WHY "CONCEPTUAL ECOLOGY" IS A GOOD IDEA

ANDREA A. DISESSA

University of California, Berkeley, USA

Abstract. This paper motivates the idea of "conceptual ecology" by critiquing the current mainstream of conceptual change research. Most research on conceptual change suffers from too little theoretical accountability concerning the nature of the mental entities involved and too little use of the details of process data to support its theoretical view. Part of the consequences of these limitations is a vast underestimate of the complexity and diversity of conceptual change phenomena. In contrast, a conceptual ecology approach involves hypothesizing that conceptual change involves a large number of diverse kinds of knowledge, organized and re-organized into complex systems. To illustrate a conceptual ecology approach, we explain two very different kinds of mental entities, p-prims and coordination classes. P-prims are small and numerous intuitive elements that are often quite context specific in their activation. Coordination classes, by contrast, are large systems whose very existence entails a high degree of coordination across diverse contexts. We claim that both p-prims and coordination classes are much more explicit and precise in their assumptions than is typically the case, and they both survive substantial empirical test in the form of analysis of process data.

1. INTRODUCTION

My aim in this chapter is to provide a critique of the current state of conceptual change research and a brief account of how I believe better progress may be achieved. In particular, much prior research in conceptual change has taken a vastly oversimplified view of the process. Figure 1 provides a graphical backdrop on which to illustrate these oversimplifications. The figure shows a naïve concept, A, and its trajectory of development into expert concept, B. What could be wrong with such a picture?

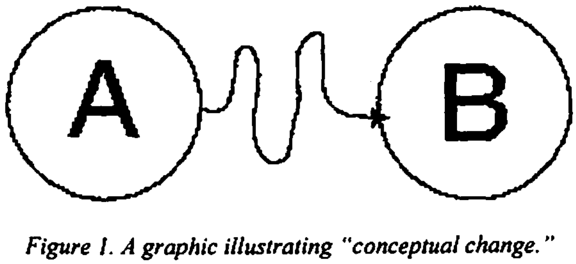

Figure 1. A graphic illustrating "conceptual change."

M. Limón & L. Mason (Eds.), *Reconsidering Conceptual Change. Issues in Theory and Practice*, 29-60.

To begin, we must ask, what are the entities, A and B? The answer most often given is "concepts," although other types of mental entities are sometimes given, say, ontologies (Chi, this volume), beliefs (Hofer & Pintrich, in press), models (Vosniadou & Brewer, 1992), or theories (McCloskey, 1983; Gopnik & Meltzoff, 1996). (To simplify exposition, for the most part I will use "concepts" as an exemplar type of mental entity, although my arguments are essentially unchanged if other types are substituted or if a few are added to a list of types.) To say A and B are concepts begs the question, what is a concept (or any of these other mental entities)? How do we know a concept when we see one? Might it not be necessary to distinguish different kinds of concepts? In this chapter I will strongly motivate the need for a significant variety of types of mental entities to replace the few listed in the literature. More significantly, I will argue that prior work has typically lacked theoretical accountability; it has, indeed, failed to tell us what concepts are, and how to distinguish them from other actual or possible types of mental entities.

Figure 1 shows only two examples of concepts, A and B. Might it not be true, however, that many mental entities contribute to the construction of B? Might it not be true that B is, in fact, a complex system consisting of many interacting parts? My belief is that it is essentially certain that scientific concepts are best considered as complex systems, and prior research has not systematically addressed this possibility seriously. For example, current practice in conceptual change research is far from being able to (and rarely attempts to) match system elements and processes against the details of student reasoning and learning data.

The logical extension of Figure 1 has exactly one naïve concept for every expert concept, and it does not make room for the distinct possibility that naïve concepts have rather different properties than expert ones. Empirical data with respect to these possibilities are easy to come by. Beginning students have many ideas that do not come close to matching expert ideas on a one-by-one basis. It well may be that the naïve conceptual ecology has no exemplar whatsoever that approaches the qualities exhibited by expert concepts.

With an impoverished view of the nature of concepts, it is no wonder that the long, winding path from naiveté to expertise has little exposed detail in the literature. Instead, one finds a variety of unhelpful, definition-begging and probably unfalsifiable terms, like "partial construction," "mixed models," and "confused ideas." And yet, in the classroom, teachers easily find rich and complex intermediate states with which they have to deal; clinical interviews of students essentially always reveal a textured mix of naiveté and learned knowledge, which, however, has had few, if any, systematic descriptions to date.

Figure 2 shows a graphic—obviously simplified—that illustrates the view of conceptual change I advocate in this paper. The naïve state consists of a large number of conceptual elements of varying types. Those elements are modified and combined in complex ways, possibly in levels and into subsystems that, together, constitute the "final" configuration of an expert concept. For reference, I call this a "complex knowledge systems" view of conceptual change—informally, "conceptual ecology."

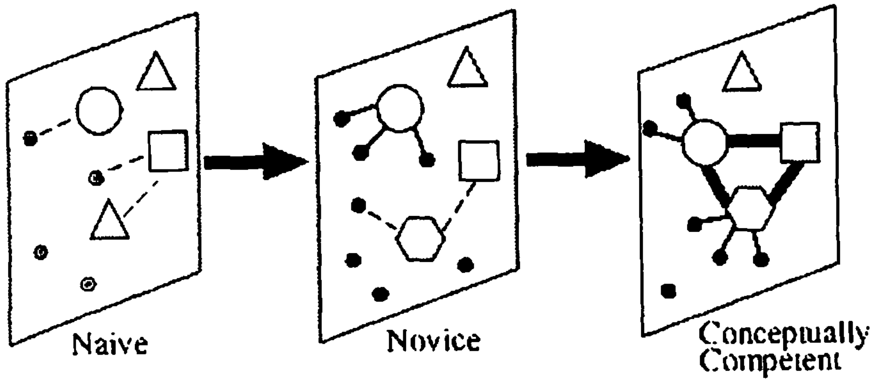

Figure 2. Illustrating a complex systems view, where many exemplars of many different types of knowledge develop and become reorganized in the process of "conceptual change." An "expert concept" draws on many different elements of naïve knowledge (some not belonging to "the naïve concept") that get gradually changed, augmented with new elements, and organized into a new configuration.

The term conceptual ecology has been used by others. In particular, in their influential work on conceptual change, Strike and Posner (1992) speak of conceptual ecology in a similar spirit.[1] More generally, there have been other advocates and allies of the complex knowledge systems view, some implicitly in the details of their analysis of learning complexities, and some more explicitly in the richness of their theoretical framing (e.g., Thagard, 1992). Indeed, in a summary chapter for a section of book on reasoning (Vosniadou & Ortony, 1989), Bill Brewer reports his synthetic conclusion that "in the long run, a proper understanding of the human mind will require that we recognize a large number of very different forms of knowledge and associated psychological processes." (p. 537) Despite sporadic recognition of the importance of a complex systems view, the mainstream of conceptual change research has persisted with vague and oversimplified assumptions about the entities and processes involved in conceptual change. In addition, of course, I intend my contribution to point out particular ways of pursuing a complex systems view that I expect will be most fruitful.

The remainder of this paper elaborates and systemizes the difficulties I see in the contemporary landscape of conceptual change research. Following, I illustrate my approach to improving on the present state.

2. DIFFICULTIES ELABORATED

Broadly, I divide the difficulties in the core of contemporary conceptual change landscape into theoretical and empirical subsets.

---

[1] *However, Strike and Posner did not seem to intend the level of detail in articulating and modeling knowledge types and architectures implicated here.*

2.1. Theoretical Considerations

A lot could be said about the lack of cogent theoretical framing for the issue of conceptual change. However, in this chapter, I underscore only one issue: the lack of well-developed technical terms.

Dictionary meanings can almost never serve the purposes of science. Instead, whenever science is successful, it refines existing terms or adds supplemental ones that can bear a stronger burden. Everyday words are known to be polysemous, combining multiple senses in useful (if ambiguous) packages. Furthermore, even the various senses of everyday words serve only everyday purposes in everyday ways. "Concept" (or one of its various senses or connotations), in particular, seems clear and useful in common usage. However, in the following section, I argue that it is hopelessly vague, covering multitudes of kinds of mental entities with a common coat.

At this point in cognitive studies, we can hope to apply high analytical (as opposed to purely empirical) standards to our technical terms. We could, for example, attempt process models that explain technical terms. Although I won't press far in this direction here, it is good to realize that the literature on conceptual change has rarely attempted process models, nor has it entertained substitute methods of making technical terms' meanings precise.

Besides "concept," other common candidates for useful technical terms suffer similar difficulties to varying degrees. We all have a vague sense of what a theory is. Yet, even if the term is sufficiently well-defined within the social conduct of professional science, in transporting the term to individual learning, a host of changes are likely necessary.[2] In particular, I believe there is convincing data that many naïve scientific ideas are inarticulate, are not easily expressible in words. This, alone, is a dramatic difference from professional science, where complex, careful and symbolically augmented expression (e.g., using algebra) is almost always evident. It seems indubitable that externalizing ideas is more than for archival purposes in professional science. Externalizing allows extraordinary reflective scrutiny and careful reformulation. In contrast, "naïve theories" are never seen directly in the words of subjects, or we might simply be able to ask students for their theories, the way we do with scientists.

"Ontology" has longstanding philosophical roots. To my knowledge, however, no researcher of conceptual change has attempted a process model of ontologies. I don't need to criticize empirical work that implicates shifting ontologies to point out that such work is weakened unless we know what an ontology is, unless we understand how such mental entities may come to exist in some detail and how they function in reasoning and learning. For ontology, as well as for concept, we want to know how we can surpass images like Figure 1 in detail and cogency.

Unless we develop theoretically well-elaborated technical terms, major empirical problems follow. Unless we know what a theory is, how do we distinguish it from a small collection of concepts, or even from a sentence one may utter and toward which one might express some commitment? Tellingly, researchers preferring one

[2] *With regard to social vs. individual perspectives on theory development, see Harris (1989).*

term of another (concept, theory, belief, ontology) essentially never use data to show students' reasoning and learning are inconsistent with other theoretical assumptions. diSessa & Sherin (1998) examine the literature of conceptual change. They argue that even the best and most widely-recognized researchers use inexplicit definitions that implicate ill-defined meanings, and they show what difficulties follow in attempting to interpret data in such vague terms.

The deep problems of conceptual change remain unscathed when we cling to vague, unelaborated terms. What aspects of "theory" are really critical in theory change, and why, after all, is conceptual change difficult when some kinds of learning proceed effortlessly?

In this chapter, I won't propose general criteria for cogency of technical terms in cognitive studies. However, I will illustrate steps toward more adequate terms with two categories among several I have developed in my own studies of conceptual change.

2.2. *Empirical and Quasi-Empirical Considerations*

This section views the current state of conceptual change research through an empirical lens. I will argue that researchers have used very weak empirical strategies that avoid the real complexity of conceptual change. In particular, I make the case that, without much effort, we can strongly motivate, if not prove, that the appropriate default approach to studying conceptual change recognizes diversity in mental entities.

Several trends that I take to be paths of improving our study of conceptual change will become evident in this section. The first, already mentioned, concerns *types*. In particular, I advocate a trend toward *multiplicity*, a greater number of (more accountable) types of mental entities. A second trend concerns *grain size*. Here, the trend should be toward *a greater number of smaller scale elements*. Concomitantly with the second trend, in investigating large-scale accomplishments like "conceptual change" we are necessarily studying *systems of interacting elements*. A final trend concerns increased care in dealing with *invariance*, that is, the issue of when two situations evoke the same conceptual elements. With a rich selection of knowledge elements, we are forced into much more specific consideration of context. If a conceptual ecology contains thousands of elements, certainly the issue of when which are activated is highlighted. In fact, we should expect a greater degree of *context dependency*. Combining trends toward increased contextual dependency, toward multiplicity and toward smaller grain size suggests that an application of a concept is likely to be better viewed as the selected activation of particular concept subcomponents, depending on context. This particular observation will become a core concern when we turn to one of my sample knowledge types, "coordination class."

2.2.1. Diversity

This subsection considers three types of diversity in conceptual change that are inadequately handled by current studies. The first is a quasi-empirical view of the diversity of *kinds* of concepts that exist. I say "quasi-empirical" to indicate that my intention is to draw on good intuitions we all should have about knowledge, rather than to suggest a particular empirical result or program of studies.

Consider the following examples of concepts and their diversity along a number of dimensions. "Dog" is a familiar concept. This is a category-like concept; its primary function is to classify entities in the world into members and non-members of some set. We shall see that other kinds of concepts very much background this particular function. Concerning the competence to recognize exemplars and non-exemplars of "dog," we note that many similar concepts must draw on some similar or identical abilities. Recognizing sheep, cats, raccoons, even people involves similar shape-determining methods; it involves recognizing families of textures (like smooth or hairy), recognizing related categories of things (like eyes, faces, etc.) and considering their systematic variation (bigger, smaller; oblate or round; etc.). Although there is no consensus that category-like concepts work by matching potential exemplars to a prototype, at least this is a plausible mechanism for doing the main work of these kinds of concepts (Rosch, 1994). We know that a fair amount of work goes into learning these concepts since, for example, children take time to get them right. On the other hand, essentially every child masters common categories without a great deal of explicit instruction.[3]

Now consider the concept "bluare," which is the artificial category of things that are both blue and square. A long tradition of psychological studies has investigated the properties of such concepts. For adults, the learning trajectory for bluare can be short, perhaps 10 seconds. It is an unproblematic combination of unproblematic other concepts (blue, square, and the logical connector "both").[4] Learning can be accomplished articulately, by explaining the concept, in contrast to the unlikelihood of learning to recognize a new kind of animal, a wombat, for example—or to acquire the ability to classify animals at all—by being given a brief description. Bluare is at an extreme end of learnability. It is less difficult than learning to recognize animals, and far less difficult that acquiring a scientific concept like force. Does it deserve the same appellation, "concept"? If it does, how much does the category tell us about

[3] *Difficulty in mastering a concept does not necessarily depend on knowledge-structural characteristics of the concept. It might be, for example, that the natural environment is simply impoverished in support for learning it. However, even in that case, difficulties in learning different concepts may certainly be systematic, even if they are in relation to existing or possible learning environments rather than intrinsic to the nature of the concept.*

[4] *The fact that bluare is a simple combination of already existing concepts probably explains its ease of learning. Wouldn't it seem important to know which already-known elements contribute to learning new scientific concepts? Surely learning most concepts can't be tabula rasa, and a profile of contributing elements is more than plausible as an important part of understanding a knowledge systems view of a "new" concept. Yet, conceptual change literature, for the most part, treats the issue at the coarsest possible grain size; as in Figure 1, the expert concepts of heat and temperature, or force, etc., grow out of naïve versions thereof.*

learning difficulties or the plausibility of various instructional strategies—such as talking as opposed to examining multiple exemplars?[5] Of course, since bluare is an extreme example, few would study its learning as related to scientific concepts. Yet, how do we know there is not a complex continuum of scientific concepts, with importantly different qualitative properties, as illustrated by dog, bluare, and (to come) force?

Consider the concept "number," cardinality. I mean to refer to the operational concept (in somewhat the Piagetian sense), enfolding the operational properties of cardinality, such as invariance on rearranging a set of objects, or invariance in the case of no additions or removals. This is much sparser knowledge territory than animal types. What categories are similar to number and might share perceptual/conceptual strategies, as recognition of dog shares with recognition of raccoon? There simply aren't any similar categories (until one gets to the rarified air of group and field theory). So, I argue, the learning of the concept number is very likely to be quite different from learning dog. Similarly, there is no plausible prototype for "number." No exemplar, say 7 as the cardinality of a set of sheep, has the right properties to be a prototype—such as a large number of typical, but neither necessary nor sufficient, properties. All cardinalities have exactly the same operational characteristics.

Finally, consider the central concept in most of the examples to follow, the physics concept of force. Quite plausibly (and a significant literature backs this up; e.g., Clement, 1983; diSessa, 1993), kinesthetic experiences of effort and accomplishment are genetic ancestors to this concept. And yet, it is clear from extended learning difficulties that the professional concept exhibits definitive differences compared to any possible naïve version of the concept. This is quite a particular regime of learning: Substantial prior resources exist (perhaps somewhat like bluare), and yet a deep gulf exists between naïve and expert concept. The nature of the gulf is as yet much in debate, which, once again, highlights the need for a refined ability to describe a learning trajectory bit-by-bit in order to understand and catalog such gulfs.

With respect to multiplicity, my program of research has attempted to identify many different types of mental entities (diSessa, 1996). For simplicity, however, I will discuss only two particular types that make evident a huge range of kinds of mental entities that may be implicated in a more refined view of conceptual change than we currently have. Fortuitously, the pair of types I discuss are both (a) among the most theoretically developed and also (b) are dramatically different in their properties. So exposition concerning these two types can do double duty in this chapter.

I wish to mention briefly two other kinds of diversity not well-respected in the literature on conceptual change. The first was already alluded to in the discussion of diversity in types. That is the diversity of states in the midst of a learning trajectory, including classifying easy and difficult parts of learning a concept and

[5] *Thomas Kuhn is a good connection here. In the "Structure of Scientific Revolutions" (Kuhn, 1970) he talks about the complex non-linguistically mediated process of learning scientific concepts in the course of studying exemplars of use of those concepts.*

understanding the reasons for their difficulty (or ease of learning). For example, learning literally the equation $F = ma$ is an easy accomplishment in the context of learning the concept of force. That may be obvious, but a theoretical accounting for ease of learning is complementary and may be part and parcel of an accounting for learning difficulties. Later, I will localize and describe the reasons for learning difficulties as part of my discussion of my two exemplar knowledge types.

Finally, accounts of conceptual change have all but ignored individual differences. Why do some people learn certain concepts almost effortlessly while others do not? While it might be true that an adequate accounting of naïve conceptual resources available for incorporation into the construction of an expert concept could account for such differences, I believe that meta-conceptual and epistemological issues are at least as prominent. See the discussion of meta-conceptual issues in coherence and consistency, later in this chapter and consult diSessa, Elby, & Hammer (in press).

All in all, I claim there is a huge diversity in conceptual learning phenomena that is not remotely accounted for in current accounts of conceptual change. Different sorts of concepts are evidently different, one from another, and may need individual accounts of relevant processes of change. Even if concepts are all the same in some deep structural sense, or differ systematically along a few dimensions, reconciling apparent differences with such hypothetical deep commonalty has not been accomplished.

In order to deal with apparent empirical diversity in knowledge types, my preferred method is straightforward—to begin developing a larger and more accountable list of types of mental entities.[6] It happens that, if one allows certain kinds of mental entities, not only does the number of kinds of entities increase, but the number of exemplars of each kind is likely to dramatically escalate. It may be that dozens or more such elements lie behind the construction of a single professional concept, in which case the class of stories about conceptual change illustrated in Figure 1 is patently hopelessly inadequate.

In addition to a diversity of types (*multiplicity*) and a proliferation in numbers of knowledge elements (*grain size*), several other trends seem strongly implicated in the above discussion. Reduction in grain size and therefore an increase in the difficulty of even listing all the cognitive elements that go into conceptual change entails a *systems approach*. If many elements go into the construction of a concept, how are they coordinated and combined to produce "a scientific concept"? Furthermore, if multiple elements are involved, then we must describe much more carefully when they work (*contextuality*). A greater accountability to contextuality also means we may have a much better chance to describe particular configurations that cause problems or lead to productive new accomplishments along an extended learning trajectory.

[6] *In principle, other solutions may exist, such as accounting for differences merely in terms of system properties of different configurations of identical elements.*

2.2.2. *Methods*

If the arguments in the above section are at all cogent, it is hard to imagine how a core portion of the relatively large literature on conceptual change has managed to ignore the implicated facts. Without entering into too much detail, I suggest that the reason is two-fold. First, as I argued above, the level of theoretical accountability generally is still very low. Hence, we simply have very fuzzy lenses with which to inspect conceptual change. It all seems a mushy soup where what counts as a theoretical perspective can never come to brass tacks in allowing a comparison of its relative effectiveness to that of another view. Vague ideas are extremely hard to hold accountable. Instead, they are mere motivators of experiments and broad interpretive frames for results. Investigators paint data with "word pictures" invoking their favorite terms without ruling out alternative interpretations, and without any strong tests for the cogency and adequacy of the terms they apply.

Concomitant to weak theory, empirical studies don't attend to details. Researchers do not attempt detailed accounts of particular applications of concepts—or descriptions of what has transpired in a segment of learning protocol—for the simple reason that there are no specific, "mechanistic" stories in the offing.[7] "Theoretical frames" are too weak to rule anything out, and they don't have enough detail even to ask good empirical questions. My basic contention is that we have said nothing falsifiable when we say force is a concept, or that impetus ideas constitute a theory (McCloskey, 1983), until we have said much more about what each of these knowledge terms entails.

Stunningly little process data is taken into account in conceptual change research. By and large, the paradigm has employed before and after snapshots. Right and wrong answers—or similar behaviors, all distanced from cogent theoretical accounts of the elements and processes of change—are counted. Protocol segments are glossed as suggestive reflections of, for example, an "underlying theory," without argument that all the elements of a theory are evident, and without competitive argument that other explanatory constructs are less adequate than "theory." Sentences are taken to represent theories, and words are taken to represent concepts, ignoring the diversity in types of concepts or theories that we should expect. There is a huge ontological gulf between, on the one hand, protocol coding categories and, on the other hand, knowledge element types or system configurations. Yet, our research techniques have yet to clearly distinguish these.

To sum up, I advocate a richer empirical accountability, parallel to one for theory. The hallmark of such accountability, I suggest, will be process analyses of concepts (or ontologies, or theories) *in use* and *in change* that rule out things that do not happen, predict things that do, and explain detailed properties of what people do in reasoning and learning—especially explaining surprising things we might not have noticed before. While I do not intend to be prescriptive or limiting, my own

---

[7] *There are exceptions. See, for example, Schoenfeld, Smith, and Arcavi (1993). Computer modeling of competence, for example, with production systems also constitutes a significant exception. However, conceptual competencies, unlike skill, is a rare target of such modeling efforts.*

empirical efforts have mainly been in terms of clinical interviews. Such interviews allow repeated episodes involving a particular mental entity to be investigated, triangulating on the properties of the entity. They allow nuanced setting of the context to investigate contextuality. They show subjects' levels of commitment, such as certainty and ambivalence. They also expose paths of reasoning that lead to answers, plainly separating answers from the ideas that generate them. In a clinical interview, we have many opportunities to look at small moves in learning that contribute, overall, to conceptual change. The data sections of this paper will show, by example, some of these properties of clinical interviews and how they can contribute to more refined empirical analysis of concepts and conceptual change.

I've set out a big agenda for the rest of this chapter. I want to illustrate:

1. What more accountable theoretical terms might look like.
2. What different exemplars of such terms are plausible and plausibly play an important part in conceptual change.
3. How different from each other knowledge types might be, and how different types can account for different aspects of reasoning and conceptual change.
4. What sort of details of subjects' behavior can be accounted for with sharpened theoretical terms. For example, what are typical easy and difficult accomplishments along the path to conceptual competence?

I don't intend to prove or demonstrate here, but can only illustrate and suggest. Adequate empirical and theoretical accountability, given the general state of the art, doesn't fit easily into a half-chapter of a book. There is simply too much theoretical groundwork to do, and too much detail in adequate empirical argument.

3. STEPS TOWARD A SHARPENED THEORETICAL AND EMPIRICAL ACCOUNTABILITY

3.1. P-prims

The first knowledge type I discuss has had an extensive history, and it has been rather thoroughly explained in other places. See, for example, diSessa (1983, 1988; 1993; 1996). What stands here is a review.

3.1.1. The nature of p-prims

P-prims, I claim, constitute the bulk (but not the totality) of intuitive physics, the precursor knowledge that gets reconstructed into schooled competence with Newtonian physics. The name, p-prim stands for "phenomenological primitive." (The relevant senses of "phenomenological" and "primitive" are explained in diSessa, 1993. However, consider the characteristics described below.) P-prims have the following properties.

- *Small and monolithic*: P-prims are small and simple knowledge elements. They are atomic in the sense that they are essentially always evoked as a whole (in contrast to scientific concepts, which I believe can be accounted for only with a systems analysis).

- *Many:* There are many, many p-prims, in the hundreds or thousands. The full collection of p-prims exhibits some mild degrees of systematicity, but p-prims are loosely coupled. They do not exhibit deductive relations or any other systematicity typically expected of, for example, theories.
- *Work by recognition*: A good candidate model for p-prims' activation and use is "recognition." One simply sees them in some situations and not in others.
- *Feelings of naturalness; judgments of plausibility*: The prototypical function accomplished by p-prims is to provide a sense of obviousness and necessity to events. If you see something pushed, you are not surprised—in fact you expect—that it moves in the direction of the push. An event or explanation feels plausible to the extent that it matches your intuitive p-prims that relate to the circumstance at issue, and it is surprising to the extent that it does not match.
- *Explanatorily primitive*: Generally, nothing can be said about why the behaviors prescribed by p-prims happen. There is no "covering theory" or articulate reasoning on tap that explains them.
- *Fluid; data driven; lack of conflict resolution*: While p-prims are sometimes strongly cued by a situation, many times they will be much less firm in their activation. In these cases, subjects might have an intuition about what might happen but then lose it as their attention shifts. In some instances, several conflicting p-prims might apply, and there is unlikely to be any way to resolve such conflicts.
- *Problematic connection to language*: P-prims are not words or word senses and are not encoded linguistically. Describing p-prims in words is difficult or impossible (for subjects, as opposed to theorists).
- *Origins in "minimal abstractions"*: Generating new p-prims is neither difficult nor particularly rare. They are frequently fairly simple abstractions of familiar event, such as the fact that a pushed object moves parallel to the push. However, p-prims' properties, especially attachment to particular circumstances, are determined typically by a long process of development.
- *Development by reorganization*: P-prims are not extinguished or replaced by learning scientific concepts. Instead, many p-prims find useful places in the complex system that is an effective scientific concept. A p-prim might come to be known as an effective special case of a scientific principle, and it will be used in place of the principle in apt circumstances. However, p-prims will no longer function as explanatorily primitive. Physics explanations need articulate accountability that p-prims can't provide. The changing function of p-prims in learning and, indeed, the natural evolution of the collection of p-prims, may be described as "shifting priorities," degrees of attachment to particular contexts of use.

These properties of p-prims are not an ad hoc collection. They are mutually dependent and mutually suggestive in many ways. For example, the fact that the elements are small suggests large numbers. Large number are reinforced by the fact that p-rims are relatively easy to generate. A single mechanism of learning (shifting priorities) accounts for naturalistic learning of p-prims and what happens in school. Lack of articulateness goes hand-in-hand with data fluidity. For details on a process model of p-prims's activation and use, which further integrates these characteristics, see diSessa (1993).

Here is a list of some mostly important p-prims. All of them will be used in empirical analyses later in the chapter.

Some basic p-prims:

- *Ohm's P-prim*: A tri-partite element with an impetus (effort), a resistance, and a result: Effort works through a resistance to achieve a result. Ohm's p-prim entails the following expectations: More effort begets more result; more resistance begets less result; and so on.
- *Force as a Mover*: An abstraction of a push or toss: Things go in the direction you push them.
- *Dying away*: Induced motion just dies away, like the sound of a struck bell.

Balance and equilibrium:

- *Dynamic balance*: Sometimes, efforts or impetuses conflict and (accidentally) cancel out, like two people of equal strength pushing against each other.
- *Overcoming*: A situation of conflicting efforts, where one wanes or increases, yields a characteristic switch from the outcome associated with one effort to the outcome associated with the other. For example, a person pushing against another increases his effort and moves the other back.
- *Return to equilibrium*: Systems that are "out of balance" tend to return to "equilibrium."[8] For example, a balance scale pushed out of level returns to level. Water levels itself in a pan.
- *Generalized springiness*: In "out of balance" systems, the displacement from equilibrium is proportional to the amount of perturbation that is applied.

A less important p-prim:

- *Contact conveys motion*: An object (typically small or light ones) in contact with another moving object (typically a large or heavy one) just moves with it. For example, a box in a wagon moves with the wagon.

3.1.2. Process data concerning p-prims

I will illustrate the characteristics and explanatory force of p-prims (and, later, coordination classes) mainly with examples taken from an extensive interviewing corpus involving one freshman college student, whom I call J. Details can be found in diSessa (1996; to be submitted) and diSessa, Elby & Hammer (in press).

The following example illustrates several of the characteristic properties of p-prims. It shows *data fluidity* in that J at first feels she sees how a situation will

[8] *Here and subsequently I use quotation marks on glosses of p-prim based descriptions in order: (a) to suggest how subjects might describe the relevant situations in words and (b) to warn that these descriptions are* not *proper physics analysis.*

behave (in this case, she seems to evoke generalized springiness), but then she loses track of that intuition (when provoked to consider "why") and changes her mind. When queried about how she had thought before changing her mind, she has essentially nothing to say except that she lost the sense she once made of the situation (that is, she loses her initial p-prim-based *sense of naturalness*). Her lack of anything to say illustrates the problematic connections between p-prims and language (*inarticulateness*).

The instigating interview probe was to ask what happens when a small weight is added to a pulley setup in which two large and equal weights are "balanced." See Figure 3. (The word balance is not used; instead, the idea is provoked by a symmetrical picture.) What J says initially is consistent with generalized springiness. The weight simply perturbs the system by some amount. However, a different intuition soon takes over, one that happens to be correct—the unbalanced system continues to move away from "equilibrium." Note that, if J had managed to keep a stable view of this situation as illustrating generalized springiness, the prediction is that she would have implicated a more extensive movement for a heavier perturbing weight, and, possibly, a counter-factual "return to equilibrium" if the weight were removed (return to equilibrium p-prim). Had she not evidently lost track of her initial conceptualization, her lack of showing proportionality of excursion from equilibrium with perturbing strength would have contradicted the proposed description of generalized springiness. Thus, matching process data provides many opportunities to contradict theoretical assumptions or prior empirical results (such as the nature of generalized springiness). See the "principle of invariance," discussed later.

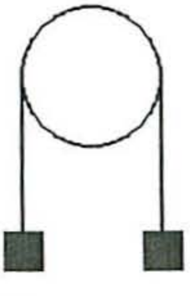

Figure 3. A "balanced" pulley situation is "disequlibrated" by a small weight on one of two larger ones.

[In protocol transcripts, ellipses denote omissions, and "//" denotes interruptions or abrupt halts, usually followed by a restart. Comments and clarifications appear in italics in brackets.]

J: I mean if that's a tiny, tiny weight, it'll probably go down, and then it'll come to rest again. And then the other one'll go up, and they'll be hanging, like, there. If that weight is as small as it's drawn. *[Note the suggestion that it might be different (e.g., greater "disequilibration") if the small weight were larger.]*

I: So but, so but why would it come to rest? That's a little funny ... why would it stop?

J: It actually// It wouldn't. It wouldn't stop. It would keep going slow, slow, slow all the way down. It would not stop.
I: So you changed your mind.
J: Changed my mind. ...
I: Can you say what made you change your mind?
J: Cause it didn't make any sense.
I: Cause it didn't make any sense?
J: No. I don't know why I said that. ... I don't know. I don't know if I didn't think about it or I just sat there thinking in my mind, but I was// I mean, I know that that wouldn't happen.

Notice that, while not very detailed, this account uses the general properties of p-prims to explain (if not predict) many details in this process segment. It uses a previously documented p-prim (generalized springiness) to explain a counter-factual prediction; it explains (momentary) sensemaking; it explains "losing track" as a general phenomenon involving p-prims (data fluidity); it explains inability to articulate.[9]

Predictions by students that have similar properties to this are incredibly easy to demonstrate. I have personally documented scores of p-prims that are at once plausibly "naturalistic" (learnable with experience in the natural world) and explain surprising non-physics predictions and explanations given by students (e.g., diSessa, 1993). Yet, no other account of conceptual change gives any status at all to "many and little" elements like p-prims. "Large," intrinsically difficult-to-change entities like concepts, theories, and ontologies provide no explanation whatsoever for commonplace cognitive phenomena.

I want to make brief reference to one of the most significant successes of p-prims in accounting for important phenomena in conceptual change. It is described in detail in diSessa (1996), and illustrates how process data can bear on issues of conceptual change. In another session, J was asked to describe what happens when a ball is tossed into the air. Initially, J provided a proper physics explanation, involving only one force, the force of gravity. However, J was then prompted to consider the peak of the toss. The point of this probe was to test the salience to J of dynamic balance and overcoming p-prims. Consistent with data fluidity and in accordance with previously documented p-prims, J began to reformulate her explanation in terms of two competing forces, where one overcame the other. After an extended bout of reasoning in which J tried alternative candidates for the force competing with gravity, J reached a stable explanation of the toss involving at least four p-prims in a natural configuration: J used force as mover, dynamic balance, overcoming, and dying away. One thing that is particularly notable is that the explanation she produced had been documented in the literature and touted as exemplifying a deeply held intuitive theory of motion (McCloskey, 1983). Thus, we see a student's description emerge on the basis of well-documented intuitive

[9] *Of course, in a short exposition we cannot rule out explanations such as that the interviewer's questioning her prediction caused her to hide (rather than lose) her initial idea. Similarly, she might have been able to explain and justify, but just chose not to. None of these alternative explanations are consistent with the broad corpus of data from J, but we cannot enter into details here.*

resources and their properties. J's description of the toss is a semi-stable—albeit constructed-on-the-fly—configuration of known entities with known properties. If the description constitutes a theory, we have seen in this analysis how that theory can arise, at least in one case. If the description does not implicate a theory, we still see how it came about and probably know more about its properties due to this process data analysis than is contributed by calling it a theory. P-prims show penetration into the details of process data that are exceptional in the conceptual change literature.

3.2. *Coordination Classes*

In this section I describe another knowledge type, which has properties almost polar opposite to p-prims. Because this type of mental entity is newer and less researched than p-prims, I make a more extended exposition of its properties.

Coordination classes are large, complex systems, rather than atomic elements. P-prims are extremely likely to constitute fragments of coordination classes. Unlike p-prims, which couldn't under the most extreme circumstances be mistaken for a full-blown scientific concept, coordination classes are, in fact, intended to constitute a model of a certain type of scientific concept (and, possibly, some non-scientific concepts as well). Other types of scientific concepts may have quite different properties, and coordination classes themselves have a range of parameters that mean the construction of a coordination class may be different in one case compared to another. For simplicity, parametric differences among coordination classes will be almost completely suppressed in the remainder of the chapter.

Unlike p-prims, which play a role in both expert and naïve thought, coordination classes may well not exist (at least with similar parameters of behavior) in naïve thinking. In any case, whether and which naïve coordination classes exist is an open empirical question.

3.2.1. *The function of coordination classes*

The functional specification for p-prims is to provide feelings of naturalness in familiar situations, surprise and possibly learning-inducing attention in situations we don't understand, and expectations that can be instrumental (e.g., we can "increase our effort" in accordance with Ohm's p-prim if we want "more effect"). In contrast, coordination classes provide a very different functionality. They provide the means for getting a certain class of information from the world. The fundamental assumption behind the idea of coordination classes is that information is not transparently available in the world. Instead, we have to learn how to access different kinds. Indeed, in different circumstances, we may need to use very different means to determine the same kind of information. In general, people must be creative in using any information that may be easily available in a situation, and then inferring the specific information they need from that.

3.2.2. The structure of coordination classes

While p-prims are nearly atomic, although embedded in a recognition system, coordination classes have a lot of internal structure. A great deal of this structure depends on specifics related to particular coordination classes. However, some large-scale partitioning of its internal parts can be made. In particular, I distinguish the set of methods by which any relevant information is gleaned from the world, which I call *readout strategies*, from the collection of possible inferences that can be drawn from available information. The latter set of inferences I call the *causal net.* "Causal" is to be understood in a very general sense. The causal net encompasses inferences that may seem more mathematical or a priori in addition to some we would instinctively describe as causal.

Let me exemplify with one of my favorite hypothetical examples of a (possible) coordination class. This example is taken from Piagetian studies of children's concepts of time (Piaget, 1969).[10] The question is, how do children learn to "see" time interval in the world? "Time interval" is the type of information whose processes of determining (coordination class) we want to explore.

Consider the following situation, illustrated in Figure 4. A blue train and a red train leave a station at the same time, A. The blue train is slower than the red train. The red train stops at a time, B, before the blue train stops, at time C. Because the red train travels much faster than the blue train, the blue train doesn't manage catch up to the red train's stop position by time it, itself, stops.

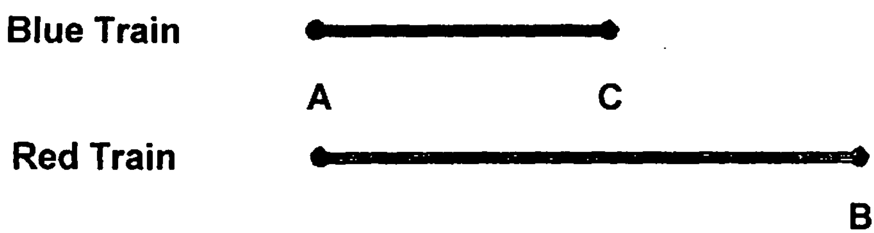

Figure 4. A red train and blue train leave the station simultaneously. The slow blue train continues to move after the red train stops, but not long enough to catch up.

There is a lot of variability in what children say in response to such a setup. However, some respond as follows. When asked which train went on for a longer time, they say that the red train did. While this may seem to entail a simple confusion of meaning of "long," that possibility can be ruled out be rephrasing the question. For example, children may be asked, "If the red train stopped at lunch time, what time did the blue train stop?" Children respond that the blue train stopped before lunch time.

[10] *I've used this example in several places. For example, it appears in diSessa and Sherin (1998).*

On further questioning, it turns out that children are accurately reading relevant information out of the situation. When queried whether the blue train was running when the red train stopped, they (sometimes) answer, accurately, "yes." Furthermore, they may also acknowledge that the red train was not running when the blue train stopped. In terms of coordination class theory, children have quite adequate readout strategies.

To an adult, the information about relative stopping compellingly and automatically suggests a conclusion. If a blue process started at the same time as a red process, and the blue process continued when the red one stopped, adults instantly conclude that the blue process went on for a longer time than the red process. They also know that no information about position (e.g., *where* the processes ran or stopped) is relevant to deciding time duration. In terms of coordination class theory, adults have a causal net that contains the appropriate inference on the basis of observations about whether one process continued when the other stopped.

It is not true that children have no causal net at all. They have a different one. They infer (sensibly enough) that if a moving object gets farther away, it has been running longer. The problem is that they don't know the applicability conditions for that inference, that it only applies to a single object or to multiple objects moving at the same speed.

The causal net is not a homogeneous subsystem. As we shall see, it might contain p-prims (like "farther implies longer duration"), articulate causal assumptions (like force—and only force—causes motion, so that if there is motion, there is a force), and even equations, like $F = ma$. We shall illustrate these possibilities shortly.

3.2.3. *The development of coordination classes*

The development of a coordination class is an extended and complex affair. The descriptor "coordination" implies that a lot of pieces must be put in place to achieve an effective coordination class. Given what we have said, above, however, we can describe different phases and different kinds of difficulties faced in this construction.

In order to read out the particular form of information of interest, one must accumulate both (1) readout strategies and also (2) inferences in the causal net that are sufficient to cover all contexts in which the information is needed. We call this achieving appropriate *span*. In each context of interest, it may be that particular readout strategies and inferences need to be combined. We call this *integration*, as in appropriately integrating several pieces of knowledge to serve the needs of particular contexts. Finally, because different readout strategies, different inferences, and different combinations may be used in different circumstances, a properly formed coordination class must be *aligned*. That is, the same information should be read out, no mater what the context, no matter which readout strategies and which inferences are used. In the data that follow, we will see, in particular, dramatic failures of alignment due to contextuality of pieces of the coordination class.

uncertain whether they are looking at a force (or an acceleration, or any other technical category). How does one "look closely" to detect whether something is a force or not? This implicates probably a very important class of meta-knowledge. Among other things, it requires an understanding of what is essential to the category "force," and what is not.

Figure 6. Do you see fiveness?

The element in Figure 6 illustrates a failure of alignment. Your coordination class knowledge sees an arc connected to a right angle and concludes that this is a five. However, it is not. It is a mirror-image five. In this circumstance, your coordination class strategies see something other than fiveness, which phenomenon we describe as a *failure of alignment* (a failure to read out the same information in different circumstances). Alignment can be restored by adding strategies that detect the difference between a true 5 and its mirror image.

3.2.4. *Process data concerning coordination classes*

"Coordination class" is intended to be a theoretically well-elaborated candidate for a model of scientific concepts. In particular, we hypothesize that the properties of conceptual development of the physics concept of force are explainable as consequences of its being a coordination class.

Force as a "concept" has many properties that suggest it is a good candidate for a coordination class. First, crude data about the difficulty of learning force suggest it may have many pieces and parts—that is, it might well be a system in the range of complexity of a coordination class. Furthermore, force has a huge range of applicability in physics. It seems very likely that students will need to coordinate many readout strategies and causal net inferences to properly see forces and their properties (strength and direction) in different circumstances, such as a ball tossed in the air, a spaceship moving in orbit about the earth, a book sitting on a table (where physicists "see" a force up on the book from the table), and so on. We expect to see typical learning problems in accomplishing appropriate span, integration, and alignment.

Before presenting data, it is worth sketching some elements of a physicist's coordination class for force. First, physicists know one really intuitive way to "see" forces. You can feel them. Physicists expect to be able to feel forces in most circumstances, and to roughly determine magnitude from its feel. This is a particularly interesting aspect of expert coordination: It is informal, and it involves sensation (rather than involving "formal" or "rational thought"). Generally, researchers equate learning physics with learning formal aspects of it. But, such a distinction has no obvious place in coordination class theory.

Two other methods of coordinating force are important to mention. These are more "formal," and relate directly to failures to coordinate properly that we will proceed to illustrate. First, physicists know the principle of "action and reaction": that when a force is applied to any object, the object applying the force always feels on itself an equal force in the opposite direction. This provides a simple inference. Any time there is a force on one object, you know there is another force present, and you can infer its direction and magnitude, "equal and opposite." Second, the equation $F = ma$ allows one to infer strength, magnitude, and perforce, existence of a force in circumstances where you can read out mass (m) and acceleration (a). $F = ma$ and "action and reaction" are elements in an expert's causal net. In the example below, the subject uses a somewhat corrupted special case of $F = ma$, that if there is motion (not necessarily acceleration) there is force.

Besides the generic challenges of span, integration and alignment associated with achieving a proper coordination class, the notion of p-prim brings a number of expectations and predictions about the development of the concept of force. Most particularly, we know (I claim) that intuitive ideas about physics come mainly in the form of p-prims. Furthermore, p-prims evidently allow inferences that might be part of a causal net. For example, force as a mover implies that the direction of a force can be determined from the direction of motion. The amount of a force is in some measure determinable by Ohm's p-prim: If the "effort" in question is a force, it must be greater in situations where there is a greater outcome. All in all, p-prims are excellent candidates for early elements in a learner's causal net. For reasons I will not go into here, it is a good guess that difficulties in learning the proper coordination class "force" reside almost exclusively in the causal net, rather than in readout strategies.

P-prims' theoretical properties lend detail to expectations about learning force. At the broadest level, knowing that intuitive ideas in physics are rich and numerous, we expect a high degree of context sensitivity. Achieving alignment across a wide range of circumstances may be very difficult. More particularly, learners start with a lot of p-prims that are evoked in circumstances that have nothing much to do with learned physics. Thus, there are many opportunities for context dependence that is directly attributable to the context dependence of p-prims. To the extent that we have documented particular p-prims, we can predict or explain behaviors in coordinating force by the use of particular p-prims in circumstances where we know those p-prims are likely to apply. Below, we will, indeed, attribute particular learning difficulties in particular circumstances to the use of particular p-prims.

In the following, we again turn to process data from the subject J. In general, we see a very high degree of context sensitivity that undermines a coherent (aligned) coordination class for force. More particularly, we will observe the following:

1. J does not expect to be able to directly feel a force. She does not use that expert coordinating strategy.

2. While she is consciously aware that "action and reaction" is a good way to "see" forces (that is, she knows it is a good element of a proper causal net for force), she does not always use it.
3. J exhibits commitment to the general principle that motion implies the existence of a force. However, in a situation where a particular p-prim is salient and explains motion, she abandons that principle to employ the p-prim, denying the existence of force in that context.

In the following segment, the interviewer tries to propose a problem to consider. If heavier things experience more gravitational force (which he assumes is obvious), then shouldn't heavier objects fall faster?

> I: ... there are people who say that, well, it sure seems like the heavier one should go faster because it's being pulled harder. I can feel that it's being pulled harder. So what would you say to that?
> J: Well, first I'd say, "how can you tell it's being pulled harder?"
> I: ... I guess I'd say, "well, you can feel it in your hand." You've got a heavy thing that's being pulled hard, you've got a light thing that's //
> J: Well, um. Gravity's uniform. So gravity won't pull any harder on something that's in the same place as it will on something else. ...
> I: So you're not feeling the force of gravity when you hold something?
> J: You're feeling the weight of the object.
> I: The weight of the object. So that's different from the force.
> J: Right.

J explicitly denies that one can feel gravitational force. Instead, she says you feel "weight." She coordinates gravitational force here by an abstract principle: "Gravity is the same on all objects." In this case, no p-prim is implicated in her failure to "see" force properly. Instead, we see that principles can serve as elements in the causal net, as well as p-prims.

J evidenced articulate commitment to "action and reaction." Below, the interviewer picks up on her apparent use of action and reaction to infer that a table is pushing up on a book placed on it from the fact that you know the book pushes down on the table. J implicates "action and reaction" by referring to the force of the table on the book as "equal and opposite." She elaborates with other examples.

> I: And you say you know the desk is pushing up because if it weren't, the book would just come down.
> J: Right. And it's the same as like equal and opposite forces. I mean, this chair right now is pushing up on me and the chair is pushing up on you. And the ground is pushing up on your feet. And that's something that's hard to think about. *[You are tempted to say:]* "No it's not; I don't feel it; I'm not moving anywhere." But it is *[pushing up]*.

A short time later, J explains this is something she learned that would not be evident to "the man in the street."

> J: ... now the chair is pushing on me as hard as I'm pushing down—130 pounds. This chair's pushing up on me. I think that's something that, once you've taken physics, that's totally normal. But if you said it to someone off the street, I think they'd say, "What are you talking about?"

In the following selection, J and the interviewer are discussing pushing a book across a table. The interviewer provides an opportunity for J to use "action and reaction" to see another force, that of the book pushing back on the hand, and she follows suit. However, directly following he provides the same opportunity to use "action and reaction" to conclude that, since the table (via friction) is pushing backward on the book, the book must be pushing forward on the table. J declines this opportunity.

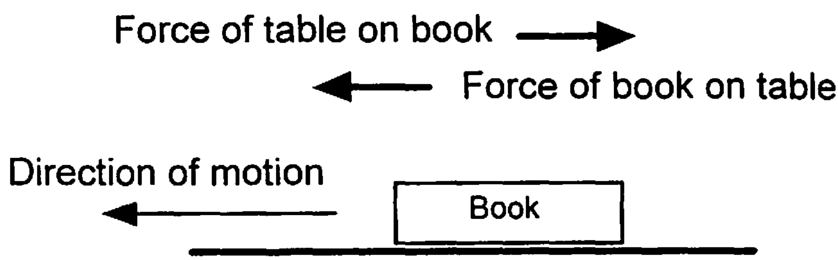

Figure 7.The table pushes backwards on the book as it slides. But J denies that the book is pushing forward on the table.

I: And now, I'm pushing on this book *[Figure 7.]*. What about the force that the book is exerting on my finger?
J: Umm. It's the same as the force you're exerting on the book.
I: Alright, what about the force of the book on the table as I'm pushing this thing along? There's a downward force that the book is exerting on the table. Is there a sideways force?
J: *[Shakes head no]*
I: No. So friction is pushing on the book that way *[against the motion]*. The book is not pushing on the table either way. Alright.

The reason for this failure to use the same principle in almost the same context (a failure of alignment due to context sensitivity in the causal net) is not evident in the available process data here. However, in other parts of the set of interviews, J shows that she thinks friction is simply a different kind of force, and, perhaps, she therefore believes it is not susceptible to the use of "action and reaction." See the description of "split concepts" in diSessa, Elby, & Hammer (in press).

Below, the interviewer tries to make the reaction force of the table salient by invoking another principle. In many contexts J exhibited articulate, reflective and deep commitment to the principle that if there is motion, there is a force. (Again, see diSessa, Elby, & Hammer, in press.) The interviewer puts a paper under the book and shows that it moves when you push the book (hence implicating a force on it). J declines the use of her own principle here. In this case, it seems clear the failure in alignment is due to the salience of a particular p-prim, one I cited earlier: contact conveys motion. If contact is creating motion, then there is no need for a force.

> I: *[continuing]* But this paper *[under the book]* moves with the book, so I have to hold it in order to keep it stopped. So, suppose I said that I have to hold it because the book is exerting a force on the paper.
>
> J: I think it's just sliding, and I think it's just bringing the paper with it. I mean, it's a really simple situation.... I would just say that it's *[the book is]* sliding against the table and bringing the paper with it.

To sum up, we observe in this process data an expectable fragmentation in context dependencies. In two cases, explicit principles that function as causal net inferences are applied in some contexts, but not others. In one case, the reason for this context dependence seems clearly the salience of a particular p-prim in a particular context. Because contact can convey motion, the principle that motion requires a force doesn't apply to this case.

4. A BRIEF REVIEW

In order to make the connection between my original critique of conceptual change research and what I have displayed concerning "conceptual ecology" more clear, I will undertake a brief review in terms of the program I set out for the chapter, following the critique.

1 .What more accountable theoretical terms might look like.
Both p-prims and coordination classes are much more specific than dictionary definitions and typical invocations of, for example, "concepts" or "theories" in conceptual change research. We have discussed the nature of elements, their origins, what happens during development, the function of the knowledge types, typical patterns of use, and the level and kinds of systematicity between elements and across contexts. (More on the latter appears in the next section.) Table 1 reviews some of the main points. Because of the level of specificity, it is not even obvious any exemplars of these knowledge types exist! Empirical investigations are necessary to verify that any hypothetical p-prim or coordination class has the necessary properties.
2. What different exemplars of such terms are plausible and plausibly play an important part in conceptual change.
Both p-prims and coordination classes are advanced here as plausible knowledge types that play distinct roles in conceptual change and account for different conceptual change phenomena. P-prims, for example, account for intuitive predictions and judgments of plausibility. Coordination classes provide a specific model of a type of full-blown concept, which entails a lot about the difficult and easy parts of conceptual change.

Table 1. Properties of P-prims and Coordination Classes

P-prims	***Coordination Classes***
Size and Structure	
• Many, small • Relatively independent (used individually) • Inarticulate	• Few, large • Complexly articulated (two major subsystems: readout strategies and causal net); integration *is* the point • May enfold articulate components, like "$F = ma$" or "you have to have a force to sustain motion"
Origins	
• Plausibly developed as a single abstraction	• Emerge slowly as a coordination of many elements
Characteristic Function	
• Sense of naturalness	• Determining a certain class of information across many contexts
Characteristics of Operation	
• Recognition, typically one by one • Sometimes quite context specific • Data fluidity	• Coordinated action across contexts entailing: (1) appropriate span, (2) integration, (3) alignment
Role in Naïve Thought	
• Simple instrumentality ("push harder") • Surprise (violation of p-prims) may evoke attention and learning	• Possibly accounts for major developmental accomplishments, such as determining time duration or "object permanence"
Role in Expert Thought	
• Small, contextually bound parts of concepts or theories (e.g., Ohm's p-prim is an intuitive gloss for Ohm's law) • Might encode contextually specific inferences in causal net	• Defines a model of the system constituting expert concepts • Supplies inferences from observations to theoretical entities

3. How different from each other knowledge types might be, and how different types can account for different aspects of reasoning and conceptual change.

Following our discussion, summarized in Table 1, p-prims and coordination classes contrast in many ways with each other. Theoretically, there is no way to mistake one for the other. Empirically, of course, any instance of a student's reasoning may require careful scrutiny to determine whether a p-prim or part of a coordination class is involved. More broadly, I fully expect that we will need many other knowledge types to fully explain the transition from naïve student to conceptually-competent physicist.

4. What sort of details of subjects' behavior can be accounted for with sharpened theoretical terms. For example, what are typical easy and difficult accomplishments along the path to conceptual competence?

Starting from the last question, establishing appropriate span and alignment is difficult, and we should expect continued examples of failures of this sort in the trajectory toward competence. Not only did we document failures of these types in a subject's process data, we even showed how a specific failure of a specific attribute of a mature coordination class (alignment) was due to the invocation of a particular previously-documented p-prim. That p-prim (contact conveys motion) "explained" a situation (a paper moving under a pushed book) and thus aborted the use of an articulate principle (for this subject), that motion requires the existence of a force.

P-prims explain many wrong expectations—and many correct ones—that students have about how the world works. They also explain phenomena like data fluidity and the general richness and detail found in intuitive thought. P-prims explain the emergence of particular macro-constructions (e.g., McCloskey's "impetus theory") as a confluence of a number of p-prims in a relatively stable configuration for a particular class of contexts.

How do p-prims and coordination classes relate to "theories" or "mental models," and so on? In general this is not a particularly good game to play precisely because of the elements of our earlier critique; most advocates of these terms say precious little about what they actually entail. Nonetheless, I can make some comments that may be at least heuristically useful. P-prims are obviously sub-conceptual, sub-theoretical, or sub-model-like. They are too small to constitute any of these macro-conceptual structures, and the most plausible developmental path between naïve p-prims and any of these structures is "incorporation as a limited part" of an emergent complex knowledge system.

Coordination classes, we have argued, are an appropriate refinement of the idea of "concept." Theories are likely to be even larger conceptual structures, encompassing several related concepts (coordination classes). For example, force, mass, and acceleration may each constitute coordination classes, and "Newton's theory" might be abbreviated in a particular relation among these coordination classes, "$F = ma$." Such a relation provides opportunities for coordination across coordination classes—for example, if you can "see" mass and acceleration, you can also "see" force. Developmentally, therefore, there are likely to be important mutual influences in coordination classes that participate in the same theory. For my own part, I would require "theories" to entail an explicit, articulable component (like $F = ma$), which requirement I did not introduce into the definition of coordination classes, (although I did not rule it out). I believe language introduces important properties that may not hold in dominantly inarticulate knowledge systems.

I expect even less consensus on what a mental model is. To my mind, mental models should (1) involve a strong, well-developed "substrate" knowledge system, such as spatial reasoning, (2) allow explicit hypothetical reasoning, and (3) involve only a small, well-defined class of causal inferences (diSessa, 1996). It is possible, for example, that a limited set of p-prims (e.g., perhaps "contact conveys motion") together with reasoning about spatial configurations (of, for example, gears) could constitute a mental model. However, other principles might define the causality in a mental model, aside from p-prims. Furthermore, the involvement of p-prims wouldn't necessarily entail the ability to support explicit hypothetical reasoning. In fact, the involvement of p-prims might even make this less likely. I see no obvious generalities about the involvement of mental models in coordination classes, or the reverse, although I also see no principle that would rule those relations out.

I know of no even remotely specific process definition of ontology in the literature. However, I believe the idea of coordination class is, in fact, a technically sufficient refinement of the general idea of ontology. This may be a contentious claim. I don't propose to defend it here, aside from noting that a coordination class may define the ability to know about a particular class of entity in the world—for example, the ability to "see" forces. If it turns out to be sensible to consider coordination classes ontological, most invocations of ontology to explain difficulties in conceptual change pale. The reason is that invocations of ontology merely *assume* that ontologies are difficult to learn, and also never explain exactly how it is possible to learn one. In contrast, the articulated definition of coordination class shows what is entailed in developing one and provides a list of focal difficulties. In certain cases we can name both the general class and the precise circumstances of conceptual difficulty, such as such when a context-specific p-prim "short-circuits" alignment.

5. COHERENCE AND CONSISTENCY

Issues of coherence and consistency in reasoning ("systematicity," speaking inclusively) become salient as one approaches a complex systems view of conceptual change.[11] For example, if a concept is, in fact, a complex system, there is likely no point in the learning trajectory where we can unequivocally decide a person "has" the concept. It may always be a matter of degree and context. With large numbers of elements and a heightened accountability concerning contextuality, it is easy to parody a complex systems view as assuming total fragmentation and inconsistency in the conceptual behavior of naïve or novice students. This is far from a sensible view. Rather, conventional assumptions about conceptual change, which I am critiquing in this chapter, are as vague and presumptuous about consistency as they are concerning the mental entities involved in conceptual change. A complex systems view may look necessarily fragmented, but that is only true because it recognizes existing complexity and takes on more accountability for details.

---

[11] The issue of coherence, systematicity, or lack thereof, is an important, unsettled issue in studies of cognition. See, for example, discussion and empirical studies in Rogers, Rutherford, and Bibby (1992).

5.1. *Entailed Elements of Systematicity*

Let me begin by sketching elements of consistency that are necessarily entailed by the view of conceptual change described in this chapter. First, although I presume there are a large number of p-prims, the existence of a p-prim actually constitutes, in itself, an important consistency in thinking. A p-prim is precisely a regularity in responding to situations in the world. Without a degree of consistency, identifying particular mental entities, particular p-prims, would make no sense. Indeed, this consistency is an explicit methodological commitment in my own work to identify p-prims. It is embodied in a fundamental idea I call "the principle of invariance." Roughly, if one has gotten the description of a p-prim correct, the p-prim should be invoked in every situation in which the description applies. Failure in invocation is a failure to describe the p-prim or its contextual specifics adequately. See the methodology section of diSessa (1993).

The mere fact that some p-prims are much more important and generally applied than others lends another kind of systematicity to intuitive thought. Sketching an important, widely applicable p-prim catches an important tendency in intuitive thought.

Another necessary element of consistency involves coordination classes. If coordination class adequately describes any scientific concept, then, to the extent that anyone learns that concept, they have achieved (a remarkable) consistency and coherence in their knowledge state—surpassing difficulties of span, integration and alignment. Obviously, I can't sensibly take the position that no one ever learns any scientific concepts. The fact that I list and demonstrate failures to achieve this consistency in several modes (e.g., failures in alignment) merely shows steps along the way to consistency and also the degree of consistency actually achieved when one learns a scientific concept.

5.2. *Allowed Elements of Systematicity*

Beyond necessary systematicity, the view of conceptual change presented in this chapter *allows* other sorts of systematicity. In particular, although space here did not allow it, I have charted in other places several kinds of systematicity in intuitive physics (see diSessa, 1993, for details). First, although elements are not tightly integrated, there are a number of loose relations. I described J's description of a toss as a (relatively stable) composition of a number of p-prims. In addition, p-prim compositions may generate new p-prims as "phenomenological syllogisms." If one knows that heavier things move slower and that bigger things tend to be heavier, it is likely one will also expect bigger things to move slower. Along similar lines, I have described a kind of systematicity that appears because of common attributes involved in many p-prims. For example, agency is important in many situations, and conflict is another fundamental attribute involved in many p-prims. Finally, a family of p-prims may engender a common abstraction—a kind of meta-p-prim—which means that a substantial range of situation may be covered, in some degree, by the same (meta) p-prim.

I mentioned previously in this chapter that coordination classes may or may not exist in naïve, intuitive thought. My guess is that some do. Piagetian object permanence may be part of a coordination class involving reading out information about physical location from the world. This conjecture is elaborated in diSessa & Sherin (1998). If naïve coordination classes exist, then all of the systematicity implied in my description of coordination classes has been achieved in physics-naïve individuals, with respect to certain ideas. The principal reason to keep the existence of naïve coordination classes hypothetical is that, given the degree of specificity in the concept of coordination class, detailed empirical support would be necessary to establish their existence. I find this far preferable to *assuming* characteristics of naïve thought in describing its contents as naïve concepts (or theories, etc.), to the extent that one is making any claim at all in making such statements.

5.3. *An Issue of Epistemology*

An illusion of fragmentation occurs because we tend to view intuitive ideas normatively (as I have, mainly, here). That is, we measure intuitive ideas against a particular standard, say, the Newtonian concept of force. This essentially guarantees a fragmented view, as we must identify all the pieces and the ways they must be coordinated in order to achieve the Newtonian concept. Scientific concepts, however, differentiate as well as integrate contexts. Galileo made a huge advance in science by concerning himself only with spatial change (motion), in contrast with Aristotle's physics, which attempted a uniform treatment of all change, including, for example, biological growth and decay.

Naïve p-prims of wide scope clearly exist. I've argued, for example, that many p-prims apply to both physical and psychological/sociological situations (diSessa, 2000). As such, they clearly extend beyond the boundaries of Newtonian concepts, which are, thus, comparatively fragmented. Making an absolute comparison between the scope of intuitive ideas and scientific ones is very difficult, even disregarding the fact that there are many kinds of ideas on both sides of the transition to scientific expertise: Naïve thinking and expert thinking draw contextual boundaries differently. Hence, we are guaranteed to find, in some instances, naïve ideas are fragmented compared to scientific ones, and, in other instances, scientific ideas are fragmented compared to intuitive ones.[12]

---

[12] *There are difficult analytical issues I am suppressing here. How does one develop a metric to be able to measure objectively the "range of circumstances" in which an idea applies? I believe it is likely that a metric can be established only relatively, that is, noting whether a concept crosses boundaries established with respect to another way of thinking. Once again, we may wind up being able to say only that one way of thinking distinguishes contexts that are treated uniformly in another way of thinking, and viceversa.*

5.4. *Intention in Systematicity*

Other issues concerning consistency and coherence are simply not touched, one way or another, in the account of naïve and expert thinking given here. For example, how much do students *strive* for coherence? These are meta-conceptual issues—issues of students' ways of conceiving of their own knowledge, issues of strategies for dealing with it, and so on. I accept the common sense that students (sometimes) strive for consistency. However, I note that this is a highly knowledge intensive activity. Clearly the world is diverse enough that insisting one must think in the same way about all situations is a foolhardy and doomed-to-failure approach. So, one has to make judgments about when it is necessary or plausible to think in the same way. More generally, it is foolish to seek more consistency than the world allows. The physical world, clearly, is quite diverse. If one takes a scientific standard, one needs to learn linear Newtonian mechanics, rotational dynamics, fluid mechanics, electricity and magnetism, optics, possibly quantum mechanics, and so on, in order to understand even everyday phenomena (bouncing balls, wobble in a spinning football, static electricity, a magnifying glass, the strength of different materials). It won't do to attempt consistency by thinking about all the implied situations in the same way.

Deciding when one can and should think in the same way involves subtle judgments. For example, even within Newtonian mechanics, one can describe a toss from the viewpoint of forces, or in terms of energy. The two viewpoints produce different sketches of the same situation. Clearly, one shouldn't conclude that one is being inconsistent in making these two different kinds of descriptions of the same situation.

In some situations the very possibility of striving for consistency is problematic. One may in principle want to be more consistent, but how can one actually make progress? How could one, for example, consider the issue of consistency with respect to p-prims? How could one consider whether dynamic balance should be applied in some circumstances where a different balancing p-prim seems to apply? Without some way of describing the knowledge that is applying—for example, a verbal expression—one can't even phrase the question of a more integrated view.

The correct way of resolving the issue of intentionality in consistency is, I believe, another version of the program described here as applied here to content-level knowledge, except applied to meta-conceptualization and epistemology. We need to describe the knowledge that people have about knowledge relevant to consistency and their strategies for dealing with it. Issues of "degree" of search for consistency are very likely at least as complicated as those concerning the relative fragmentation of content-level knowledge.[13]

---

[13] *This precise task is taken on, in a rough and ready form, in diSessa, Elby, and Hammer (in press). We show that the subject whose thinking is displayed here, J, has a very interesting meta-conceptual profile that influences how much fragmentation or consistency she feels is possible. In broad strokes, she feels that the scope of scientific ideas, like* F = ma, *is generally smaller than is true. She also simply does not judge that she is being inconsistent if she says there is both one and two forces acting in a situation.*

Finally, there is one more open dimension of possible systematicity that I have simply had no space to touch upon in this chapter. As I mentioned, p-prims and coordination classes are only two exemplars of knowledge types that I feel are necessary to distinguish in thinking about conceptual change. Were we to delve into others, the issue of consistency and coherence becomes yet more complicated (and my position less specified by what has been said in this chapter).

6. SUMMARY

In this chapter, I have argued that the mainstream in the study of conceptual change is flawed in several respects. Overall, I feel the theoretical accountability of the research community is far from adequate. Instead, researchers use commonsense or dictionary terms, or terms imported from other disciplines, that are much too vague or implicitly defined to allow good progress. Adequate theoretical terms may be developed and validated in many ways, but, first, must come a commitment to clarity and cogency. "Concept," "belief," "theory," and "ontology" have all been used to describe the elements of mind involved in conceptual change. Yet none are elaborated to the point that we know exactly what we are talking about, and to the point that we can even in principle empirically determine whether one or another of these is more adequate than others (or whether they are all necessary) to describe difficulties and accomplishments in conceptual change.

Theoretical vagueness and imprecision trickle down and reinforce a tendency to use data impressionistically or merely statistically, without putting strong hypotheses about what is involved in conceptual change to strong tests. I advocate moving from before/after studies, and studies that use only constructs (like coding categories) distanced from cogent theoretical terms, to the use of process data to test and illustrate theoretical commitments about concepts, or other theoretical elements of mind, *in use* and *in change.*

The greatest casualty of weak theoretical and empirical accountability is a widespread and dramatic over-simplification of the complexity, diversity and nuance in conceptual change data. This complexity is less evident in conceptual change research than in common sense about types of concepts and in expert teachers' detailed reactions and judgments concerning students' partial states of development in the process of conceptual change. Clinical data is rich in possibilities for more detailed accountability, if it is used for that purpose.

I have argued that, in order to recognize and deal with the diversity and complexity involved in conceptual change, we need to move simultaneously in several directions: With respect to types of knowledge, we need to move toward multiplicity. With respect to "size" and number of elements of mind, we need to move toward smaller grain size and greater numbers. Concomitantly with this move to smaller grain size, we need to deal much more effectively with evident contextual dependence in the way students think. With respect to details in describing both change, and even what constitutes a concept, we need to move toward a systems view that describes scientific concepts as complex, finely configured systems involving named parts and relations.

To make the case that such progress is possible (and to point out directions I feel will be most fruitful), I have tried to illustrate what more theoretically accountable replacements for "concepts," or "theories" might look like. I defined and illustrated the idea of p-prims, which are numerous, small, recognition-driven elements that define significant properties of naïve physics knowledge, and which become involved in learning scientific concepts. P-prims exhibit a detailed contextuality and sometimes cause unstable, data-driven conceptualization. I also defined and illustrated coordination classes, which are complex systems that consist of many parts (of types I delineated), in particular ways (that I named and classified). Coordination classes are a proposed model for a particular class of scientific concepts.

Finally, I tried to show how an improved theoretical view may be held accountable to details in process data. I showed hypothesized behaviors of p-prim-driven cognition, such as data-driven instability. I argued that we can see in process data that intuitive "theories" or other large-scale constructs may actually be composed of p-prims. I also showed classes of learning difficulties predicted by coordination class theory in examples of student thinking. I showed a great deal of context dependence in the ways that students read out and infer information about forces. I showed that p-prims (and other kinds of ideas, like articulate principles) play specific roles in mistakes students make in learning to "see" forces.

ACKNOWLEDGMENTS

Much of the work reported in this chapter was sponsored at various times by the Spencer Foundation. Its help is gratefully acknowledged. The chapter also benefited from comments by David Kaufman, Joe Wagner, and the editors of this volume.

REFERENCES

Clement, J. (1983). A conceptual model discussed by Galileo and used intuitively by physics students. In D. Gentner & A. Stevens (Eds.), *Mental models* (pp. 325-340). Mahwah, NJ: Lawrence Erlbaum Associates.

diSessa, A. A. (1983). Phenomenology and the evolution of intuition. In D. Gentner and A. Stevens (Eds.), *Mental models* (pp. 15-33). Mahwah, NJ: Lawrence Erlbaum.

diSessa, A. A. (1988). Knowledge in pieces. In G. Forman and P. Pufall (Eds.), *Constructivism in the computer age* (pp. 49-70). Mahwah, NJ: Lawrence Erlbaum Associates.

diSessa, A. A. (1993). Toward an epistemology of physics. *Cognition and Instruction, 10*, 105-225.

diSessa, A. A. (1994). Speculations on the foundations of knowledge and intelligence. In D. Tirosh (Ed.), *Implicit and explicit knowledge: An educational approach* (pp. 1-54). Norwood, NJ: Ablex.

diSessa, A. A. (1996). What do "just plain folk" know about physics? In D. R. Olson and 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. (2000). Does the mind know the difference between the physical and social worlds? In L. Nucci, G. Saxe, and E. Turiel (Eds.), *Culture, development and knowledge*. Mahwah, NJ: Lawrence Erlbaum Associates.

diSessa, A. A. (to be submitted). An interactional analysis of clinical interviewing.

diSessa, Elby, & Hammer (in press). J's epistemological stance and strategies. In G. Sinatra and P. Pintrich (Eds.), *Intentional conceptual change*. Mahwah, NJ: Lawrence Erlbaum Associates.

diSessa, A. A., & Sherin, B. (1998). What changes in conceptual change? *International Journal of Science Education, 20*, 1155-1191.

Gopnik, A., & Meltzoff, A. N., (1996). *Words, thoughts, and theories: Learning, development, and conceptual change*. Cambridge, MA: MIT Press.

Harris, P. L. (1994). Thinking by children and scientists: False analogies and neglected similarities. In L. Hirschfeld & S. Gelman (Eds.), *Mapping the mind: Domain specificity in cognition and culture* (pp. 294-315). Cambridge, UK; NY: Cambridge University Press.

Hofer, B. K., & Pintrich, P. R. (in press). *Personal epistemology: The psychology of beliefs about knowledge and knowing*. Mahwah, NJ: Lawrence Erlbaum Associates.

Kuhn, T. (1970). *The structure of scientific revolutions*. Second edition. Chicago: University of Chicago Press.

McCloskey, M. (1983). Naïve theories of motion. In D. Gentner & A. Stevens (Eds.), *Mental models* (pp. 299-324). Mahwah, NJ: Lawrence Erlbaum Associates.

Piaget, J. (1969). *The child's conception of time*. (A. J. Pomerans, Trans.). New York: Ballantine Books.

Rogers, Y., Rutherford, A., & Bibby, P. A. (1992). *Models in the mind: Theory, perspective and applications*. NY: Academic Press.

Rosch, E. (1994). Categorization. In V. S. Ramachandran (Ed.), *Encyclopedia of human behavior* (Vol. 1, pp. 513-523). San Diego: Academic Press.

Schoenfeld, A., Smith, J., & Arcavi, A. (1993). Learning: The microgenetic analysis of one student's understanding of a complex subject matter domain. In R. Glaser (Ed.), *Advances in instructional psychology* (Vol. 4, pp. 55-175). Mahwah, NJ: Lawrence Erlbaum Associates.

Strike, K. A., & Posner, G. J. (1992). A revisionist theory of conceptual change. In R. A. Duschl, & R. J. Hamilton (Eds.), *Philosophy of science, cognitive psychology, and educational theory and practice*. Albany, NY: SUNY Press.

Thagard, P. (1992). *Conceptual revolutions*. Princeton, NJ: Princeton University Press.

Vosniadou, S. & Brewer, W. F. (1992). Mental models of the earth: A study of conceptual change in childhood. *Cognitive Psychology, 24*, 535-585.

Vosniadou, S., & Ortony, A. (1989). *Similarity and analogical reasoning*. Cambridge, UK; NY: Cambridge University Press.