INT. J. SCI. EDUC., 1998, VOL. 20, NO. 10, 1155–1191

What changes in conceptual change?

Andrea A. diSessa, Graduate School of Education, University of California–Berkeley and Bruce L. Sherin, School of Education and Social Policy, Northwestern University, Evanston, Illinois, USA

This paper has two aims. First, it reviews literature about conceptual change and about the study of concepts more broadly. The principal claim is that much prior work has suffered from inexplicitness and imprecision in terms of what constitutes a concept. Second, we introduce a theory of one particular type of concept. A *coordination class* is a systematic collection of strategies for reading a certain type of information out from the world. We identify both structural components and performance properties of coordination classes. Using this theory, we analyse protocol data from a student with respect to the difficult concept of force in Newtonian mechanics.

Introduction

The intent of this paper is to add to theory about the fundamentals of conceptual change. We want to show where imprecision in prior descriptions has marred progress and we want to take some steps towards a more sound account of the nature of conceptual change. Our basic contention is that it is essential for researchers to answer the question 'What changes in conceptual change?' with substantially greater precision than has heretofore been attempted. We will, of course, argue for this contention as the paper proceeds, but the central observation is that how one understands and describes the *processes* of conceptual change must flow from an account of the entities that are changing. For example, suppose we begin with the image of a network of nodes, each of which corresponds to a concept, with the nodes connected by links of multiple types. This is called a semantic net and its purpose is to model meaning by situating a concept in a larger field. What are the relevant types of change? Some possibilities are:

- We can add or delete nodes (concepts).
- We can add or delete links between nodes, or change the type of an existing link.
- If the network has a characteristic global organization (perhaps a hierarchical structure), we can perform more dramatic alterations to this structure.

These are just a few of the possibilities. Thagard (1992), for example, begins with a network of this type and lists nine degrees of change.

Alternatively, suppose we start with the image of a concept as a set consisting of elements which are instances of the concept. Given such a model, changes to the concept correspond to changes in the constituting elements, which means adding or deleting instances. We might also identify some special types of changes to these sets. For example, the set could be broken into two parts (differentiation), or we could form the union of two sets (coalescence). The point, once again, is that the possible types of change are largely dictated by the nature of what is changing.

As this paper proceeds, we will be working towards our own answers to the 'what changes' question. We are, of course, not the first people to propose to answer this question. However, we will argue that many proposed answers only end up begging the question. Furthermore, we believe that, even where answers are given, greater precision is needed.

In what follows, we begin by presenting what we call the 'standard model' of conceptual change. This model is, we will argue, typical of the way that conceptual change is understood in educational research. Within this model, the answer that is typically given to the 'what changes' question is 'concepts'. This observation motivates us to turn our attention to the notion of concept. We will examine how this notion has been treated in the literature dealing with conceptual change, and we will try to pinpoint some difficulties. Ultimately, we will argue that it is necessary to replace the notion of concept with a variety of more carefully defined theoretical constructs.

In the second major portion of this paper we will present our own candidate for one of these new theoretical constructs, what we call a *coordination class*. After presenting a general introduction to coordination classes, we will focus our attention on physical or mathematical quantities as examples of coordination classes. We will then illustrate these ideas with excerpts from some clinical interviews, providing an analysis focused around the concept of force.

The standard model of conceptual change

The purpose of introducing the notion of conceptual change, especially in educationally-oriented work, is to separate learning into at least two types: we want to have a vocabulary that distinguishes deep varieties of change from more routine sorts of learning, such as fact memorization and the gradual accretion of beliefs. Any account of conceptual change should draw a line between these two types of learning. The standard model does this; it is an image that is invoked by a wide array of researchers, both implicitly and explicitly, to describe what counts as 'conceptual change', which represents deep, more difficult to accomplish learning, in opposition to more mundane sorts of learning. For us, the standard model serves as a jumping-off point, but we will ultimately react against it.

We begin with a version of the standard model that appears in the work of Susan Carey (1988). In order to distinguish more and less fundamental varieties of change, Carey contrasts the accounts that appear in two bodies of literature. First, she begins with a discussion of literature concerning the novice/expert shift, especially as it appears in accounts of physics learning. She describes how some novice misconceptions may be replaced by more expert beliefs. For example, the misconception 'there is no motion without a force' should be replaced with the expert belief 'there is no acceleration without a force'.

What sort of change is this? According to Carey, one possibility is that we should understand this as a change in the *relations* among the concepts that appear in these beliefs. The point is that concepts such as force, motion, and acceleration may remain unchanged in the move from novice to expert; they are simply related in new ways. Students develop new beliefs, and follow those which are a type of relation. But they may be beliefs about the same concepts.

For contrast, Carey describes another body of literature that emphasizes the possible existence of more dramatic varieties of change. Here, Carey turns to the history of science, and especially to the work of Thomas Kuhn. Carey states that, although there has been much argument and counter-argument, 'a strong view of restructuring has survived, one that allows for true conceptual change among core concepts of successive theories'. In discussing the historical development of theories, historians have argued that there may be changes 'even in the individual concepts in the center of each system' and 'because of these differences, some terms in one theory may not even be translatable into the terms of the other' (Carey 1988: 6).

Carey thus provides a contrast between more and less dramatic varieties of change. If the less dramatic variety is accurate for describing what happens when a student learns physics, then it may be good enough to understand students as simply acquiring new beliefs about (that is, new relations among) concepts, such as force and mass, that they already possess. However, if the stronger version is correct, it may not be accurate to think of students as even talking about the same things when they say 'force' or 'mass'; the very concepts may differ in some fundamental manner. The concepts may be incommensurable, to use the term that is used by historians of science.

Carey calls these two varieties of change 'weak' and 'strong' knowledge restructuring. In weak restructuring 'new relations among concepts are represented, and new schemata come into being which allow the solution of new problems and which change the solutions to old problems'. In contrast, strong restructuring involves 'change at the level of individual concepts at the core of the successive systems' (p. 7). The latter, stronger variety of restructuring we call conceptual change.

This is the standard model of conceptual change. We can see that it attempts to do the job of distinguishing between conceptual change and more mundane varieties of learning. Other researchers use different language than Carey, but present very similar images. For example, Gentner *et al.* (1997: 31–2), writing in a recent special issue on conceptual change, distinguish three grades of change: belief revision, theory change, and conceptual change.

> Belief revision is a change in facts. Theory change is a change in the global knowledge structure. Conceptual change, in some sense the most drastic, is a change in the fundamental concepts that compose the belief structure. Conceptual change thus requires at least locally nonalignable or incommensurable beliefs.

Although Gentner *et al.* (1997) distinguish two degrees of change in the relations among concepts, it is clear that their notion of what counts as conceptual change is very similar to that of Carey.

It is worth emphasizing that the standard model originated in literature from the history of science. We saw that Carey drew on this literature to develop her own distinction. And, although not all researchers subscribe to many of the specifics – such as whether it is ever accurate to speak of concepts as being 'incommensurable' – terminology from the history of science has been central to discussions of conceptual change. For example, in an influential article that has been described as 'seminal' (Magnusson *et al.* 1997), Posner *et al.* (1982: 212) develop their version of the standard model precisely by making analogies to the history of science. In particular, they make analogy to the Kuhnian notions of 'normal science' and 'scientific revolution':

> There are two distinguishing phases of conceptual change in science. Usually scientific work is done against the background of central commitments which organize research . . . The second phase of conceptual change occurs when these central commitments require modification. Here the scientist is faced with a challenge to his basic assumptions. If inquiry is to proceed, the scientist must acquire new concepts and a new way of seeing the world. Kuhn terms this kind of conceptual change a 'scientific revolution'

With Kuhn's version laid out, Posner and colleagues (1982: 212) make their analogy to individual learning:

> We believe there are analogous patterns of conceptual change in learning. Sometimes students use existing concepts to deal with new phenomena . . . Often, however, the student's current concepts are inadequate to allow him to grasp some new phenomenon successfully. Then the student must replace or reorganize his central concepts.

Not only are ideas borrowed from the history of science, much of the conceptual change literature uses historical examples for illustration. For example, the article by Gentner *et al.* (1997) cited above, looks in detail at the work of Kepler. Similarly, Wiser (1988) discusses the historical development of the concepts of heat and temperature. Thagard (1992) uses a number of historical examples, including the replacement of the phlogiston theory by the oxygen theory.

Not all conceptual change researchers would subscribe to the standard model, at least not in the particular way that we have presented it. Furthermore, even those that do subscribe to a version of this model might reasonably complain that what we have presented is overly simplified. There are complications and theoretical positions that relate to the standard model in ways that are not trivial to work out. In addition, most writers who describe conceptual change in this manner also count other types of change as conceptual change. To cite what is perhaps an extreme example, Thagard (1992) counts six of his nine degrees of change, mentioned above, as conceptual change.

Nonetheless, the above exposition of the standard model is sufficient to establish the jumping-off point that we need for the remainder of this paper. What we have tried to do, thus far, is document that conceptual change is very often understood as involving changes in 'the very concepts' at the 'core' of a conceptual system, the very 'terms' in which the world is understood. When this foundation of terms changes, everyone seems to agree that this is difficult, and we call it conceptual change.

Thus, we must turn our attention more seriously to *concepts* and how they change. In asking what changes in conceptual change, we will now answer 'concepts'. But, since we have yet to say much about what a concept is, this only goes so far towards answering our question. Indeed, one of the most important difficulties of the standard model, in most of the relevant presentations in the literature, is its failure to unpack what 'the very concepts' are in sufficiently rigorous terms. Without a clear notion of what a concept is, the standard model really only begs the question of what counts as conceptual change.

The concept of 'concept' and research on conceptual change: the problems identified

In this section we begin our focus on concepts. We work on this first by taking a look at how the notion of concept has been treated in the literature on conceptual change, both implicitly and explicitly. For a first example, we return to the article by Carey (1988), discussed above.

After laying out the notions of strong and weak restructuring, Carey applies these notions to the restructuring of biological knowledge that occurs between the ages of 4 and 10. In doing so, she makes a strong case that there are dramatic changes in biological knowledge during these years. Among her more important contentions is that the acquisition of detailed biological knowledge is the key to these broad changes.

Let us look briefly at one example from the paper to give a flavour of Carey's argument and to see the role that the notion of concept plays within that argument. Some of the key developments addressed by Carey relate to *childhood animism*. A well-known and well-reproduced result is that young children tend to describe many things as 'alive' – such as the moon and cars – that would not be judged to be alive by adults. For Carey, these observations are relevant to human understanding of the concept *living thing*, and she describes an experiment that is designed to explore whether and in what way young children possess this concept.

Six-year-old children were shown pictures of dogs and flowers and told 'here are two things in the world that have golgi'. Carey reasons that if 6-year-olds possess the concept *living things*, then they should be able to induce that other livings things also possess golgi. Since they do not possess this concept, they will sometimes say that inanimate objects have golgi. In Carey's words 'By hypothesis, the 6-year-old does not have the concept *living thing* available as an inductive base, and so should sometimes include inanimate objects as golgi-havers' (p. 22). This hypothesis was validated by the experiment.

Carey ultimately addresses the question of whether the changes in biological knowledge that occur between the ages of 4 and 10, such as those described above, should count as strong restructuring. In the present paper, we are especially interested in whether Carey concludes that there is a fundamental change at the level of individual concepts. Near the end of her paper, Carey concludes (speculatively) that there is a fundamental change of this particular sort:

> Are important biological concepts such as *dog, animal, alive, death, growth, eating*, and so on different in the two systems in the same ways that the concept *motion* differs in successive mechanics? Attribution of strong restructuring hangs on the answer to this question, for it is change at the level of individual concepts, more than anything else, that distinguishes the two types of restructuring. I believe it is misleading to describe the 10-year-old merely as having represented different relations among the same concepts... Learning the role of eating in maintaining the function of the body is part of coming to distinguish dead from animate. (p. 23)

We are using Carey's work here not because we believe that it is especially problematic. On the contrary, we believe it is an outstanding example in the field. Carey makes her points clearly and, more importantly, she has assembled substantial evidence of deep learning within a domain. Nonetheless, we will now point to what we feel are important difficulties in this work, difficulties that are common across work in conceptual change. These problems, we believe, need to be addressed before we can profitably continue to make claims concerning conceptual change.

Furthermore, we should emphasize that we are using Carey (1988) as an extended reference point for purposes of specificity and concreteness. We do not mean to imply that it summarizes her or anyone else's point of view, nor that Carey's work has not evolved. Nonetheless, we do not take it that more recent work changes the basic issues about which we speak. For example, in Carey and Spelke (1994) concepts – indeed, core concepts – and conceptual change are still central ideas. Although the article conjectures about mechanisms, these are not connected in any detailed way to a model of concepts or how pieces of that model are changed by the proposed mechanisms.

The items on the following list of difficulties are not independent. Rather, they constitute approaches to seeing a core underlying problem. We do not present the list in order to provide a careful analysis of the problem, but, instead, simply to help the reader to begin to see the dramatic nature of the difficulties involved.

1. What should we count as a concept?

First, we ask the question: How does Carey know a concept when she sees one? One of her research questions is to ask whether children of a certain age possess a particular concept such as *living thing*. But how does she know that anyone, even adults, can accurately be described as possessing this concept?

The central problem here is that just because we, as researchers, can name a particular cognitive task (such as the task of determining whether something is 'alive' or not), this does not tell us anything directly or obviously about how that task is accomplished. Even if we can name the task with a word of phrase ('living thing'), this does not mean that any unified mental structure, such as a concept, is responsible for an individual's ability to perform the task. To cite a dramatic possibility, we might decide, on various grounds, that it is more accurate to describe our apparent understanding of a word such as 'dog' as being due to some completely different, hidden set of concepts, ones for which we have no words or even any explicit awareness. Some researchers have believed that there is a language of the mind, mentalese, but that the vocabulary of that language has nothing much to do with common words (Fodor 1975). In mentalese, 'dog' might correspond to 'nog fruk chuk'. So, learning 'dog' might mean learning any subset of the (let us say) concepts 'nog', 'fruk', or 'chuk'. Or, in mentalese, 'nog fruk chuk' might implicate a theory, so learning it might not even be conceptual. How do we know? We do not intend to argue for mentalese, but much research assumes and seems to validate that tacit knowledge, which is not evident in language, does key work in learning (Tirosh 1994).

The temptation to assume that words or phrases correspond to concepts has other difficulties. Notably, there are tens of thousands of words. If every such is a concept, then the concept of concept almost certainly cannot do the work we need to do. That is, conceptual change (learning a new word) cannot separate difficult, deep learning from easy learning. Consider that children learn many words per week during their peak language-learning period, not to mention that some words are almost certainly learnable simply by looking them up in the dictionary. We might recover our research goals by determining empirically which word-concepts are difficult and which are easy. But this begs the question of why, and provides no insight on how to improve learning.

To be fair to Carey, she does do some of the triangulating that would be necessary to establish her candidate concepts as solid elements of a descriptive theory. We cannot argue the specific cases here and, indeed, Carey is, in our estimation, among the best in this regard. Instead, we only want to note that there is extremely difficult work to be done. A significant amount of argument and evidence is needed before one should venture to assert that any individual or group possesses a particular concept, such as *force* or *mammal*, much less that acquiring the concept explains difficulties in learning.

2. Lack of attention to manifest variety

In the above quote from Carey, notice the long list of things that are called concepts: *dog, animal, alive, death, growth, eating*. To this list we can add *motion, force, velocity* and *acceleration*, which are referred to as concepts earlier in her paper. In calling all of the above things concepts, Carey is putting together quite a variety of things. This is not necessarily a problem; indeed, science is partly about looking for unifying categories. But it is entirely possible that knowing *force* is very different from knowing *dog*. The reader is asked to ponder (as we will more carefully later in this paper) the difference between recognizing dogs and forces in the world, and the very different types of reasoning that one might do involving each. Furthermore, how should we handle possibly even more exotic concepts, such as the concept of *number*? Numbers are very different than dogs.

In sum, it is very easy to imagine that we will need more than one theoretical construct to deal with this variety. If there are differences, then this may well matter for how we understand conceptual change in each case.

3. Theoretical imprecision

Even if it turns out that it makes sense to call all of *force, dog* and *number* concepts – even if we can usefully deal with all of these with a single theoretical construct – at rock bottom, we are still left with the problem of saying what a concept is. If we accept the standard model, we can say that concepts are at the 'core' of our understanding, that they are the entities about which we have beliefs, but they are nonetheless left as something of a black box. The precise form that an account of concepts takes will depend, of course, on the research tradition that one adopts, and we certainly do not want to rule out multiple traditions. In the present paper, we adopt the tradition of cognitive science, and from that point of view we want to have more precision about the processes and form of the mental representations associated with concepts. However, whatever the tradition, we need enough precision on the nature of concepts to be able to say something about what in essence a concept is, what concepts people actually have, and how these entities can change.

What might a concept be

At this point, we want to pause in our examination of the notion of concept in conceptual change research and look a little more broadly at how concepts have been treated in the research literature. It is possible that other fields have provided empirically grounded theoretical accounts of the notion of concept that can serve as the necessary basis for conceptual change research. This discussion will also help give a feel for what a 'more precise' model of concept might be like.

'Categories' and 'concepts'

We first turn our attention to a literature most of which falls within the realm of experimental psychology. This research is typically described as being about the psychological nature of categories, but, in some cases, the researchers involved frame their work more generally as being about the psychological nature of concepts (Medin 1989, Rosch 1994). Although this research is rich and fascinating, we can only scratch the surface here. Our primary job will be to show why this work cannot provide the answers we need, that is, why we will not be able to look to the category research for an account of the nature of concepts that can provide the sole – or even primary – basis for further study of conceptual change.

We begin with a little of the history of this work. This story has been told many times by more authoritative sources, and the interested reader should look to existing summaries for a more detailed account (Lakoff 1987, Medin 1989, Rosch 1994). The story usually begins with the rejection of the 'classical' view of categories. In the classical view, a category is defined by a list of features – sometimes also called 'properties' or 'attributes' – that are both necessary and sufficient to determine category membership. For example, a bird might be defined as a thing that has wings, a beak and feathers. If something has all of these features, then it is a bird; if it does not, then it is not a bird.

The rejection of the classical view of categories was fuelled in large part by the recognition of what have been called 'typicality effects' (Medin 1989). It was found that, psychologically, humans judge some instances of a concept to be better examples of the concept than others. For example, robins are judged to be better examples of birds than chickens. And there are fuzzy cases – cases in which it is not clear whether or not the case should count as an instance of the concept.

These observations led to the development of what have been called 'probabilistic' theories of concepts. In one version of this theory, it is argued that a concept is represented by a prototype. So, for the concept of bird, the mental representation would be a prototypical bird – something like an idealized bird. Then category membership is determined by measuring the similarity of some instance to the prototype. In contrast, in an alternative class of probabilistic theories called 'exemplar' theories, the concept is represented by means of specific examples of the concept.

All of the above models are at least moderately specific about the mental representation of concepts. If they do not describe the precise form of the representation and process details, at least they describe content – what the representation represents. For the classical view, the representation is a list of features; for prototype theories, the mental representation is simply some form of the prototype. In some cases, there have been attempts to specify form and process, for example, how the prototype is represented and how similarity is measured for prototype theories. In any case, all of these models constitute steps beyond 'concept as black box'. If we believed that they provided a useful account of the psychological nature of concepts, we would probably be in a relatively good position to say something about possible types of conceptual change. For example, if concepts were feature lists, then changes in a concept would correspond to various sorts of changes to the associated feature list.

However, this history of category research is not yet done. The formerly more promising theories – prototype and exemplar – have now been rejected by many researchers, even by some of the original proponents of these theories. In fact, one of the most prominent of these, Eleanor Rosch, has altered her view. The evolution in Rosch's thinking is summed up concisely by Lakoff (1987: 43):

> Rosch considered the possibility that prototype effects, as operationalized by the experiments cited above, might provide a characterization of the internal structure of the category. Thus, for example, the goodness-of-example ratings might directly reflect the internal structure of the category in mental representation ... Rosch eventually gave up on these interpretations of her experimental results ... She came to the conclusion that prototype effects, defined operationally by experiment, underdetermined mental representations. The effects constrained the possibilities for what representations might be, but there was no one-to-one correspondence between the effects and mental representations.

The problem is that, although prototype (typicality) effects are very robust, the interpretation of these effects in terms of mental structures and processes is far from clear. In particular, there is no reason to believe that these effects are necessarily generated by a simple mental structure (e.g. a prototype) that embodies the typicality distribution in a straightforward manner. It is possible, for example, that the effects are generated by a complex ensemble of mental structures. This realization was partly driven by some observations that apparently 'classical' categories, such as odd and even number, also exhibit typicality effects (Gleitman *et al.* 1983).

These problems are closely related to the first difficulty we mentioned in relation to Carey's work – what should count as a concept? In most cases, the category researchers framed their experiments as asking questions about specific concepts. They asked: How can we characterize human understanding of Category X? They then went on to map out the structure of this category in the response given by subjects. However, the assumption that there is anything like a mental representation of Category X is a big assumption; it presumes that there is an entity in our mental ecology that is both worthy of reification and also straightforwardly associated with Category X. Again, how do we know, a priori, what concepts to look for? If every idea one might trouble to name is a concept and concepts are responsible for extraordinary difficulties in learning, then surely we are in an untenable position.

The notion that more elaborate mental structures may ultimately be responsible for generating prototype effects has led to a new class of 'theory-based' models. For example, Lakoff (1987: 70) understands concepts as defined relative to what he calls 'idealized cognitive models' (ICM), and he uses the concept of *bachelor* for illustration. To explain why priests and homosexuals are not typical examples of bachelors, even though they are unmarried males, Lakoff states that *bachelor* is defined with respect to an ICM in which there is a human society with (typically monogamous) marriage, and a typical marriageable age.

The move to theory-based models may very well be a move in the right direction; we believe that loosening the presumed relation between concepts – as delineated a priori – and mental representations is appropriate. However, although we cannot begin to analyse and critique this work here, we agree with Rosch (1994:

522) that existing theory-based models do not provide much precision on the nature of the 'theories' that are at the heart of these models.

> Is the theory of categories as theories a new claim of substance, or is it only a battle cry? On the one hand, it comes as a refreshing recognition of the larger context in which categorization always occurs – which it is hoped will provide an antidote to the minimodel mentality. However, the theories view is remarkably silent on all positive issues one might expect it to address. What is meant by a theory?... It is hard to escape the impression that presently absolutely anything can count as a theory and that the word *theory* can be invoked as an explanation of any finding about similarity or attributes.

If we accept this appraisal, we must conclude that we cannot look to current theory-based models to solve the problem of theoretical precision for us.

We have one final point to make about the category literature. This is perhaps, from the point of view of this paper, the most important point we will make, and it relates closely to our Difficulty 2. Although the category researchers often describe their work as being about the nature of human concepts, the term concept tends to be used much more broadly within educational research. When we say concept, we are not only talking about *bird* and *bachelor*, we also intend to include *force* and *number*. Thus, even if the models generated by the category literature work well for concepts like *bird* and *bachelor* – perhaps call them 'category-like concepts' – it is extremely difficult to see how the same models could possibly work for all of the things that have been labelled concepts by educational researchers. For example, it does not seem conceivable that one understands the concept of *number* only by holding some representation of a prototype number. Would all numbers be represented by a single prototype, or do we need a prototype for every number? We believe that this inquiry is *prima facie* absurd. There are many different places that we need to see numbers in the world, so even one prototype per number seems unlikely. Furthermore, knowing the concept of *number* means much more than deciding if a particular entity in the world is an instance of this concept. What might the ICM for number be?

What else might a concept be?

The discussion above of the category literature captures an important fraction of the noteworthy theorizing concerning the psychological nature of concepts. As we argued, although the category literature is relevant to our understanding of concepts and the models described may be among the best currently available, we believe that this literature does not provide the big answers we need to press forward our study of conceptual change. The early and more precise models, such as prototype theories, are proving to be unsatisfactory. Successor models have back-pedalled on precision. Furthermore, it seems evident that the existing models could not possibly cover the range of uses of the term 'concept'.

To fill out this discussion, we briefly mention additional noteworthy classes of models.

- Relational theories of concepts. In relational theories, concepts get their meaning by participating in a web of relations with other concepts. Given this model, changes correspond to the adding and deletion of nodes, as well as changes in the existing relations between these nodes (although we might not count the latter as conceptual change). Models in which concepts correspond to terms, and these terms are defined by the beliefs and propositions in which they participate, may be considered to be relational theories. It is possible that we might ultimately decide to include relational theories among theory-based models.

- Patterns of activation in a neural net. Models of concepts have also been proposed by connectionists. In these models, there may be no simple, localized 'mental representation' that corresponds to the concept; instead, concepts can correspond to patterns of activation that emerge during the activity of the neural net (e.g. Rumelhart *et al.* 1986). If we decide to adopt this viewpoint, we are left with two options. First, we could abandon any attempts to provide accounts of learning at the level of concepts. (This is a real option!) Alternatively, we could work on describing how and when the patterns corresponding to concepts arise. The latter path, of course, is not trivial, and still leaves us with all of the difficult issues pertaining to understanding the nature of concepts. For example, we still need to know: How can we determine which items in our vocabulary, if any, correspond to concepts?
- Actional/situated perspectives. Situated perspectives go, in a sense, a step farther than the distributed connectionist viewpoint. These perspectives not only reject a simple localization of a concept in a particular mental representation, they reject any sort of localization of concepts in the mind of an individual. (See, for example, Magnusson *et al.* 1997). In these views, concepts are taken to be abstractions over people acting in settings. Thus, rather than being a pattern of activation in neurons, they are somehow a pattern of people, things in the world and neuronal activations. If this notion of concept is retained, many tasks remain; we still must decide what to reify as concepts, and we must come up with a precise notion of concept that fits within this perspective.

Models of concept in conceptual change research

The list of models of concept just presented is far from exhausive. Nonetheless, we believe that the preceding discussion can give a good feel for the range of views concerning the nature of concepts, and for what a more precise model of a concept might be like. With this discussion in mind, we now return, once again, to Carey (1988). To which, if any, of the above models does Carey subscribe? There is no place in Carey's article in which she provides any approximation of a definition of the term concept. So, to determine how she understands the nature of concepts, we must resort to looking at the various statements she makes about concepts. Here, we will use just one representative passage. In arguing that the core biological concepts must change in the development that occurs between ages 4 and 10, Carey states:

> At least two important kinds of conceptual change that are found in theory changes are found in this case. Both differentiation and coalescence must be achieved (such as the just-mentioned differentiation of two senses of *not alive* and the 'coalescence' of *animal* and *plant*). The coalescence of *animal* and *plant* into a single category *living*

> *things* involves more than a simple disjunction *animal* or *plant*. By age 10 the two are seen as fundamentally alike in a way that involves having reconceptualized both. (p. 24)

First, we note that Carey is using the word category here, rather than the word concept. This may suggest that she has in mind the category literature described above. Indeed, the concepts mentioned in this last quote seem plausibly to be of the 'category-like' variety; unlike the concepts of *force* and *number*, it does not seem, prima facie, absurd to think of the concept of *animal* as being represented as a mental prototype. However, whatever model we attribute to Carey must be consistent with her statements concerning how concepts change. This poses a difficulty for prototype models, for example, since it is difficult to understand what it might mean to coalesce the prototypes of *animal* and *plant* to get a new prototype for *living thing*. How might we fill in some details of a prototype representation so that such a coalescence makes sense? Furthermore, it is hard to see how a single prototype could cover a category as large as *living thing*. Likely, we would want to resort to using multiple prototypes, or a large set of exemplars.

It is possible that Carey would be more comfortable within a theory-based or relational model. In fact, there are hints of this in the paper. But the notions of differentiation and coalescence are harder to understand in models that give meaning to concepts through non-localized representations. In the case of a relational model, for example, the coalescence of two concepts would have to involve a complicated rearrangement of relations in a manner that corresponds roughly to coalescing.[1]

The problem is not that any of these possibilities is untenable. Rather, the difficulty is that we simply cannot be sure what sort of model Carey has in mind. If we do not know what a concept is, statements about changes in concepts are questionable utility; saying that *animal* and *plant* coalesce is less helpful if we are not very clear on the general nature of *animal* and *plant* before and after coalescence.

To make these judgements more pointed, we look at work by David Trowbridge and Lillian McDermott that is described in two papers published in 1980 and 1981. In these two papers, Trowbridge and McDermott provide empirical results that they describe as relevant to students' understanding of the physical concepts of *velocity* and *acceleration*. Although our comments here are relevant to both of these papers, we will focus on the first, which describes itself as 'systematic investigation of the understanding of the concept of velocity among students enrolled in a wide variety of introductory physics courses' (1980: 1920).

Trowbridge and McDermott employed modified versions of some traditional Piagetian tasks. One of these tasks is illustrated in figure 1. In this task, two balls

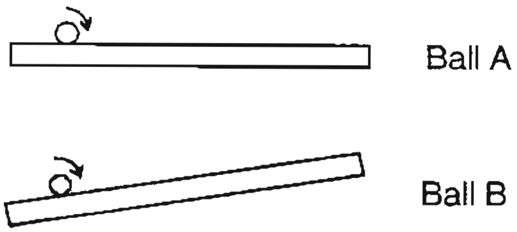

Figure 1. Trowbridge and McDermott's task.

roll simultaneously along U-shaped channels. The top ball, Ball A, travels at a constant speed from left to right. In contrast, Ball B starts off moving more quickly than Ball A, but slows down as it goes up the incline. The apparatus is set up so that Ball B passes Ball A near the start of the motion, but is then itself passed by Ball A.

The question that the researchers asked their subjects was: Do these balls ever have the same speed? In response, many students stated that the balls have the same speed at the two passing points, that is, when they were at the same position.

Trowbridge and McDermott attributed these incorrect responses to what they called the 'position–speed confusion', though the authors themselves were somewhat tentative in the use of this phrase. They stated that 'The use of the word "confusion" here should not be misconstrued to mean the mistaking of one fully developed concept for another. We are using the expression "confusion between speed (or velocity) and position" to refer to the indiscriminate use of nondifferentiated protoconcepts' (p. 1027).

Trowbridge and McDermott presents us with some quite clear empirical results, which they interpret in terms of the concepts of *velocity* and *position*. The authors' interpretation is important, since they prescribe instruction that is specifically directed at helping students 'resolve' (differentiate) these concepts. But Trowbridge and McDermott's conclusions need to be reexamined in the light of the three difficulties discussed above:

1. What should we count as a concept?

Trowbridge and McDermott's whole programme is framed in terms of questions about the concepts of *position* and *velocity*. But how do we know, a priori, whether there is any psychological reality to these concepts? To use Lakoff's language, the observations presented in this work greatly 'underdetermine mental representations' – there is a big step from subjects; responses in these Piagetian tasks to any claims about mental representations. As noted above, Trowbridge and McDermott are somewhat sophisticated in this regard; they realize that it may not be accurate to describe student behaviour in terms of the concepts of *position* and *velocity*, and they allow that students may have some undifferentiated mixture. Nonetheless, the presumption is that, in experts, the non-differentiated protoconcepts are ultimately differentiated into the concepts of *position* and *velocity*.

Trowbridge and McDermott's analysis actually illustrates a special case of Difficulty 1. The presumption that we can use expert concepts as a lens through which to view novice or everyday concepts is clearly, given our above discussions, questionable.[2] Other researchers have been sensitive to this particular difficulty (e.g. Wiser 1988). But to this we add the following: As of yet, we are not in a position to identify *expert* concepts! We must not make the mistake of presuming that there is a simple alignment between expert concepts (psychologically understood) and the technical vocabulary of a domain.

2. Lack of attention to manifest variety

There is less for us to say about this difficulty. We only note that Trowbridge and McDermott are using the term 'concept' for *position, velocity* and *acceleration*. This is the same term that Carey and others have used for *mammal* and *number*.

As we have discussed, some of the best developed theories of concepts do not appear to be readily applicable to non-category-like concepts of this sort.

3. Theoretical imprecision

Although Trowbridge and McDermott do try to be careful, phrases such as 'non-differentiated protoconcepts' cannot be helpful without further elaboration, and the authors fall far short of any model of concept. It is worth nothing that Trowbridge and McDermott use language that implicates processes of use as well as representation. 'Indiscriminate use' suggests something like carelessness. But no details of this characterization of process are forthcoming.

For a view of what they understand a concept to be, we can, as with Carey, look at some of their statements about concepts:

> Physics instructors generally share a common interpretation of the kinematical concepts based on operational definitions and precise verbal and mathematical articulation. On the other hand, students in an introductory physics course are likely to have a wide variety of somewhat vague and undifferentiated ideas about motion based on intuition, experience, and their perception of previous instruction. (p. 1020)
>
> Although students who were unsuccessful could generally give an acceptable definition for velocity, they did not understand the concept well enough to be able to determine a procedure they could use in a real physical situation for deciding if and when two objects have the same speed. (p. 1027)

For experts, Trowbridge and McDermott state that concept understanding is based on 'operational definitions and precise verbal and mathematical articulation'. This seems closest to a classical model of a concept since it appears to bound understanding within a simple and statable definition, even if that definition is not based on features and attributes. In any case, this statement implies a model of concept that is very localized in a particular mental representation. This inference is supported by the latter passage in which the authors speak of a concept being used to 'determine a procedure'. If concepts determine procedures, then the procedures are themselves not part of the concept. In contrast, the model of concept that we introduce below incorporates procedures into the heart of 'concept'.

Drawing as many possible inferences as we can, we still cannot be definitive about what 'concept' means to Trowbridge and McDermott. For the most part, they and others treat concept as a black box. Without a more precise notion of concept, serious difficulties follow. We do not know whether the terms position and velocity adequately cover novice or expert conceptual competence. Even if these terms do apply in some manner, without a better model of concept we cannot conclude anything general about conceptual change from these experiments, nor can we say anything reasonably precise about the change that happens in this case, aside from remarking on the different performances of novices and experts.

Finally, we note that a number of other terms have been used in conceptual change research. For example, some researchers use conception in place of concept and framework or theory to implicate a relational view. There are many nuances that may be explored comparing the meanings of these terms. However, to keep this paper manageable, we remark only that, in our estimation, none of these escape the difficulties listed here and in previous sections.

In the following section, we build on the above critiques towards a positive contribution to understanding the nature of concepts and conceptual change. The first section prepares for later work by setting the standards to which we aspire in presenting a model of concept. These standards flow, for the most part, directly from what was said above. We then introduce a particular model of concept – heuristically first, then in more detail, including particular expectations and issues that deserve empirical investigation. Last, we will present several episodes in a case study of conceptual change to illustrate and make plausible the value of the proposed model of concept.

Tenets of accountability

What should we expect from a model of concepts?

1. Standards of what counts as a concept

We should be in a much better position than the work cited above for determining what is a concept, and what is not. These standards should be both theoretical and empirical. They should have analytic clarity, and such clarity should lead to empirical focus and determination. We can expect an articulation of the pieces of 'knowing a concept'.

2. An account of different varieties of concepts

An articulated model of concept should lead to an understanding of parameters in the nature of different concepts, or towards an understanding of qualitatively different kinds of concepts. Indeed, the term 'concept' may not survive as a technical term and may instead be replaced by a variety of theoretical constructs. If there are warrants for treating concepts in a unitary manner, these should be explicitly given – theoretically and/or empirically.

3. Theoretical accounts of a variety of processes involved in concept use and conceptual change

Given a better articulation of the nature of concept, we should be in a much better position to understand the operation, emergence and change of such entities. More particularly, we should have a better understanding of which are the most difficult parts of assembling a concept and which parts are easy, at least in some circumstances. These processes need to specify what aspects of prior understanding are involved in new concepts and in what ways. This 'prior understanding' might include something akin to a prior version of the target concept (such as notions of 'alive' that precede a proper biological concept), but, more generally, other aspects of prior competence will be implicated. Because of this, conceptual change involves more than the change of prior related concepts. Ancillary knowledge may well specify contextual effects that determine in significant measure how easy or difficult learning a new concept might be.

4. The ability to match real-time data from students' thinking with the theoretical accounts of processes described above.

It is remarkable how much study of conceptual change has been simply description of before and after states, without watching in any detail how concepts work and develop, fitting such observations into theoretical patterns. Simply put, a good articulation of the nature of concepts should allow us, for example, to 'see concepts and conceptual change at work'.

Overall, having an articulated model of concept may allow us to organize and differentiate the processes involved in the application of a concept and in conceptual change. We should be able to localize difficulties students have with a concept and in conceptual change, and we may be able to theoretically identify types or components of change, possibly making connections that might otherwise be surprising. In addition, Item 4 reminds us that this ability to localize and be precise in our accounts must be empirical as well as theoretical; it is important that we are able to account for process data relating to the use and change of concepts.

Coordination classes: a heuristic introduction

Preamble

In the particular model of concept we will present, readers' expectations about what concepts are may well be violated. This is exactly as it should be. If expectations are not violated in developing a scientifically sharp and useful concept of concept, then it is almost certain we have not produced anything significantly new. So, our first task is to sketch briefly what can be and what should not be expected.

Our model of concept will not include a simple account of a unitary mental structure ('concept nodes'), or even a small number of mental structures, associated with a concept. Instead, what we will describe is more like a knowledge system (diSessa 1996). This is a fundamental break from much prior work, and we believe it is an important one. In particular, a knowledge system approach is exactly what allows us to analyse real-time performance in terms of particular parts of the system. As a consequence of adopting a system view, however, the boundaries of any concept become somewhat fuzzy, since many system parts are involved. Our contention is that this is not a disadvantage, but reflects a fundamental reality. Instead of stating that one either has or does not have a concept, we believe it is necessary to describe specific ways in which a learner's concept system behaves like an expert's – and the ways and circumstances in which it behaves differently.[3]

In accord with this shift to a system view, we describe our model partly in terms of performance. We will still attempt to provide structural descriptions of the parts of the system; however, we believe it is equally important to understand the function of the system. Performance specifications can also enhance our ability to make connections with empirical data

Another difficulty that may hinder readers in understanding our model is that when many people hear the term 'concept', they may think of category-like concepts. However, our model is not clearly about category-like concepts. It turns out that, in some instances, the model of concept that we will present behaves in a category-like manner: it defines a class of entities that are cleanly distinguished from other entities. But this need not be the case, or it may not be obviously the case. Furthermore, as we will see, there are instances in which our model behaves in a way that is manifestly not category-like.

A start

Our model concerns a class of concepts that we believe is important in science learning. We call them *coordination classes*. To motivate the idea, we begin with the contention that 'seeing things' in the world – gaining information about them – is a complex cognitive accomplishment. What we want to describe is what lies behind this ability to 'see things'. In this regard, the comment is often made that every observation is theoretically based, as well as empirical. The notion of pure and indubitable 'data' is no longer regarded as a serious possibility in the philosophy of science, nor is it commonly used in the avowed warrants of professional science. While we do not believe it is appropriate to describe what lies behind observation necessarily as a 'theory', this nonetheless provides a good beginning orientation. In this way, our concept of concept might be considered a refinement of theory-based models mentioned in the subsection on categories and concepts.

Coordination classes, then, are systematically connected ways of getting information from the world. Note that this is a performance specification and it stipulates that the prototypical task for coordination classes is getting information. The prototypical task for category-like concepts, in contrast, is to determine whether something is or is not a member of the category. For illustration, consider the two candidate concepts *velocity* and *bird*. The primary task that a concept of velocity must do is to determine the amount of velocity. Is the velocity of that car high or low? Which of those two cars has the higher velocity? In contrast, the primary task that a concept of bird must accomplish, if we follow the lead of category researchers, is determining whether some particular entity in the world is a bird.

Of course, getting information and determining class membership may be related. But one can use information for many purposes in addition to determining category membership. Our implicit assertion is that getting information is a better characterization of the work that must be done by scientific concepts – at least of the sort we consider – as compared to sorting entities by category.

In a certain sense, this programme goes back to Kant. He saw that an essential contribution of the individual was the fusion of diverse perceptions into an integrated reality. Kant believed this was an a priori necessity for any form of cognition and that the forms for accomplishing it were universal. In the present case, our interest is in understanding how 'seeing in different situations' can constitute the core function of concepts (coordination classes). Shifting the means of seeing, a fortiori, is the core problem of conceptual change.

The difficult job of a coordination class is to penetrate the diversity and richness of varied situations to accomplish a reliable 'readout' of a particular class of information. This information may be, for example, seeing the defining attributes of instances of the concept in question, if the concept is category-like. 'Seeing', as used here, is at least partially metaphorical. In many instances this seeing is a substantial accomplishment of learning and will depend only very partially on basic perceptual capabilities. In addition, this form of seeing sometimes involves explicit strategies and extended reasoning, as later examples will show. We will use *coordinate* as a verb in place of 'see' or 'determine information' to emphasize a coordination class perspective.

Think of an individual faced with a rich phenomenological world, where things present themselves in very many different contexts. That individual is faced with the problem of opportunistically picking out features in the current context that relate to the critical information required. These features may vary from situation to situation. Indeed, even within one situation, one may have to combine several contingently related observations via some form of reasoning to infer the necessary information. Coordination classes include strategies of selecting attention and strategies of determining and integrating observations into the requisite information.

Coordination has a double meaning. First, it refers to the fact that, within a given situation, multiple observations or aspects may need to be coordinated to determine the necessary information. This version of coordination might be described more precisely as *integration*. Second, it refers to the fact that, across instances and situations, the knowledge that accomplishes readout of information must reliably determine the *same* information. Otherwise we might count the concept as confused or incoherent. This latter sense of coordination might be called *invariance*.

Imagine you have gone to a party and meet someone named John. Consider two different ways of getting information about John. First, you may be interested in him as a personality. Is he witty? Is he trustworthy? Is he aggressive? To determine these things, you look at him and listen to him, but you focus on very particular things, and not others. You listen to his words, but 'tone' may be more important than meaning. You watch his face for signs of sincerity (however one does that). You may even attend to feelings he engenders in you, like fear. Often, you combine several 'observations'. Perhaps you are afraid of him, but you are certain he didn't intend this. Notice that causality is part of your judgments – in this case, psychological causality of intentions, means of carrying them out, and their consequences. Notice how subtle and critical attention must be, and how multiple modalities, sight and sound might be combined. In sum, our 'readout' of personality in this case is a complex task, requiring integration of very diverse observations and extended reasoning based on psychological causality.

Cues available to determine John's personality might depend on the situation. In a formal situation, for example, it might be hopelessly difficult, and part of your coordination knowledge might be the ability to appreciate this difficulty. On the other hand, you might find John on the tennis court, and believe how he handles losing is a dead giveaway for 'good sports'. In general, having multiple strategies to see the same thing in multiple situations and the ability to judge the reliability of various strategies in various situations is prototypical coordination knowledge.

Notice that it is really a significant stretch to think of personality-determining strategies as defining a category-like concept. The task of determining which things in the world are 'personalities' and which are not is much backgrounded. Instead, information gathering is the primary function. 'What is a person's personality?'[4] Personality determining illustrates the idea of a systematically connected set of strategies to get a class of information, which is the essence of coordination classes.

Now, consider that you want to get information about John simply as a physical object, rather than as a personality. Suppose, for example, John has become drunk, and he is no longer interesting to talk to, or to interact with as a personality. Furthermore, if he is sufficiently drunk, he may even have lost the ability to get from place to place on his own. Intellectually boring and without autonomous locomotion, John might as well be a sack of potatoes, for most purposes. Indeed, finding and moving a sack of potatoes might involve essentially the same set of skills as finding and moving John to get him home. This is a very different way of coordinating John compared to determining his personality. You might still attend to the sounds he makes. For example, you might listen for his muttering in order to find him if you forgot where you last saw him. But you attend to these sounds not for personality or meaning, but to locate the source of a sound, just as a child might listen for the quacking of a mechanical duck in order to find it. And if you can locate a quacking toy in a room full of people, you likely have the skills to find John by sound, even if you are naive about human personality. Coordinating location doesn't have much to do with any of the categories people, 'John', or toy duck.

Location is less category-like even than personality. The reasons are fairly clear. Every physical object has a location, so it does not make sense to ask if a scene contains a location, as it does for dogs and cups. Furthermore, there just aren't many location-like things, so that distinguishing location from nearby categories may not be important. In general, complex attributes like location and personality don't make prototypical categories, although they are fine candidates for coordination classes.

Determining personality is probably not very closely related to the sort of coordination classes we have in mind, those involving concepts from physics. However, coordinating 'as a physical object' (including location) is probably more closely related. Below we discuss the development of this capability, in part to illustrate its relevance and in part to continue developing the model for coordination classes.

Piagetian examples

Piagetian psychology provides a wealth of examples that may very well be amenable to analysis as learning particular coordination classes. Early on, babies must learn to coordinate various observations in order to determine the position of an object. 'Object permanence' may be defined roughly as having the 'concept' of spatially-localized and permanently existing entities, such as those you may choose to employ in dealing with poor John. Empirical studies show that this is an extended accomplishment. An accumulation of a complex and broad set of strategies and understandings is characteristic of coordination classes, as opposed to, say, learning a rule or definition. In particular, babies give little indication very early on of even continued attention when an object disappears from view. Maintaining attention and instigating search strategies are accomplishments. Later, children may search, but do so ineffectively. They may look where they last saw an object even though they are told or are given enough information to determine that the object has been moved. When children have the capability to track and determine the location of most objects under most conditions they are likely to meet, we say they have the 'concept of object permanence'. An analysis of the knowledge necessary to do this would constitute an analysis of the coordination class 'location of physical objects'.

Figure 2. The blue and red trains start at the same place and time (A); the blue train is slower, but travels for a longer time – until time C, after time B.

'Giving children enough information to infer that an object has moved' implicates a critical pool of knowledge relevant to defining a coordination class. This is the set of inferences that lead from observable information to the determination of things that may not be directly or easily observable. We call this class of information *the causal net*. In many instances, the causal net corresponds roughly to what people intuitively expect of causality, especially in the empirical analyses we discuss later. Inferring something about John's personality – by inferring motives, means and intentions – involves psychological causality. In other instances, 'causal' must be understood in the very general sense of relating observations and determinations to other plausible or necessary determinations. Causal nets are, roughly, our replacement for the 'theories that lie behind observations', or the theories implicated in theory-based notions of categories.

One of our favourite examples of coordination, which highlights the causal net and how critical and frequently invisible it is, comes from Piaget's studies of children's conception of time. Consider two trains that start from the same place at the same time (labelled A in figure 2). Suppose that the blue train is slower than the red train. However, the blue one continues running for a longer period of time (until C), while not long enough to overtake the red train spatially (which stopped at time B). Some children seem entirely capable of making all the relevant observations about relative time duration in order to conclude that the blue train went on for a longer duration than the red train. After observing the trains, they will assert that both trains started at the same time, and that neither started before the other. They will also assert that the blue train was still going when the red train stopped, and that the red train was stopped when the blue train stopped. In terminology that we will explain more fully later, such children possess in-principle sufficient *readout strategies*. But then they will claim that the red train went on for a longer time than the blue train.[5] To an adult, making the former observations about when the trains started and what was happening when each stopped is an invisible step away from determining that the train that stopped after the other went on for more time. But, some children appear simply not to make this inference. It is missing from their causal net.[6]

To take a different Piagetian example, acquiring the concept of 'conservation of volume' as an issue of coordination highlights the two senses of the term coordination. Early on, children will fairly systematically claim that an amount of liquid that has just been poured from a wide container into a narrow one is now 'more'. But this is not as simple as confusing height with 'more'. On other occasions, they may use width as a determinant of amount. This is (lack of) invariance, coordinating across different situations (as in creating a coherent concept).

In the case of a child who 'has conservation of volume', he or she must coordinate, in the sense of jointly using, width and height to estimate volume; this is integration. In the case in question, the child may well judge that 'the same liquid, with none removed or added' is a more precise and reliable determinant of more or less than coordinating width and height. So, integration might be, in some cases, selecting from different determining possibilities.[7] This last case again shows that judgements of reliability and contextual effects that determine use of one method or another, are central in defining coordination class knowledge.

A final Piagetian example of (hypothetical) coordination class knowledge suggests that, sometimes, coordination is complex, extended in time and strategic in nature, compared to being quick, perceptual and invisible. Reading out the information about cardinality of a set involves a complex knowledge system consisting of many coordinated parts, responsive to a wide range of situations (e.g. arrangements of objects, kinds of objects, visibility, and so on). The most obvious part of cardinality competence, counting, is itself complex and requires a high degree of coordination (for example, coordinating the 'number song' with sequentially touching objects). One decides to count, it takes time to do so, and then one may reflect on parts of this strategy – for example, when to stop. Still, invisible and quick components are central parts of this hypothetical coordination class – for example, judging that if no objects join or leave a group, it has the same cardinality.

We have not defined coordination class precisely. However, even at this stage we believe it is evident that the concept has some specifcity and 'bite'. We can see that some examples that might be referred to as concepts are clearly not coordination classes. Other cases are strongly suggestive that they are, while a third class of 'concepts' might take extended empirical analysis to determine their status as coordination classes. In the first case, non-coordination class concepts, it seems clear that certain categories piggyback on preexisting coordinations. For example, learning categories such as 'a glurb is blue and square' piggybacks on coordinating, individually, shape and colour information, specifically using logical connectives. It is implausible that people learn how to see 'square' or 'blue' when learning about glurbs. And it is also implausible that people learn logical connectives in the context of learning about glurbs. In net, a glurb may be a category, but learning it is not learning a coordination class. In contrast, an example that is prototypical of coordination class learning is object permanence. Finally, whether there is any component of coordination class learning involved in scientific concepts, such as force or acceleration, is, at this point, unclear. Since the relevance of coordination classes to science education depends on encompassing ideas like force and acceleration, we turn directly to this issue.

Quantities as coordination classes: refining the definition of coordination class[8]

We suggested above, with little argument, that cardinality might constitute a coordination class. In this section we argue, more generally, that physical quantities might constitute coordination classes. We review what we have said about coordination classes and present a slightly elaborated model of coordination class knowledge for the case of quantities. At this point, we will view the issue of whether quantities constitute coordination classes as hypothetical. Later we will argue – briefly from a theoretical point of view and more extensively using empirical data – that there is good reason to believe some quantities do constitute coordination classes, and that viewing them as such is insightful. To aid clarity, we italicize the key theoretical terms in this section.

In our heuristic introduction to coordination classes, we emphasized *readout strategies* that deal with the diversity of presentations of information to determine, for example, characteristic attributes of a concept exemplar in different situations. In the case of quantities, this amounts mainly to determining the value of the quantity in a particular situations. Readout strategies constitute the first of two main structural components of coordination classes.

We suggested that two different kinds of coordination are central to readout. First, in a given situation, one must frequently collect, select or combine diverse 'observations' to determine what we wish to 'see' (*integration*). Indeed, when each strategy will be used, and how reliable these strategies are, is prototypical coordination class knowledge. The second sense of coordination deals specifically with how observations in different circumstances can manage to determine the same information (*invariance*)[9]. Integration and invariance are the two primary performance specifications of coordination classes, beyond 'getting information'.

The general class of knowledge and reasoning strategies that determines when and how some observations are related to the information at issue (that is, related to 'determining the value of the quantity') is called the *causal net*. In the case of physical quantities, the term causal may be particularly apt. This is because quantities are involved in regularities that connect circumstances (preconditions) to outcomes. If certain preconditions causally determine particular outcomes, then the relations between preconditions and outcomes can be used to determine outcome information using only precondition information. Similarly, if we are trying to determine quantities involved in certain preconditions, causal links may allow us to observe outcomes and infer preconditions. The causal net is the second primary structural component of coordination classes.

To illustrate the causal net, consider what is a particularly important case for this paper, determining force. The existence of a force 'causes' acceleration, which is the essence of Newton's second law, the equation $F = ma$. (F denotes force; a denotes acceleration, and m denotes the mass of the object that is accelerating and on which the force acts.) So, if you want to determine force, you can sometimes look first to determine acceleration. Conversely, the same equation allows you to determine effects from preconditions. If you happen to know the force, then you can determine the acceleration.

This example illustrates the potentially pivotal role that equations may play in the causal net for quantities, especially for experts. However, for many reasons, some of which will become evident, we caution that identifying the causal net with equations, even equations that are defining, simply will not do. In particular, students' non-quantitative assumptions about what things cause or determine other things are quite evident in their thinking about and determining quantities. We will provide some vivid empirical examples later where a student abandons equations in deference to more intuitive assumptions in acts of coordination. Although we won't argue the case here, we believe that, even for experts, non-quantitative connections among quantities and other aspects of the physical situation are critical in coordinating.

The relations between readout strategies and the causal net are intimate. One looks for things that are related (via the causal net) in order to determine some quantity. Indeed, even seeing those secondary features may involve additional inferences from the same or another causal net. In general, readout strategies and the causal net should co-evolve as learning occurs. There should be episodes of 'conceptual bootstrapping', where causal assumptions drive the learning of new readout strategies. On other occasions, 'noticings' – for example, that something surprisingly affects something else – may drive reformulations in the causal net. In general, characteristics of one will have important influences on how the other behaves and develops.

On the basis of these distinctions, we can be more refined in describing the process of change in acquiring a coordination class. In particular, the separate changes in readout strategies and in the causal net constitute parameters of conceptual change. For example, it is plausible that no new readout strategies are necessary in learning a new coordination class, but existing ones come to be organized and used differently. On the other hand, sometimes new strategies may be necessary. On the causal net side, the issue is similarly differentiated. In principle, one may have essentially no prior causal net and a new net may need to be built from whole cloth. Or an old causal net may need to be developed and reorganized to varying degrees.

In the following, we will not systematically replace 'concept' with 'coordination class'. We avoid this for two reasons. First, coordination classes are only one type of concept, and we still need a more general and open term. Second, we want to emphasize that we are working on the same issues as those concerned with concepts and conceptual change.

Assumptions and predictions

We have already, implicitly or explicitly, provided significant possibility for differentiating aspects of the accomplishment of coordination class competence into different categories. In particular, we distinguished acquiring new readout strategies from altering the causal net. We have also distinguished coordination across situations (invariance) from coordination within particular situations (integration). All of these distinctions should allow us to be more precise in describing difficulties, easy accomplishments and partial success in conceptual change. In this section, we elaborate a few specific expectations about how things will turn out empirically in the case of physics. This is based mainly on prior work about intuitive conceptions in physics and their relation to learning schooled concepts (diSessa 1993).

Probably the most important observation in this prior work is that naive understandings of the physical world constitute a rich knowledge system. First, this means that perceptual (low-level readout) strategies may need to be organized and somewhat extended. However, on the whole, this is not a particular difficulty in achieving conceptual change in physics. What may be more surprising is that a rich and flexible causality already exists for understanding physics prior to instruction. diSessa (1983, 1993) called this causality the naive 'sense of mechanism'. In brief, the claim is that naive physical causality consists of a rich system of elements that are organized only in a limited degree. The elements themselves, called p-prims, are relatively simple and usually abstracted from common experiences. For example, people expect the greater effort is accompanied by greater results. Push harder, and things go faster or farther. The extension of this prior work that we make here is to project that the naive sense of mechanism *is* the causal net beginners use in bootstrapping their way toward understanding physical concepts. In short, p-prims coordinate and are especially prominent at early stages of learning. Recalling that the causal net is roughly akin to 'theory' in other accounts of concepts, the properties of p-prims – including their individual simplicity, their large number, limited organization and their phenomenological nature – are stark contrasts to what usually comes to mind as properties of theories.

Since students start with a different sense of mechanism, we project that the causal net is, in fact, the primary locus of difficulty in learning concepts of schooled physics. Among other things, it needs to become more systematically organized (diSessa 1993). In the terms of this discussion, both issues of invariance and issues of integration are relevant potential difficulties in 'recruiting and organizing' prior causal net knowledge. In fact, because naive physical causality is so rich and diverse, invariance in particular is extremely problematic. In different situations, different p-prims may be evoked, and hence students may behave as if their concepts varied from context to context. Indeed, the diversity of naive phenomenology may mean that even achieving a sense of invariance – a sense that one is getting at the same kind of thing in different circumstances – let alone full coordination, may be extremely difficult.

Finally, although we will not justify it theoretically (but see later data, and consult diSessa 1991), we project that the details of calculation are minimally relevant to the coordination that equations accomplish. Instead, rough qualitative relationships seen via equations do the most work. Sherin (1996) provides a theory of equation interpretation that names and empirically determines a critical set of schemata that lie at the base of the skill of 'equation interpretation'. These schemata are thus prime candidates for the sort of relationship that *do* do the critical work in how equations coordinate.

In sum, the observed richness of intuitive phenomenology suggests that:

(a) There is little shortage of capability to see physically relevant things in the world. Students don't need to lean new sensory strategies to coordinate properly. Instead, the causal net will constitute a more serious problem.
(b) Because of the richness of intuitive causality, the prototypical problem is that *too much* will be seen as relevant, or the wrong things. Thus, invariance will be a particular issue.
(c) P-prims coordinate. That is, the beginning causal net is, substantially, the naive sense of mechanisms, which we identify as the set of p-prims a student has.
(d) Qualitative interpretations of equations are more important than precise calculation in equations' role in coordination. This may have important educational implications for a curriculum full of calculation and relatively devoid of both qualitative analysis and explanation as learning activities.

It is worth remarking that p-prims manifestly cannot count, themselves, as coordination classes. While they may constitute concepts in some sense, individually a p-prim is simply too small and isolated to constitute a coordination class. This fact is particularly notable as expectations extremely similar to p-prims have been identified as 'the native concepts' (for example, see 'motion implies a force' in Clement 1982) or even 'the naive theory' (McCloskey 1983) involved in learning physics. If naive coordination classes exist, undoubtedly multiple p-prims are involved in the causal net, but these p-prims are not sensibly considered to be concepts themselves.

J coordinates force and acceleration

We turn to empirical work in this section. We look at several episodes from an extended clinical interviewing study conducted (by AdS) with a freshman at the University of California, Berkeley. The subject, called J, had done well in high school physics and managed introductory university physics reasonably, if not without some difficulties. She was articulate and spoke easily about her thinking about physics.

Our intent is not to demonstrate or prove that coordination classes exist and have the properties we say. That would be much too much to expect for anything as complex and elusive as a model of concepts. Instead, we want to illustrate coordination class ideas in an empirical context to help readers tune up their own coordinations for coordination class. Beyond this, we want to show that viewing concepts as coordination classes at least plausibly applies to the phenomenology of learning physics, and to suggest that coordination classes' existence is not only empirically tractable, but that it may well have significant consequences in how we understand concepts and conceptual change.

The main focus in these episodes is the concept of force. Force is an excellent focus for several reasons. First, it is important – arguably the most important concept in Newtonian physics. Second, relying on prior work (such as diSessa 1993), we know that the naive sense of mechanism is very rich in the vicinity of force-related phenomena. Hence, especially since force-related p-prims have already been investigated, we should be in a good position to see their role in learning the schooled concept of force. Finally, the concept of force embraces a very wide range of situations and phenomenon. It applies to phenomenologically very different situations, such as the force that supports an object sitting on a table or the motion of an object tossed into the air. Given the fact that this diversity is covered by a single concept, force, we should expect issues of invariance particularly to show in the development of the schooled concept. We will see that J has difficulties with the concept of force that seem explainable in terms of force's constituting a coordination class. Not only does she show these difficulties precisely in the context of reading out the amount of force, but we can compare in detail how she does this to how experts do it. Furthermore, how she coordinates force is consistent with the assumptions and predictions of the previous section.

In addition to force, we will look briefly at acceleration. We consider a second concept primarily to provide a contrast to the difficulties we find J has with force, and to understand both concepts as variations comprehensible within the framework of coordination classes.

In what follows we make our points as directly as possible. Some limitations and difficulties will be discussed in the section, 'How have we done?'

Episode 1: J denies one can feel the force of gravity

In this segment, J and the interviewer are discussing the fact that all objects, no matter how heavy or light, fall at the same rate – Galileo's famous discovery. The interviewer proposes the problem: If gravity pulls harder on different objects, how is it that they still fall at the same rate?[10] He presumes the following coordination analysis, which is ordinarily well within the reach of first-year students. The force of gravity is easy to feel (read out) in the situation where you hold a body. The force of gravity is pulling down, and the supporting force of your hand is pushing up. If the body is not moving, these forces must be balanced, so the force your hand exerts – which you can feel – is equal to the force of gravity. Indeed, this analysis is so transparent to competent physics students that they might say one can directly feel the force of gravity. In contrast, J simply denies that one can feel the force of gravity. Evidently, she does not coordinate 'force of gravity' in the usual way.

In protocol transcriptions, comments or explanations are in italics. Ellipses denote that portions of the transcript have been omitted for clarity. The symbol // denotes an abrupt stop or interruption.

> *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?'
> [J *has focused directly on the critical coordination problem*.]
> *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 ... Weight equals mass times gravity.... So when you talk about the weight you do incorporate gravity. But whether you have something heavy or light, gravity's gonna be the same.... You don't feel the force of gravity.

J considers the situation where you hold a body in your hand, and 'observes' the force of gravity in a nonstandard way. She agrees with part of the standard coordination inference, that one feels something more or less directly. But she describes this as the 'weight' of an object, and distinguishes it from the force of gravity. She coordinates gravity indirectly via two things. First, she coordinates via the property 'gravity acts the same on all bodies'.[11] And she coordinates via the correct equation $F = mg$, incorrectly identifying g, rather than F, as the force of gravity. (This g needs to be coordinated as an acceleration, not as a force). These two elements – 'gravity (gravitational force) is constant' and 'the g in $F = mg$ is gravitational force' – constitute the part of J's causal net that we can see in this situation.

This seems a strikingly clear example of a student having a different coordination for force than a physicist. It is patently a situation where the issue is one of readout, how one sees a force in a particular situation. Rather than following the expert causal net to coordinate gravitational force directly from felt force, she introduces an intervening concept, weight, which can be felt, but which is related only abstractly, via an equation, to gravitational force. Her problem with using the equation is neither that the equation is incorrect nor that it is inappropriate.

Instead, she incorrectly coordinates a term in the equation, g, as a force rather than as an acceleration.

Episode 2: 'It's just bringing the paper with it'

In the following brief segment, J again does not use the expert way of coordinating force. In this case, the expert causal net inference is that 'action and reaction' forces are equal and opposite. Thus, if you know one, you know the other. In place of action and reaction, J uses a familiar p-prim: *contact conveys motion*. A prototypical situation in which one may abstract this p-prim is a child pulling a wagon with an object inside. Because the object rests on the wagon, it is unproblematically moved with the wagon. The critical feature of this p-prim is that it explains motion with no need for intervention or force.

The situation at issue in this episode is simply pushing a book along the surface of a desk. Strikingly, J begins by identifying a correct action and reaction pair, the force your fingers exert in pushing the book, and the 'resistance', the pushing back of the book against your fingers.

I: And now, I'm pushing on this book. 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.

Below, the interviewer attempts to get J to identify another action and reaction pair, which is the frictional force pushing against the book's motion and the complementary force the book exerts on the desktop in the direction of its motion (refer to figure 3).

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.

J accepts that the frictional force of the table on the book acts backward on the book, resisting the motion, but she rejects the existence of a complementary force of the book pushing forward on the table. Next, the interviewer puts a sheet of paper under the book, demonstrating that it moves when the book is moved. He suggests, as a physicist would, that there must be a force to move the paper. Indeed, one can 'feel the force' by holding the paper, preventing it from moving.

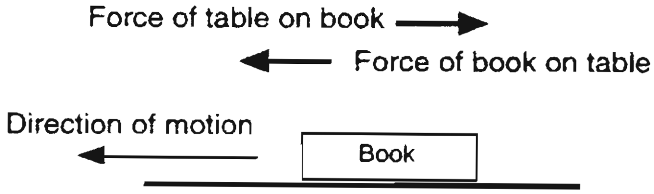

Figure 3. When a book is pushed along a table with friction, the book exerts a force on the table and the table exerts an equal and opposite force on the book.

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 think you could start saying there's all these millions and trillions of forces. [*Several sentences are omitted in which J describes complex situations with many forces.*] So I would just say that it's [*the book*] sliding against the table and bringing the paper with it.[12]

J understands, in some degree, the principle of action equalling reaction. She will use it to 'see' the amount of force involved in one situation (the amount of force pushing back on fingers that are pushing a book). But she encounters problems of invariance. She just doesn't see action and reaction as applying in another situation, to the pair consisting of the friction that the table exerts on the sliding book (which she acknowledges) and the force the book exerts on the table (which she does not acknowledge). While this protocol does not show the reasons for this lack of invariance, it shows why she rejects an argument of the interviewer. She does not see that a force is necessary to account for a moving paper precisely because of a specific p-prim, *contact conveys motion*, applies and, in that situation, the p-prim negates the need for force. This is a second example of non-invariance because, in other circumstances (see below) J articulately uses a principle that motion requires a force. In this case, it seems clear that the context specficity of a particular p-prim explains the lack of invariance. This is precisely one of our expected occurrences: because the naive causal net is rich and diverse, it sometimes overrides principles that seem logical in other situations (e.g. that an unbalanced force is needed to create motion).

Episode 3: J denies $F = ma$

In the following episode, J finds a fundamental contradiction between her naive causal net and the equation $F = ma$. In general, J sees rest as a special state usually caused by a balance of forces (diSessa 1996). In her view, motion is the result of an imbalance, when one force *overcomes* another. In other words, J uses the p-prims *balance implies stability (rest)*, and *imbalance* (or *overcoming*) *implies motion*. The contradiction is that an object moving at constant speed, for J, must be the result of unbalanced forces; but, $F = ma$ says that unbalanced forces result in acceleration – not constant speed.

Just before Episode 2, the interviewer had introduced the problem of pushing a book along on a table. J had announced that, in that situation, the force of the hand pushing the book had to be greater than the force of friction pushing backward on the book.[13] After Episode 2, the interviewer returns to focusing on the force of the hand and friction. He reminds J that she said the hand force is greater, and he sows the seeds of contradiction by raising the issue of acceleration.

I: And what about the situation where I'm just pushing along like that, after I get it going?

J: After you get it going, it's going at a constant velocity.

I: Right. In that situation is the force of my finger greater than the force of friction?

J: [*shakes head yes*]

I: But if I have an unbalanced force, isn't it going to accelerate, continue to accelerate?

J: No.... Because if you want to keep it accelerating, you have to keep applying, you know, a greater force.... I mean if you said that every time forces were unbalanced, something accelerated, that would change everything.

This is the beginning of a 15-minute inquiry into the conflict between her naive causal net, which implies that the force of friction must be less than the force of the hand pushing the book, and $F = ma$, which makes it impossible to have a situation with both unbalanced forces and constant velocity (no acceleration). We reproduce here a few critical sections.

Before, J reviews the contradiction she has run into, and she considers carefully whether it is possible that she just hasn't seen the acceleration of the book. For example, perhaps it may be too small to see. But J is confident and competent in her coordination knowledge for acceleration (and gets a little support from the interviewer), and she concludes that there simply is no acceleration.

J: Because of $F = ma$, if you have a constant acceleration, then you should be able to have constant force.
I: Sounds right. If $F = ma$, constant mass.
J: Well, then, I don't know.
I: This becomes puzzling.
J: Yes. See, to me, if you're applying this constant force like this, that doesn't look to me like the book is accelerating. At all.
I: No, it doesn't.
J: Maybe it's just accelerating at the same rate. It seems to me that, see, if it has constant acceleration, the velocity is still increasing. It's just increasing at a constant rate, right?
I: Say that again.
J: If it has constant acceleration, the velocity is increasing, but it's at the same rate.
I: Yeah, that's right.
J: So if you're pushing this, it has constant acceleration. It still has to be getting faster and faster and faster. I don't see it getting faster and faster, unless you're applying greater and greater force.

Shortly thereafter:

J: See, I would think that when you're pushing it, it has a constant velocity, not a constant acceleration. I can't imagine someone telling me something that would convince me otherwise. Because// It's, like, you see accelerations, you feel acceleration. It's not like this book is really accelerating, and we just don't see it.
I: I see. It's something in physics you can see.
J: Yeah.
I: Forces are hard, but acceleration you can see.

This is canonical coordination behaviour. Because the amount of acceleration has come into question, J focuses her attention on seeing acceleration. With that increased care, J still (correctly) cannot see acceleration, and the contradiction remains. She sees the force of the hand overcoming friction just as clearly as she sees acceleration. Only, in that case, she is using an inappropriate coordination from her naive causal net, motion implies imbalance.

J's relative success with acceleration, we believe, is a general phenomenon among learners. It is accounted for by the fact that the naive causal net for acceleration is more like the experts' than is true in the case of force. Still, we know non-trivial changes in the causal net must be accomplished to coordinate acceleration adequately. For example, the naive distinction between acceleration and deceleration is irrelevant to experts. For physics-naive people, deceleration is a natural phenomenon that needs no explanation (diSessa 1993). However, acceleration always needs a forceful intervention.

Before resolving the conflict, J states it once more.

> *J*: No, I'm just thinking about if that's true, let's say, that every time there's a force something's moving, then it's going to be accelerating. Because anytime something moves, the forces aren't equal. And, to me, it's going to be very mind boggling to think that every time something's moving, it's accelerating.
>
> *I*: That just does not sound possible.
>
> *J*: That just is not slightly possible.

J's resolution to this problem is stunning. She decides $F = ma$ simply does not apply to pushing a book across a table! She rejects the equation as relevant to coordinating force in this situation.

> *J*: Like, to me, you look at $F = ma$, and there's a force, and that has to mean acceleration. But, then, it's easy to say, 'that's true'. But, I mean, there's no way it is. I guess you can just say that, you know, those darn equations aren't applicable to every single thing. They're not always true. You can't live by them.

Later:

> *J*: I mean, to me, it just makes so much sense that it's crazy to even debate it: that if you're pushing this with a constant force, and you see it moving, that it's not accelerating.

Still later:

> *J*: I mean, you learn these formulas in school and you can't// You know, half the time they only apply to certain perfect models, especially like in chemistry. You know, certain situations. . . . I just thought that $F = ma$ was one of those that was universal. You know, it wasn't, like, specific.

We see that, at least for J, equations are not at the center of her causal net. Here, p-prims hold sway by virtue of their salience in this situation. Once again (as in the case of sliding paper), lack of invariance is accounted for by the context specificity of intuitive inference.

There are three ancillary but notable observations one can make about this episode. First, J's difficulties with physics, at least in Episodes 2 and 3, have nothing to do with categorization. She does not misclassify any non-force phenomenon as a force, or a force as some other entity. Instead, her problems are prototypically coordination – selecting attributes, like motion or rest, relevant to determining force, and reasoning about magnitudes of various forces via a naive or expert causal net.

Second, there does not seem to be any sense in which J is being 'indiscriminate' (in Trowbridge and McDermott's sense) about her use of $F = ma$ and action and reaction. She shows signs of being interested, attentive and careful. More centrally, she is being highly discriminating – deciding that, in this case, a p-prim is more causally relevant than other means of coordinating. Lest this seem naive, we hasten to add that $F = ma$ does not even come close to explaining all of the rich phenomenology of the physical world. For example, it is not possible to explain the behaviour of light (electromagnetic radiation) or why some objects are brittle and others are durable by reference to $F = ma$. Discriminating what ideas apply where is critical, and J is working on these issues.

Third, this might have been an excellent learning opportunity for J. Had she had more conviction that $F = ma$ *is* general and *does* apply to this case, she might well have questioned her own causal assumptions more thoroughly. About how that might have been managed, unfortunately, we have no data. More generally, although we will not argue it here, J's epistemological stance and strategies concerning the nature of physical knowledge seemed to interfere with many learning opportunities. This is the sort of knowledge, not specifically related to the concept in question, that we mentioned as potentially influencing conceptual development in Tenet of Accountability number 3. (See page 15.)

To sum up, in this situation J has been placed in a predicament. Her naive causal net leads to one determination – that the force of the hand is overcoming friction in moving a book – which is clear and incontrovertible to her. However, this directly conflicts with $F = ma$, which implies that when there are unbalanced forces, there must be an acceleration. After checking her coordination for acceleration and finding it adequate, J resolves the conflict by asserting that $F = ma$ does not apply to pushing books across tables. She will certainly coordinate, sometimes, with this critical equation. But in this instance, she has chosen her naive causal net. This episode contrasts with the others in that J knew, here, that she had two methods to coordinate. So she was forced into a more or less articulate choice.

How have we done

In this section, we briefly review what we have accomplished in comparison to the standards and expectations we set. Of course, in a single paper, one can't really expect to present a complete theory of concepts plus experimental demonstration. Theoretically, we are clearly short of computational precision. Empirically, we can't pretend to have shown that, for example, the difficulties we identify account for widespread, systematic problems in learning physics, despite evidence from J. But, we hope we have made some progress, reviewed below.

1. Standards of what counts as a concept

We haven't answered the question, 'What is a concept?' in general. However, we have said quite a lot about a particular kind of concept. A coordination class is the competence to see a particular class of information in the world. This core function is at least as specific as 'distinguishing members from non-members' in category research, and much better adapted to understanding how scientific concepts function in explanation and problem-solving. While it may be surprising that 'how one sees it in the world' is a sufficient characterization of a scientific concept, this is the hypothesis behind the definition of coordination class, and none of our empirical work so far is inconsistent with it.[14]

We have provided an articulation of the pieces of 'having a coordination class'. We have defined two structural components: the set of readout strategies and the causal net. If one claims to have identified a coordination class, one must further specify these other components.

In addition, we identified two performance functions of coordination classes: coordinating observations within a single situation (integration), and coordinating different features in different situations to find the same information (invariance).

One learns about a coordination class, in part, by identifying the range of invariance and the particular ways that people integrate.

In the case of quantities, we delimited a role for equations. Equations are not coordination classes, but one can use an equation to coordinate, to relate some observations to particular needed information. We suggested that specific quantitative use of equations is less important than qualitative uses of the connections evident in equations.

In evaluating how well we have done in distinguishing what are concepts (coordination classes) from what are not, it is important to remember that there is always an empirical component of such determination. Psychological reality is not a purely theoretical issue. Even with a perfectly precise and clear theory, cases may require extensive and difficult experiment to resolve. Before data such as we obtained with J, it may have been difficult to determine whether force would admit of a coordination class analysis. To take a case in point, we do not assume that the proposition 'force is a coordination class' is a priori valid. It requires the kind of empirical work introduced here to validate.

2. An account of different varieties of concepts

It seems completely clear that coordination classes are not any old concept. In particular, many category-like concepts are almost certainly not coordination classes. The concept 'dog' may be extremely important, but it is far from determining a coordination class. Most of the strategies for observing physical objects, or, more particularly, animals, overlap what it takes to see 'dogness'.[15]

Any particular description of structural or performance attributes constitutes a refinement of the broad category, coordination class. For example, we can imagine a type of coordination class that involves a localized causal net, perhaps embodied in an explicit definition. This type of coordination class might behave very differently from one with a diffuse causal net, like the many, loosely related p-prims in the naive sense of mechanism.

As we differentiate various versions of concept, each one covers less. That means, even if we understand coordination classes completely, other forms of conceptual change involving other kinds of concepts probably need to be understood. It is intriguing, if speculative, to note that all concepts may have a component of coordination to them – how does one identify exemplars and their characteristic attributes? On the other hand, coordination may not be definitive of the nature of other types of concepts, and, for example, learning trajectories and difficulties might be very different from coordination classes.[16]

3. Theoretical accounts of a variety of processes involved in concept use and conceptual change

Separate changes in the causal net and in readout strategies localize problems with conceptual change. In particular, we hypothesized that the causal net is where most difficulties with conceptual change in physics are. We can also consider separately how people achieve invariance and how they achieve integration. Learners may fail to have the relevant coordination class in either of these ways, and each failure may need to be specified as to context and reason for failure.

In terms of specific mechanisms and mental representation, this paper has been comparatively light. On the other hand, we made excellent contact with another theory – that of p-prims – which does attempt fairly explicit treatment of representation, origins and development of the intuitive knowledge hypothesized here to constitute the naive causal net. For details, see diSessa (1993).

4. The ability to match real-time data from student's thinking with the theoretical accounts of process described above

We have interpreted several episodes of clinical interviewing data in terms of coordination classes. In particular, J seems to have difficulty precisely with the central function of a coordination class. She just doesn't see forces and determine their values in the same way that a physicist does. In addition, we have (tentatively) localized some of her difficulties. Consistent with our general expectations (prediction a), J did not appear to have any shortage of readout strategies; she could see motion and feel the weight of an object in her hand. Instead, she had difficulties with her causal net. She showed that previously documented naive knowledge elements, p-prims, served to coordinate (prediction c). More specifically, equations as a means of coordinating may well be rejected in particular cases in favour of coordinating via p-prims. She coordinated force appropriately in some cases and not in others. J's problems also seemed all to be in the category of invariance, consistent with prediction b. Finally, consistent with prediction d, when J used equations, she nonetheless never computed a single number.

We haven't shown very much learning in these segments, although there appear to be good opportunities for J to learn. Tracing learning data with coordination class theory remains a project for future work.

Among our tenets of accountability and the various criticisms we brought to bear against prior work, this presentation of coordination classes probably fares worst with respect to the issues of whether there is any well-defined mental entity that is referred to by 'force' or 'acceleration'. In our defence, we note:

- Because of our systems approach, establishing integrity and relative independence of a coordination class may be among the most difficult empirical accomplishments.
- We did not claim or expect that naive or novice students have well-defined coordination classes that change during 'conceptual change'. It may be that only experts have coordination classes for these things; novices merely coordinate as best they can. On the other hand, we provide data that prior knowledge, p-prims, were implicated in J's coordination, exactly as predicted.
- Even if experts don't have coordination classes either, our data strongly suggests that *how* they coordinate, and how students come to that, is important and insightful of learning difficulties. The theory may not lose much if the assumption of integrity turns out to be false.
- Finally and most generally, we have presented details in our theory and empirical expectations that can be checked far beyond typical specifications of 'concept'. Those details allowed us to analyse J's use of force in real-time data, showing at least examples of phenomena predicted by the theory and even allowing us to note which aspects of the theory do not show in that

data. (For example, we have no explanation for the non-invariance of J's use of action and reaction). This is qualitatively different from assuming, a priori, that a concept is at issue and then using data in which no process details show to determine properties of the concept (e.g. whether people at various ages 'have it').

Summary

We have had two overarching goals in this paper. The first was to provide an overview and critique of research literature on the nature of concepts and conceptual change. The point is to sensitize readers to a plethora of difficulties and unclarities. Most notably, research has been strikingly inexplicit about what it means by concepts, thus begging the central question, "What changes in conceptual change?" In addition, we believe that researchers have uncritically applied the term concept to what may be very different entities – dog, force, or number. Underlying both of these difficulties is a theoretical imprecision that is insuperable.

Against this backdrop, we have sought to take modest positive steps. We outlined the beginnings of a theory of one type of concept that we call a coordination class, and which we believe to be important in learning science. Coordination classes, in a nutshell, are the complex set of ways that people use to read particular classes of information out of the world. We have tried: (a) to provide sufficient analysis so that what counts as a coordination class is clear enough to allow determination, at least in principle, of whether something is a coordination class or not; (b) to use that analysis to distinguish not only coordination classes from other concepts, but also to map a comprehensible variety of coordination classes based on, for example, parameters like the richness of the causal net and its relation to prior knowledge; (c) following the analyses above, to give some accounting of the processes involved in learning a coordination class, such as typical phenomena (e.g. lack of invariance) and more or less difficult aspects of learning; (d) to analyse real-time data illustrating all of the above.

Acknowledgements

This work was supported, in part, by grant number B-1393 to A. diSessa from the Spencer Foundation. Both authors express their appreciation. Writing was completed while the first author was a Fellow at the Center for Advanced Study in the Behavioral Sciences. Financial support for this was also provided by the Spencer Foundation, grant number 19940032. In addition, we gratefully acknowledge helpful comments by Reed Stevens, Marie-Françoise Legendre-Bergeron and two insightful anonymous reviewers.

Notes

1. We believe 'coalesce' and 'differentiate' are better understood as a priori descriptions of the behaviour of category-like concepts. They merely describe changes over time in the results of many 'Is this an X?' questions. In general, these terms will say essentially nothing about the structure of concepts or change in that structure.
2. We are not saying that one shouldn't compare novices and experts, only that expert ideas can't be the only reference point. Instead, if we want a theory of conceptual change, we need to be able to describe novice (and expert!) concepts *in their own*

terms. We want to describe in detail what students *do* do, not just what they don't. Empirical accounts later in this paper are dedicated to showing that we can do better than 'nondifferentiated protoconcept' in this regard.

3. Again, as in footnote 2, we want to characterize novice and expert concepts in their own terms. But comparison is also enlightening, especially in theoretical terms that apply equally (and insightfully) to characterize both novices and experts.
4. It is remotely possible that determining personality could reduce to determining membership in a finite set of categories, like 'good sport', etc. Given a specification of categorizing processes, this would provide one model for coordination. This does not appear to be the way scientific concepts work.
5. The situation is more complicated than 'confusing' duration with distance. In many circumstances, it is clear that these same children can distinguish time from distance, and even that, in this questioning, they believe they are talking about issues such as 'before lunch' or 'after lunch', rather than stopping 'in front of' or 'after' (in the spatial sense) the other train. Indeed, understanding when people use one indicator or another to determine some information is a canonical task of understanding their coordination knowledge.
6. Although beyond the scope of this paper, it is useful to distinguish *relevance* from *strategy* as stages in the acquisition of an inference in a causal net. Many examples in the developmental literature show children know or conjecture relevance of some factor before they have strategic competence to use that relevance (Metz 1993).
7. In still other circumstances, it might mean that conflicting information requires yet other perspectives to decide.
8. Historical note: The original theoretical motivation for coordination classes appeared in diSessa (1994). In diSessa (1991) the conjecture that quantities may be coordination classes was elaborated. The theory given here is slightly different in some respects, to suit present circumstances.
9. An additional subspecification is possible. Achieving invariance involves appropriate *span* (covering readout in the breadth of situations necessary) and *alignment* (making sure that, in different situations, the same information is read out). This is more detail than we will attempt to elaborate in this paper.
10. This is a perfectly well-formed question, and all its presuppositions are valid. Heavier objects are pulled harder by gravity. The answer to why heavier objects don't fall faster is that heavier objects also 'resist' motion in an amount proportional to the extra force they feel. Hence, increased gravitational force is exactly compensated by increased inertial resistance.
11. Gravitational acceleration is the same on all bodies at the same distance from the center of the earth. However, gravitational force varies with the mass of the object according to $F = mg$, where g is the constant gravitational acceleration.
12. An alternate interpretation that does not involve *contact conveys motion* is that the paper is just 'sticking' to the book. Either way, the coordination implications are identical.
13. In fact, a correct analysis is that these forces are equal, even though they are not an action/reaction pair.
14. The causal net, by itself, may make a more intuitive definition of a concept – the set of inferences that one can make with and about a quantity. With that stipulation, the definition of coordination class used in this paper might work equally well because it happens that every aspect of a causal net may show up in some act of determining. And the paper's definition might have methodological advantages in defining a focused but sufficient set of inferences, those involved in empirical determining.
15. On the other hand, if one uses behaviour as a key indicator, a causal net concerning animal behaviour and how to recognize different characteristic behaviours may be implicated. However, this causal net is unlikely to be restricted to dogs.
16. This observation highlights an implicit assumption that we have made. That is, the information-gathering strategies associated with the 'concept' in question are both relatively integrated and characteristic of the concept. If gathering the relevant information uses an *ad hoc* collection of strategies that are variously characteristic of other

concepts as well, then we can speak of coordination, but coordination class may be inappropriate.

References

Carey, S. (1988) Reorganization of knowledge in the course of acquisition. In S. Strauss (ed.), *Ontogeny, Phylogeny, and Historical Development* (Norwood, NJ: Ablex), 1–27.

Carey, S. and Spelke, E. (1994) Domain-specific knowledge and conceptual change. In L. A. Hirschfeld and S. A. Gelman (eds), *Mapping the Mind: Domain Specificity in Cognition and Culture* (Cambridge: Cambridge University Press).

Clement, J. (1982) Students' preconceptions in introductory physics. *American Journal of Physics,* 50, 66–70.

diSessa, A. A. (1983) Phenomenology and the evolution of intuition. In D. Gentner and A. Stevens (eds), *Mental Models* (Hillsdale, NJ: Erlbaum), 15–33.

diSessa, A. A. (1991) Epistemological micromodels: The case of coordination and quantities. In J. Montangero and A. Tryphon (eds), *Phychologie génétique et sciences cognitives.* (Volume from the Eleventh Advanced Course) (Geneva: Archives Jean Piaget), 169–194.

diSessa, A. A. (1993) Toward an epistemology of physics. *Cognition and Instruction, 10* (2–3), 105–225; Responses to commentary, 261–280 (*Cognition and Instruction*, Monograph no. 1).

diSessa, A. A. (1994) Speculations on the foundations of knowledge and intelligence. In D. Tirosh (ed.), *Implicit and Explicit Knowledge: An Educational Approach* (Norwood, NJ: Ablex), 1–54.

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* (Oxford: Blackwell), 709–730.

Fodor, J. A. (1975) *The Language of Thought* (Cambridge, MA: Harvard University Press).

Gentner, D., Brem, S., Ferguson, R. W., Markman, A. B., Levidow, B. B., Wolff, P. and Forbus, K. D. (1997) Analogical reasoning and conceptual change: a case study of Johannes Kepler. *Journal of the Learning Sciences,* 6 (1), 3–40.

Gleitman, L. R., Armstrong, S. L. and Gleitman, H. (1983) On doubting the concept 'concept'. In S. E. Kofsky (ed.), *New Trends in Conceptual Representation: Challenges to Piaget's Theory?* (Hillsdale, NJ: Erlbaum), 87–109.

Lakoff, G. (1987) *Women, Fire, and Dangerous Things* (Chicago: University of Chicago Press).

Magnusson, S. J., Templin, M. and Boyle, R. A. (1997) Dynamic science assessment: a new approach for investigating conceptual change. *Journal of the Learning Sciences,* 6 (1), 91–142.

McCloskey, M. (1983) Naive theories of motion. In D. Gentner and A. Stevens (eds), *Mental Models* (Hillsdale, NJ: Lawrence Erlbaum Associates), 299–324.

Medin, D. (1989) Concepts and conceptual structure. *American Psychologist*, 44 (12), 1469–1481.

Metz, K. E. (1993) Preschoolers' developing knowledge of the pan balance: from new representation to transformed problem solving. *Cognition and Instruction,* 11 (1), 31–93.

Posner, G. J., Strike, K. A., Hewson, P. W. and A., G. W. (1982) Accommodation of a scientific conceptions: toward a theory of conceptual change. *Science Education*, 66, 211–227.

Rosch, E. (1994) Categorization. In V. S. Ramachandran (ed.), *Encyclopedia of Human Behavior* (vol. 1) (San Diego, CA: Academic Press), 513–523.

Rumelhart, D., Smolensky, P., Hinton, G. and McClelland, J. (1986) Schemata and sequential thought processes in PDP models. In J. McClelland and D: Rumelhart (eds), *Parallel Distributed Processing: Explorating in the Microstructure of Cognition* (vol. 2) (Cambridge, MA: MIT Press/Bradford Books), 7–57.

Sherin, B. L. (1996) The symbolic basis of physical intuition: A study of two symbol systems in physics instruction. Unpublished doctoral dissertation, University of California, Berkeley.

THAGARD, P. (1992) *Conceptual Revolutions* (Princeton, NJ: Princeton University Press).

TIROSH, D. (1994) *Implicit and Explicit Knowledge: An Educational Approach* (Norwood, NJ: Ablex).

TROWBRIDGE, D. E. and MCDERMOTT, L. C. (1980) Investigation of student understanding of the concept of velocity in one dimension. *American Journal of Physics,* 48 (12), 1020–1028.

TROWBRIDGE, D. E. and MCDERMOTT, L. C. (1981) Investigation of student understanding of the concept of acceleration in one dimension. *American Journal of Physics,* 49 (3), 242-253.

WISER, M. (1988) The differentiation of heat and temperature: history of science and novice-expert shift. In S. Strauss (ed.), *Ontogeny, Phylogeny, and Historical Development* (Norwood, NJ: Ablex), 28–48.