular intuitions will be identified in examples.

I use the terms “intuitive” and “intuition” here loosely and informally to describe students’ commonsense, everyday “prior conceptions.” However, consistent with the larger program, I will introduce a technical model of a very particular class of such ideas that has proven important in KiP studies.

A modeling approach—The learning sciences are still far from knowing exactly how learning works. It is more productive to recognize this fact explicitly and to keep track of how our ideas fail as well as how they succeed. Concomitantly, KiP builds models, typically models of different types of knowledge, not a singular and complete “theory of knowledge and learning,” and the limits of those models are as important (e.g., in determining next steps) as demonstrated successes.

Continuous improvement—A concomitant of the modeling approach is a constant focus on improving existing models, and, sometimes, developing new models. In fact, the central models of KiP have had an extended history of extensions and improvements (diSessa et al. 2016). It is a positive sign that the core of existing models has remained in tact, while details have been filled in and extensions have been produced to account for new phenomena.

I call the themes above “macro” because they are characteristic of the larger program, and they are best seen in the sweep of the KiP program as a whole.

In contrast, the “micro” themes, below, can be relatively easily illustrated in many different contexts, which will be seen in the example work presented below.

A multi-scaled approach—I already briefly called out the commitment to both short-term and long-term scales of learning and performance phenomena, a *temporally multi-scaled approach*. Most conceptual change research, and, indeed, a lot of educational research, is limited to before-and-after studies, and there is almost no accountability to process data, to change as it occurs in moments of thinking.

A systems orientation also entails a second dimensional scale. Complex systems are built from “smaller” elements, and indeed, system change is likely best understood at the level of transformation and re-organization of system constituents. So, for example, the battery of “little” ideas, intuitions, which constitute “prior conceptions,” can be selected from, refined, and integrated in order to produce normative complex systems, normative concepts. Since normative concepts are viewed as systems, their properties as such—both pieces and wholes—are empirically tracked. I describe a focus on both elements and system-level properties as *structurally multi-scaled*.

Richness and productivity—This theme is not so much a built-in assumption of KiP, but it is one of the most powerful and consistent empirical results. Naïve knowledge is, in general, rich and escapes simple characterizations (e.g., as isolated “misconceptions,” simple false beliefs). Furthermore, learning very often—or always—involves recruiting many “old” elements into new configurations to produce normative understanding. This is the essence of KiP as a strongly constructivist framework, and it is one of its most distinctive properties in comparison to many competitor frameworks for understanding knowing, learning, and conceptual change. diSessa (2017) systematically describes differences compared to some contrasting theories of conceptual change. In my reading, assuming richness and productivity of naïve knowledge is comparatively rare, but certainly not unheard of, in mathematics, just as it is in science.

Diversity—An immediate consequence of the existence of rich, small-scaled knowledge is that there are many dimensions of potential difference among learners. Each learner may have a different subset of the whole pool of “little” intuitions, and might treat common elements rather differently. KiP may be unique among modern theories of conceptual change in its capacity to handle diversity across learners.

Contextuality—“Little” ideas often appear in some contexts, and not others. Furthermore, as they change to become incorporated into normative systems of knowledge, the contexts in which they operate may change. So, understanding *how knowledge depends on context* is core to KiP, while it is marginally important or invisible in competing theories. This focus binds KiP with situative approaches to learning (“situated cognition”). See Brown et al. (1989) for an early exposition, and continuing work by such authors as Jean Lave and Jim Greeno.

5.1.2 Empirical Methods

KiP is not doctrinaire about methods, and many different ones have been used.

Two modes of work are, however, more distinctive. First, KiP has *the development and continuous improvement of theory* (models) at its core. We in the community articulate limits of current models, encourage the refinement of old models and the development of new ones, when necessary.

Theory development, in turn, usually requires the richest data sources possible in order to synthesize and achieve the fullest possible accountability to the details of process. This is opposed to data that is quickly filtered and reduced to a priori codes or categories. In practice, *microgenetic* or *micro-analytic* study of rich data sources of students thinking (e.g., in clinical interviews) or learning (full-on corpora of individual or classroom learning) have been systematically used in KiP not only to validate, but also to generate new theory. See Parnafes and diSessa (2013) and the methodology section of diSessa et al. (2016). This kind of data collection and analysis is strongly synergistic with design-based research (diSessa and Cobb 2004), and iterative design and implementation of curricula—along with rich real-world tracking of data in concert with more cloistered and careful “break-out” studies of individuals—have been common.

I now proceed to concretize and exemplify the generalizations above with respect both to theory development and empirical work. I will boldface themes from the above list, as they are relevant. As mentioned, I start with examples having to do with physics, but then proceed to mathematics.

5.2 Two Models: Illustrative Data and Analysis

In this section I sketch the two best-developed and best-known KiP models of knowledge types. As such, the section illustrates KiP as a **modeling approach**. While both models are both **temporally** and **structurally multi-scaled**, the first model, p-prims, emphasizes smaller scales in time and structure. The second, coordination classes gives more prominence to larger scales.

5.2.1 Intuitive Knowledge

P-prims are elements of intuitive knowledge that constitute people’s “sense of mechanism,” their sense of which happenings are obvious, which are plausible, which are implausible, and how one can explain or refute real or imagined possibilities. Example p-prims are (roughly described): increased effort begets greater results; the world is full of competing influences for which the greater “gets its way,” even if accidental or natural “balance” sometimes exists; the shape of a

situation determines the shape of action within it (e.g., orbits around square planets are recognizably square). Comparable ideas in mathematics are that “multiplication makes numbers bigger” (untrue for multipliers less than one); a default assumption that a change in a given quantity generally implies a similar change in a related quantity (more implies more; less implies less, whereas, in fact, “denting” a shape may decrease area but increase circumference); and “negative numbers cannot apply to the real world” (what could a negative cow mean?). In the rest of this section, I will discuss physics examples only.

We must develop a new model for this kind of knowledge because, empirically, it violates presumptions of standard knowledge types, such as beliefs or principles. First, classifying p-prims as true or false (as one may do for beliefs or principles) is a category error. P-prims *work*—prescribe verifiable outcomes—in typical situations of use, but always fail in other circumstances. Indeed, when they will even be brought to mind is a delicate consequence of context (**contextuality**, both internal: “frame of mind”; or external: the particular sensory presentation of the phenomenon). So, for example, it is inappropriate to say that a person “believes” a p-prim, as if it would always be brought to mind when relevant, and as if it would always be used in preference to other ways of thinking (e.g., other p-prims, or even learned concepts). Furthermore, students simply cannot consider and reject p-prims (a commonly prescribed learning strategy for dealing with “misconceptions”). Impediments to explicit consideration are severe: There is no common lexicon for p-prims, and people may not even be aware that they have such ideas. Furthermore, “rejection” does not make sense for ideas that usually work, nor for ideas that may have very productive futures in learning (see upcoming examples).

Example data and analysis: J, a subject in an extended interview study (diSessa 1996), was asked to explain what happens when you toss a ball into the air. J responded fluently with a completely normative response: After leaving your hand, there is only one force in the situation, gravity, which slows the ball down, eventually to reverse its motion and bring it back down.

Then the interviewer asked a seemingly innocuous question, “What happens at the peak of the toss?” Rather than responding directly, J began to reformulate her model of the toss. She added another force, air resistance, which is changing, “gets stronger and stronger [as if to anticipate an impending balance and overcoming; see continuing commentary] to the point where when [sic] it stops.” But then, she introduced yet another force, an upward one, which is equal to gravity, “in equilibrium for a second” at the top, before yielding to gravity. Starting anew, she provided a source for the upward force: It comes from your hand, and it “can only last so long against air and against gravity.” In steps, she further decided that it’s just gravity that is opposing the upward force, not air resistance, and gradually she reformulated the whole toss as a competition where the upward force initially overbalances gravity, reaching an equilibrium at the top, and then gravity takes over.

The key to understanding these events is that the interviewer “tempted” J to apply intuitive ideas of balancing and overcoming; he asked about the peak because the change of direction there looks like overcoming, one influence is getting weaker, or another is getting stronger. J “took the bait” and reformulated her ideas

to include conflicting influences: The downward influence is gravity, but she struggled a bit to find another one, first trying air resistance, getting “stronger and stronger,” but then introducing an upward force that is changing, getting weaker and weaker. This is a striking example of **contextuality**: J changed her model entirely after focusing attention on a particular part of the toss that suggested balancing. However, more surprises were to come.

Over the next four sessions, the interviewer continually returned to the tossed ball, providing increasingly direct criticism. “But you said the upward force is gone at the peak of the toss, and also that it balances gravity there. How can it both be zero and also balance gravity?” Over the last two sessions, the interviewer broke clinical neutrality and provided a computer-based instructional sequence on how force affects motion, including the physicist’s one-force model of the toss. At the end of the instructional sequence, J was asked again to describe what happens in the toss. Mirroring her initial interview but with greater precision and care, she gave a pitch-perfect physics explanation. But, when asked to avoid an incidental part of her explanation (energy conservation), J reverted to her two-force model. So, we know that J exhibits not only surprising **contextuality** in terms of what explanation of a toss she would give, but that contextuality, itself, seems strongly persistent, a core part of her conceptual system.

After the completion of interviewing sessions, J reflected that she knew that it would appear to others that she described the toss in two different ways, and the “balancing” one might be judged wrong. But she felt both were really the same explanation.

Salient points: The dominant description of intuitive physics in the 1990s was that it constituted a coherent theory (see diSessa 2014, for a review and references), and the two-force explanation of the toss was a perfect example. External agents (the hand) supply a force that overcomes gravity, but is eventually balanced by it, and finally overcome. The KiP view, however, is that the “theory” only appears in particular situations (e.g., when overcoming is salient). Indeed, J did not seem to have the theory to start, but constructed it gradually, over a few minutes. **Contextuality** is missing from the then “conventional” view; “theories” comparable to Newton’s laws don’t come and go depending on what you emphasize in a visual scene. J’s case is particularly dramatic since she never relinquished her intuitive ideas, even while she improved her normative ones. Instead, situation-specific saliences continued to cue one or the other “theory” of the toss. The long-term stability of an instability (the shift between two models of a toss) shows an attention to **multiple temporal scales** that is unusual in conceptual change studies but critical to understanding J’s frame of mind. What happened in a moment each time it happened (shifting attention and corresponding shift in model of the toss), nonetheless continued to happen regularly over months of interviewing. Such critical phenomena test the limits of observational and analytical methods. For example, before and after tests are very unlikely even to observe the phenomenon. Attributing “misconceptions” categorically to a subject—“J has the non-normative dual force model of a toss”—fails to enfold this essentially **multi-scaled** and highly contextual analysis of J.

Another subject in the same study, K, started by asserting the two-force model of the toss. However, this subject reacted to similar re-directions of her attention concerning her explanation by completely reformulating her description to the normative model. She then observed that she had changed her mind and explained the reasons for doing so. The two-force model was then gone from the remainder of her interviews.

Ironically, a standard assessment employing first responses would classify J as normative, and K as “maintaining the naïve theory.” Rather, K was a very different individual who could autonomously correct and stabilize her own understanding. J, in contrast, alternated one- and two-force explanations, and didn’t really feel they were different. KiP methodologies did not assume simple characterization of either student’s state of mind (**richness**), and they could also therefore better document and understand their differences (**diversity**). Neither J nor K would be well characterized by their initial responses. J, and not K, was deeply committed to a balancing view of many aspects of physics, even if both found balancing salient and significant in some cases.

Some lessons learned: The knowledge state of individuals is complex, and assessments cannot presume first responses will coherently differentiate them. The assumption of coherence in students’ understanding is plainly suspect; J consistently maintained both the correct view and the “misconception,” even in the face of direct instruction. The interviewer, knowing that fragile knowledge elements like p-prims are important, primed one (balancing, at the peak), and saw its dramatic influence. P-prims explain a lot about the differences and similarities between J and K (both used balancing, but J had a much greater commitment to it), but not everything. In continuing study (diSessa et al. 2002), we discovered that J showed an unusual and often counterproductive view of the nature of physics knowledge, which K did not. Modesty is the best policy: The complex conceptual ecology of students needs continuing work (**continuous improvement**).

One lesson learned here is that p-prims behave very differently than normative concepts. In terms that might be familiar to mathematics education researchers, p-prims provide a highly articulated version (specific elements whose use and contextuality can be examined across many circumstances) of a student’s “concept image” (Tall and Vinner 1981). We need a different model to understand substantial, articulate and context-stable ideas, something roughly akin to “concept definition,” but something that, in my view, uses KiP to better approach the cognitive and learning roots of expertise.

5.2.2 *Scientific Concepts*

Coordination classes constitute a model aimed at capturing central properties of expert concepts.

According to the coordination class model, the core function of concepts is to read out particular concept-relevant information reliably across a wide range of circumstances, unlike the slip-sliding activation of p-prims. Figure 5.1 explains.

Figure 5.2 shows the primary difficulty in creating a coherent concept. All possible paths from world (or imagined world) to concept attributes must result in the same determination. This is called *alignment*, and it is a property of the whole system, not of any part of it.

A physics example of lack of alignment is that students will sometimes determine forces by using intuitive inferences (“An object is moving; there must be a force on it.”), and sometimes by “formal methods” (“An object is moving at a constant speed; according to Newton’s third law, there is no net force on it.”). A mathematical example is that students may deny that an equator on a sphere with three points marked on it is a triangle, even if they have agreed that any part of a great circle is a “straight line,” and that a triangle is any three connected straight line segments.

Coordination classes are large and complex systems. This is structurally unlike p-prims, which are “small,” simple, and relatively independent from one another. Alignment poses a strict constraint on all possible noticings (e.g., noticing F_1 or F_2 in Fig. 5.2) and all possible inferences (e.g., I_1 and I_2): All paths should lead to the same determination. That is, there is a *global* constraint on all the pieces of a

Fig. 5.1 Coordination classes allow reading out information relevant to concepts, here illustrated by “location,” from the world. The readout happens in two stages. (1) “See” or “notice” involves extracting *any* concept-relevant information: “The cat is *above* the mat,” and “The cat is *touching* the mat.” (2) “Infer” draws conclusions specifically about the relevant information (location) using what has been seen: “The cat is *on* the mat.”

Fig. 5.2 In situations where multiple features (F_1 , F_2) are available, different choices of what to observe may lead to different inferences (I_1 , I_2) and potentially contradictory determinations (D_1 , D_2) of the “same” information

coordination class, which makes the model essentially **multi-scaled**. In this case, multi-scaled refers to the **structure** of the knowledge system—pieces and the whole system—rather than to its **temporal** properties, which were emphasized with J.

I will not belabor a full taxonomy of parts of coordination classes, but, because it is relevant to an example from mathematics (Sect. 3.1), I note that a coordination class needs to include *relevance*, in addition to *noticings* and *inferences*. Relevance means that a coordination class needs to “know” when a concept applies and when information about it *must* be available. If you are asked about slope, there *must* be some available information about “rise” and “run,” and it behooves one to attend to that information.

Dufresne et al. (2005) provided an accessible example of core coordination class phenomena. They showed two groups of university students, engineering and social science majors, various simulated motions of a ball rolling along a track that dipped down, but ended at its original height. They asked which motion looked most realistic. Subjects saw the motions in two contexts: one that showed only the focal ball, and another that also showed a simultaneous and constant ball motion in a parallel, non-dipping path. The social scientists’ judgments of the realism of the focal motion remained nearly the same from the one- to two-ball situation. But, the engineers showed a dramatic shift, from preferring the correct motion to preferring another motion that literally no one initially believed to be realistic. In the two-ball case, engineers performed much worse than social scientists!

Using clinical interviews, the researchers confirmed that the engineers were looking at (“noticing”) different things in the different situations. Relative motion became salient with two balls, changing the aspects of the focal motion that were attended to. In the two-ball presentation, a kind of balancing, “coming out even” dominated their inferences about realism. The very same motion that they had resoundingly rejected as least natural became viewed as most realistic.

Lessons learned: Scientific concepts are liable to shifts of attention during learning, and thus different (incoherent) determinations of their attributes. This is an easily documentable feature of learning concepts such as “force,” and there is every reason (and some documentation) to believe this is also true for mathematical concepts. So, people must learn a variety of ways to construe particular concepts in various contexts, ways that are differentially salient in various conditions, yet all determinations must “align.” Again, this local/global coherence principle shows KiP’s attention to **multiple scales of conceptual structure**.

It is only mildly surprising that the “culprit” inference here is a kind of balancing, as implicated in J’s case. So, once again, a relatively small-scaled element, similar to balancing p-prims, plays a critical role. Balancing is a core intuitive idea, but it also becomes a powerful principle in scientific understanding (**productivity**). Changes in kinetic and potential energy do always *balance out*. In this case, engineering students have elevated the importance and salience of balancing compared to social scientists, but have not yet learned very well what exactly balances out, and when balancing is appropriate (relevance). Certain p-prims are thus learned to be powerful, but they have not yet taken their proper place in understanding physics. Incidentally, this analysis

also accounts for a very surprising difference (**diversity**) between different classes of students—engineers and social scientists.

P-prim and coordination classes are nicely complementary models. Within coordination class theory, p-prim turn out to account for certain problems (mainly in terms of inappropriate inferences), but they also can lie on good trajectories of learning, in constructing the overall system. Balancing is a superb physical idea, but naïve versions of balancing need to be developed precisely and not overgeneralized. Linearity is a comparable idea in mathematics. It is a wonderful and powerful idea, but it does not work, for example, for functions in general. $\sin(a + b)$ is not $\sin(a) + \sin(b)$. As balancing and linearity develop, they both need to be properly coordinated with checks and other ways of thinking.

5.3 Examples in Mathematics

This section displays some mathematical examples. The field of KiP analyses in mathematics is less rich than for physics, and overall trends are less well scouted out. But, to give a sense of what KiP looks like in mathematics and to encourage further such work is a primary goal of this article.

5.3.1 *The Law of Large Numbers*

Joseph Wagner (2006) used the main ideas of coordination class theory to study the learning of the statistical “law of large numbers”: The distribution of average values in larger samples of random events hews more closely to the expected value (long-term average) than for smaller samples. In complementary manner, smaller samples show a greater dispersion; a greater proportion of their averages will be far from the expected value. So, if one uses a sample of 1000 coin tosses, one is nearly assured that the sample will have an average close to 50% heads and 50% tails. A sample of 10 tosses can easily lead to averages of, say, 70% heads and 30% tails. In the extreme case, a single toss, one is guaranteed of “averages” that are as far as possible from the long-term average: one always gets 100% heads, or 100% tails.

Wagner discovered that students often showed canonical coordination class difficulties during learning. Many had exceedingly long trajectories of learning, corresponding to learning in different contexts of use of the law of large numbers. In more technical detail, thinking in different contexts typically involves different knowledge (different noticings and different inferences), which may need to be acquired separately for different contexts. Furthermore, reasoning about the law in each context must *align* in terms of “conceptual output” (e.g., what is the relevant expected value) across all contexts. In short, **contextuality** is a dramatic problem for the law of large numbers, and systematic integrity (a **large-scale structural property**—in fact, the central-most large-scale property of coordination classes) is

hard won in view of the **richness** of intuitive perspectives that may be adopted local to particular contexts (**small-scale structure**; think p-prims).

I present an abbreviated description one of Wagner's case studies to illustrate. Similar to the case of J, this is a fairly extreme case, but one in which characteristic phenomena of coordination class theory are easy to see. In particular, we shall see that learning across a wide range of situations appears necessary. The law of large numbers might not even appear to the learner as relevant to some situations, or it might be applied in a non-aligned way, owing to intuitive particulars of the situations. I sketch the subject's learning according to diSessa (2004), although a fuller analysis on most points and a more extensive empirical analysis appear in Wagner (2006).

The subject, called M ("Maria" in Wagner 2006), was a college freshman taking her first course in statistics. Wagner interviewed her on multiple occasions throughout the term (methodologically similar to J's study), and used a variety of near isomorphic questions involving the law of large numbers. The questions asked whether a small or large sample would be more likely to produce an average within particular bands of values, bands that include the expected value, or bands that are near or far from it. Would you choose a small or large sample if you wanted to get an average percentage of heads in coin tosses between 60 and 80% of the tosses? The law of large numbers says you would want a smaller number of tosses; in contrast, a very large number of tosses is almost certain to come out near 50% heads.

We pick up M's saga after she learned, with some difficulty, to apply the law of large numbers to coin tosses. Just after an extensive discussion of the coin situation, the interviewer (Jo) showed M a game "spinner," where a spun arrow points to one of 10 equal angular segments. Seven of the segments are blue, and three are green. Jo proceeded to ask M whether one would want a greater or lesser number of spins if one wanted to get an average of blues between 40 and 60 percent of the time.

M: OK. ... Land on blue? ... Well, 70% of the // of that circle is blue. Yeah. Seventy percent of it is blue, so, for it to land between 40 and 60 percent on blue, then, I would say there really is no difference. [She means it doesn't make a difference whether one does few or a lot of spins.]

Jo: Why?

M: Because if 70% of the // the circle, or, yeah, the spinner is blue, so ... it's most likely going to land in a blue area, regardless of how many times I spin it. It kinda really doesn't matter. It's not like the coins...

M is saying that she does not see the spinner situation as one in which the law of large number applies. The coordination class issue of *relevance* defines one of her problems. The larger data corpus suggests that a significant part of the problem is that M does not see that the concept of expected value applies to the spinner. She knows that in one spin, 70% of the time you will get blue, and 30% of the time you will get green. She reasons pretty well about "chances" for individual spins. But she simply does not believe that the long-term average, the expected percentage of blues or greens, exists. She "sees" chances, but does not infer from them a long-term average, nor even appear to know that a long-term average exists in this case.

Jo showed M a computer simulation of the spinner situation and proposed to do an experiment of plotting the result (histogram) of many samples of a certain number of spins. Would the percentages of blue pile up around any value, the way coin tosses always pile up around 50%? M was reluctant to make any prediction at all. But she very hesitantly suggested that the results might pile up around 70%. When the simulation was run, M was evidently surprised. “It *does* peak [pile up] around 70!!”

Here, we are at a disadvantage because we know much less about the relevant p-prims (or similar knowledge elements) that are controlling M’s judgments, unlike the fact that, for J, the interviewer suspected balancing might provoke a different way of thinking about the toss, or that Dufresne et al. found that “balancing out” also sometimes controlled engineers’ judgments about the realism of depicted motions of rolling balls. A good coordination class analysis demands a better analysis than the data here allow. However, a hint was offered earlier in the conversation when Jo pressed M to explain how the spinner differed from coins. M reported, “The difference, uh, between the coins and this [spinner] is that, in every toss, in the coin, I know that there’s a ... 50% chance of getting a head, 50% chance of getting a tail.” But with a spinner, “It’s just not the same.” Although M cannot put her finger on the difference, it seems plausible that she sees the 50–50 split of a coin flip to be *inherent in the coin*, “in every toss...,” while the spinner arrow, per se, does not visibly (to her) have 70–30 in its very nature. An alternative or contributing factor involves the well-known fact concerning fractions that students seem conceptually competent first with simple ones, like $\frac{1}{2}$. But, again, there is not enough data to distinguish possibilities.

Independent of the reason, the big picture relevant to coordination classes is that M simply does not see the spinner as essentially similar to coins. The *relevance* part of her developing coordination class is the most obvious problem. In particular, she doesn’t naturally see an expected value as relevant to (nor determinable for) spinners. This case has a happy ending because the empirical (computer simulation) result was enough to convince M that expected value existed in the spinner case, and she began to reason more normatively about Jo’s questions. To summarize, there was a conceptual **contextuality** that prevented using the same pattern of reasoning, the law of large numbers, in different situations. M needed to learn that expected value existed for spinners, and that it related to the “chances” concerning a single case in the same way as for coins: The long-term expected average is the same as the “chances” for a single case.

The final case of contextuality I report (there are many others!) concerns the average height of samples of men, corresponding to men in the U.S. registering for the military draft at small or large post offices. If the average height in the U.S. is 5 ft 9 in., would a small or large post office (small or large sample) be more likely to find an average height for one day of more than 6 ft? At first, M had no idea how to answer the question. Pressed, she offered an uncertain reference to larger sets of numbers having smaller averages. The law of large numbers was, again, invisible to her in this context.

Jo improvised yet another context. Would you rather take a big or small sample of men at a university in order to find the average height? M was quick and confident in her answer. A larger sample would be “more representative,”² “more accurate.” Arguably, the sampling context evoked a memory or intuition that larger samples are “better.” Having made the connection to this intuition, M applied it relatively fluently to the post office problem.

The reason “representativeness” and “accuracy” were cued in the university sampling situation and not previously might not be clear. But M did not mention these intuitive ideas in any previous problems, and, once cued, she took those ideas productively into new contexts. The combination of **contextuality** and **productivity**, shown here, is highly distinctive of KiP analyses. Some intuitions, even if they are not usually evoked, can be useful if, somehow, they are brought to the learner’s attention.

The next example is among the first applications of KiP to mathematics (a decade earlier than Wagner’s work), and the final one is among the latest (a decade later than Wagner).

5.3.2 *Understanding Fractions*

Smith (1995) did an investigation of student understanding of rational numbers and their representation as fractions according to broad KiP principles. He began by critiquing earlier work as (a) using a priori analysis of dimensions of mathematical competence, and also (b) systematically assessing competence according to success on tests. Instead, he proposed to look at competence directly in the specific strategies students use to solve a variety of problems. In particular, he did an exhaustive analysis of strategies used by students during clinical interviews on a set of fractions problems that was carefully chosen to display core ideas in both routine and novel circumstances. Smith looked most carefully at the strategies used by students who could be classified as “masters” of the subject matter. So, his intent was to describe the nature of achievable, high-level competence by looking directly at the details of students’ performance.

The results were surprising in ways that typify KiP work. Masters used a remarkable range of strategies adapted rather precisely to particulars of the problems posed. While they did occasionally use the general methods that they had been taught (methods like converting to common denominators or converting to decimals), general methods appeared almost exclusively when none of their other

²Kahneman and Tversky (1972) provide a now-canonical treatment of statistical “misconceptions,” including representativeness. However, their theoretical frame is very different from KiP. **Productivity**, in particular, is missing, unlike the cited role of representativeness in M’s learning. These authors maintain that, to learn, intuitions must be excluded, and formal rules must be followed without question. Pratt and Noss (2002) provide a KiP-friendly treatment of statistical intuitions.

methods worked. A careful look at textbooks suggested that it was unlikely that many, if any, of the particular strategies had been instructed. Student mastery seems to transcend success in learning what is instructed.

In net, observable expertise is: (a) “fragmented” (**contextual**) in that it is highly adapted to problem particulars; (b) **rich**, composed of a wide variety of strategies; and (c) significantly based on invention, rather than instruction. The latter two points suggest **productivity**, the use of rich intuitive, self-developed ideas, and that that richness is maintained into expertise, in contrast to what conventional instruction seems to assume.

One can summarize Smith’s orientation so as to highlight typical KiP strategies, which contrast with those of other approaches:

- avoiding a priori or “rational” views of competence in favor of directly empirical approaches: Look at what students do and say about what they do.
- couching analysis in terms of knowledge systems (**a complex systems approach**) of elements and relations among them (e.g., particular strategies were often, but not always, defended by students by reference to more general, instructed ways of thinking).
- discovering that the best student understanding, not just intuitive precursors, is rich (many elements), diverse, and involves a lot of highly particular and contextually adapted ideas (**contextuality**). Thus it is in some ways more similar to pre-instructional ideas than might be expected.

Smith did not use the models (p-prims, etc.) that later became the recognizable core of KiP. But, still, the distinctiveness of a KiP orientation proved productive. I believe this is an important lesson, that, independent of technical models and details, KiP’s general principles and orientations can provide key insights into learning that are not available in other perspectives. Newcomers to KiP might do well to start their work at this level, and move to more technical levels when those details come to seem sensible, and when and if the value of technicalities becomes palpable.

5.3.3 Conceptual and Procedural Knowledge in Strategy Innovation

The relationship of procedural to conceptual knowledge is a long-standing, important topic in mathematics education. There is a general agreement that one should strike a balance between these modes. However, at a more intimate level, the detailed relations should be important. What conceptual knowledge is important, when, and how? It is known that students can (e.g., Kamii and Housman 2000) and do (e.g., Smith’s work, above) spontaneously innovate procedures. How might conceptual knowledge be important to innovation, specifically what knowledge is important, and what is the nature and origin of those resources?

Levin (2012) studied strategy innovation in early algebra. Her study involved a student who started with an instructed guess-and-check method of solving problems like: “The length of a rectangle is six more than three times the width. If the perimeter is 148 ft., find the length and width.” Over repeated problem solving, this student moved iteratively, without direct instruction, from guess-and-check to a categorically different method: a fluent algorithmic method that mathematicians would identify as linear interpolation/extrapolation. One of the interesting features of the development was that intuitive “co-variation schemes,” more similar to calculus (related rates) than anything instructed in school, rooted his development (**productivity**). Indeed, his development could be traced through six distinct levels of co-variation schemes, progressively moving from qualitative (the “more implies more” intuition, but in a circumstance where it is productive), toward more quantitatively precise, general, and “mathematical-looking” principles.

In order to optimally track and generalize this student’s progress, Levin extended the coordination class model to what she calls a “strategy system” model, demonstrating the generative and evolving nature of KiP (**continuous improvement**). Her model maintained a focus on perceptual categories (“seeing” in Fig. 5.1), and inferential relations (e.g., co-variation schemes). But there were also theoretical innovations: Typically more than one coordination class is involved in strategy systems. General conceptions (inferences) specifically supported procedural actions in particular ways.

In addition to the core co-variational idea, a cluster of intuitive categories, such as “controller,” “result,” “target,” and “error” played strongly into the student’s development. All in all, Levin’s study showed the surprising power of intuitive roots—ones that may never be invoked in school—and provided a systematic framework for understanding their use in the development of procedural/conceptual systems.

5.3.4 *Other Examples*

In addition to what was presented above, I recommend a few other examples of KiP work that will be helpful for mathematics education researchers with different specialties in order to understand the KiP perspective. Andrew Izsák’s has developed an extensive body of work using KiP to think about learning concerning, for example, area (Izsák 2005), and early algebra (Izsák 2000). Similarly Adiredja (2014) treated the concept of limit from a KiP perspective. Adiredja’s analysis is important in the narrative of this article in that it takes steps to comprehend learning of the topic, limits, at a fine grain-size, including the productivity and not just learning difficulties that emerge from prior intuitive ideas. The work may be profitably contrasted with that of Sierpiska (1990) and Tall and Vinner (1981) on similar topics.

5.4 Cross-Cutting Themes

In this final section, I identify KiP’s position and potential contributions to two large-scale themes in the study of learning in mathematics and science.

5.4.1 *Continuity or Discontinuity in Learning*

I believe that one of the central-most and still unsettled issues in learning concerns whether one views learning as a continuous process or a discontinuous one. In particular, how do we interpret persistent learning problems that appear to afflict students for extended periods of time? In science education, so-called “misconceptions” or “intuitive theories” views treat intuitive ideas as both entrenched and unproductive. They are assumed to be unhelpful—blocking, in fact—because they are simply wrong (Smith et al. 1993). In mathematics education, one also finds a lot of discussion about misconceptions (e.g., concerning graphing, Leinhardt et al. 1990) and also about the essentially problematic nature of “intuitive rules” such as “more implies more” (Stavy and Tirosh 2000). But, more often than in science, researchers implicate discontinuities of form, rather than just content. For example, Sierpinska (1990) talks about basic “epistemological obstacles,” large-scale changes in “ways of knowing.” Vinner (1997) talks about “pseudo-concepts” as bedeviling learners, and some interpretations of the distinction between process and object conceptualizations in mathematics (Sfard 1991) put process forms as inferior to conceptions that are at the level of objects (not necessarily Sfard’s contention). Or, the transition from process to object modes of thinking is always intrinsically difficult. Tall (2002) emphasizes the existence of discontinuities possibly due to deep-seated brain processes (“the limbic brain;” sensory-motor thinking). Along similar lines (as anticipated in footnote 2), Kahneman and Tversky’s view of difficulties in learning about chance and statistics relies on so-called “dual process” theories of mind. (See Glöckner and Witteman 2010, for a review and critical assessment.) Instinctive (intuitive) thinking must be *replaced* with a categorically different kind of thinking based on a conscious and explicit rule following.

On the reverse side, mathematics education researchers sometimes have supported the productivity of intuitive ideas (e.g., Fischbein 1987), and, most particularly, constructivist researchers have pursued important lines of continuity between naïve and expert ideas (Moss and Case 1999, is, in my view, an exceptional example from a large literature). However, very few studies approach the detail and security of documentation of elements, systems of knowledge, and processes of transformation of the best KiP analyses.

The issues are too complex and unresolved for a discussion here, but KiP offers a view and accomplishments to support a more continuous view of learning and to critique discontinuous views. For example, both experts and learners use intuitive ideas, even if their knowledge is different at larger scales of organization.

Gradual organization and building of a new system need not have any essential discontinuities: There may not be any chasm separating the beginning from the end of a long journey. It is just that, before and after, things may look quite different. A core difficulty in learning might simply involve (a) a mismatch between our instructional expectation concerning how long learning should take and the realities of the transformation, and (b) a lack of understanding of the details of relevant processes. KiP offers unusual but tractable and detailed models of small-scale, intuitive knowledge that can support its incorporation into expertise, and methodologies capable of discovering and carefully describing particular elements. These issues are treated in more detail in Gueudet et al. (2016).

5.4.2 Understanding Representations

To conclude, I wish to mention two KiP-styled studies concerning the general nature of representational competence—central to mathematical competence—and the roles of intuitive resources in learning about representations.

Sherin (2001) undertook a detailed study of how students use and learn with different representational systems (algebra vs. computer programs) in physics. One of Sherin’s key findings was that p-prim-like knowledge mediates between real-world structure (“causality”) and representational templates. For example, the idea of “the more X, the more Y” (e.g., more acceleration means greater force) translates into the representational form “ $Y = kX$ ” (e.g., $F = ma$). Sherin’s work will be most interesting to mathematics education researchers interested in how representations become meaningful in thinking about real-world situations (modeling), how such situations bootstrap understanding of mathematical structure, and the detailed role that intuitive knowledge plays in these processes. This work builds on similar earlier work by Vergnaud (1983), but in distinctly KiP directions.

Finally, diSessa et al. (1991) studied young students’ naïve resources for thinking about representations. In contrast to misconceptions-styled work, we uncovered very substantial expertise concerning representations. However, the expertise was different than what is normally expected in school. It had more to do with the generative aspects of representation (e.g., design and judgments of adequacy) and less to do with the details of instructed representations. This repository of intuitive competence is essentially ignored in school instruction, an insight shared with a few (e.g., Kamii and Housman 2000), but not many, mathematics education researchers.

References

- Adiredja, A. (2014). *Leveraging students’ intuitive knowledge about the formal definition of a limit* (Unpublished doctoral dissertation). Berkeley, CA, University of California.

- Brown, J. S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher*, 18(1), 32–42.
- diSessa, A. A. (1996). What do “just plain folk” know about physics? In D. R. Olson & N. Torrance (Eds.), *The handbook of education and human development: New models of learning, teaching, and schooling* (pp. 709–730). Oxford, UK: Blackwell Publishers, Ltd.
- diSessa, A. A. (2004). Contextuality and coordination in conceptual change. In E. Redish & M. Vicentini (Eds.), *Proceedings of the International School of Physics “Enrico Fermi”: Research on Physics Education* (pp. 137–156). Amsterdam: ISO Press/Italian Physics Society.
- diSessa, A. A. (2014). A history of conceptual change research: Threads and fault lines. In K. Sawyer (Ed.), *Cambridge handbook of the learning sciences* (2nd ed., pp. 88–108). Cambridge, UK: Cambridge University Press.
- diSessa, A. A. (2017). Conceptual change in a microcosm: Comparative analysis of a learning event. *Human Development*, 60(1), 1–37.
- diSessa, A. A., & Cobb, P. (2004). Ontological innovation and the role of theory in design experiments. *Journal of the Learning Sciences*, 13(1), 77–103.
- diSessa, A. A., Elby, A., & Hammer, D. (2002). J’s epistemological stance and strategies. In G. Sinatra & P. Pintrich (Eds.), *Intentional conceptual change* (pp. 237–290). Mahwah, NJ: Lawrence Erlbaum Associates.
- diSessa, A. A., Hammer, D., Sherin, B., & Kolpakowski, T. (1991). Inventing graphing: Meta-representational expertise in children. *Journal of Mathematical Behavior*, 10(2), 117–160.
- diSessa, A. A., Sherin, B., & Levin, M. (2016). Knowledge analysis: An introduction. In A. diSessa, M. Levin, & N. Brown (Eds.), *Knowledge and interaction: A synthetic agenda for the learning sciences* (pp. 30–71). New York, NY: Routledge.
- Dufresne, R., Mestre, J., Thaden-Koch, T., Gerace, W., & Leonard, W. (2005). When transfer fails: Effect of knowledge, expectations and observations on transfer in physics. In J. Mestre (Ed.), *Transfer of learning: Research and perspectives* (pp. 155–215). Greenwich, CT: Information Age Publishing.
- Fischbein, E. (1987). *Intuition in science and mathematics: An educational approach*. Dordrecht, The Netherlands: Kluwer Academic.
- Glöckner, A., & Wittenman, C. (2010). Beyond dual-process models: A categorisation of processes underlying intuitive judgement and decision making. *Thinking & Reasoning*, 16(1), 1–25.
- Gueudet, G., Bosch, M., diSessa, A., Kwon, O. N., & Verschaffel, L. (2016). *Transitions in mathematics education*. In G. Kaiser (Series Ed.), ICME-13 topical surveys. Switzerland: Springer International.
- Izsák, A. (2000). Inscribing the winch: Mechanisms by which students develop knowledge structures for representing the physical world with algebra. *Journal of the Learning Sciences*, 9(1), 31–74.
- Izsák, A. (2005). “You have to count the squares”: Applying knowledge in pieces to learning rectangular area. *Journal of the Learning Sciences*, 14(3), 361–403.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3, 430–454.
- Kamii, C., & Housman, L. (2000). *Young children reinvent arithmetic* (2nd ed.). New York, NY: Teachers College Press.
- Leinhardt, G., Zaslavsky, O., & Stein, M. K. (1990). Functions, graphs, and graphing: Tasks, learning, and teaching. *Review of Educational Research*, 60(1), 1–64.
- Levin (Campbell), M. E. (2012). *Modeling the co-development of strategic and conceptual knowledge during mathematical problem solving* (Unpublished doctoral dissertation). Berkeley, CA, University of California.
- Moss, J., & Case, R. (1999). Developing children’s understanding of the rational numbers: A new model and an experimental curriculum. *Journal for Research in Mathematics Education*, 30(2), 122–147.
- Parnafes, O., & diSessa, A. A. (2013). Microgenetic learning analysis: A methodology for studying knowledge in transition. *Human Development*, 56(5), 5–37.

- Philip, T. (2011). A “knowledge in pieces” approach to studying ideological change in teachers’ reasoning about race, racism and racial justice. *Cognition and Instruction*, 11(3), 297–329.
- Pratt, D., & Noss, R. (2002). The microevolution of mathematical knowledge: The case of randomness. *The Journal of the Learning Sciences*, 11(4), 453–488.
- Sfard, A. (1991). On the dual nature of mathematical conceptions: Reflections on processes and objects as different sides of the same coin. *Educational Studies in Mathematics*, 22(1), 1–36.
- Sierpinska, A. (1990). Some remarks on understanding in mathematics. *For the Learning of Mathematics*, 10(3), 24–36.
- Sherin, B. (2001). A comparison of programming languages and algebraic notation as expressive languages for physics. *International Journal of Computers for Mathematical Learning*, 6(1), 1–61.
- Smith, J. P. (1995). Competent reasoning with rational numbers. *Cognition and Instruction*, 13(1), 3–50.
- Smith, J. P., diSessa, A. A., & Roschelle, J. (1993). Misconceptions reconceived: A constructivist analysis of knowledge in transition. *Journal of the Learning Sciences*, 3(2), 115–163.
- Stavy, R., & Tirosh, D. (2000). *How students (mis-)understand science and mathematics*. New York, NY: Teachers College Press.
- Tall, D. (2002). Continuities and discontinuities in long-term learning schemas. In D. Tall & M. Thomas (Eds.), *Intelligence, learning and understanding: A tribute to Richard Skemp* (pp. 151–177). Flaxton QLD, Australia: Post Pressed.
- Tall, D., & Vinner, S. (1981). Concept image and concept definition in mathematics with particular reference to limits and continuity. *Educational Studies in Mathematics*, 12(2), 151–169.
- Vergnaud, G. (1983). Multiplicative structures. In R. Lesh & M. Landau (Eds.), *Acquisition of mathematics concepts and processes* (pp. 127–174). New York, NY: Academic Press.
- Vinner, S. (1997). The pseudo-conceptual and the pseudo-analytical thought processes in mathematics learning. *Educational Studies in Mathematics*, 34(2), 97–129.
- Vosniadou, S. (Ed.). (2013). *International handbook of research on conceptual change* (2nd ed.). New York, NY: Routledge.
- Wagner, J. (2006). Transfer in pieces. *Cognition and Instruction*, 24(1), 1–71.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.