

Towards a General Model of Applying Science

RENS BOD

To appear in *International Studies in the Philosophy of Science*, Vol. 20(1), March 2006, pp.5-25.
Symposium on "Applying Science"

How is scientific knowledge used, adapted and extended in deriving phenomena and real-world systems? This paper aims at developing a general account of "applying science" within the exemplar-based framework of Data-Oriented Processing (DOP), which is also known as Exemplar-Based Explanation (EBE). According to the exemplar-based paradigm, phenomena are explained not by deriving them all the way down from theoretical laws and boundary conditions but by modelling them on previously derived phenomena that function as exemplars. To accomplish this, DOP proposes to maintain a corpus of derivation trees of previous phenomena together with a matching algorithm that combines subtrees from the corpus to derive new phenomena. By using a notion of derivational similarity, a new phenomenon can be modelled as closely as possible on previously explained phenomena. I will propose an instantiation of DOP which integrates theoretical and phenomenological modelling and which generalises over various disciplines, from fluid mechanics to language technology. I argue that DOP provides a solution for what I call Kuhn's problem and that it redresses Kitcher's account of explanation.

1. Introduction

How do we get from theory to the real world? That is to say, how is scientific knowledge used, adapted and extended in deriving concrete phenomena and real-world systems? It has for a long time been assumed that phenomena and systems are derived by solving the laws from pure science for specific boundary conditions (see Boon 2006). Yet it has become increasingly clear that in deriving a phenomenon or real-world system we also add non-theoretical elements, such as corrections, normalisations and other adjustments, that stand in no deductive relation to laws (see Cartwright 1983, 1999; Boumans 1999; Morrison and Morgan 1999). Applying a scientific theory to a concrete situation is a matter of intricate approximation and de-idealisation for which no general rules are known. How then do we derive a real-world system or phenomenon in science and engineering? According to Ronald Giere, Thomas Nickles and others, scientists work in an exemplar-based way. In deriving a new phenomenon they look for known phenomena that are in various ways similar to the new phenomenon. The derivations and techniques that successfully accounted for the known phenomena are extended and adapted to the new phenomena (see Giere 1988, 1999; Nickles 2003). Such known phenomena function as exemplars on which new phenomena are modelled.

The notion of exemplar is usually attributed to Thomas Kuhn in his account of normal science (Kuhn 1970). Kuhn urged 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). Scientists possess what Kuhn called "acquired similarity relations" that allow them "to see situations as like each other, as subjects for the application of the same scientific law or law-sketch" (ibid. p. 190). Instead of explaining a phenomenon from scratch (i.e. all the way down from laws), Kuhn maintained that scientists try to match the new phenomenon to one or more previous phenomena-plus-explanations.

In similar vein, Philip Kitcher argues that new phenomena can be 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). Unlike Kuhn, Kitcher proposes a rather concrete account of explanation, known as the "unificationist view", which he explicitly links to Kuhn's view by interpreting exemplars as argument patterns (ibid., p. 437-8). However, neither Kitcher's account is entirely formalised, leaving most intermediate steps in derivations and argument patterns undiscussed.

Thomas Nickles relates Kuhn's view to Case-Based Reasoning (Nickles 2003, p. 161). Case-Based Reasoning (CBR) is an artificial intelligence (AI) 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. Falkenhainer et al. 1989; Kolodner 1993; Veloso and Carbonell 1993; VanLehn 1998). Yet, to the best of my knowledge, CBR has never been employed to develop a formal model of applying science.

In this paper, I intend to show that an instantiation of CBR which is known as *Data-Oriented Processing (DOP)* (Bod 1998), and which has also been referred to as *Exemplar-Based Explanation (EBE)* (Bod 2004), can be used to develop a formalised model of applying science. While DOP/EBE was originally developed in AI, in particular for natural language processing (cf. Bod 1998; Kaplan 1996; Collins and Duffy 2002), it can be rather straightforwardly extended to exemplar-based explanation in science and engineering. The key idea of DOP is to represent derivations of phenomena by *trees*. Given a corpus of previously derived phenomena that function as exemplars, DOP proposes that new phenomena can be derived by combining *subtrees* from the corpus. In Scha et al. (1999) it is demonstrated that DOP can be defined as a probabilistic, recursive instantiation of CBR.

In the following section, I will first show how derivations in physics can be represented by trees that describe each step from laws to a description of a phenomenon. Next, I show how explanations of new phenomena can be constructed by combining subtrees from previously explained phenomena. In section 3, I demonstrate that an (extremely) large number of different derivations exists for virtually every phenomenon, even though these derivations are all subsumed under the same general laws. In other words, derivational explanation is *massively redundant*, a problem which has been largely underestimated in the philosophy of science. I argue that redundancy can be dealt with by maximising a notion of "derivational similarity" which mimics as closely as possible prior exemplars and which constructs new derivations out of *largest* possible derivational chunks. In section 4, I will show how the resulting DOP model can be extended to real-world systems and technological devices from fluid mechanics and hydraulics. In section 5, I will provide an excursion into a field at the other end of the technological spectrum, discussing some examples from language technology. What counts for fluid mechanics also counts for language technology: (linguistic) phenomena are derived not all the way down from theory, but from concrete derivations of previous phenomena. I contend that DOP provides a general model of "applying science" across different disciplines.

2. The DOP view of scientific explanation

What do derivational explanations in physics look like? Let's start with a simple textbook example. Consider the derivation in figure 1 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

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

Figure 1. Derivation of the Earth's mass according to Alonso and Finn (1996)

By substituting the data for the Moon, $r = 3.84 \cdot 10^8$ m and $P = 2.36 \cdot 10^6$ s, Alonso and Finn compute the mass of the Earth: $M = 5.98 \cdot 10^{24}$ 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 idealised, the derivation provides a concrete problem solution on which various other (idealised) phenomena can be modelled. 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.

To create a formal model that reuses derivational patterns, we first need a formal representation of derivations. Analogous to proof trees in formal logic, DOP proposes to represent derivations by tree structures which indicate how a mathematical description of a phenomenon (or problem) is compositionally derived from theoretical laws and conditions. Figure 2 shows how the derivation for the Earth's mass in figure 1 can be turned into a tree.

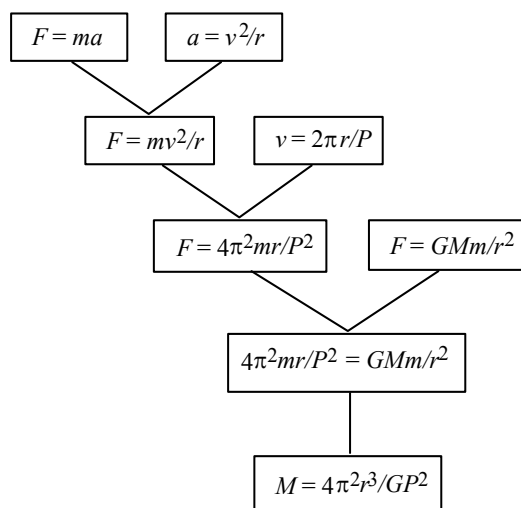


Figure 2. Derivation tree for the derivation in figure 1

The derivation tree in figure 2 represents the various derivation steps, insofar as they are carried out in figure 1, from higher-level laws to an equation for the mass of a planet. In general, a derivation tree is a finite tree in which each node is labelled with a formula; the boxes are only convenient representations of these labels. The formulas at the top of each "vee" (i.e. each pair of connected 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 in this tree 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 this case, for the mass M). Thus, a binary branch refers to a physical derivation step which introduces and combines physical laws or conditions, while a unary branch refers to a mathematical derivation step which solves an equation for (a) certain variable(s).

The notion of derivation tree is clearly reminiscent of the deductive-nomological (D-N) account of explanation by Hempel and Oppenheim (1948). In the D-N account, a phenomenon is explained if it can be deduced from general laws and antecedent conditions. But while the D-N model usually focuses on the initial premises (laws and antecedent conditions) and the final conclusion, a derivation tree describes each step in getting from general laws to a description of the phenomenon. DOP thus interprets Kitcher's "switch in conception from premise-conclusion pairs to derivations" (Kitcher 1989, p. 432) by tree representations. Although there may be other representations of derivations, trees are exceedingly flexible structures that can be decomposed into parts (*subtrees*) and be recomposed to derive new phenomena.

For example, consider the subtree in figure 3 which is extracted from the derivation tree in figure 2 by leaving out the last derivation step (i.e. the solution for the mass M).

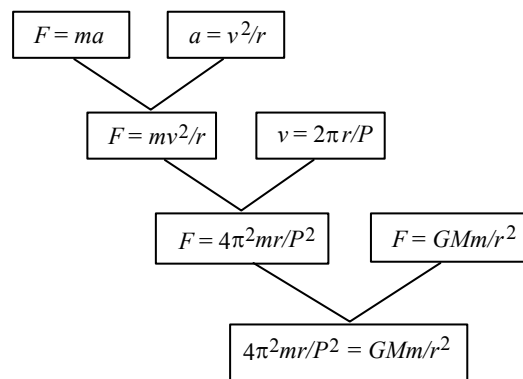


Figure 3. A subtree from figure 2 reflecting a theoretical model of a planet-satellite system

This subtree can be applied to a range of 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 3 needs only to be extended with a derivation step that solves the last equation for r^3/P^2 , as represented in figure 4.

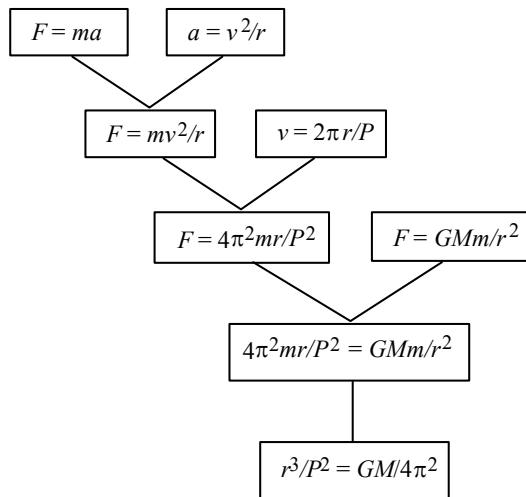


Figure 4. Derivation tree for Kepler's third law from the subtree in figure 3

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 3 for r and taking P as the rotation period of the Earth.

However, it is 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 3, but need to extract two smaller subtrees from figure 2 that are first combined by term substitution (represented by the operation " \circ ") and then solved for v in figure 5.

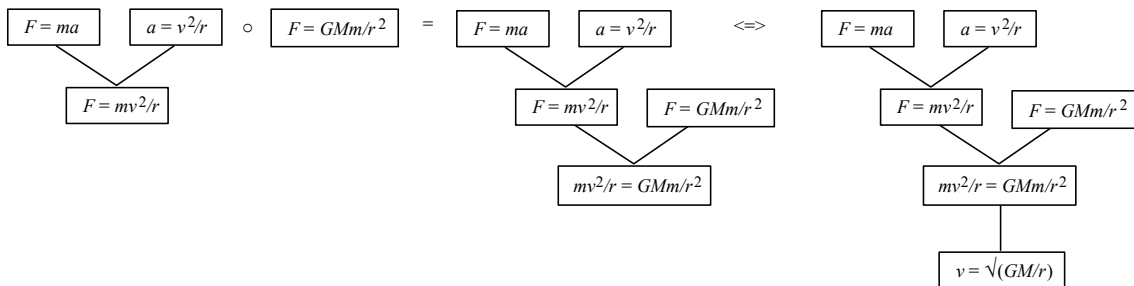


Figure 5. Constructing a derivation tree for a satellite's velocity by combining two subtrees from figure 2

Figure 5 shows that we can create new derivation trees by combining subtrees from previous derivation trees that function as exemplars. The notion of term substitution, though widely used in rewriting systems, may need some further specification. The combination of tree t and tree u , written as $t \circ u$, yields a tree that expands the root nodes of t and u to a new common root node where the righthandside of the equation at the root node of u is substituted in the corresponding term 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. This DOP model employs (1) a *corpus of derivation trees* representing exemplars and (2) a *matching procedure* that explains a phenomenon by combining subtrees from the

corpus into a derivation tree such that the root of the tree represents the phenomenon and the leaves represent laws or antecedent conditions. Note that subtrees can be of any size: from single equations to any combination of laws up to entire derivations. This reflects the continuum between laws, law-schemes and derivations in DOP.

3. Derivational redundancy and derivational similarity

Given a corpus of derivation trees and a mathematical description of a phenomenon, existing equational reasoning systems, such as *TKSolver* or *Mathematica*, can be employed to derive the phenomenon from the equations in the subtree-roots (see Baader and Nipkow 1998 for an overview on equational reasoning). However, there is a problem we have not considered so far: there can be many, sometimes extremely many, different derivation trees for the same phenomenon. In general, the number of derivation trees grows exponentially with the number of terms in the mathematical description of the phenomenon. To give a simple example, let's enlarge our corpus 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 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 idealised: no centrifugal force is taken into account, let alone influences from the Sun or other planets.) Thus the problem of the Earth's mass is derivationally redundant in that it can be solved in at least two different ways. And *both* derivations are used in Alonso and Finn's textbook as exemplars for deriving solutions to other problems.

When we add the tree in figure 6 to our corpus of exemplars, we can also derive Kepler's third law from this exemplar, resulting in the alternative derivation in figure 7, which uses a large subtree from figure 6 in combination with two small subtrees from the exemplar in figure 2.

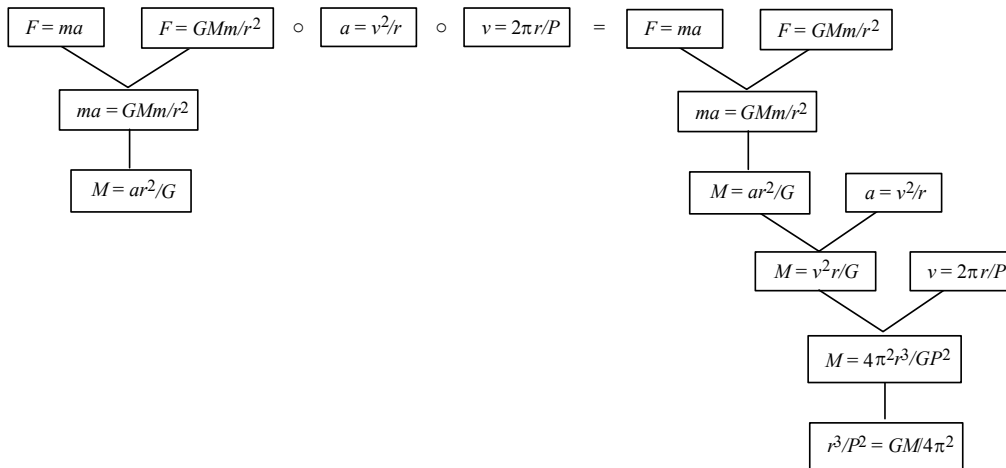


Figure 7. An alternative derivation tree for Kepler's third law

There is nothing wrong with this alternative derivation tree: there are no spurious non-explanatory laws that are irrelevant (as would be e.g. Hooke's or Boyle's law): Kepler's third law is subsumed under the same general laws in both cases. The main difference is that the derivation in figure 7 is modelled on a *different* exemplar, i.e., different from the exemplar the derivation in figure 4 is modelled on. The alternative derivation in figure 7 is insightful as it refers to the conceptual equivalence between terrestrial and celestial mechanics in Newtonian dynamics. The fact that Kepler's third law can be derived from figure 6 suggests that if we bring a satellite down to the Earth's surface it still follows the same law. Note that there are even more derivation trees for Kepler's law. By combining subtrees from the two exemplars in figures 2 and 7 in different ways, we get a combinatorial explosion of possible derivations.

Yet, experiments with advanced physics students show that no student comes up with the alternative derivation tree in figure 7 (see Bod 2004). Why? Apart from the fact that the derivation tree in figure 4 is slightly smaller, the tree in figure 4 is more "derivationally similar" to an exemplar in the corpus. That is, the tree in figure 4 can be constructed by just one large subtree from the corpus, whereas the tree in figure 7 needs at least 3 smaller subtrees to be constructed from the corpus.

This suggests that if we want to model the way humans derive phenomena, we should mimic as closely as possible the exemplary derivations that were learned from textbooks. The distinctive feature between different derivations of a phenomenon is that *some derivations are more similar to exemplary derivations than others*. The larger the partial match between a derivation and an exemplar the more "derivationally similar" they are. Since students learn physics not just by memorising laws, but also by studying exemplary problem solutions, I conjecture that they derive a phenomenon by maximising derivational similarity, or equivalently, by *minimising derivation length* where the length of a derivation is defined as the number of corpus-subtrees it consists of. I will refer to the derivation of minimal length also as the "shortest derivation". Since subtrees 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*.

DOP embodies the hypothesis that scientists try to explain a new phenomenon by maximising derivational similarity between the new phenomenon and previously derived phenomena. And the shortest derivation provides a possible way to attain this goal. The rationale behind maximising derivational similarity is that it favours

derivation trees which maximally overlap with previous derivation trees, such that only minimal recourse to additional derivational steps needs to be made.

Rather than trying to get hold on similarity between phenomena (see Sterrett 2006), I thus focus on the similarity between *derivations* of phenomena. This brings me back to Kuhn's problem, that is, how do we know on which exemplar a new phenomenon can be modelled (Kuhn 1970, p. 190)? Kitcher's account does not help us here. His "unificationist" view does not tell us whether we can best model, for example, the gravitational acceleration at a planet's surface on the exemplar in figure 2 or on the exemplar in figure 6. DOP's solution is to model a phenomenon on the exemplar from which the largest subtree can be reused to derive it. And hypotheses of the largest reusable subtree for deriving a phenomenon can be generated by best-first heuristics even before the problem gets actually solved (see Baader and Nipkow 1998). Thus DOP suggests a modification of Kitcher's account of explanation: rather than "using the *same* patterns of derivations again and again" as Kitcher argues, we use the *largest possible* patterns of derivations again and again.

We should keep in mind that the phenomena discussed so far are highly idealised and limited to textbook examples. There is *no* historical analogue of using parts from the derivation of the Earth's mass to derive Kepler's third law (it rather happened the other way round). Yet in science education the two problems are treated as closely interconnected, and with good reason: the two problems can be solved by using the same partial derivations. In the next section I will investigate how the DOP approach can be extended to real-world systems and technological devices. As an intermediate step, I could also have dealt with idealised phenomena that are *not* exactly solvable. A typical example is the three-body problem in Newtonian dynamics. Even if we make the problem unrealistically simple (e.g. by assuming that the bodies are perfect spheres that lie in the same plane), the motion of three bodies due to their gravitational interaction can only be approximated by techniques such as perturbation calculus. However, in perturbation calculus every derivation step still follows numerically from higher-level laws. Our main interest is in phenomena for which there are derivation steps that are *not* dictated by any higher-level law.

4. Extending DOP to real-world systems and technological devices

Derivations of real-world systems and technological devices are strikingly absent in physics textbooks. But they are abundant in engineering practice. As an example I will discuss a concrete system from hydraulics: (the velocity of) a jet emerging from a small orifice in a tank, which I will refer to as an *orifice system*. I have chosen this system because it functions as a shared example in hydraulics on which several other systems are modelled, and yet it has no rigorous solution from higher-level laws but involves additional coefficients. I will extend the original DOP model and show how a "derivation" of the orifice system can be used as an exemplar by this model for deriving other real-world systems such as weirs, notches and water breaks.

The orifice system is usually derived from Daniel Bernoulli's famous equation, which in turn is derived from the Principle of Conservation of Energy.¹ According to the Principle of Conservation of Energy the total energy of a system of particles remains constant. The total energy is the sum of kinetic energy (E_k), internal potential energy ($E_{p,int}$) and external potential energy ($E_{p,ext}$):

¹ Bernoulli used a precursor of this principle which was known as "Equality between the Potential Ascent and Actual Descent" (see Mikhailov 2002, p. 70).

$$\Sigma E = E_k + E_{p,int} + E_{p,ext} = \text{constant}$$

Applied to an incompressible fluid, the principle comes down to saying that the total energy per unit volume of a fluid in motion remains constant, which is expressed by Bernoulli's equation:

$$\rho gz + \rho v^2/2 + p = \text{constant}$$

The term ρgz is the external potential energy per unit volume due to gravity, where ρ is the fluid's density and z the height of the unit (note the analogy with mgh in classical mechanics). The term $\rho v^2/2$ is the kinetic energy per unit volume (which is analogous to $mv^2/2$ in classical mechanics). And p is the potential energy per unit volume associated with pressure. Bernoulli's equation is also written as

$$\rho gz_1 + \rho v_1^2/2 + p_1 = \rho gz_2 + \rho v_2^2/2 + p_2$$

which says that the total energy of a fluid in motion is the same at any two unit volumes along its path.

Figure 8 illustrates how the engineering textbook *Advanced Design and Technology* derives the orifice system (together with Torricelli's theorem) from Bernoulli's equation (Norman et al. 1990, p. 497):

We can use Bernoulli's equation to estimate the velocity of a jet emerging from a small circular hole or orifice in a tank, Fig. 12.12a. 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.

Figure 8. Derivation of Torricelli's theorem in Norman et al. (1990)

Thus the theoretically derived discharge of the system differs substantially from the actual discharge and is corrected by a coefficient of discharge, C_d . This is mainly due to an additional phenomenon which occurs in any orifice system: the *vena contracta*. Although this phenomenon has been known for more than three and a half centuries (cf. Torricelli 1644), no rigorous derivation exists for it and it is taken care of by a correction factor. The correction factor is not an adjustment of a few percent, but of almost 40%. The value of the factor varies however with the profile of the orifice and can range from 0.5 (the so-called Borda mouthpiece) to 0.97 (a rounded orifice).

Introductory engineering textbooks tell us that coefficients of discharge are experimentally derived corrections that need to be established for each orifice separately (see Norman et al. 1990; Douglas and Matthews 1996). While this is true for real-world three-dimensional orifices, there are analytical solutions for idealised two-dimensional orifice models by using free-streamline theory (see Batchelor 1967, p. 497). Moreover, Sadri and Floryan (2002) have shown that the *vena contracta* can also be simulated by a numerical solution of the general Navier-Stokes equations which is, however, *again* based on a two-dimensional model. For three-dimensional orifice models there are no analytical or numerical solutions (Munson 2002; Graebel 2002). The coefficients of discharge are then derived by physical modelling, i.e. by experiment. This explains perhaps why *physics* textbooks usually neglect the *vena contracta* in dealing with Torricelli's theorem. And some physics textbooks don't deal with Torricelli's theorem at all. To the best of my knowledge, all engineering textbooks that cover Torricelli's theorem also deal with the coefficient of discharge. (One may claim that the *vena contracta* can still be qualitatively explained: because the liquid converges on the orifice, the area of the issuing jet is less than the area of the orifice. But there exists no quantitative explanation of C_d for a three-dimensional jet.)

Although no analytical or numerical derivations exist for real-world orifice systems, engineering textbooks still link such systems via experimentally derived corrections to the theoretical law of Bernoulli, as if there were some deductive scheme. One reason for enforcing such a link is that theory does explain some important features of orifice systems: the derivation in figure 6, albeit not fully rigorous, explains why the discharge of the system is proportional to the square-root of the height h of the tank, and it also generalises over different heights h and orifice areas A . Another reason for enforcing a link to higher-level laws is that the resulting derivation can be used as an exemplar for solving new problems and systems. To show this, I will first turn the derivation in figure 8 into its corresponding derivation tree. But how can we create such a derivation tree if the coefficient of discharge is not derived from any higher-level equation? The orifice system indicates that there can be phenomenological models that are not derived from the theoretical model of the system. Yet, when we write the

coefficient of discharge as the empirical generalisation $Q(actual) = C_d Q(theoretical)$, which is in fact implicit in the derivation in figure 8, we can again create a derivation tree and "save" the phenomenon. This is shown in figure 9 (where we also added the principle of conservation of energy).

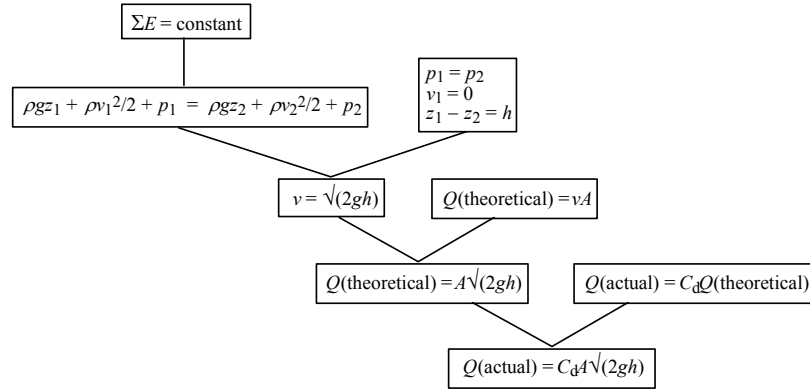


Figure 9. Derivation tree for the derivation in figure 8

The tree in figure 9 closely follows the derivation given in figure 8, 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 equation $Q(actual) = C_d Q(theoretical)$. Although this equation does not follow from any higher-level law or principle, we can use it *as if* it were a law. Of course it is not a law in the universal sense; it is a correction, a rule of thumb, but it can be reused for a range of other hydraulic systems, ranging from nozzles, notches, weirs, open channels and many pipeline problems -- see Douglas and Matthews (1996).

What does this mean for DOP? By using the derivation tree in figure 9 as an exemplar and by using the same substitution mechanism for combining subtrees from exemplars as in section 2, together with a mathematical procedure that can solve an equation (see Baader and Nipkow 1998), we obtain an exemplar-based model for fluid mechanics that can explain a range of new real-world systems. For example, the three subtrees in figure 10 can be extracted from the derivation tree in figure 9 and be reused in deriving the rate of flow of a rectangular *weir* of width B and height h (see e.g. Norman et al. 1990, p. 498).




Figure 10. Three subtrees from figure 9 that can be reused to derive a description of a weir system

By adding the mathematical equivalence $vA = \int v dA$ and the equation $dA = B dh$, which follows from the definition of a rectangular weir, we can create the derivation tree in figure 11 for the discharge of a weir.

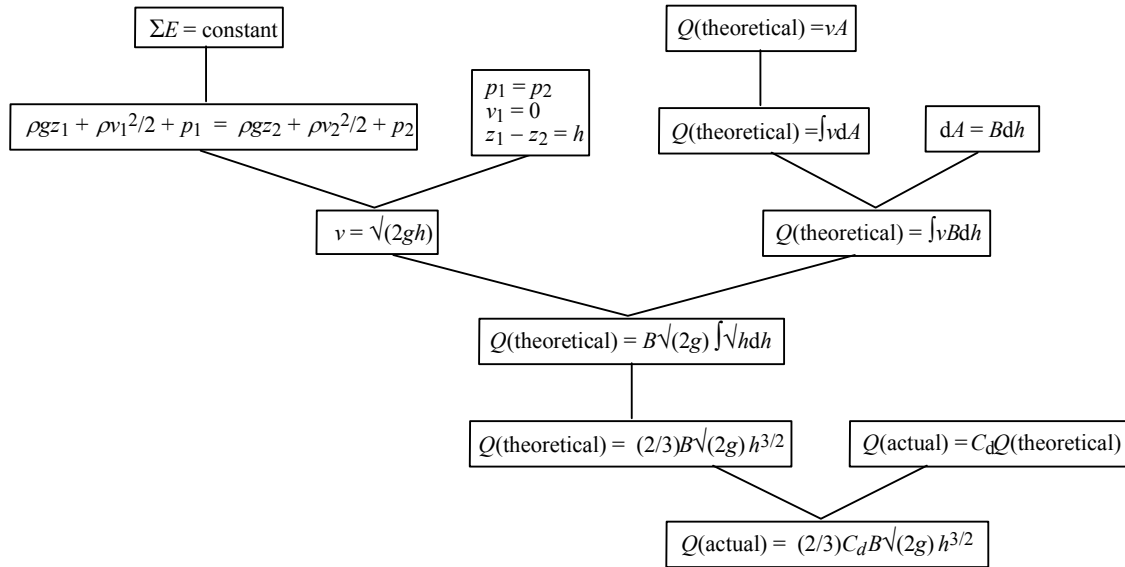


Figure 11. Derivation tree resulting from the shortest derivation for a description of a weir by combining the subtrees from figure 10

The derivation tree in figure 11 closely follows the derivations given in Norman et al. (1990, p. 498) and Douglas and Matthews (1996, p. 117), where a description of a weir system is derived by modelling it as closely as possible on the orifice system. This corresponds to engineering practice where new systems are almost literally built upon or constructed out of similar previous systems. To give an historical example, the earliest known derivation of the equation for the weir system by Jean-Baptiste Bélanger (1828, p. 37) takes the orifice system as given and reuses it to derive the formula in figure 11. Bélanger reused not only Bernoulli's derivation of Torricelli's theorem (as reflected by the leftmost subtree in figure 10) but also the empirical coefficient of discharge (the rightmost subtree in figure 10) and the equation for discharge itself (the intermediate subtree in figure 10). Bélanger thus modelled the weir system by reusing and extending (parts of) a previously derived system in such a way that only minimal recourse to additional derivational steps was needed.

DOP can simulate this exemplar-based modelling by combining those derivational chunks that maximise derivational similarity or, equivalently, minimise derivation length. The three subtrees in figure 10 indeed correspond to the smallest number of subtrees needed to construct, via some intermediate mathematical steps, a derivation for the equation of the weir system. There are many other possible derivations for this equation (the system is massively derivationally redundant), but they all result in considerably longer derivations than the one given above.

The derivation tree in figure 11 has effectively become an exemplar itself in hydraulics (which in DOP means that it is added to the corpus). The derivation has been reused and extended to derive a so-called *V-notch*. The *V-notch*, in turn, has been extended to derive a so-called *trapezoidal notch*, which has again been further extended to derive a *Cipolletti weir*, etc. (see Chanson 2002 for an overview). Modelling in engineering is highly cumulative: new systems are built upon or constructed out of previous systems and their derivations form increasingly complex wholes. We can handle this complexity by taking large(st) partial derivations from (descriptions of) previous systems as "given" (as in DOP) and work from there.

Note that the DOP approach naturally allows for using corrections, normalisations, fudge factors and other adjustments. Morrison and Morgan (1999, p. 11) refer to such

adjustments as "additional 'outside' elements" but in the exemplar-based approach they have the same status as other equations in a derivation tree. As long as "additional 'outside' elements" can be stated in terms of mathematical equations they can be integrated by a derivation tree, and be reused to solve new problems by maximising derivational similarity. We only need to slightly extend our definition of DOP given at the end of section 2, where leaf nodes in a derivation tree referred to either general laws or antecedent conditions. In the new DOP model, the leaf nodes may also be empirical rules or any other equations that are not deduced from higher-level laws. We may lump these three kinds of knowledge (laws, antecedent conditions and empirical rules) together as *knowledge that is not derived from higher-level knowledge*. The definitions of derivational similarity and shortest derivation remain the same.

But how far does DOP stretch for technological devices and real-world systems? Can we also represent more complex systems by means of derivation trees? DOP hinges on the fact that all "theory-external" knowledge must be represented by equations. But what if such knowledge involves an intermediate model, as in the case with Prandtl's boundary layer model discussed by Heidelberger (2006)? Morrison (2006) rightly notes that the boundary layer model is autonomous in that it was not created by some approximation of the Navier-Stokes equations. Yet, as Heidelberger shows, 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. 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 reused as an exemplar by DOP to derive a range of new engineering problems.

Note that DOP does not demand that every phenomenon or system be linked to universal laws. A phenomenon may be derived from a phenomenological model only, without any derivational relation to high-level theory. This is the case in, for example, quantum-chromodynamics (QCD), 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, and they were even taken themselves from exemplars in other fields, albeit 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 (see next section). As long as there is a model or regularity representing the phenomenon, we can construct a derivation tree for it. In our new definition of 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.

5. DOP in other disciplines: an excursion into language technology

What counts for fluid mechanics and hydraulics also counts for many other disciplines: real-world phenomena are derived not from general laws, but from (largest possible) parts of derivations of previous phenomena. As an example from the other end of the technological spectrum, I will give a brief excursion into language technology.

While language *theory* is permeated by the idea that a language is aptly described by a formal grammar, i.e. a finite and succinct set of rules which can derive an infinite set of well-formed utterances, language *technology* does not work that way. As soon as natural language processing systems need to deal with a non-trivial fragment of a language, say English, formal grammars turn out to be severely inadequate. Grammars either *undergenerate*, which means that they provide no derivation for otherwise well-formed utterances, or they *overgenerate*, which means that they provide too many derivations for well-formed utterances (cf. Manning and Schütze 1999). "All grammars

leak", is the well-known dictum of Edward Sapir (Sapir 1921, p. 38). In fact, there are so many idiosyncratic and idiomatic phenomena in natural language that only an approach which takes into account a stock of previously produced sentences can accurately model a language. After unsuccessful attempts to apply formal grammars to automatic linguistic analysis, a different paradigm has been developed since the 1980s in language technology: new sentences are derived not by using a concise set of rules, but by using a large corpus of previously derived sentences together with a matching algorithm (see Manning and Schütze 1999 for an historical overview).²

Before going into the details of this matching algorithm for linguistic analysis, let me first explain what derivations of sentences look like. It is by now widely acknowledged that sentence derivations can be represented by tree structures, similar to the derivation trees in physics in the previous sections. The first linguistic tree structure was (most likely) proposed by Wilhelm Wundt in his *Logik* (Wundt 1880). But it was Noam Chomsky who made the notion of *syntactic phrase-structure tree* more widely accepted (Chomsky 1957). Although richer structures have also been proposed in the meantime, there is ample agreement that tree structures form the backbone of sentence-analyses, sometimes enriched with phonological, morphological and semantic representations (see Sag et al. 2003; Bresnan 2000; Goldberg 1995). In this section, I will focus on syntactic representations only.

So what does a syntactic phrase-structure tree look like? Figure 12 gives two tree structures for respectively the sentences *She wanted the dress on the rack* and *She saw the dog with the telescope*.

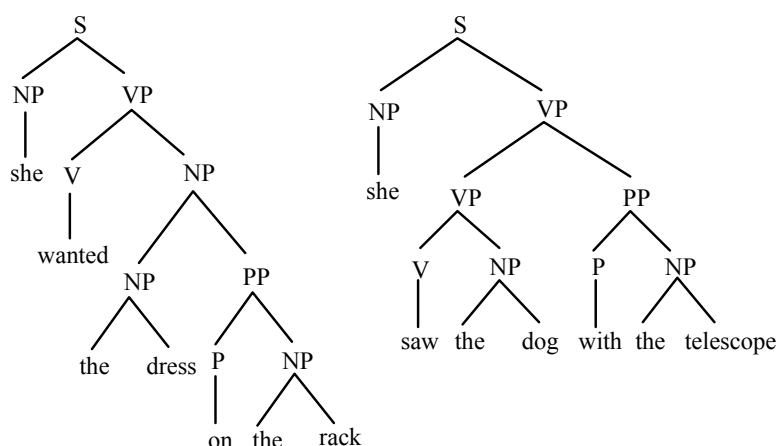


Figure 12. Two sentences with their phrase-structure trees

A phrase-structure tree describes how parts of a sentence combine into constituents and how these constituents combine into a representation for the whole sentence. The constituents in a phrase-structure tree are labelled with syntactic categories such as *NP* for noun phrase, *PP* for prepositional phrase, *VP* for verb phrase and *S* for the whole sentence. To keep the example simple, we have left out some low-level labels for Noun and Article. The two trees in figure 12 are structurally different in that in the first sentence the prepositional phrase *on the rack* forms a noun phrase with *the dress*, whereas in the second sentence the prepositional phrase *with the telescope* forms a verb phrase with *saw the dog*.

² Even if the notion of "grammar" is still used by many systems it is not succinct but consists of (tens of) thousands of rules that are derived from actual language corpora (see e.g. Knuutila and Voutilainen 2003).

Although phrase-structure trees are not labelled with equations, they are compositionally built up as in physics derivation trees: each category is defined in terms of its underlying subcategories (and if we enrich each syntactic label with its logical-semantic interpretation, we would again obtain derivation trees with equations -- see Bod 1998). Note that phrase-structure trees are represented upside down: the root is at the top instead of at the bottom. This is pure convention.

How can these sentences be used to derive new sentences, i.e. what does a "matching algorithm" for language look like? There is no single (or unique) way to do this. One straightforward but not very successful method is to read off the "grammar rules" that are implicit in the trees, such as $S \Rightarrow NP VP$, $VP \Rightarrow V NP$, $NP \Rightarrow NP PP$, $NP \Rightarrow the\ dress$, etc. in figure 12 (Charniak 1996).³ Another, more successful method is by reading off for every single word a subtree including that word (Chiang 2000). Yet another and still more successful method is by first enriching each syntactic label with its so-called "headword" and by next reading off the rules from the trees (Collins 1997). We know the relative success of these methods as they have been evaluated on the same benchmark, the so-called Penn Treebank corpus consisting of 50,000+ manually analysed sentences (Marcus et al. 1993). We will not go into further details of these different methods (but see Bod et al. 2003a, 2003b).

While these methods may seem rather disparate, they are based on the same underlying idea: new sentences are derived by parts of previously derived sentences. The distinctive feature of each method is their definition of what are to be considered the *relevant* parts. Yet it is also possible to generalise over these different methods by taking *all* partial trees as "relevant" parts. Interestingly, this generalised model is captured by the DOP approach. If we put restrictions on the partial trees, all aforementioned other models can be instantiated by DOP (see Charniak 1997).

The example in figure 13 illustrates how DOP is used for natural language processing. If we take the sentences in figure 12 as our (unrealistically small) corpus, we can derive the new sentence *She saw the dress with the telescope* by extracting two subtrees from the trees in figure 12 and by combining them by means of *label substitution*.

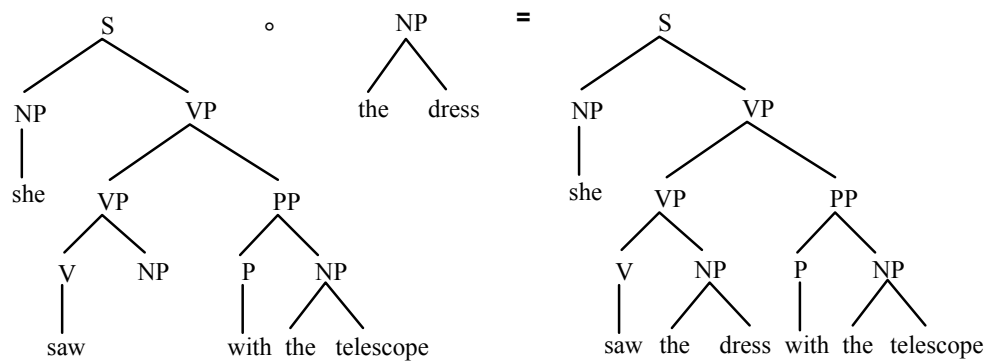


Figure 13. Deriving a new sentence by combining subtrees from figure 12

Note the similarity of label substitution with term substitution in the previous sections. Analogous to term substitution, the label substitution of two trees, again written as $t \circ u$, yields a tree where u is substituted on the leftmost syntactic leaf node of t . As in the DOP model for physics, the linguistic DOP model constructs new trees by combining

³ A grammar rule like $S \Rightarrow NP VP$ says that a sentence (S) consists of a noun phrase (NP) followed by a verb phrase (VP).

subtrees from prior trees. The analogy may be even stronger for hydraulics: just as engineers build new systems almost literally out of previous systems, language users construct new sentences almost literally out of previous sentences. Like in engineering, this can result in increasingly complex wholes (for example, by iteratively combining subtrees from the corpus in figure 12, we can construct sentences of arbitrary length, such as *She saw the dress on the rack with the telescope* and *She saw the dress with the dog on the rack with the telescope*, etc. -- see Bod 2006).

However, the most important commonality between the two disciplines is probably the way derivational redundancy is resolved. Turning back to figure 13, the new sentence *She saw the dress with the telescope* is interpreted analogous to the corpus sentence *She saw the dog with the telescope*: both sentences receive roughly the same phrase structure. But we can also derive an alternative phrase structure for this new sentence, namely by combining three (rather than two) subtrees from figure 12, as shown in figure 14.

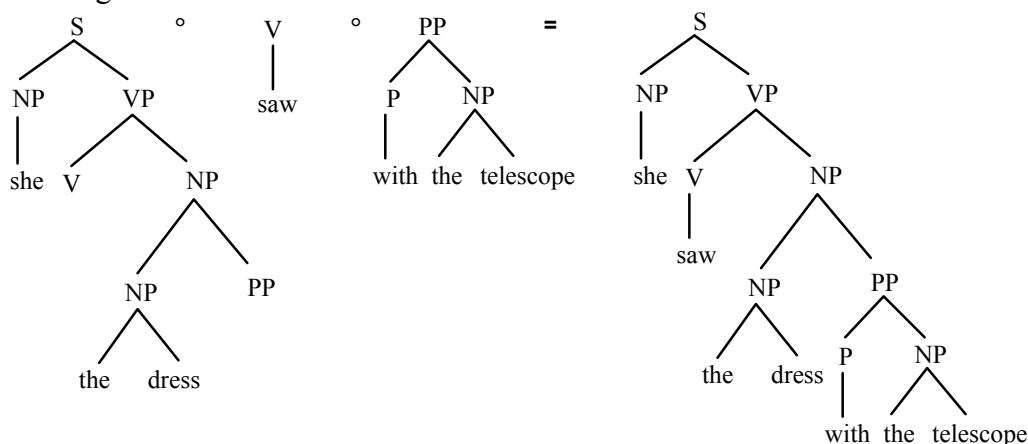


Figure 14. A different derivation for *She saw the dress with the telescope*

Thus the sentence *She saw the dress with the telescope* can be derived in (at least) two different ways: either analogous to the first tree in figure 12 or analogous to the second tree in figure 12. That is, the sentence is "structurally ambiguous" or "derivationally redundant", and the two different structures represent two different meanings for the sentence (due to the different prepositional phrase attachments). Which tree should be chosen? Experiments with human language users show that people have a strong preference for understanding the sentence as conveyed by the syntactic structure in figure 13. The reason for this seems to be that, as in physics, the structure in figure 13 is more derivationally similar to previously perceived sentence-structures (although our small corpus in figure 12 is of course not representative). Humans tend to understand a new sentence by modelling it as closely as possible on previously understood sentences. This can be accomplished by maximising derivational similarity -- or equivalently, by minimising derivation length -- with respect to previously derived sentences.⁴ For our example sentence *She saw the dress with the telescope*, the shortest derivation (which maximises derivational similarity) is represented by figure 13: only two subtrees from the corpus are needed to construct this tree, while at least three corpus-subtrees are needed to construct the tree in figure 14.

Of course, the corpus in figure 12 is much too small for simulating actual language processing. More realistic experiments use corpora of tens of thousands of

⁴ Most models also take into account the frequency of occurrence of derivational chunks in the corpus (see Manning and Schütze 1999), but I will not go into this here.

(manually constructed) phrase-structure trees (cf. Bod et al. 2003b). Yet it is noteworthy that the use of largest possible derivational chunks is important not only for solving derivational redundancy, but also for covering idiomatic expressions and multi-word units such as "to take advantage of", "more pricks than kicks" and "everything you always wanted to know about X but were afraid to ask". Natural language is replete with idiosyncratic and idiomatic expressions, which can only be taken into account by units larger than single grammar rules or lexical entries.

Note that DOP allows in principle for corpora of any sort of trees -- be they physical, linguistic, musical or of any other kind (see Bod 2002a, 2002b). As long as we can construct a corpus of exemplary derivations for a certain discipline, we can create a DOP model for it and use it to derive new phenomena without making recourse to an axiomatic system of rules. Sure enough, the exemplary derivations in the corpus do include general laws or rules, such as $F = ma$ in physics or $S \Rightarrow NP VP$ in linguistics, but they also include very particularist information ranging from empirical coefficients in hydraulics to idiomatic expressions in natural language that do not follow from these laws or rules. The world may be full of nomothetic elements like laws. But it is also full of idiographic elements such as empirical coefficients. DOP does justice to both.

6. Conclusion

I have argued for a general model of "applying science", known as DOP or EBE. According to this model, we explain new phenomena by recombining largest possible fragments or chunks from a corpus of previously derived phenomena. Examples from hydraulics and natural language processing suggest an underlying methodology across different disciplines. As long as we can construct a corpus of exemplary derivations for a certain discipline, we can create a DOP model for it and use it to derive descriptions of new phenomena.

References

- Alonso, Marcelo and Edward Finn (1996), *Physics*, Addison-Wesley, Harlow.
- Baader, Franz and Tobias Nipkow (1998), *Term Rewriting and All That*, Cambridge University Press, Cambridge.
- Batchelor, G. (1967), *An Introduction to Fluid Dynamics*, Cambridge University Press, Cambridge.
- Bélangier, Jean-Baptiste (1828), *Essai sur la Solution Numérique de quelques Problèmes Relatifs au Mouvement Permanent des Eaux Courantes*, Carilian-Goeury, Paris.
- Bod, Rens (1998), *Beyond Grammar*, CSLI Publications, Stanford.
- Bod, Rens (2002a). "A unified model of structural organization in language and music", *Journal of Artificial Intelligence Research*, 17(2002), 289-308.
- Bod, Rens (2002b). "Memory-based models of melodic analysis: challenging the Gestalt principles", *Journal of New Music Research*, 31(1), 27-37.
- Bod, Rens (2004), "Exemplar-Based Explanation", *Proceedings E_CAP'2004*, Pavia, Italy.
- Bod, Rens (2006), "Exemplar-Based Syntax: How to Get Productivity from Examples?", *The Linguistic Review (Special Issue on Exemplar-Based Linguistics)*, in press.
- Bod, Rens, Jennifer Hay and Stefanie Jannedy (2003a), *Probabilistic Linguistics*, The MIT Press, Cambridge.
- Bod, Rens, Remko Scha and Khalil Sima'an (2003b), *Data-Oriented Parsing*, University of Chicago Press, Chicago.
- Boon, Mieke (2006), "How Science is Applied in Technology", *International Studies in the Philosophy of Science*, this issue.
- Boumans, Marcel (1999), "Built-in justification", in M. Morgan and M. Morrison (eds.), *Models as Mediators*, Cambridge University Press, Cambridge, 66-96.
- Bresnan, Joan (2000), *Lexical-Functional Syntax*, Blackwell, Oxford.
- Cartwright, Nancy (1983), *How the Laws of Physics Lie*, Oxford University Press, Oxford.
- Cartwright, Nancy (1999), *The Dappled World*, Cambridge University Press, Cambridge.
- Chanson, Hubert (2002), *The Hydraulics of Open Channel Flow*, Butterworth-Heinemann, Oxford.

- Chiang, David, (2000), "Statistical parsing with an automatically extracted tree adjoining grammar", *Proceedings ACL'2000*, Hong Kong, China.
- Charniak, Eugene (1996), "Tree-bank Grammars", *Proceedings AAAI'96*, Menlo Park, California.
- Charniak, Eugene (1997), "Statistical Techniques for Natural Language Parsing", *AI Magazine*, 1997(4), 16-25.
- Chomsky, Noam (1957), *Syntactic Structures*, Mouton, The Hague.
- Collins, Michael (1997), "Three Generative Lexicalised Models for Statistical Parsing", *Proceedings ACL'97*, Madrid, Spain.
- Collins, Michael and Nigel Duffy (2002), "New Ranking Algorithms for Parsing and Tagging: Kernels over Discrete Structures, and the Voted Perceptron", *Proceedings ACL'2002*, Philadelphia.
- Douglas, J. and R. Matthews (1996), *Fluid Mechanics*, Vol. 1, 3rd edition, Longman, Harlow.
- Falkenhainer, B., K. Forbus and D. Gentner (1989) "The Structure-Mapping Engine: Algorithm and Examples", *Artificial Intelligence*, 41, 1-63.
- Giere, Ronald (1988), *Explaining Science: A Cognitive Approach*, University of Chicago Press, Chicago.
- Giere, Ronald (1999), *Science without Laws*, The University of Chicago Press, Chicago.
- Goldberg, Adele (1995), *Constructions*, The University of Chicago Press, Chicago.
- Graebel, W. (2001), *Engineering Fluid Mechanics*, Taylor & Francis, New York.
- Hartmann, Stephan (1999), "Models and stories in hadron physics", in M. Morgan and M. Morrison (eds.), *Models as Mediators*, Cambridge University Press, Cambridge, 326-346.
- Heidelberger, Michael (2006), "Models in Fluid Dynamics", *International Studies in the Philosophy of Science*, this issue.
- Hempel, Carl and Paul Oppenheim (1948), "Studies in the Logic of Explanation", *Philosophy of Science*, 15, 135-175.
- Kaplan, R. (1996), "A Probabilistic Approach to Lexical-Functional Analysis", *Proceedings of the 1996 LFG Conference and Workshops*, CSLI Publications, Stanford.
- Kitcher, Philip (1989), "Explanatory unification and the causal structure of the world", in Kitcher, Philip, and Salmon, Wesley (eds.), *Scientific Explanation*, University of Minnesota Press, Minneapolis, 410-505.
- Kolodner, J. (1993), *Case-Based Reasoning*, Morgan Kaufmann, Menlo Park.
- Knuutila, Tarja and Atro Voutilainen (2003), "A parser as an epistemic artefact: a material view on models", *Philosophy of Science*, 70(5), 1484-1495.
- Kuhn, Thomas (1970), *The Structure of Scientific Revolutions*, 2nd edition, University of Chicago Press, Chicago.
- Manning, Christopher and Hinrich Schütze (1999), *Foundations of Statistical Natural Language Processing*, The MIT Press, Cambridge.
- Marcus, M., B. Santorini and M. Marcinkiewicz, (1993), "Building a Large Annotated Corpus of English: the Penn Treebank", *Computational Linguistics*, 19, 313-330.
- Mikhailov, Gleb (2002), *Die Werke von Daniel Bernoulli*, Band 5, Birkhäuser Verlag, Basel.
- Morgan, Mary and Margaret Morrison (eds.) (1999), *Models as Mediators*, Cambridge University Press, Cambridge.
- Morrison, Margaret (2006), "Applying Science and Applied Science: What's the Difference?", *International Studies in the Philosophy of Science*, this issue.
- Morrison, Margaret and Mary Morgan (1999), "Models as mediating instruments", in M. Morgan and M. Morrison (eds.), *Models as Mediators*, Cambridge University Press, Cambridge, 10-37.
- Munson, Bruce (2002). *Fundamentals of Fluid Mechanics*, Wiley, London.
- Nickles, Thomas (2003), "Normal science: from logic to case-based and model-based reasoning", in Thomas Nickles (ed.), *Thomas Kuhn*, Cambridge University Press, Cambridge, 142-177.
- Norman, Eddie, Joyce Riley and Mike Whittaker (1990), *Advanced Design and Technology*, Longman, Harlow.
- Sadri, R. and Floryan, J. M. (2002), "Entry flow in a channel", *Computers and Fluids*, 31, 133-157.
- Sag, Ivan, Thomas Wasow and Emily Bender (2003), *Syntactic Theory: A Formal Introduction*, CSLI Publications, Stanford.
- Sapir, Edward (1921), *Language*, Hartcourt Brace & Company, San Diego.
- Scha, Remko, Rens Bod and Khalil Sima'an (1999). "Memory-Based Syntactic Analysis", *Journal of Experimental and Theoretical Artificial Intelligence*, 11(3), 409-440.
- Sterrett, Susan (2006), "Models of Machines and Models of Phenomena", *International Studies in the Philosophy of Science*, this issue.
- Toricelli, Evangelista (1644), *De Motu Gravium Naturaliter Descendentium et Proiectorum Libri duo*, Florentiae, Florence.
- 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.
- Wundt, Wilhelm (1880), *Logik. Eine Untersuchung der Prinzipien der Erkenntnis und der Methoden Wissenschaftlicher Forschung*, Enke, Stuttgart.