

Getting Rid of Derivational Redundancy or How to Solve Kuhn’s Problem

Rens Bod

School of Computer Science
University of St Andrews, and
Institute for Logic, Language and Computation
University of Amsterdam
rens@science.uva.nl

Abstract

This paper deals with the problem of derivational redundancy in scientific explanation, i.e. the problem that there can be extremely many different explanatory derivations for a natural phenomenon while students and experts mostly come up with one and the same derivation for a phenomenon (modulo the order of applying laws). Given this agreement among humans, we need to have a story of how to select from the space of possible derivations of a phenomenon the derivation that humans come up with. In this paper we argue that the problem of derivational redundancy can be solved by a new notion of “shortest derivation”, by which we mean the derivation that can be constructed by the fewest (and therefore largest) partial derivations of previously derived phenomena that function as “exemplars”. We show how the exemplar-based framework known as “Data-Oriented Parsing” or “DOP” can be employed to select the shortest derivation in scientific explanation. DOP’s shortest derivation of a phenomenon maximizes what is called the “derivational similarity” between a phenomenon and a corpus of exemplars. A preliminary investigation with exemplars from classical and fluid mechanics shows that the shortest derivation closely corresponds to the derivations that humans construct. Our approach also proposes a concrete solution to Kuhn’s problem of how we know on which exemplar a phenomenon can be modeled. We argue that humans model a phenomenon on the exemplar that is derivationally most similar to the phenomenon, i.e. the exemplar from which the largest subtree(s) can be used to derive the phenomenon.

Keywords: Scientific explanation, derivational redundancy, exemplar-based reasoning, data-oriented parsing, shortest derivation, derivational similarity, Kuhn’s problem

1 Introduction

This paper deals with the problem of derivational redundancy in scientific explanation, i.e. the problem that there can be extremely many different explanatory derivations for a phenomenon while students and experts tend to come up with only one and the same derivation for a phenomenon (modulo the order of applying the various laws, as we shall see in section 4). There is already an extensive literature on redundancy in deductive explanation, especially with respect

to spurious laws that appear as extra premises in a deductive argument but that do not contribute anything to the explanation itself (see Salmon 1990 for an overview). Yet, the current paper focuses on another kind of redundancy that is still underestimated: a phenomenon may have many different derivations even if they are subsumed under the same general laws and even if they do not contain spurious laws that are non-explanatory or irrelevant. We will show that there is a combinatorial explosion of derivations for virtually every phenomenon. In other words, derivational explanation is inherently *massively redundant*. This massive redundancy stands in strong contrast with the fact that physics students and experts tend to come up with only *one and the same* derivation for a certain phenomenon (see section 4). Given this remarkable agreement among students in deriving phenomena, we need to have a story of why humans focus on one derivation and not on others.

In this paper we argue that the problem of massive derivational redundancy can be solved by a notion of “shortest derivation”. By the shortest derivation of a phenomenon we mean the derivation that can be constructed by the fewest (and therefore largest) partial derivations of previously derived phenomena that function as “exemplars”. We show how a model of exemplar-based reasoning, as instantiated by DOP (Data-Oriented Parsing), can be applied to selecting the shortest derivation (Bod 1998). DOP represents derivations of phenomena by trees, just like proof trees in formal logic. When a phenomenon has more than one derivation, DOP returns the tree that can be constructed by the fewest subtrees from previously derived phenomena, i.e. exemplars. We show that DOP’s notion of shortest derivation closely corresponds to the derivations that humans construct. We argue that other exemplar-based accounts do not solve the redundancy problem.

The idea that natural phenomena can be explained by modeling them on exemplars is usually attributed to Thomas Kuhn in his account on “normal science” (Kuhn 1970). Kuhn urges that exemplars are “concrete problem solutions that students encounter from the start of their scientific education” (ibid. p. 187) and that “scientists solve puzzles by modeling them on previous puzzle-solutions” (ibid. p. 189). Instead of explaining a phenomenon from scratch (i.e. all the way down from laws), Kuhn contends that a scientist has acquired “a way to see his problem as *like* a problem he has already encountered” (ibid. p. 189). These “acquired similarity relations” allow a scientist to match a new problem or phenomenon to one or more previous phenomena-plus-explanations.

In similar vein, Philip Kitcher argues that new phenomena are derived by using the same patterns of derivations (“argument patterns”) as used in previously explained phenomena: “Science advances our understanding of nature by showing us how to derive descriptions of many phenomena, using the same patterns of derivation again and again” (Kitcher 1989, p. 432). Different from Kuhn, Kitcher proposes a rather concrete account of explanation, known as the “unificationist view”, which he still links to Kuhn's view by interpreting exemplars as derivations (ibid. pp. 437-8). Yet, we will argue in section 3 that Kitcher’s account does not solve the problem of derivational redundancy.

Thomas Nickles relates Kuhn's view to Case-Based Reasoning (Nickles 2003, p. 161). Case-Based Reasoning (CBR) is an artificial intelligence technique that stands in contrast to rule-based problem solving. Instead of solving each new problem from scratch, CBR stores previous problem-solutions in memory as cases. When CBR begins to solve a new problem, it retrieves from memory a case whose problem is similar to the problem being solved. It then adapts the example's solution and thereby solves the problem. CBR has been instantiated in many different ways and has been used in various applications such as reasoning, learning, perception and understanding (cf. Carbonell 1986; Falkenhainer et al. 1989; Kolodner 1993; Veloso and Carbonell 1993; VanLehn 1998). However, none of these instantiations specifically addresses the problem of massive derivational redundancy.

An instantiation of CBR that does address the problem of derivational redundancy, albeit in a different domain, is Data-Oriented Parsing (DOP). DOP is a natural language processing technique that provides an alternative to rule-based language processing. It analyzes new sentences by modeling them on analysis-trees of previous sentences (Scha 1990; Bod 1998; Goodman 2003). DOP operates by decomposing the given analysis-trees into "subtrees" and recomposing those pieces to build new trees. When a sentence has more than one possible analysis or logical interpretation -- which is the typical case in natural language -- DOP selects the analysis-tree that is constructed by the "shortest derivation", which is the tree consisting of the fewest (and therefore largest) subtrees from previous trees (Bod 2000). DOP has been highly successful in solving syntactic and semantic redundancy ("ambiguity") in natural language understanding (see Manning and Schütze 1999; Bod et al. 2003). In Scha et al. (1999) it is shown how DOP can be defined as an instantiation of CBR.

In the current paper we argue that DOP can also be used for solving derivational redundancy in physics. The DOP approach may be particularly suitable to tackle the redundancy problem because of the analogy between explanatory derivations in physics and tree structures in linguistics and logic. If we can convert explanatory derivations into trees, we can directly apply the DOP approach to the redundancy problem. That is, when a phenomenon has more than one derivation tree, DOP proposes to select the tree that can be constructed by the fewest subtrees from trees of previously derived phenomena. We argue that DOP also proposes a solution to Kuhn's problem, i.e. the problem of how we know on which exemplar a phenomenon can be modeled. DOP's answer is: the exemplar from which the largest possible subtree(s) can be reused.

We will first show in section 2 how derivations in physics can be interpreted as trees, and how explanations of new phenomena can be constructed by combining subtrees from previously explained phenomena. In section 3 we give an in-depth exploration of the problem of derivational redundancy and argue that DOP's notion of shortest derivation can solve this problem. We urge that the redundancy problem is not an artifact of DOP and that neither Kitcher's nor Kuhn's account solves it. The resulting DOP model, which we may term "Data-Oriented Physics", is evaluated in section 4 on a corpus of phenomena from classical and fluid

mechanics that were derived by third-year physics students. It turns out that there is a very close correspondence between the derivations constructed by humans and DOP's notion of shortest derivation. We end with a discussion and a conclusion.

2 Extending the DOP Model to Scientific Explanation

What do derivational explanations in physics look like? Let's start with a simple textbook example. Consider the following derivation of the Earth's mass from the Moon's orbit in the textbook by Alonso and Finn (1996, p. 247):

Suppose that a satellite of mass m describes, with a period P , a circular orbit of radius r around a planet of mass M . The force of attraction between the planet and the satellite is $F = GMm/r^2$. This force must be equal to m times the centripetal acceleration $v^2/r = 4\pi^2r/P^2$ of the satellite. Thus,

$$4\pi^2mr/P^2 = GMm/r^2$$

Canceling the common factor m and solving for M gives

$$M = 4\pi^2r^3/GP^2.$$

By substituting the data for the Moon, $r = 3.84 \times 10^8$ m and $P = 2.36 \times 10^6$ s, Alonso and Finn compute the mass of the Earth: $M = 5.98 \times 10^{24}$ kg. In doing so, Alonso and Finn abstract from many features of the actual Earth-Moon system, such as the gravitational forces of the Sun and other planets, the magnetic fields, the solar wind, etc. Albeit heavily idealized, the derivation provides a concrete problem solution on which various other (idealized) phenomena can be modeled. In fact, Alonso and Finn reuse parts of this derivation to solve problems such as the velocity of a satellite and the escape velocity from the Earth.

In order to create a formal model for derivational explanation that reuses derivational patterns, we first need a formal representation of derivations. Analogous to proof trees in formal logic, DOP proposes to represent derivational explanations by tree structures which indicate how a mathematical description of a phenomenon (or problem) is compositionally derived from theoretical laws, antecedent conditions and other knowledge. Figure 1 shows how the derivation for the Earth's mass above may be turned into a tree.




Figure 1. Derivation tree for the Earth's mass

The tree in figure 1 represents the various derivation steps (insofar as they are carried out in the example above) from higher-level laws to an equation of the mass M . We will refer to such a tree as a *derivation tree*. A derivation tree is a labeled tree in which each node is annotated with a formula; the boxes are only meant as convenient representations of these labels. The formulas at the top of each “vee” (i.e. each pair of binary branches) in the tree can be viewed as premises, and the formula at the bottom of each “vee” can be viewed as a conclusion which is arrived at by simple term substitution. The last derivation step in the tree is not formed by a “vee” but consists in a unary branch which solves the directly preceding formula for a certain variable (in the tree above, for the mass M). Thus, in general, a unary branch refers to a mathematical derivation step that solves an equation for (a) certain variable(s), while a binary branch refers to a physical derivation step which introduces and combines physical laws or conditions (or other knowledge such as phenomenological corrections and coefficients).

Note that a derivation tree corresponds to the notion of covering-law explanation or deductive-nomological (D-N) explanation of Hempel and Oppenheim (1948). In the D-N account, a phenomenon is explained by deducing it from general laws and antecedent conditions. D-N explanations usually focus on the initial premises (laws and conditions) and the final conclusion (the phenomenon). But they can just as well represent the intermediate steps as derivation trees do. For every derivation tree there is a corresponding D-N explanation and vice versa. Although the D-N account is known to suffer from various shortcomings and is nowadays superseded by other approaches, such as the unificationist account (cf. Friedman 1974; Kitcher 1981, 1989), most derivations in textbooks basically follow this scheme. Our main reason for representing derivations by trees is of course that we can then apply the DOP approach to the problem of derivational redundancy, as we shall see in the next section.

But before we can do this, we will need to demonstrate how DOP builds new explanations out of previous ones (since our solution to the redundancy problem is defined in terms of subtrees of previous trees). Consider the following subtree in figure 2 which is obtained

from the derivation tree in figure 1 by leaving out the last derivation step (i.e. the solution for the mass M).

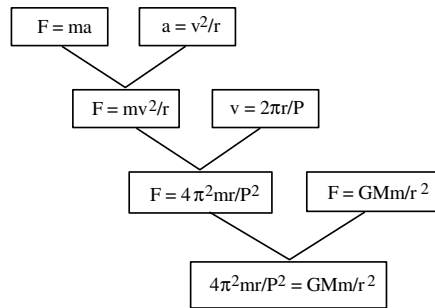


Figure 2. A subtree from the tree in figure 1

This subtree can be applied to various other situations. For instance, in deriving the regularity known as Kepler's third law (which states that r^3/P^2 is constant for all planets orbiting around the Sun, or satellites around the Earth if you wish) the subtree in figure 2 needs only to be extended with a mathematical derivation step that solves the last equation for r^3/P^2 , as represented in figure 3.

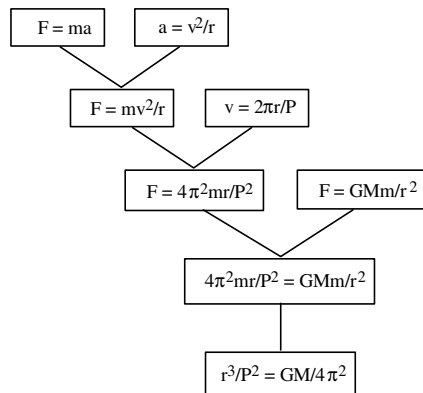


Figure 3. Derivation tree for Kepler's third law

Thus instead of starting each time from scratch, we learn from previous derivations and can reuse them for solving new problems. In a similar way we can derive the distance of a geostationary satellite, namely by solving the subtree in figure 2 for r and taking P as the rotation period of the Earth.

It is of course not typically the case that derivations involve only one subtree. In deriving the velocity of a satellite at a certain distance from a planet, we cannot directly use the large subtree in figure 2. Instead, analogous to DOP models for natural language, we *decompose* the tree in figure 1 into two smaller subtrees and *recompose* them by term substitution (represented by the operation ‘ \circ ’) and finally solve for the velocity v in figure 4.

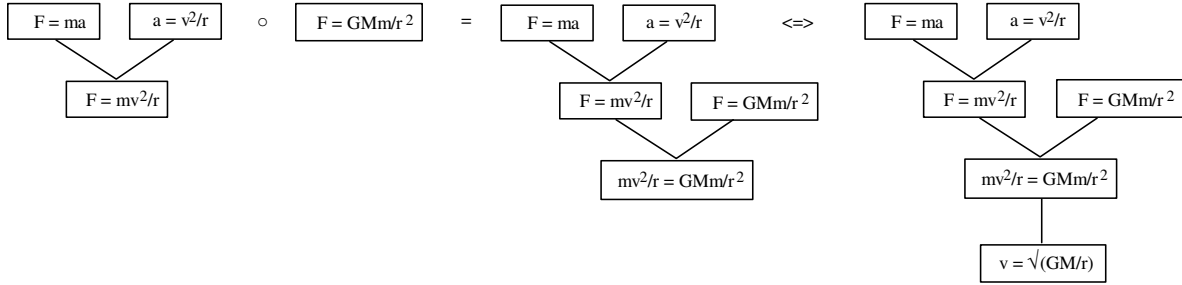


Figure 4. Deriving a phenomenon by combining two subtrees

Figure 4 shows that we can create new derivation trees by combining subtrees from previous derivation trees. Note that subtrees can be of arbitrary size: from single equations to combinations of laws, up to entire derivations.

The notion of term substitution, though widely used in rewriting systems, may need some further specification. The term-substitution operation ‘ \circ ’ is a partial function on pairs of labeled trees; its range is the set of labeled trees. The combination of tree t and tree u , written as $t \circ u$, is defined iff the equation at the root node of u can be substituted in the equation at the root node of t (i.e. iff the lefthandside of the equation at the root node of u literally appears in the equation at the root node of t). If $t \circ u$ is defined, it yields a tree that expands the root nodes of copies of t and u to a new root node where the righthandside of the equation at the root node of u is substituted in the equation at the root node of t . Note that the substitution operation can be iteratively applied to a sequence of trees, with the convention that \circ is left-associative.

We now have the basic ingredients for a DOP model of derivational explanation, which we may term “data-oriented physics”. This DOP model employs (1) a *corpus of derivation trees* representing exemplars and (2) a *matching procedure* that combines subtrees from the corpus into new derivation trees. This brings us to the following definition for an explanation of a phenomenon with respect to a corpus.

Definition 1 Given a corpus C of derivation trees T_1, T_2, \dots, T_n representing exemplars and a term substitution operation \circ , an explanation of a phenomenon P with respect to C is a derivation tree T such that (1) there are subtrees t_1, t_2, \dots, t_k in T_1, T_2, \dots, T_n for which $t_1 \circ t_2 \circ \dots \circ t_k = T$, (2) the root node of T is mathematically equivalent to P and (3) the leaf nodes of T are either laws or antecedent conditions or equations that cannot be derived from higher-level equations.

In our examples above, the mathematical derivation steps all occur at the end of a derivation (figures 1, 3 and 4). But they may of course just as well occur in the course of a derivation between two subtrees, as we will see in the following sections.

So far, we have only dealt with explanations of *idealized* phenomena that can be constructed from theoretical laws and antecedent conditions only. It is well known, however,

that explanations of *real-world* phenomena may involve *non*-deductive elements such as corrections, normalizations and other adjustments that stand in no deductive relation to laws (Cartwright 1983, 1999; Giere 1988, 1999). Yet, it is easy to see that trees can integrate theoretical laws and phenomenological adjustments as long as such adjustments can be described in terms of mathematical formulas. Definition 1 captures in fact non-deductive elements by referring to them as “equations that cannot be derived from higher-level equations”.

It is convenient to give an example of a phenomenological adjustment in a derivation tree, as we will need this example in section 4. Consider the problem of deriving the discharge of a jet from a tank in fluid mechanics, which is derived by Norman et al. (1990, p. 497) from Bernoulli's law as follows:

Suppose the subscripts 1 and 2 refer to a point in the surface of the liquid in the tank, and a section of the jet just outside the orifice. If the orifice is small we can assume that the velocity of the jet is v at all points in this section.




Figure 12.12

The pressure is atmospheric at points 1 and 2 and therefore $p_1 = p_2$. In addition the velocity v_1 is negligible, provided the liquid in the tank has a large surface area. Let the difference in level between 1 and 2 be h as shown, so that $z_1 - z_2 = h$. With these values, Bernoulli's equation becomes:

$$h = v^2/2g \quad \text{from which} \quad v = \sqrt{(2gh)}$$

This result is known as Torricelli's theorem. If the area of the orifice is A the theoretical discharge is:

$$Q(\text{theoretical}) = vA = A\sqrt{(2gh)}$$

The actual discharge will be less than this. In practice the liquid in the tank converges on the orifice as shown in Fig. 12.12b. The flow does not become parallel until it is a short distance away from the orifice. The section at which this occurs has the Latin name *vena contracta* (*vena* = vein) and the diameter of the jet there is less than that of the orifice. The actual discharge can be written:

$$Q(\text{actual}) = C_d A \sqrt{2gh}$$

where C_d is the coefficient of discharge. Its value depends on the profile of the orifice. For a sharp-edged orifice, as shown in Fig. 12.12b, it is about 0.62.

Note that the coefficient of discharge C_d is not derived from higher-level laws in the derivation but is introduced ad hoc. Yet we can still create a derivation tree for this system (and thus also a D-N argument) if we write the coefficient by the rule $Q(\text{actual}) = C_d Q(\text{theoretical})$, which is shown in figure 5 (in which we also added Bernoulli's law in full):

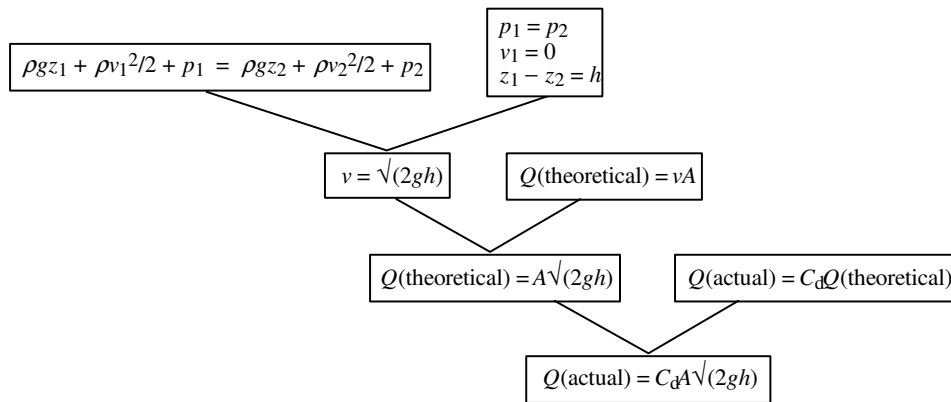


Figure 5. Derivation tree for the discharge of a jet from a tank

The tree in figure 5 closely follows the derivation above, where the initial conditions for p_1 , p_2 , v_1 , z_1 and z_2 are represented by a separate label in the tree. The coefficient of discharge is introduced in the tree by the rule $Q(\text{actual}) = C_d Q(\text{theoretical})$. Although this empirical rule does not follow from any higher-level law, it is widely used in fluid mechanics to solve a large number of other systems, ranging from nozzles, notches, weirs, open channel flows and many pipeline problems -- see Douglas and Matthews (1996). The rule is employed in all hydraulic systems with a fluid discharge. Without using it, the predicted discharge of a system can be up to 50% off the mark. By writing the rule in terms of a mathematical equation we fit it into a derivation tree and can reuse it wherever needed.

3 The Redundancy Problem and Its Proposed Solution

Now that we have extended the DOP model to derivational explanation, we can go into the main problem of this paper, and show how DOP may solve it. This is the problem that there can be many, often extremely many, different derivations for the same phenomenon, even if they are subsumed under the same general laws and even if they do not contain spurious laws that are non-explanatory or irrelevant.

In order to show this, we will first enlarge our tiny corpus used in section 2 (which consisted only of the derivation in figure 1) with another derivation from Alonso and Finn's textbook. This derivation again provides an exemplary problem solution for the Earth's mass but this time by computing it from the acceleration of an object at the Earth's surface (Alonso and Finn 1996, p. 246). This second exemplar can be represented by the derivation tree in figure 6.




Figure 6. An additional exemplar in the corpus for deriving the Earth's mass

By substituting the values for g (the acceleration at the Earth's surface), R (the Earth's radius) and G (the gravitational constant), Alonso and Finn obtain roughly the same value for the Earth's mass as in the previous derivation in figure 1. They argue that this agreement is “a proof of the consistency of the theory” (ibid. p. 247). (Note that the derivation is again idealized: no centrifugal force is taken into account, let alone influences from the Sun or other planets.) Thus the problem of the Earth's mass can be solved in at least two different ways. And both derivations are used in Alonso and Finn's textbook as exemplars for deriving other phenomena.

When we add the tree in figure 6 to our corpus, we can model Kepler's regularity also on this exemplar, resulting in an alternative derivation tree which is constructed by using a large subtree from figure 6 in combination with two small subtrees from the exemplar in figure 1, as shown in figure 7.

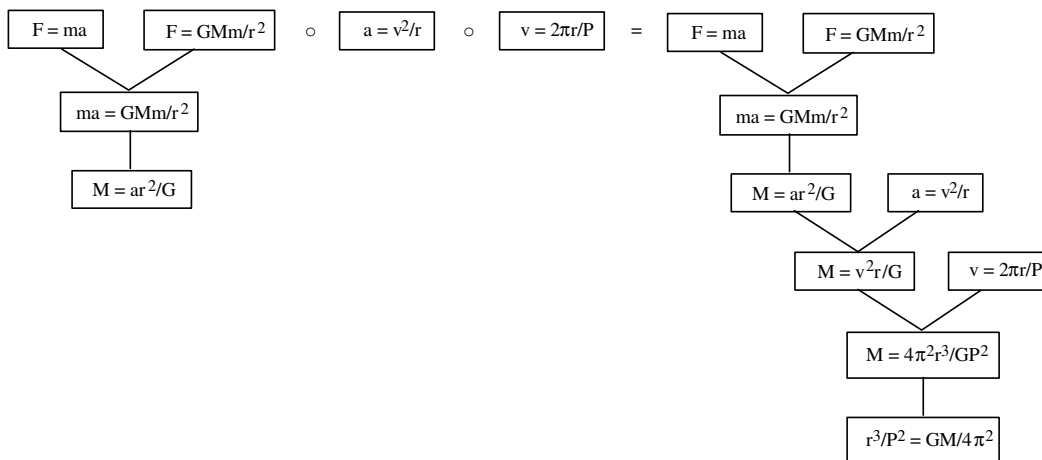


Figure 7. An alternative derivation and resulting derivation tree for Kepler's regularity

Thus there are at least two different derivation trees for Kepler's regularity, represented by figures 3 and 7, and it is easy to see that there are many more trees for Kepler's regularity. By combining subtrees from the two exemplars in figures 1 and 6 in different ways, we get a combinatorial explosion of possible derivation trees of Kepler's law. In other words, derivational explanation is *massively redundant*. Figure 8 shows two more derivation trees of Kepler's law.

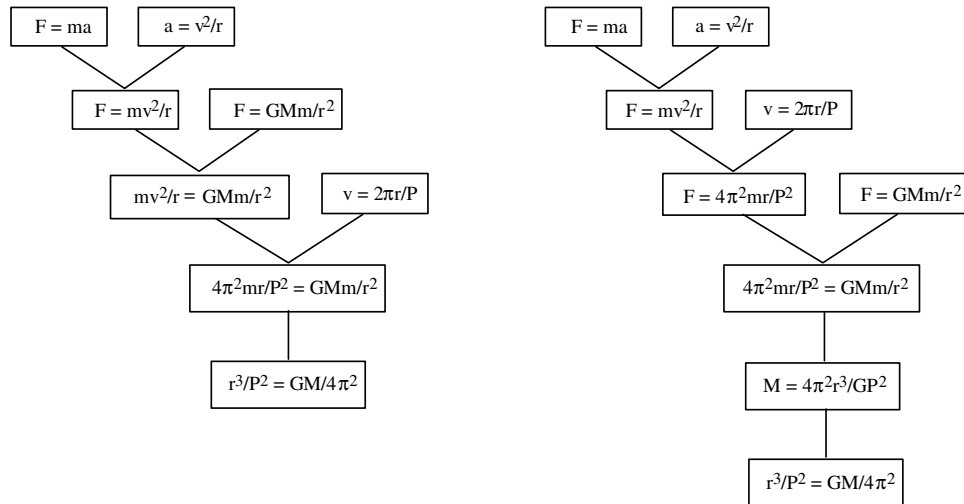


Figure 8. Two more derivation trees for Kepler's regularity

The leftmost tree in figure 8 differs from the tree in figure 3 only in the order of combining the laws. The rightmost tree in figure 8 can be obtained by reusing the entire derivation tree from figure 1 (rather than the partial tree in figure 2) which is next solved for Kepler's regularity. We will see that these two derivation trees can be easily ruled out by our mechanism described below.

It may of course also occur that a phenomenon can be derived in different ways from *different* laws. For example, we can derive the velocity that an object attains in free fall from height h in at least two ways: either by applying the law of energy conservation or by applying the Galilean equations of motion. In the first case, we equate potential energy mgh with kinetic energy $mv^2/2$ and solve for v . In the second case, we combine the Galilean equations $h = (1/2)gt^2$ and $v = gt$, and also solve for v . Both derivations lead to the same result for the velocity, i.e. $v = \sqrt{2gh}$. Yet it turns out that students have a clear preference for the first derivation which equates potential and kinetic energy (see section 4).

These examples show that the redundancy problem is *not* just an artifact of the DOP model. Any representation that describes how a phenomenon is derived from laws and conditions will exhibit derivational redundancy. Even if we do not consider the order by which laws are combined in a derivation, there may still be a multitude of derivations (as illustrated by

the derivations of Kepler's law in figure 3 and the rightmost derivation in figure 8). And since humans display a remarkable agreement in deriving phenomena -- as we will show in section 4 - - we need to have a story of how to select from the space of possible derivations the derivation that humans come up with. To this effect, let's go back to Kepler's third law and discuss more closely the differences between the two derivation trees in figures 3 and 7 that were modeled on two different exemplars.

First note that there is nothing "wrong" with the longer derivation tree in figure 7: there are no spurious non-explanatory laws that are irrelevant, as would be e.g. Hooke's or Boyle's law (see Friedman 1974 for a discussion of non-explanatory laws in derivations). The main difference is that the derivation in figure 7 is modeled on a *different* exemplar than the derivation in figure 3. The alternative derivation in figure 7 is even insightful as it refers to the conceptual equivalence between terrestrial and celestial mechanics in Newtonian dynamics. The fact that Kepler's regularity can be derived not only from the exemplar in figure 1 but also from figure 6 suggests that if we bring a satellite down to the Earth's surface it still follows the same regularity.

Yet, it turns out that no physics student comes up with the derivation tree in figure 7. Why? Apart from the fact that the derivation tree in figure 3 is smaller, the tree in figure 3 is more "derivationally similar" to an exemplar in the corpus. That is, the tree in figure 3 can be constructed by just one large subtree from the corpus -- i.e. from the exemplar in figure 1 -- whereas the tree in figure 7 needs at least 3 subtrees to be constructed from the corpus -- one from the exemplar in figure 6 and two from figure 1. Of course, for another phenomenon it may be the exemplar in figure 6 rather than in figure 1 that can derive the phenomenon in one go. For example, to derive a formula for the gravitational acceleration at the Earth's surface we can use one large subtree from figure 6 and not from figure 1. Different problems may be modeled on different exemplars.

Thus in solving the problem of derivational redundancy, it seems that we need to determine on which exemplar a phenomenon can best be modeled. Note that Kitcher's account of explanation does not help us here. According to Kitcher (1989, p. 432), "Science advances our understanding of nature by showing us how to derive descriptions of many phenomena, using the same patterns of derivation again and again". But his "unificationist" account does not tell us whether humans model a phenomenon like the gravitational acceleration at a planet's surface on the exemplar in figure 1 or on the exemplar in figure 6. This is what should perhaps be called "Kuhn's problem", i.e. the problem of how we know on which exemplar a phenomenon can be modeled. DOP's answer is: the exemplar from which the largest subtree can be reused. This finally brings us to our notion of "shortest derivation" and to a solution of the problem of derivational redundancy.

Let's get more concrete. We have seen that there can be different derivation trees for a phenomenon. A distinctive feature between different derivation trees is that *some trees are more similar to exemplars than others*. The larger the partial match between a derivation tree and an

exemplar, in terms of their largest common subtree, the more “derivationally similar” they are. Since students learn physics not just by memorizing laws, but also by studying exemplary problem solutions, they try to derive a phenomenon by maximizing derivational similarity with previously derived phenomena, or equivalently, by *minimizing derivation length* where the length of a derivation is defined as the number of corpus-subtrees it consists of. We will refer to the derivation of minimal length as the “shortest derivation”. Since subtrees in DOP can be of arbitrary size, *the shortest derivation corresponds to the derivation tree which consists of largest partial match(es) with previous derivation trees in the corpus*. This brings us to the following definition of the “best derivation tree” for a phenomenon derived by DOP:

Definition 2 Let $L(d)$ be the length of derivation d in terms of its number of subtrees, that is, if $d = t_1 \circ \dots \circ t_k$ then $L(d) = k$. Let d_T be a derivation which results in tree T . Then the best tree, T_{best} , derived by DOP is the tree which is produced by a derivation of minimal length:

$$T_{best} = \operatorname{argmin}_T L(d_T)$$

It is important to understand the difference between a tree produced by the smallest number of subtrees and an absolute smallest tree. While the tree in figure 3 is produced by the shortest possible combination of corpus-subtrees, it does not correspond to the smallest possible tree, i.e. the tree with the smallest possible number of nodes (or labels). There exists a smaller tree that simply applies all laws at once to arrive at the formula for Kepler's regularity. However, it turns out that no student constructs such minimal derivations, and we therefore believe that our notion of shortest derivation consisting of the smallest number of (corpus-)subtrees is more appropriate than a notion of shortest derivation defined as the smallest number of *nodes*. Only in case our notion of shortest derivation does not lead to a unique result, i.e. if a phenomenon can be derived by the same smallest number of subtrees, it seems reasonable to choose the tree with the fewest nodes from among the shortest derivations, reflecting a preference for economy if DOP does not break ties.

This means that the tree in figure 3 is preferred to the rightmost tree in figure 8: both trees can be constructed by reusing one subtree only (from the exemplar in figure 1), but the tree in figure 3 has fewer nodes. Note that the tree in figure 3 is also preferred to the leftmost tree in figure 8. This is due to our notion of shortest derivation: the leftmost tree in figure 8 cannot be constructed by one subtree only. It can only be constructed by combining smaller subtrees (like the tree in figure 7). Thus our mechanism indeed seems to eliminate derivational redundancy. But in how far does it correspond to the way humans eliminate derivational redundancy? This will be explored in the next section.

4 Evaluating DOP's Solution for the Redundancy Problem

How can we evaluate our DOP model? Since we propose DOP to be a cognitive model of human problem solving, it seems appropriate to evaluate the model on derivations constructed by humans. To this end, we developed a *test corpus* of manually solved problems by third-year physics students and a *training corpus* of exemplary problem solutions taken from textbooks. Next, we developed a simple implementation of DOP which computed T_{best} for each test problem by means of subtrees from the training corpus. The performance of DOP on the test problems was compared with the derivations provided by the students.

4.1 Method and Procedure

The procedure was as follows. 19 third-year physics students were paid to solve 6 elementary problems from classical mechanics and 5 elementary problems from fluid mechanics. The students had previously followed courses in classical mechanics and fluid mechanics. The 11 problems given to them consisted in deriving a phenomenon from laws, initial conditions and, in the case of fluid mechanics, empirical coefficients. The students were given no other instructions than that they should solve the problems by paper and pencil in class. The first two and the last two of the problems are given below:

Problem nr. 1

Show that the period of the Earth's rotation for which an object at the equator would become weightless is given by $P = 2\pi\sqrt{R/g}$ where R is the Earth's radius and g is the gravitational acceleration at the Earth's surface.

Problem nr. 2

Show that the theoretical velocity which an object attains in free fall from height h is given by $v = \sqrt{2gh}$ where g is the gravitational acceleration at the Earth's surface.

Problem nr. 10

When water flows through a right-angled V-notch, show that the discharge is given by $Q = KH^{5/2}$ in which K is a constant and H is the height of the surface of the water above the bottom of the notch.

Problem nr. 11

Show that the theoretical rate of flow through a rectangular notch is given by $Q = (2/3)B\sqrt{2g}H^{3/2}$ where B is the width of the notch and H is the height of the water level above the bottom of the notch.

After the students had solved the problems on paper, they were given a short, fifteen-minutes tutorial on the concept of derivation tree, especially on the difference between binary branches in a tree (used for combining laws, conditions etc.), and unary branches (used for mathematical

derivation steps). The students were told that the exact order of combining laws in a tree was not important as long as these laws could be properly combined by term substitution to solve the problem. What *was* important was to represent in the tree the derivation steps they had used to get from laws to the phenomenon, so we told them. Thus we did not distinguish between trees whose only difference was the order of the applied laws, as we found out in a pilot experiment that this was neither done by students. We will see below that even with this abstraction there were still many different derivations. After this brief tutorial, the students were asked to draw derivation trees for their problem solutions.

There was a high agreement among the derivation trees constructed by the students: on average 91.4% (SD=1.2) of the derivation trees per problem matched (modulo law order). To create a gold standard, only the most voted derivation tree for each problem was put in the test corpus. It is important to mention that the students had no difficulties with drawing trees for their problem solutions, and there were no questions during this task. This suggests that derivation trees are suitable structures for representing problem solutions by humans.

Next, the students were asked to draw derivation trees for 9 exemplary problem solutions that are used as exemplars in the textbooks by Alonso and Finn (1996, chapter 11) and Douglas and Matthews (1996, chapter 7). The three example problems in figures 1, 5 and 6 were among these exemplars. It should be said, however, that none of the students derived $F=mv^2/r$ from $F=ma$ and $a=v^2/r$, as we did in figure 1. Instead, all students used the equation for centripetal force $F=mv^2/r$ directly as a law. The agreement among the derivation trees for the exemplary solutions was very high: 97.7% (SD=0.3). The most voted (i.e. most frequent) tree for each exemplary solution was put in the training corpus.

All test problems could be solved out of subtrees from the training corpus, but this fact was not told to the students: they first had to solve the test problems after which they were asked to draw trees for the exemplary problem solutions from the textbooks. Each student accomplished the task in less than 2.5 hours (including the tutorial).

The goal for DOP was to construct T_{best} for each of the 11 problems from the test corpus by means of the subtrees from the training corpus of 9 exemplars. To accomplish this, we developed a simple implementation of DOP that used *TK Solver* as a backbone (release 5.0, Universal Technical Systems Inc.). *TK* solves an equation given a list of other equations -- provided that there is a solution. To make *TK* suitable for DOP, we programmed a shell around *TK* which first converted each derivation tree from the training corpus into all its subtrees and next extracted the equations from the subtree-roots. Each equation was indexed to remember the subtree it was extracted from. This resulted in a list L of 148 equations. For each test problem (i.e. equation to be solved), *TK* derived a set of solutions given the list L . All problems received more than 60 different solutions, even *after* abstracting from the order of the equations used in the solution, which gives an idea of the derivational redundancy if we do not have any mechanism to break ties.

From *TK*'s output, our program selected the shortest solution for each problem that used the fewest equations. The equations of the shortest solution were converted back to their corresponding subtrees, which were combined into the derivation tree T_{best} . Note that in this way T_{best} consisted of the largest partial matches with trees in the training corpus. In case T_{best} was not unique the program chose the tree with fewest nodes among the best trees. An advantage of *TK* is that it hides the algebra, which is good as this was also asked from the students and which corresponds to our use of unary branches in trees. (We will not go into the equational reasoning techniques used in *TK*, as these are nowadays rather standard -- see Baader and Nipkow 1998).

Of course *TK* is by itself not comparable to human problem solving. But what we want to evaluate is not *TK*, but the claim that (DOP's notion of) the shortest derivation matches the derivations constructed by humans. Thus *TK* is only taken as a tool and should be seen as a "provider" of the set of *possible* solutions from which DOP should predict the *best* solution that humans come up with.

4.2 Results and Discussion: Solving Kuhn's Problem

The best trees computed by DOP were compared with the derivation trees constructed by the students in the test corpus. Abstracting from the order of the laws in the trees, the accuracy of DOP was 91%. That is, for 10 out of 11 phenomena, the derivation trees predicted by DOP matched (modulo law order) the derivation trees assigned by (the vast majority of) the students.

To put our 91% accuracy into perspective, we also computed the accuracy by choosing a *random* derivation tree, T_{random} , from among *all* possible trees that were constructed by DOP (i.e. trees that did not necessarily correspond to the shortest derivation but that could be constructed from *TK*'s output). In this case, the accuracy was only 9% (again, modulo law order). Although our test set consists of only 11 trees, the difference between 91% accuracy obtained by T_{best} and the 9% accuracy obtained by T_{random} -- which is however still a "correct" derivation -- suggests that T_{best} mimics human problem solving more closely than T_{random} . We also computed the accuracy by choosing the *absolute* smallest tree containing fewest nodes among all proposed trees in *TK*'s output. This resulted in 18% accuracy (2 out of 11).

These results suggest that if we want to predict the derivations that humans construct for phenomena, we should not search for the smallest tree in terms of nodes, let alone select a random tree, but search for the tree which consists of the fewest subtrees from previous derivations, as generated by our notion of shortest derivation.

One may claim that our results are not very surprising since they reconfirm Kuhn's insight that scientists explain phenomena by modeling them on previously explained phenomena with only minimal recourse to additional derivation steps (Kuhn 1970, p. 189). However, Kuhn does not provide an exact procedure that tells us on which exemplar a new phenomenon can be modeled. DOP does provide such a procedure, albeit indirectly by solving the problem of derivational redundancy. In doing so, DOP also suggests a precise notion of similarity between a

phenomenon and a set of exemplars: the most similar exemplar is the one from which the largest subtree can be reused to derive the phenomenon. Moreover, even when we know on which exemplar a phenomenon can be modeled, there may still be different derivations for the phenomenon, also if we abstract from the order of law application -- as for instance illustrated by the rightmost tree in figure 8. Thus we need an additional notion to break ties, as given by DOP's shortest derivation.

This brings us to another interesting result: DOP correctly predicted for each problem whether it could be solved by problem solutions from classical or from fluid mechanics. It happened only once that T_{best} was not unique. This occurred with problem nr 2 given above: "Show that the theoretical velocity which an object attains in free fall from height h is given by $v = \sqrt{2gh}$ where g is the gravitational acceleration at the Earth's surface". Although this problem is exceedingly simple, DOP constructed two equally short derivation trees for it with equally many nodes. It is interesting therefore to go into this problem with somewhat more detail.

All students but one derived the formula $v = \sqrt{2gh}$ by equating potential energy mgh with kinetic energy $mv^2/2$ (one student derived it from the Galilean equations of motion, resulting in a somewhat longer derivation, which was also discussed in section 3). DOP, on the other hand, returned two derivations that were equally short. One of these derivations was the same as the (majority of the) students came up with, i.e. by equating potential and kinetic energy -- which in fact literally appeared as a subtree in a larger derivation tree in the training set and could thus be constructed by reusing exactly that subtree. The other derivation was taken from the exemplary solution from fluid mechanics for the jet through an orifice, and also consisted of just one subtree, namely the subtree in figure 5 that has the formula $v = \sqrt{2gh}$ as its root and Bernoulli's equation and the initial conditions as its leaves. Both trees could thus be constructed by exactly one subtree from the corpus and both consisted of three nodes. Thus even after choosing the tree with the fewest nodes from among the shortest derivations, T_{best} was not unique. (Of course, this all hinges on the contents of our training corpus, but remember that we did actually use the exemplary derivations from two widely used textbooks.)

To get out of this impasse, we could add yet another selection criterion to DOP, for example one which chooses the tree with fewest *terms* (variables and constants) in the labels. Because the derivation from classical mechanics contains fewer terms than the derivation from fluid mechanics, this selection criterion would let DOP predict the same derivation as the vast majority of the students did for this problem. But this additional criterion does not solve our impasse in general, since DOP would then predict the wrong derivation if the problem was framed in terms of fluid mechanics, i.e.: "Show that the theoretical velocity of a jet through a small orifice from a tank of height h is $v = \sqrt{2gh}$ where g is the gravitational acceleration at the Earth's surface".

What is actually happening is that DOP constructs derivation trees by syntactic reasoning only. It does not distinguish between a point mass and a unit-volume of fluid. In fact, in reading problem nr 2, we may just as well interpret the word "object" as a part of fluid rather than as a

point mass, resulting into two different phenomena. This is not as far-fetched as it seems, since historically Daniel Bernoulli solved the problem of the velocity of water by analogically treating a flow in terms of Newtonian-like particles, which makes the two phenomena “equivalent”.

Of course, if we want to avoid DOP mixing up derivations from different domains, we could either create different corpora for different domains, or introduce different variables for point mass velocity and fluid velocity. The latter can be accomplished by using subcategorizations, e.g. v_p for the velocity of a particle and v_f for the velocity of a fluid. This way, the two velocities cannot be substituted, and the phenomena $v_p = \sqrt{2gh}$ and $v_f = \sqrt{2gh}$ get each a different derivation. However, as problem nr 2 is stated above, it is inherently ambiguous and DOP rightly comes up with the somewhat unexpected result of two different best trees. In terms of Kuhn’s problem: the phenomenon can be modeled equally well on two different exemplars: one from classical mechanics and one from fluid mechanics.

While DOP may thus capture the notion of derivational similarity between a phenomenon and an exemplar quite well, even with respect to exemplars that contain theory-external knowledge like ad hoc corrections or approximations (see figure 5), the model hinges on the fact that all knowledge must be represented by equations. But what if such knowledge involves an intermediate model, as is e.g. the case with Prandtl’s *boundary layer* model discussed by Morrison (1999)? Morrison rightly notes that the boundary layer model is autonomous in that it was not created by some approximation of the general Navier-Stokes equations for fluid dynamics. Yet, Prandtl’s model does represent a mathematical structure which approximates the Navier-Stokes equations, which means that it can be represented by a derivation tree (cf. Bod 2006). Such a derivation tree does of course not represent the creative act of inventing the boundary-layer model, but once it has been invented it can be fit into a tree and can be reused as an exemplar by DOP to derive a range of new engineering problems.

Another important feature of DOP is that it does *not* demand that a phenomenon or system be linked to universal laws (see definition 1). A phenomenon may be derived from a phenomenological model only, without any deductive relation to high-level theory. This is e.g. the case in quantum-chromodynamics, where phenomenological models such as the MIT-bag model are used to describe certain features of quarks (see Hartmann 1999). Such phenomenological models do serve as exemplars, even though there is no deductive relation with theory. This kind of situation also occurs in disciplines where universal laws are difficult to come by or where they are not present at all, such as in biology, economics and linguistics. As long as there is a model or regularity representing the phenomenon, we can construct a derivation tree for it. In DOP there is no need to link a phenomenon to general laws, except if the phenomenon can be derived from them. In the “worst” case a derivation tree consists only of the empirical regularity describing the phenomenon. Whether DOP can still solve derivational redundancy in such cases will have to await further research.

5 Conclusion

In this paper we proposed a solution to the problem of derivational redundancy in scientific explanation, which also suggested a solution to Kuhn's problem. We showed that the DOP approach provides a way to reduce the combinatorial explosion of different derivations of a phenomenon by selecting the shortest derivation consisting of the fewest partial derivations from previously derived phenomena. A preliminary investigation with a corpus of phenomena from classical and fluid mechanics showed that DOP accurately predicts the derivations humans come up with. To the best of our knowledge, this paper provides the first concrete proposal to tackle the problem of massive derivational redundancy in scientific explanation, providing a concrete solution to Kuhn's problem of how we know on which exemplar a problem can be modeled. While we have only shown how Kuhn's problem can be solved for routine problem solving, we argued that our approach can be extended to frontier science where intermediate models play an important role.

References

- Alonso, M. and E. Finn (1996). *Physics*, Addison-Wesley.
- Baader, F. and T. Nipkow (1998). *Term Rewriting and All That*, Cambridge University Press.
- Bod, R. (1998). *Beyond Grammar: An Experience-Based Theory of Language*. CSLI Publications/Cambridge University Press.
- Bod, R. (2000). Parsing with the Shortest Derivation. *Proceedings COLING 2000*, Saarbruecken, Germany. 69-75.
- Bod, R., R. Scha and K. Sima'an (eds.) (2003). *Data-Oriented Parsing*. University of Chicago Press.
- Bod, R. (2006). Towards a General Model of Applying Science. *International Studies in the Philosophy of Science* 20(1), 5-25.
- Carbonell, J. (1986). Derivational Analogy: A Theory of Reconstructive Problem Solving and Expertise Acquisition. In Michalski et al. (eds.), *Machine Learning*, Vol. II, Morgan Kaufmann, 371-392.
- Cartwright, N. (1983). *How the Laws of Physics Lie*, Oxford University Press.
- Cartwright, N. (1999). *The Dappled World*. Cambridge University Press.
- Douglas, J. and R. Matthews (1996). *Fluid Mechanics*. Vol. 1, 3rd edition, Longman.
- Falkenhainer, B., K. Forbus and D. Gentner (1989). The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligenc*, 41, 1-63.
- Friedman, M. (1974). Explanation and scientific understanding, *Journal of Philosophy*, 71: 5-19.
- Giere, R. (1988). *Explaining Science: A Cognitive Approach*, University of Chicago Press.
- Giere, R. (1999). *Science without Laws*, The University of Chicago Press.
- Goodman, J. (2003). Efficient algorithms for the DOP model. In R. Bod, R. Scha and K. Sima'an (eds.). *Data-Oriented Parsing*, CSLI Publications, 125-146.
- Hartmann, S. (1999). Models and stories in hadron physics, in M. Morgan and M. Morrison (eds.), *Models as Mediators*, Cambridge University Press, 326-346.
- Hempel, C. and P. Oppenheim (1948). Studies in the Logic of Explanation, *Philosophy of Science*, 15: 135-175.

- Kitcher, P. (1981). Explanatory unification, *Philosophy of Science*, 48: 507-531.
- Kitcher, P. (1989). Explanatory unification and the causal structure of the world. In P. Kitcher and W. Salmon (eds.), *Scientific Explanation*, University of Minnesota Press, 410-505.
- Kolodner, Janet (1993), *Case-Based Reasoning*, Morgan Kaufmann.
- Kuhn, T. (1970). *The Structure of Scientific Revolutions*, 2nd edition, University of Chicago Press.
- Manning, C. and H. Schütze (1999). *Foundations of Statistical Natural Language Processing*. The MIT Press.
- Morrison, M. (1999). Models as autonomous agents, in M. Morgan and M. Morrison (eds.), *Models as Mediators*, Cambridge University Press, 38-65.
- Nickles, T. (2003). Normal science: from logic to case-based and model-based reasoning, In Thomas Nickles (ed.), *Thomas Kuhn*, Cambridge University Press, 142-177.
- Salmon, W. (1990). *Four Decades of Scientific Explanation*. University of Minneapolis Press.
- Scha, R. (1990). Language Theory and Language Technology; Competence and Performance. In Q.A.M. de Kort and G.L.J. Leerdam (eds), *Computertoepassingen in de Neerlandistiek*, Almere.
- Scha, R., R. Bod and K. Sima'an (1999). A Memory-Based Model of Syntactic Analysis: Data-Oriented Parsing. *Journal of Experimental and Theoretical Artificial Intelligence*, 11(3), 409-440.
- VanLehn, K. (1998). Analogy Events: How Examples are Used During Problem Solving, *Cognitive Science*, 22(3): 347-388.
- Veloso, M. and J. Carbonell (1993). Derivational Analogy in PRODIGY: Automating Case Acquisition, Storage, and Utilization. *Machine Learning*, 10(3), 249-278.