DOCUMENT RESUME

ED 355 028 PS 021 215

AUTHOR	Halford, Graeme S.
TITLE	Experience and Processing Capacity in Cognitive Development: A PDP Approach.
PUB DATE	Mar 93
NOTE	30p.; Paper presented at the Annual Conference of the Society for Research in Child Development (New Orleans, LA, March 25-28, 1993).
PUB TYPE	Speeches/Conference Papers (150) -- Reports - Descriptive (141)
EDRS PRICE	MF01/PC02 Plus Postage.
DESCRIPTORS	*Cognitive Development; Cognitive Mapping; *Cognitive Processes; *Concept Formation; Models; Schemata (Cognition); Thinking Skills; Young Children
IDENTIFIERS	Chunking; Knowledge Acquisition; *Parallel Processes; *Process Models; Transitive Inferences

ABSTRACT

Cognitive development is driven by experience, but is mediated by domain general processes, which include learning, induction, and analogy. The concepts children understand, and the strategies they develop based on that understanding, depend on the complexity of the representation they can construct. Conceptual complexity can be defined in terms of the number of independent dimensions that need to be represented. Parallel Distributed Processing (PDP) models of the way in which information is represented help to explain why the number of dimensions that can be processed in parallel is limited. This explanation leads to a new definition of processing capacity, which seems to account for many phenomena, including some that have traditionally been attributed to stages. The definition implies that cognitive development is an interaction of domain specific and domain general processes. An overarching goal of research is to define the nature of this interaction. An important result of the interaction is the growth in the capacity to represent concepts of increasing structural complexity. This capacity to represent information controls the concepts that are acquired as a function of experience. (MM)

Experience and Processing Capacity in Cognitive Development: A PDP Approach

Graeme S. Halford

University of Queensland

Paper presented to conference of the Society for Research in Child Development, New Orleans, La., March 25-28, 1993

Author's address: Graeme Halford, Psychology Department, University of Queensland, 4072, Australia. Fax: 61 7 365 4466.

Email: gsh@psych.psy.uq.oz.au.

Abstract

Tnis model proposes that concepts children understand, and the strategies they develop, depend on the complexity of their representations, defined in terms of number of independent dimensions. A PDP model explains why concepts of high dimensionality impose high processing loads, and suggests that representations become differentiated with age into more vectors, so more dimensions can be represented in parallel. It is suggested that cognitive development also depends on induction of schemas that can be used as mental models, and can guide development of strategies. Processing loads can be reduced by conceptual chunking, and by acquisition of serial processing strategies.

Experience and Processing Capacity in Cognitive Development: A PDP Approach

Graeme S. Halford

The theory of cognitive development which I propose can be summarized in the following propositions:

1. Cognitive development is experience driven, but is mediated by domain general processes, which include learning, induction, and analogy.

2. The concepts children understand, and the strategies they develop based on that understanding, depend on the complexity of the representations they can construct.

3. Conceptual complexity can be defined in terms of the number of independent dimensions that need to be represented. Parallel Distributed Processing models of the way information is represented help to explain why the number of dimensions that can be processed in parallel is limited. This leads to a new definition of processing capacity, which appears capable of accounting for many phenomena, including some that have been attributed traditionally to stages.

These propositions imply that cognitive development is an interaction of domain specific and domain general processes, and so an overarching goal of research is to define the nature of this interaction. An important component of it is the growth in capacity to represent concepts of increasing structural complexity. This capacity to represent information controls the concepts that are acquired as a function of experience.

The key methodological features of the approach are:

1. Explicit definition of the nature of processing capacity. This has been done by developing a parallel distributed processing (PDP) model of the way information is processed.

2. Objective assignment of concepts to levels of complexity, by using computer models of the representations required.

3. Experimental demonstration of capacity limitations, using the easy-to-hard paradigm (Hunt & Lansman, 1982). This entails using a secondary task indicator to determine whether the criterion task is resource dependent. It is not subject to the ambiguities of some alternate dual-task paradigms.

4. Empirical assessment of the nature of processes entailed in cognitive tasks. Competence is only interpretable if we know how a task is performed.

One major goal of cognitive development theory is to predict children's cognitive capabilities. Part of this is to explain why some tasks are typically mastered at later ages than others. Much debate has concerned whether certain pivotal tasks, such as transitive inference and class inclusion, are normally mastered by children under five years.

Our desire to accelerate cognitive development has often caused us to try and explain these difficulties away, on the grounds that they result from flawed tests, or inadequate knowledge. However, many of the claims that children succeed with alternative tests are flawed due to either false positives (e.g. reporting chance results as success), or failure to consider alternative bases for the performance (Halford, 1989). Furthermore, many of the improvements have been with children *over* five years, and therfore do not account for the finding that these tasks are specially difficult for children below this age.

Another problem is that lack of process models makes it difficult to define test validity, resulting in circularity; "good" tests tend to be those that children pass. Therefore it seems appropriate to conclude that while a lot of important causes of failure have been discovered, there are still sources of difficulty for young children that remain to be explained.
A theory of cognitive development should be able to account for this. I will address the problem by considering the case of transitivity.

Transitive inference

Consider a transitive inference task such as that shown in figure 1.

Insert figure 1 here

There is a reasonable concensus in the literature that such tasks are performed by arranging the terms in order (Sternberg, 1980; Trabasso, 1977; Thayer & Collyer, 1978). Given the analogical character of human reasoning, we can conceptualize this process as mapping the premises into a schema, which is used as an analog.

A common ordering schema, the left-right or top-down arrangement, is used as an analog. In effect it serves as a kind of template, or "mental model" for imposing order on the premises. Once the premises are ordered in this way transitive inferences are easily made by accessing the ordered representation; e.g. we can easily see that John is fairer than Tom.

There is some difficulty in performing the mapping however. This is because both premises must be processed to map any premise term into a slot in the ordering schema; e.g. we need both premises to know that John must go in first position. This decision requires cognitive effort, and it illustrates the operation of processing load in the theory.

A mental model of the task, in the form of an ordering schema, can be used to guide the development of strategies. We have developed a self-modifying production system model which acquires strategies through experience, guided by a specific example of an ordered set which is used as an analog, as shown in the previous figure. Once such a strategy is developed there is no further need for analogical reasoning, except where the strategy must be modified, or transferred to a new domain.

Once developed the strategies become autonomous in most familiar applications, but their initial development depends on ability to represent the structure of the concept. That is, children cannot autonomously develop their own strategies for transitive inference unless they can represent the concept of order adequately. This means they must have a mental model of an ordered set of at least three elements, with an asymmetric, transitive, binary relation between them. This is an instance of a principle that I believe to be of general validity, which is that autonomous strategy development depends on adequate representation of a concept of the task.

The capacity required to represent a concept will depend on its structural complexity. In order to explore the reasons for this, we must consider the way concepts are represented. We have gained considerable

insights into the problem by considering how to represent concepts in parallel distributed processing (PDP) architectures.

PDP implications for processing capacity

We will examine how concepts of varying complexities are represented. The representation of a binary relation, such as LARGER THAN is shown in figure 2.

Insert figure 2 here

A vector is used to represent the predicate, LARGER THAN, and another vector is used to represent each argument. In this example, there is a vector representing arguments *elephant* and *dog*. The predicate-argument binding, that is, the fact that elephant is larger than a dog, is represented by the tensor product of the three vectors, as shown in figure 2. Because LARGER-THAN is a binary relation, with two arguments, it is represented by a rank 3 tensor product, that is, one with three vectors.

Dimensionality

The rank of a tensor product can be shown to relate to a conceptual complexity metric originally devised by Halford & Wilson (1980). The complexity of a concept is defined in terms of its *dimensionality*; i.e. the number of independent items of information required for the computations the concept entails. Dimensionality is quite similar to the idea of degrees of freedom. The idea is that all aspects of a task that enter into a particular computation must be represented in parallel, and aspects which are free to vary independently must be represented as separate dimensions. The number of vectors required for a representation is one more than the number of dimensions. Hence a binary relation, which is 2 dimensional, is represented by a tensor product of rank 3.

Insert figure 3 here

There are four levels of relations, unary, binary, ternary, and quaternary, represented by tensor products. Each level of relation corresponds to a level of dimensionality, because each argument of a relation corresponds to an independent source of variation. Higher dimensional representations permit more complex associations to be computed; e.g. with a ternary relation R(a,b,c), we can compute how *a*

varies as a function of *b*, how *a* varies as a function of *c*, how *b* varies as a function of *c*, how *a* varies as a function of *b* and *c*, etc. With a binary relation only one type of association, *a* as a function of *b*, is possible.

It has been shown elsewhere (Halford, 1993) that the four levels of dimensionality bear a broad correspondence to Piaget's (1950) major stages, as shown. It has also been shown that representations of higher dimensionality impose higher processing loads (Halford et al., 1986; Maybery et al., 1986).

A tensor product of higher rank imposes a higher computational cost, because the number of tensor product units increases exponentially with the number of vectors, and the number of connections increases accordingly. The PDP model therefore provides a natural basis for the increase in processing load that has been observed empirically.

Insert figure 4 here

Now we will consider the levels of representation in more detail. Unary relations include simple categories, defined by one attribute, such as the category of large things. They also include categories defined by a collection of attributes that can be represented as a single chunk, such as the category of dogs. One vector (shown vertically) would represent the category label DOG. The other vector would represent the instances. Representations of different dogs would be superimposed on this set of units. Thus vectors representing each known dog would be superimposed, so the resulting vector would represent the central tendency of the person's experience of dogs. It would represent the person's prototype dog. However the representations of the individual dogs can still be recovered. Questions such as "are chihuahuas dogs", or "tell me the dogs you know" can be answered by accessing the representation. Note that the representation is one dimensional because if one component is known, the other is determined. Thus if the argument vector represents a labrador, the other vector must be "dog". Similarly, if the predicate vector represents "dog", the argument vector must represent one or more dogs.

Unary relations also include ability to represent variable-constant bindings. The well-known A not-B error in infant object constancy research can be thought of as requiring ability to treat hiding place as a

variable. That is, when an infant has repeatedly retrieved an object from hiding place A, then continues to search for it at A despite having just seen it hidden at B, the infant is treating the hiding place as a constant. However if hiding place were represented as a variable this perseveration would be overcome. This requires a rank 2 tensor to represent the binding between the object and its location.

The fact that children can represent category membership at about one year, and the A not-B error disappears about the same time, is consistent with ability to represent rank-2 tensor products at that age. Thus Piaget's preconceptual stage appears to require this level of representation.

At the next level binary relations, and univariate functions can be represented. These are all 2 dimensional concepts (given any two components, the third is determined), and they entail tensor products of rank 3. Based on an assessment of the cognitive development literature Halford (1982, 1993) suggests they develop at approximately two years of age. They correspond to Piaget's observation that in the intuitive stage children process one binary relation at a time.

At the next level concepts based on ternary relations, binary operations, and bivariate functions, are represented. These are 3-dimensional, and require tensor products of rank 4. Well known examples include transitivity and class inclusion, but there are many other concepts that belong to this level, including conditional discrimination, the transverse pattern task, the negative pattern task, dimension checking in blank trials task, and many more (Halford, 1993). The familiar binary operations of addition and subtraction belong to this level. One vector represents the operation (+ or x) while two others represent the addends (multiplicands), and the fourth vector represents the sum (product). Note that if you know three of these, the fourth is determined; e.g. if you know the numbers are 2,3,5 you know the operation is addition; if you know the numbers 2, ?, 5, and the operation is addition, you know the missing number is 3, and so on. (Readers interested in PDP might note that there is no catastrophic forgetting when addition and multiplication are superimposed on a rank 4 tensor product).

More complex concepts are represented by structures with more vectors. The representation of transitivity requires a rank 4 tensor product, as shown in figure 5.

Insert figure 5 here.

Given that transitive inferences are made by organizing premise information into an ordered set of three elements, as shown in Figure 1, the core of the transitivity concept is a ternary relation. That is, transitivity is a relation with three arguments, corresponding to *a,b,c* or top, middle, bottom, depending on the particular instantiation.

Consequently, it has to be represented by a tensor product of higher rank than a binary relation, such as LARGER-THAN.

Insert figure 6 here

Class inclusion will be represented as shown in figure 9. There is a vector representing the concept, and three vectors representing its arguments, the superordinate, the first subordinate, and its complement.

All of these tasks are performed by about five years of age, but cause considerable difficulty below this age. In a broad sense, this level of processing corresponds to Piaget's concrete operational stage, which can be conceptualized as ability to process binary operations, or compositions of binary relations (Halford, 1982, 1993; Sheppard, 1978).

At the fourth level concepts based on quaternary relations, and compositions of binary operations, can be represented. These include understanding proportion and concepts such as distributivity, that are based on compositions of binary opersations. In a broad sense this level of processing corresponds to Piaget's formal operations stage, which entails relations between binary operations (Halford, 1993). The representation of proportion is shown in figure 7.

Insert figure 7

The representation of the balance scale is shown in figure 8.

Insert figure 8

A PDP model has been developed which shows that these representations are capable of carrying out the computations relevant to each concept.

Chunks and dimensions

We have argued (Halford 1993; Halford et al., in press) that the number of dimensions can be identified with the number of chunks. An attribute on a dimension, like a chunk (Miller', 1956) is an independent unit of information that can vary in size. For example letters, digits, and words vary considerably in the amount of information they contain, but each is a chunk because it is an independent unit. Similarly, an attribute on a dimension can represent varying amounts of information, and attributes on different dimensions are independent.

Working memory research suggests that the number of chunks that adults process in parallel is about four (Schneider & Detweiler, 1987; Halford et al., in press). Therefore we would predict that adults can process a maximum of four dimensions in parallel. We have also produced some empirical evidence supporting this prediction (Halford et al., in press). This would mean that the most complex tensor product representations that can be processed would be rank 5, i.e. with five vectors.

Age and dimensionality ... presentations

This argument enables us to reformu..te the longstanding question of whether processing capacity changes with age. The question becomes, not whether overall capacity changes, but whether representations become more differentiated so that tensor products of higher rank can be processed. This would mean that concepts of higher dimensionality would be represented, enabling higher-order relations to be understood.

Our developmental work suggests that the dimensionality of representations does increase with age: Children can represent one dimension in parallel at a median age of one year, two dimensions at 2 years, 3 at five years, and 4 at 11 years. There are indications that this factor is at least partly maturational (Halford, 1983, Chapter 3), but more data are needed.

Chunking and segmentation

Concepts more complex than four dimensions can be processed by either *conceptual chunking* or *segmentation*. Conceptual chunking entails recoding concepts of higher dimensionality into fewer dimensions, most commonly into one dimension; i.e. it entails reducing multiple chunks to a single chunk. An example would be the concept of velocity, defined as $v = s/t$. It is 3 dimensional, and requires a tensor product of rank 4. However it is also possible to think of velocity as a single dimension, such as the position of a pointer on a dial.

Insert figure 9

When velocity is chunked as a single dimension, it can be represented by a single vector, and combined with up to three other dimensions. Thus velocity can now be used to define acceleration, $a = (v_2 - v_1)t^{-1}$.

Acceleration in turn can be chunked, and combined with up to three other dimensions. Thus force, $F = ma$ can be defined as the product of mass and acceleration. Conceptual chunking enables us to bootstrap our way up to concepts of higher and higher dimensionality, without exceeding the number of dimensions that can be processed in parallel.

If the number of dimensions can be reduced by chunking, is the limit in processing capacity meaningful? It is meaningful because when representations are chunked, we lose the ability to recognize relations within the representation. When velocity is represented as a single dimension, we can no longer compute the way velocity changes as a function of time or distance, or both. Similarly, we cannot compute what happens to time if distance is held constant, and velocity varies, and so on. This example illustrates the point that any computation requires a minimum number of dimensions to be represented.

Segmentation entails developing serial processing strategies. In this case tasks are segmented into steps, each of which is small enough not to exceed the capacity to process information. Only that part of a concept that is the focus of attention is represented at any one time.

However autonomous development of strategies requires a concept of the task, and this requires that there be sufficient processing capacity to represent the dimensions of the concept. Where children cannot represent sufficient dimensions for a particular concept, they will default to lower dimensionality representations, which will result in strategies that are partly correct, but which lead to errors on some variants of the task (Halford et al., 1992).

Capacity overload

A child (or adult) who was unable to construct a representation of the dimensionality required for a task would have three options:

1. Chunk the task to a lower dimensional representation. However this requires the ability to "unpack" the chunks to represent the relations they contain, and it also depends on previous experience with mapping components into chunks.

2. The task can be segmented into smaller components that are processed serially. However this requires a strategy the development of which depends on ability to represent the concept of the task, so there is a catch 22 involved here. This difficulty can be overcome by instruction, but generalization will be limited if the child cannot represent the task concept.

3. The child can default to a lower level representation. This typically results in performance which is partly correct, but will be invalid on telltale aspects of the tasks that depend on representing more complex relations.

The performance of a child who cannot construct representations of adequate dimensionality is analogous to analysing (say) a three-factor experiment as a series of two-way ANOVAS. Most findings will be a correct account of the data, just as the hypothetical child's performance will be mostly correct. There will be however, at least in certain telltale cases, higher order interactions that will be missed. Similarly, the child who deals with an N-dimensional concept using representations of dimensionality less than N is really looking at the task through restricted

windows. Sooner or later telltale performances will occur which show that the representation was not really adequate.

Much of controversy that has dogged cognitive development may be attributable to this situation. Advocates of precocious development can always point to those aspects of young children's performance that appear adequate. However advoctes of capacity limitations can point to what they regard as telltale failures on more complex features of the tasks. Resolution of this polemic depends on more precise definition of competence in each domain.

Learning, and strategy development

If we accept that knowledge acquisition is a major component of cognitive development, it follows that learning, defined as acquisition of knowledge through experience, must play a significant role. The conspicuous lack of attention to the role of learning is probably because it is associated with behavioristic learning theories which have not been found to offer many solutions. However there are contemporary learning theories which do have the potential to explain how children acquire important concepts, and are worthy of further study by cognitive developmentalists.

A theory of learning is needed to explain how children acquire knowledge about the structure of the world. A reinterpretation of some established learning phenomena, including classical conditioning (Rescorla, 1988) and discrimination learning (Halford, 1993, Chapter 4) shows that humans and (other) animals possess very basic and effective learning mechanisms for this purpose. Theories of this process have been proposed by Holland et al. (1986) and by Holyoak, Koh & Nisbett (1989). Furthermore PDP theory provides powerful explanations for our ability to extract regularities from experiences which include a lot of randomness.

The second aspect of learning is acquisition of skills and strategies. There are well-substantiated computational models of skill acquisition (Anderson, 1987) which can be applied to showing how children acquire reasoning strategies. One such model (Halford, et al. 1992; Halford, er al. on contract) shows how transitive inference strategies can be acquired. These models recognize the active, constructive role of the child in

building its own knowledge base, and are a far cry from passive associationistic theories of the past.

Domain-General versus domain-specific acquisitions

Processing capacity, defined as the number of dimensions that can be represented in parallel, is a domain-general factor. It would affect any performance which depended on central representations. This would include all strategies and cognitive skills which develop under the guidance of a concept of the task. Most intellectual activities such as reasoning, mathematics and understanding of concepts, would be subject to this factor. We have not yet attempted to model language processing explicitly in this architecture, but there are indications that it also would be affected in this way. Specifically, no more than four dimensions would be processed in parallel. The difficulty of understanding complex centre-embedded sentences appears to be amenable to explanations in these terms; the sentence "The boy the man the girl saw met slept" exceeds human processing capacity because it requires five dimensions to be processed in parallel.

Learning, induction, and the mechanisms underlying strategy development, such as analogy and means-end analysis (the "weak" methods) are domain general, in that they appear to operate with more or less equivalent efficiency indendently of domain. However experience necessarily occurs within some domain. Given that cognitive development is experience driven, it will therefore be domain dependent. This means that domain-general factors which relate to the core cognitive processes must interact with domain-specific experience to produce the cognitive skills and concepts that children acquire.

Cognitive growth

Cognitive growth depends therefore on four main factors:

The first is *learning and induction*, which enables the child to build up an extremely rich store of world knowledge. This is the "raw material" of the schemas which can be used as mental models in reasoning and problem solving.

The second factor is *conceptual chunking*, which entails recoding representations into fewer vectors, so they can be combined into more complex representations, without overloading processing capacity.

The third factor is the development of *serial processing strategies* which permit tasks to be performed in smaller steps, timesharing the available representational capacity.

The fourth factor is the *development of ability to represent concepts of higher dimensionality*. The first three factors are essentially experiential, but the fourth is probably at least partly maturational. The actual mechanism is not yet known, but it probably entails differentiating distributed representations into more vectors. This entails rearranging the connections, to make the representations equivalent to higher rank tensor products. It would not increase overall processing capacity, but would enable higher orders of relationship to be represented.

The type of change that is envisaged here is analogous to splitting an experimental design into more independent variables. The total number of conditions represented might not change, but the orders of interaction that can occur do change; e.g. if we take a two-way ANOVA with four levels of one factor and two levels of another, and convert it into a three factor design with two levels on each factor, we still have the same number of conditions (8), but now we have added three-way interactions. Thus the most important change is in the orders of relations that can be represented. Similarly, growth in processing capacity through development is more likely to mean that higher order relations can be represented, rather than that more information can be stored. This ability to represent more dimensions in parallel enables children to conceptualize tasks more adequately, and thereby to construct more effective strategies

References

Anderson, J.R. (1987). Skill acquisition: compilation of weak-method problem solutions. Psychological Review, 94,192-210.

Halford, G.S. (1982). The development of thought. Hillsdale, N.J.: Erlbaum.

Halford, G.S. (1989) Reflections on 25 years of Piagetian cognitive developmental psychology, 1963-1988. Human Development, 32, 325-387.

Halford, G.S. (1993). Children's understanding: The development of mental models. Hillsdale, N.J.: Erlbaum.

Halford, G.S., Maybery, M.T., & Bain, J.D. (1986). Capacity limitations in children's reasoning: A dual task approach. Child Development, 57, 616-627.

Halford, G.S., Smith, S.B., Maybery, M.T., Stewart, J.E.M., & Dickson, J.C. (1992) A computer simulation model of acquisition of transitive inference. Technical Report 92/1, Centre for Human Information Processing and Problem Solving, University of Queensland, Australia.

Halford, G.S., Smith, S.B., Dickson, J.C., Maybery, M.T., Kelly, M. E., Bain, J. D. & Stewart, J.E.M. (on contract). Modelling the development of reasoning strategies: The role of analogy, knowledge and capacity. In Developing Cognitive Competence: New Approaches to Process Modelling. Hillsdale, NJ: Erlbaum.

Halford, G.S. & Wilson, W.H. (1980) A category theory approach to cognitive development. Cognitive Psychology, 12, 356-411.

Halford, G.S., Wilson, W.H., Guo, J., Gayler, R., Wiles, J., & Stewart, J.E.M. (in press). Connectionist implications for processing capacity limitations in analogies. In K.J. Holyoak & J. Barnden (Eds.), Advances in connectionist and neural computation theory, Vol. 2: Analogical Connections. Norwood, N.J.: Ablex.

Holland, J.H., Holyoak, K.J., Nisbett, R.E., & Thagard, P.R. (1986). Induction: Processes of inference, learning and discovery. Cambridge, MA: MIT Press.

Holyoak, K. J., Koh, K., & Nisbett, R. E. (1989). A theory of conditioning: Inductive learning within rule-based hierarchies. Psychological Review, 96(2), 315-340.

Hunt, E., & Lansman, M. (1982) Individual differences in attention. In R.J. Sternberg (Ed.) Advances in the psychology of human intelligence. (vol. 1., pp. 207-254). Hillsdale, N.J.: Erlbaum.

Maybery, M.T., Bain, J.D., & Halford, G.S. (1986). Information processing demands of transitive inference. Journal of Experimental Psychology: Learning, Memory and Cognition, 12, 600-613.

Miller, G.A. (1956) The magical number seven, plus or minus two: Some limits on our capacity for processing information. Psychological review, 63, 81-97.

Piaget, J. (1950). The psychology of intelligence. (M. Piercy & D.E. Berlyne, Trans.) London: Routledge & Kegan Paul, (Original work published 1947).

Rescorla, R.A. (1988). Pavlovian conditioning: it's not what you think it is. American Psychologist, 43, 151-160.

Sheppard, J.L. (1978). A structural analysis of concrete operations. In J.A. Keats, K.F. Collis, & G.S. Halford, (Eds.) Cognitive Development: Research based on a neo-Piagetian approach. London: Wiley.

Schneider, W., & Detweiler, M. (1987). A connectionist/control architecture for working memory. The psychology of learning and motivation, 21, 53-119.

Sternberg, R.J. (1980). The development of linear syllogistic reasoning. Journal of Experimental child psychology, 29, 340-356.

Thayer, E.S. & Collyer, C.E. (1978). The development of transitive inference: A review of recent approaches. Psychological Bulletin, 85, 1327-1343.

Trabasso, T. (1977). The role of memory as a system in making transitive inferences. In R.V. Kail, Jr. & J.W. Hagen, Perspectives on the development of memory and cognition, (pp. 333-356). Hillsdale, NJ: Erlbaum.

Figure Captions

Figure 1. Transitive inference problem mapped into an ordering schema.

Figure 2. Tensor product representation of predicate-argument binding.

Figure 3. Four levels of relations, with dimensionality, tensor product representation, and equivalent Piagetian stage.

Figure 4. Tensor product representation of Dog category.

Figure 5. Tensor product representation of transitivity.

Figure 6. Tensor product representation of class inclusion.

Figure 7. Tensor product representation of proportion.

Figure 8. Tensor product representation of balance scale.

Figure 9. Unchunked and chunked representation of velocity concept.

Premises: Peter is fairer than Tom
John is fairer than Peter

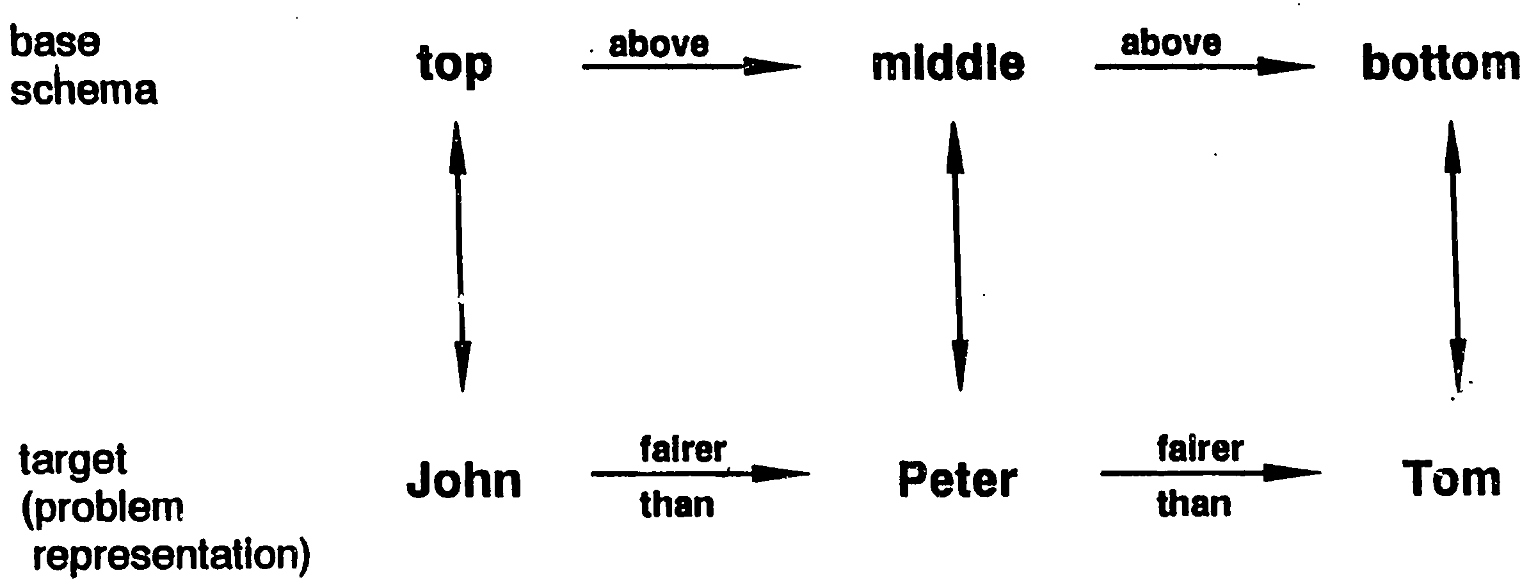

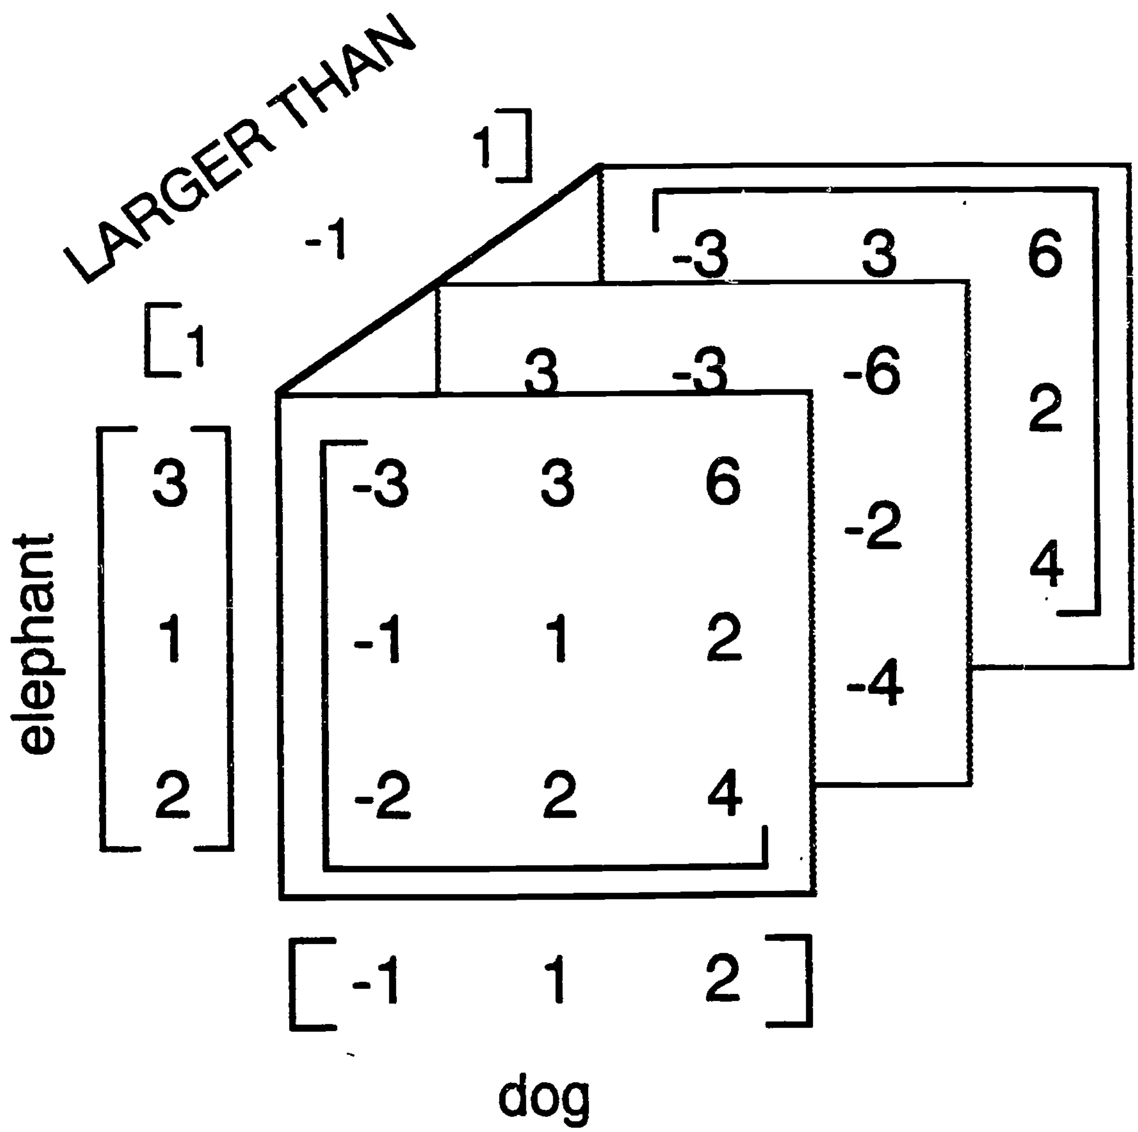
LARGER THAN
elephant
dog

Piagetian Stage	Formal Specification	Dimensionality	PDP Implementation
Preconceptual	Unary relations	1	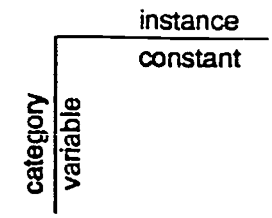
Intuitive	Binary relations, univariate functions.	2	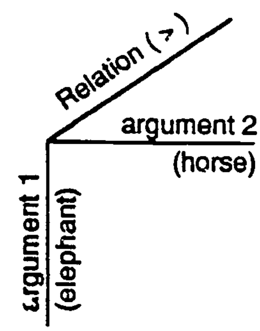
Concrete Operational	Ternary relations, binary operations, bivariate functions.	3	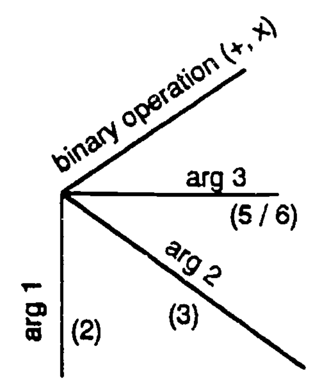
Formal Operations	Quaternary relations, compositions of binary operations.	4	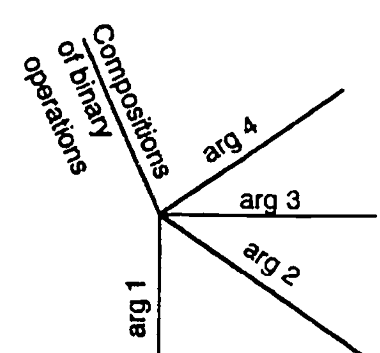

$2(2+3) = 10$

DOG

Fido
Lady
Stella
Cleo
Smoky

rank-4 tensor product
representing transitivity

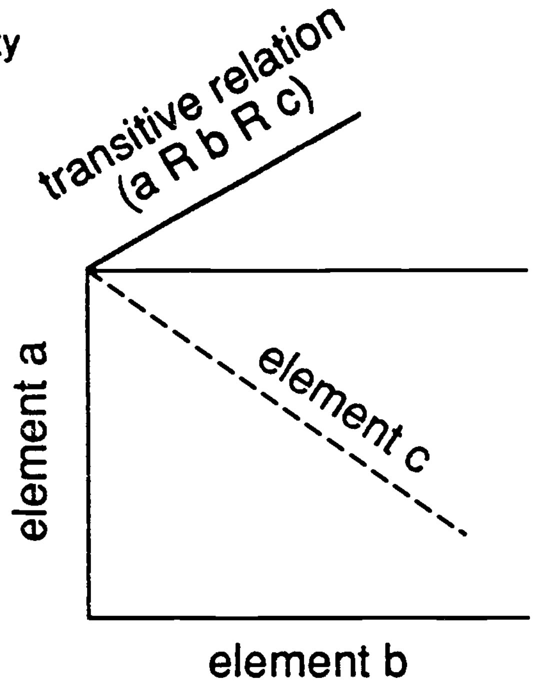

rank-3 tensor products
representing premises

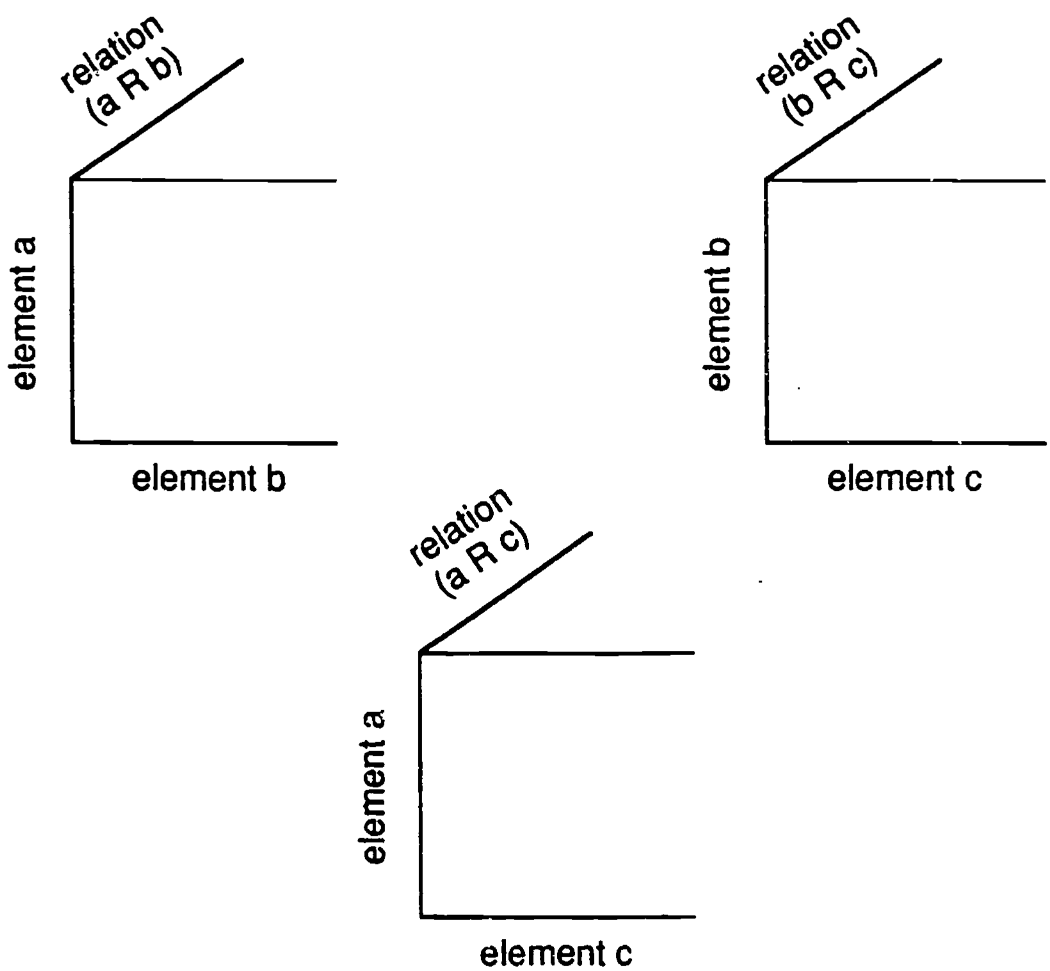

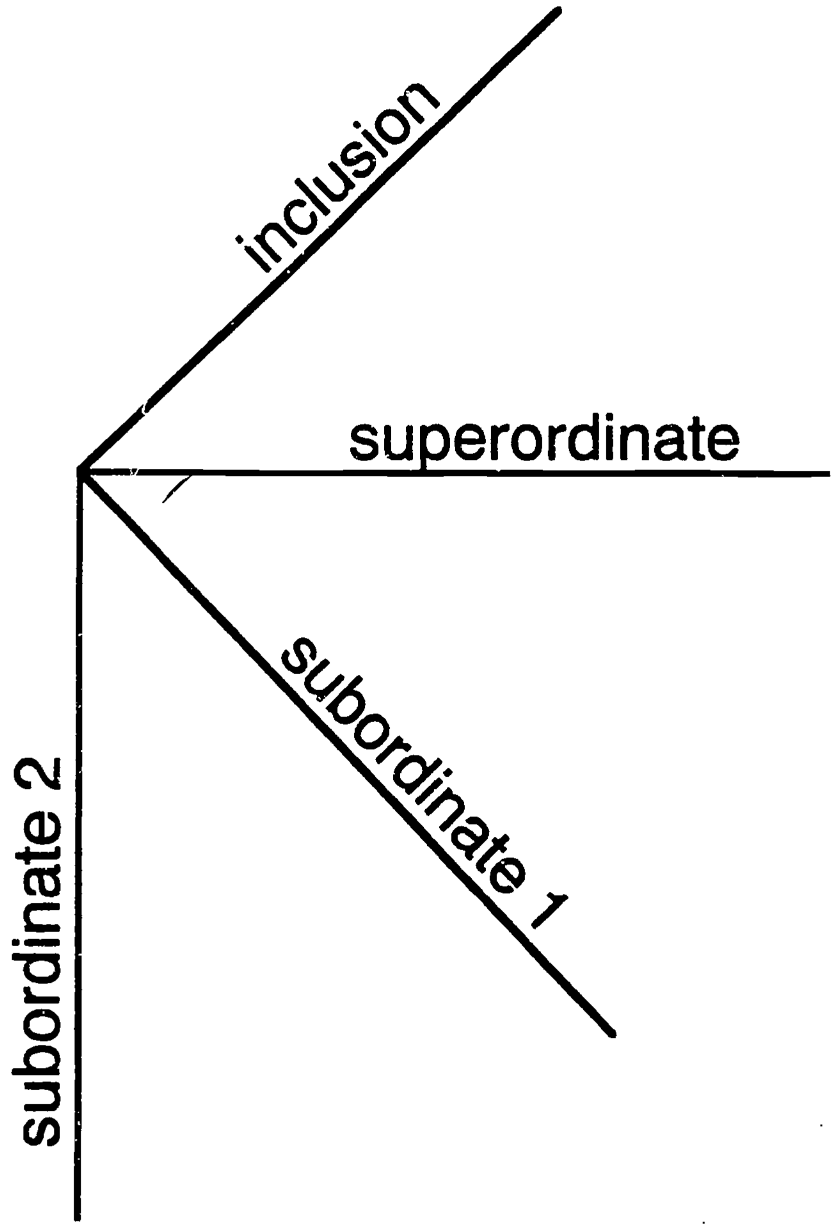
inclusion
superordinate
subordinate 1
subordinate 2

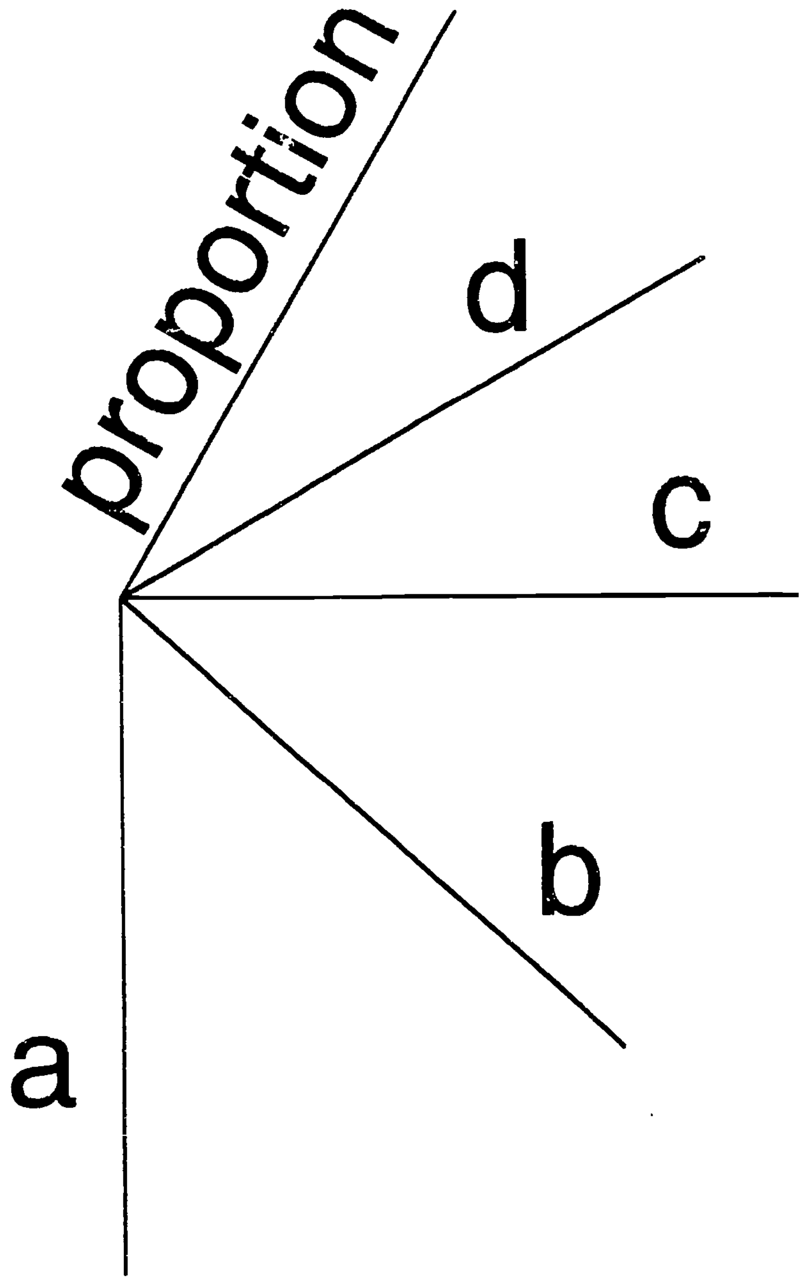
proportion
d
c
b
a

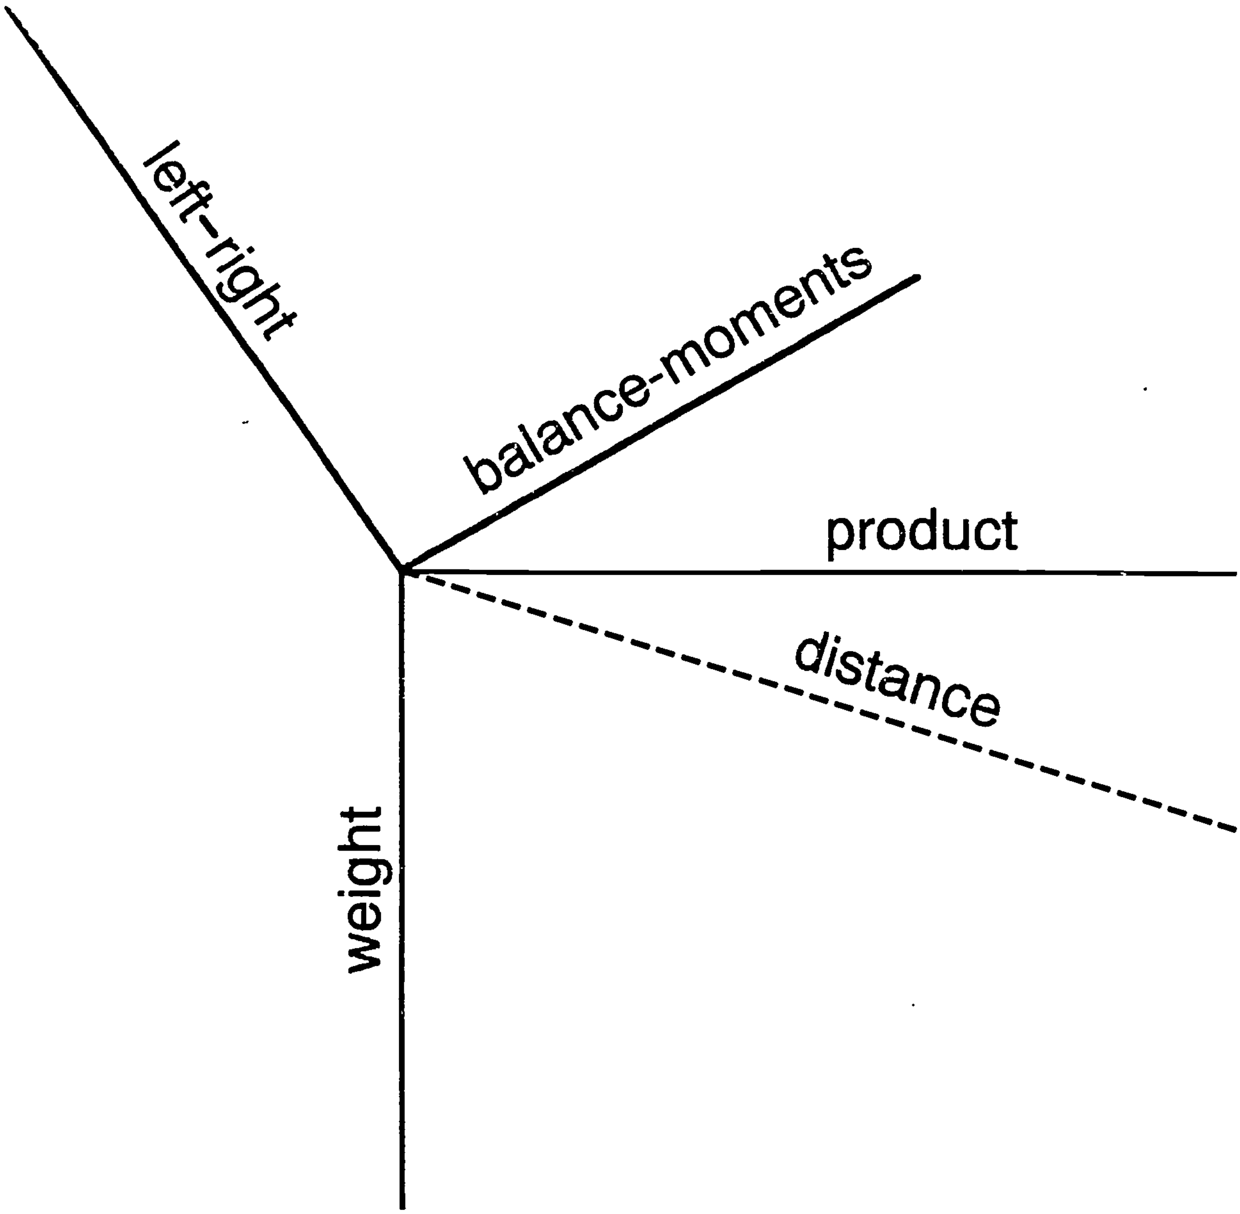
left-right
balance-moments
product
distance
weight

$$v = s\,t^{-1}$$

$$a = (v_2 - v_1)\,t^{-1}$$

$$f = m\,a$$