BAC0N.5: THE DISCOVERY OF CONSERVATION LAWS[1]

Pat Langley
Gary L. Bradshaw
Herbert A. Simon
Department of Psychology
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213 USA

ABSTRACT

BACON.5 is a program that discovers empirical laws. The program represents information at varying levels of description, with higher levels summarizing the levels below them. The system applies a small set of data-driven heuristics to detect regularities in numeric and nominal data. These heuristics note constancies and trends, leading BACONS to formulate hypotheses, define theoretical terms, and postulate intrinsic properties. Once the program has formulated an hypothesis, it' uses this to reduce the amount of data it must consider at later times. A simple type of reasoning by analogy also simplifies the discovery of laws containing symmetric forms. These techniques have allowed the system to rediscover Snail's law of refraction, conservation of momentum, Black's specific heat law, and Joule's formulation of conservation of energy. Thus, BACON.S'S heuristics appear to be general mechanisms applicable to discovery in diverse domains.

Introduction

The dual notions of symmetry and conservation lie at the very heart of physics, and the assumption of symmetry has been a powerful aid in the discovery of many physical laws. In this paper we describe BACON.5, a program that discovers conservation laws through a simple form of analogical reasoning, BACON.5 is not intended as a detailed model of historical discoveries. Rather, the system presents one of many ways in which historical discoveries could have occurred. The system is named after Sir Francis Bacon (1561-1626), the early philosopher of science. The data-driven heuristics of the BACON 5 program are in much the same spirit as those proposed by Bacon the philosopher over three hundred years ago.

Previous discovery systems such as AM [1] and METADENDRAL [2] relied strongly on theory-driven methods of discovery. In contrast, BACON.s employs data-driven heuristics to direct its search for interesting laws. A major goal of our research has been the identification of general discovery principles, and we feel that generality is more likely to reside in data-driven approaches than in theory driven ones, BACON.s is the fifth in a line of discovery systems developed by the authors [3, 4, 5]. The current system differs from its precursors in its data collection strategies, and in its ability to make symmetry assumptions, note simple analogies, and deal with noise. As an introduction, we consider BACON.S'S discovery of the ideal gas law in the following section. After this, we discuss the program's heuristics and control structure. Next we consider in detail BACON.S'S performance on some additional tasks. Finally, we draw conclusions from these results, and propose some directions for future efforts.

Thto research was supported by Grant IST-7918266 from the National Science Foundation and Grant MH 07722 from the National Institutes of Mental Health.

BACON.S, like its predecessors in the BACON series, addresses itself to that part of the scientific process that starts with a set of data and seeks to discover regularities in them. Of course, this is only one of several ways in which scientific discovery takes place, and only one of the steps (albeit a crucial one) in data-driven discovery. Scientists must also determine what data to gather, must invent and use appropriate instruments for making observations, must draw out potentially observable consequences from theories, and must analogize from one phenomenon to another. Thus, BACON.S'S capabilities provide an account of one important part of the scientific enterprise, while leaving other equally important components to future investigation.

An Example: The Ideal Gas Law

The ideal gas law may be stated as pV = nRT, where p is the pressure on a quantity of gas, V is the volume, n is the amount of gas in moles, T is its temperature, and R is the constant 8.32. BACON.5 discovers this law in a straightforward manner. The system begins by holding the values of n and T constant (say at n = 1 and T s 300), while varying the values of the pressure p. When the resulting values of V are examined, BACON.S notes that the values of V decrease as those of p increase. This trend suggests that the product pV should be considered; when the values of this term are computed, they are found to have the constant value 2496.0. Now the program examines this phenomenon when n s 1 and T = 310; in this case, pV is 2579.2. Similarly, pV is2662.4 when n =1 and T= 320.

At this point, BACON.S has enough information to relate the values of pV and the temperature T. These terms are linearly related with an intercept of 0, making the ratio pV/T equal to 8.32. Now the discovery system can vary its third independent term by setting n to 2. Upon gathering the relevant data, pV/T is found to be 16.64, while when n is set to 3, it becomes 24.96. When it compares the values of n and pV/T, BACON.S finds another linear relation with a zero intercept. The resulting equation, pV/nT = 8.32, is equivalent to the ideal gas law.

Standard analyses of the scientific method partition the work) into data or observations, and hypotheses or laws that explain those data, BACON.S replaces this dichotomy with information represented at varying levels of description. A description at one level acts as an hypotheses with respect to the descriptions below it, and as a datum for those above it. In the case of the ideal gas law, the direct observations of n, T, p, and V represent the first level of description, BACON.S treats these as data, generating the hypothesis that pV is 2496.0 when n is 1 and T is 300. Next the program treats these values and others at the second level as data, leading to the rule that pV/T is 8.32 when n is 1 at the third level. This datum and others like it produce the final fourth level law, that pV/nT = 8.32 for all cases. As we shall see below, much of BACON.S'S power derives from its ability to recursively apply the same heuristics at every level of description.

The Heuristics of BAC0N.5

One of the defining features of intelligence is the ability to search. In this section we discuss the heuristics used to direct BACON.S'S attention. We focus initially on the discovery system's techniques for gathering data. Next we consider the central heuristic - looking for constancies • followed by the strategies responsible for defining theoretical terms. We then examine the processes which lead to the proposal of intrinsic properties. Finally, we discuss BACON.S'S heuristics for reasoning by analogy and dealing with symmetrical laws.

Gathering Data

BACONS begins a task by asking the user for the independent terms it should consider, the values of these terms it should examine, and the dependent variables it should inspect. Having gathered this information, it starts to run experiments by varying *one* of the terms and holding all others constant. If this leads to a law, the program stores the term found to be constant with the term which was varied. Now BACON.S alters the value of another independent term. However, since it has a strong expectation about the form of the law it will find at the level below, it does not vary[2] the first term again. Instead, K selects a single independent value, examines the associated dependent value, and computes the expected constant from these. Similarly, once it has discovered a law at the second level, BACON.S need not vary the second term again. This method requires significantly less data than the complete factorial design technique used in earlier versions of BACON [3]. Figure 1 presents the data search trees for the two methods; the factorial design approach is shown by the light lines, while BACON.S'S search is superimposed in darker lines.

Noting Constancies

After manipulating an independent term which it has not varied before, BACONS examines each of the dependent terms at the current level of description. If the values of a dependent term are constant, it passes this variable and its value to the next higher level, where they are treated as data. In determining whether a set of values is constant, BACON.S first takes the mean of those values. It then multiplies this mean by a user-modifiable fraction (say 0.05), and creates an interval by adding/subtracting the resulting product to the mean. Next, each of the values is tested to see if it falls within the interval, BACON.S also computes the number of allowable exceptions by multiplying a second parameter by the number of values. If this product exceeds the number of exceptions, then the test succeeds and an hypothesis is formulated; otherwise the system must look elsewhere.

If a dependent term is not constant, then processing passes on to other heuristics. Some of these define arithmetic combinations of the independent and *dependent* terms. When this occurs, the program computes the values of the new term and sees if they are constant. Since BACON.S'S central goal is to find invariance, this process of defining terms and testing for recurrence continues until a constant is found or the terms become too complex. In addition to applying to terms of different complexity, the same constancy detector applies to ail levels of description, whether the first or the final independent term has been varied.

[2]This to not entirely true: W 1*00*6 has previously found a linear edition with a nonzero Intercept and a atopo other then 1.0, it requires at least two points to determine the slope end 'intercept under the new conditions. Also, if significant noise is present, one may use mutiple observations to achieve a reliable estimate.

Finding Linear Relations

If a term fails the test for invariance described above, BACON.S looks for a linear relation between that term and other variables. This test uses a standard linear regression technique to determine the slope of the best line relating two variables. Using this slope, it then computes the intercepts predicted by each observation. If these intercepts pass the test for constancy discussed above, then the system creates names for the slope and intercept, and returns their values to the next higher to be treated as data.

Certain values for slopes and intercepts are handled as special cases. If x and y are linearly related, then the slope s may be expressed as (y • i)/x, where i is the intercept. However, if the intercept is 0, this simplifies to s s y/x; in such cases, BACONS assumes that all future intercepts relating x and y will be 0, and defines the ratio y/x. Similarly, the intercept can in general be expressed as i = y * sx, but this reduces to I * y • x when the slope is 1. In this case, the program assumes all future slopes will be 1 and defines the difference y • x.

As with constants, the system may not find a linear relation between the first terms it considers. In such cases, the heuristics discussed below may generate new terms which are then examined. For example, consider BACON.S'S discovery of Ohm's law for electric circuits. Here the program manipulates the length I of the wire connected to a battery, and measures the resulting current I. The values of I are not constant, nor are I and L linearly related. However other factors lead BACON.S to consider the product IL, and this term is linearly related to I. The intercept of this line corresponds to the voltage of the battery, while the intercept corresponds to its internal resistance.

Defining Theoretical Terms

The slopes, intercepts, ratios and differences discussed above may be viewed as types of theoretical terms. They are not directly observable, but their values can be computed from observables. Although a term like IL or pV/nT may be replaced by its definition at any time, its use can greatly simplify the statement of a law. BACON.S invokes a third heuristic when both of the above tests fail, which also proposes useful theoretical terms. If a monotonic increasing relation is found between two terms and the signs of their values are the same, then the ratio of those terms is defined. If a monotonic decreasing trend occurs and the signs agree, then a product is defined. When the signs differ, the reverse events occur. The test for monotonicity allows for a certain number of exceptions before the test fails. Once a term has been defined, its values are computed and passed to the tests for constancy and linearity.

In its discovery of Kepler's third law of planetary motion, BACON.5 observes[3] the period p of a planet, along with its distance d from the sun. Upon comparing the values of these terms for a group of planets, the program finds that d and p increase together. Their ratio d/p is defined and its values are calculated. Now BACON.S notes that d/p decreases as p increases; the product p(d/p) is considered but abandoned, since it is equivalent to the distance d. However, d/p also decreases as d increases, leading to the product d^2/p. When the values of this theoretical term are computed, the program notes that it decreases as d/p increases. The result is the product d^3/p^2, which has a constant value; this invariance is equivalent to Kepler's third law.

Both terms are actually dependent, being functions of the independent nominal variable planet. which takes on values such as mercury, venus, and earth The intrinsic property heuristic discussed below lets the values of d and p be compered as if one were independent.

Postulating Intrinsic Properties

Although the above heuristics are useful for relating numeric variables, they cannot be applied when an independent nominal or symbolic variable influences the values of a numeric dependent term. For example, one finds that replacing one battery (e.g., battery A) in an electric circuit with another (e.g., battery B) produces changes in the slope and intercept of the line relating IL and I. In such cases, BACON.5 postulates an intrinsic property associated with the nominal term. The values of this property are borrowed from one[4] of the dependent variables. As a result, the ratio of that dependent variable and the new term is guaranteed to be 1 by definition.

Fortunately, BACONS is often able to move beyond such tautological reasoning through a process of generalization. Initially, the intrinsic values associated with a set of nominal values are retrieved only under the exact conditions in which they first occurred. But later, as the program alters other variables, a situation will arise which differs from the original in only one respect (e.g., x 2 instead of x = 1). When this situation occurs, BACON.5 compares its set of intrinsic values to the set of dependent values observed under the new conditions. If the two sets of values are linearly related, a generalization results. In future cases, the intrinsic values will be retrieved regardless of the value of the varied term. Once this has happened, the ratio of the dependent term and the intrinsic property can take on values other than 1, which can be related to other terms at higher levels of description.

In experiments where the same nominal term is associated with different objects, a similar process can occur. For example, *in* its discovery of Snell's law (discussed at length in the next section), BACON.5 associates values of an intrinsic property called the index of refraction with different gases and liquids (such as air and water). The assignment occurs when one of the media involved in the experiment is varied. Later, when the other medium is varied at a higher level, the program retrieves the values it generated earlier. This assumption that the same values apply to both cases lets the system break out of the tautological circle that can result from the introduction of intrinsic properties.

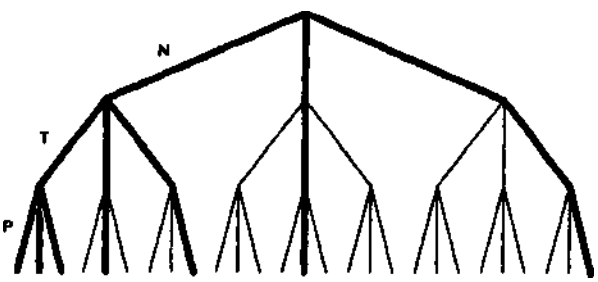

Figure 1. Search trees for the ideal gas law.

If multiple dependent terms exitt, then BACON 5 attempts to relate each 0* them to the new property. When no relation can be found, the program proposes a second property baaed on the unrelated variable This occurs in the discovery of Ohm's law, leading to separate intrinsic properties for voltage and internal resistance.

Reasoning by Analogy

We have already discussed BACON.5's technique for using its expectations to minimize the data it must gather. In addition, the program incorporates a heuristic for reducing its search still further in special cases. When BACONS is informed of a variable, it requests, the name of the object with which that term is associated. The program generates a unique name for this variable-object combination, and moves on to collect more information. However, before it plans an experiment, the system reorganizes the known variables, grouping together all those associated with the same object.

In some situations, two objects have exactly the same associated terms. In this case BACON.S assumes that the final law will be symmetrical, though it cannot predict the actual form. It orders the terms so that those attached to one of the objects are varied first. Once it finds a constant theoretical term incorporating these variables, the program assumes that an analogous term should be used for the other object. Thus, BACON.S must systematically vary only one of the second set of terms. When this is done, the two higher level terms take on different values, and can be related by other heuristics. In this way, the symmetrical laws so common in physics can be easily discovered.

In other cases, BACON.S must use its standard techniques to relate two entirely different sets of variables, each associated with an object. Later, a new task may be presented in which one set of terms is identical with that from the previous problem. When this occurs, the system assumes that the theoretical term which proved useful before will do so again. None of the analogous variables *need* be systematically varied, nor must the system tediously search for the appropriate combination of terms. Thus, this heuristic enhances the savings which normally occurs, reducing the data which must be examined and directing search through the space of theoretical terms. We discuss examples of each situation in the following section.

The Discoveries of BACON.5

In the previous section we discussed BACON 5's strategies in the abstract. Below we trace some of the system's discoveries in detail, focusing on laws in which symmetry and conservation play a role. These include Snell's law, conservation of momentum, Black's specific heat law, and Joule's formulation of conservation of energy. Finally, we examine the generality of the program's heuristics with respect to these and other tasks.

$MEDIUM_1$	H_1	O_1	$MEDIUM_2$	H_2	O_2	O_2/H_2
VACUUM	40	10	WATER	40	13.33	0.33
VACUUM	40	10	WATER	50	16.67	0.33
VACUUM	40	10	WATER	60	20.00	0.33

Table 1. First level data obeying Snell's law.

Snell's Law of Refraction

As a ray of light passes from one medium to another, the passage alters the direction in which the ray travels. Snell's law of refraction relates the initial and final directions, as well as the media involved. If we let 0_1 be the incoming angle of incidence and O_2 be the outgoing angle of refraction, then Snell's law may be stated as sine $0_1/n_1$ = sine 0_2/V in thisequation, n_1 the value of an intrinsic property of the first medium called the index of refraction, while n_2 is the analogous value for the second medium. The index of refraction for vacuum is normally assigned as the standard 1.0, while the values for water, oil, and glass are 1.33,1.47, and 1.66. respectively.

In rediscovering Snell's law, BACON. 5 is presented with two objects, each with three associated variables. The first of these terms is the medium, which takes nominal values such as air and water. The second is the distance h from the intersection of the media at which the light is measured. The final term is more complicated. Draw a perpendicular line through the point at which the ray crosses from one medium into the other. Now find a point on the ray which has distance h from this intersection. The variable o is defined as the distance of this point from the perpendicular. Thus, any combination of values for h and o determine a particular angle, BACONS is given experimental control of the first medium, h_1 o_1, the second medium, and h_2. To find o_2 for a given experimental situation, one moves down the ray of light for the distance h_2; the value of o_2 is simply the distance of this point from the perpendicular. This formulation[5] of the task allows the program to discover Snell's law without knowledge of trigonometry.

$MEDIUM_1$	H_1	O_1	$MEDIUM_2$	O_2/H_2	N_2	O_2/H_2N_2
VACUUM	40	10	WATER	0.33	0.33	1.0
VACUUM	40	10	OIL	0.37	0.37	1.0
VACUUM	40	10	GLASS	0.42	0.42	1.0

Table 2. Second level data obeying Snell's law.

When presented with these terms, BACONS first plans its experiment so the h_2 will be varied first, followed by $medium_2$ o_r h_1, and $medium_1$, in that order. The program makes a special effort to group variables associated with the same object together. Table 1 presents some data collected when the value of h_2 was varied. Comparing the values of o_2 and h_2, the system finds a linear relation with a zero intercept. Accordingly, the ratio o_2/h_2 is defined, which has the constant value 0.33. Note that this term is equivalent to sine O_2

Next BACONS alters the value of the second *medium,* using oil instead of water. Since the system expects o_2/h_2 to be constant, it computes its value when o_2 is 20, giving 0.37 as the result. Using the same approach, o_2/h_2 is found to be 0.42 when the second medium is glass. These values are summarized in Table 2. At this point, BACONS has a set of nominal independent values associated with a set of numeric dependent values. Since no more progress is possible until numbers are associated with these symbols, the program postulates an intrinsic property n_2 to be connected with the second medium. The values of this term are taken to be those of o_2/h_2, and the ratio O_2/h_2n_2 is defined. This new property has (by definition) the constant value 1.0. However, when the values of n_2 are retrieved during later periods of the discovery process, this ratio may take on different values.

O_1/H_1N_1	O_2/H_2N_2	$O_2H_1N_1/O_1H_2N_2$
1.00	1.00	1.0
1.20	1.20	1.0
1.35	1.35	1.0

Table 3. Highest level data obeying Snell's law.

[5] We would like to thank Robert Akscyn for suggesting this approach.

Having related all variables associated with the second object, BACON 5 locuses on the first set of terms. Assuming symmetry, the system does not bother to systematically vary o, and h_1 Instead, the values of the first medium, h_1 and o_1 are varied in conjunction, and the corresponding values of o_1/h_1n_1 are immediately computed. In addition, the values of o_2/h_2n_2 are found for each situation, and the higher level terms are compared, as shown in Table 3. A linear relation with a zero intercept is uncovered, and the ratio $o_2h_1n_1/o_1h_2n_2$ is found to be the constant 1.0. Replacing o/h with sine *8,* we see that this law becomes $n_1$since $0_2/n_2$sine O_1 = 1.0, which is equivalent to Snells law. The intrinsic properties n_1 and n_2 are proportional to the indices of refraction of the two media. Figure 2 presents BACON.5's search through the data space for this task.

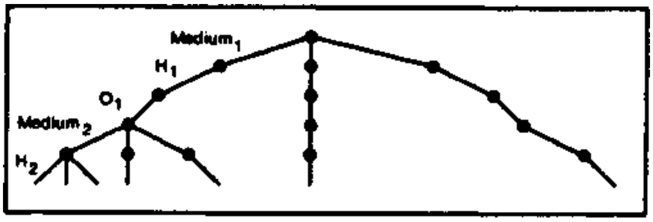

Figure 2. Search tree for Snell's law.

Conservation of Momentum

Physicists define the momentum of an object as the product of its mass m and its velocity v. Similarly, the total momentum of a group of objects is the sum of their individual momenta. One of the most basic laws of physics states that when two or more objects collide, the total momentum of the system is conserved. For two-object collisions, this may be stated as m_1v_1 m_2v_2 = $m_1u_1 + m_2u_2$, where m_1 and m_2 are the masses of the objects, v_1 and v_2 are their initial velocities, and u, and u_2 are their final velocities. This equation may be rewritten as $m_1(v_1 - u,)$ = $m_2(u_2 \cdot v_2)$, which says that the loss in momentum for one object is compensated by a gain in the momentum of the other object.

$NAME_1$	V_1	U_1	$NAME_2$	V_2	U_2	$U_2 \cdot V_2$
A	40	10	B	40	55	15.0
A	40	10	B	50	65	15.0
A	40	10	B	60	75	15.0

Table 4. First level data obeying conservation of momentum.

In discovering conservation of momentum, BACON.5 is informed of three terms, each associated with two objects. The terms name, and $name_2$ take on nominal values such as A, B, and C; these serve to distinguish specific objects from one another. The terms v_r v_2, u_r and u_2 take on the same meanings as above. BACON.5 is not directly presented with the masses of the objects. The first five of these variables are under the system's experimental control, while u_2 is viewed as dependent. The treatment of u, as independent requires some justification. In the normal interpretation of conservation of momentum, both u1 and u_2 are considered dependent on the other four terms. Moreover, the law cannot uniquely predict the values of u1 and u2; usually another equation such as conservation of energy must be invoked to achieve predictive power. However, since velocities are relative to the motion of the observer, one can simply move at a velocity which gives an arbitrary value to u1 Once this is done, the value of u_2 is uniquely determined.

BACON.5 begins by arranging these variables so that v_2 is varied (irst, followed by $name_2$, u1 V1 and name,. Upon varying the values of v_2, the program notes a linear relation between u_2 and v_2 with a slope of 1.0. Accordingly, the difference u_2 - v_2 is defined, which has the constant value 15.0. When $name_2$ is varied next, BACONS considers only enough values of u_2 and v_2 to find the differences it needs. Table 5 shows the results of these calculations for three different objects, where the relative masses are A = 1, B = 2, C = 3, and D = 4. Confronted with nominal independent values, the system defines the intrinsic property k_2 and assigns it values based on those of u_2 • v_2. The ratio (u_2 • v_2)/k_2 is also defined and found (not surprisingly) to be 1.0. Note that k_2 is proportional to the inverse of m_2, the mass of the second object.

$NAME_1$	V_1	U_1	$NAME_2$	$U_2 \cdot V_2$	K_2	$(U_2 \cdot V_2)/K_2$
A	40	10	B	15.0	15.0	1.0
A	40	10	C	10.0	10.0	1.0
A	40	10	D	7.5	7.5	1.0

Table 5. Second level data obeying conservation of momentum.

After relating k_2, v_2, and u_2, BACON.5 turns its attention to k,, v_v and u_v By analogy, it decides to relate the values of (u1 - v1)/k1 to those of (u_2 • v_2)/k_2, gathering only the data it requires. Table 6 gives some values of these theoretical terms. The program immediately finds a linear relation between these higher level variables, this time with a zero intercept. The ratio k,(u_2 * v_2)/k_2(u_1 • v^ is defined and found to have the constant value • 1.0. This expression is easily transformed into the second statement of conservation of momentum given above, BACON.S'S search through the data space is identical in form to that for Snell's law.

$(U_1 - V_1)/K_1$	$(U_2 \cdot V_2)/K_2$	$K_1(U_2 \cdot V_2)/K_2(U_1 \cdot V_1)$
-1.0	1.0	-1.0
-2.0	2.0	-1.0
-3.0	3.0	-1.0

Table 6. Highest level data obeying conservation of momentum.

Black's Specific Heat Law

If one brings two containers of liquid having different temperatures into contact, their temperatures will eventually come into equilibrium. The quantitative relation is $c_1m_1t_1 + c_2m_2t_2 = c_1m_1f_1 + c_2m_2f_2$, where t_1 and t_2 are the initial temperatures, f_1 and f_2 the final temperatures, m, and m_2 the masses of the two liquids, and c, and c_2 are intrinsic properties associated with various types of liquids called their specific heats. Joseph Black discovered this relation, usually called the specific heat law, in the 1860s. The equation may be recast as $c_1m_1(t_1$ •f,) = $c_2m_2(f_2 - t_2)$, which states that any heat lost from one body is gained by the other.

In this case, BACONS is told about two objects, each with an associated liquid (taking nominal values sgch as water and oil), mass m, initial temperature t, and final temperature f. All variables are *under* experimental control except f_2, which is the single dependent term. The final temperature of the first liquid, f,, can be manipulated observationally by continuously watching the temperature of the first liquid. When the temperature reaches the desired value, then the temperature of the second liquid is sampled as well. Thus, one need not wait until the liquids have reached equilibrium to discover Black's law.

Joule's Law

In the 1840's James Prescott Joule set forth his law of energy conservation. Joule tested his hypothesis in a variety of ways. In one experiment, he dropped heavy objects to turn a wheel, whose motion in turn heated a large container of water by friction. By comparing the potential energy lost by the dropped objects to the heat energy gained by the liquid, he found that the total energy remained constant. In a second experiment, the scientist sent an electric current through water and measured the resulting change in temperature; the total energy was conserved in this case as well.

BACON.5 can discover both forms of Joule's law. In the first situation, the program is informed about two objects a heavy ball and a container of liquid. Three independent terms are associated with the ball • its weight w_1 the height h_1 from which it is dropped, and n_1 the number of times it is dropped. The liquid is also associated with three controllable variables - the type of liquid L_2 occupying the container, its mass m_2, and its initial temperature t_2. In addition, a single dependent term, the final temperature f_2, is associated with the liquid.

As in its discovery of Black's law (itself a special case of Joule's law), BACON.s begins by varying t_2 and examining the resulting values of f_2 The difference f_2 - t_2 is found to be constant, followed by $m_2(f_2 - 1_2)$ at the next level of description. Still later, the intrinsic property d_2 is defined, followed by the ratio $m_2(f_2 - t_2)/d_2$. At this point the program shifts to manipulating characteristics of the first object. Since different variables are involved, it must rely on its data-driven heuristics instead of jumping ahead through analogy. When the number n, is varied, a linear relation with intercept 0 is noted; the ratio $m_2(f_2$ - t) / d_2n_1 is defined and found constant. This process repeats itself for the remaining terms, leading first to the ratio m_2 (f_2 • t_2)/$d_2n_1w_1$ and then to m_2 (f_2 - t_2)/$d_2n_1w_1h_1$ This last expression has a constant value, and the resulting hypothesis is equivalent to Joules statement that the loss in potential energy (11,111,11,) is gained in heat energy (m_2 (f_2 - t_2)/d_2).

The form of this law is quite similar to that for conservation of momentum, and BACON.5's path of discovery is correspondingly similar. At the outset, the program groups terms associated with each object together, to simplify its analogical reasoning at a later stage. Upon varying t_2 and examining the resulting values of f_2, the system finds a linear relation with a slope of 1.0. This leads it to consider the difference f_2 • t_2, which has a constant value. When BACONS varies m_2 at the next level, it finds that f_2 • t_2 decreases as m_2 increases; their product $m_2(f_2 - 1_2)$ is computed and also found to be a constant. At the third level of description, the program finds itself attempting to relate nominal values to numeric ones. Its solution is to define the intrinsic property d_2 based on the values of $m_2(f_2$ • $t_2)$; this new property is equivalent to the inverse of c_2, the specific heat of the second liquid. The ratio $m_2(f_2 - t_2)/c_2$ is immediately defined and found to be the constant 1.0.

Since it has discovered the relations between the four terms associated with the first body, BACON.S next considers the second liquid. The term $m_2(f_2$ • $t_2)/d_2$ was useful before, so it computes the values of $m_1(f_1 - t_1)/d_1$ for a number of cases and attempts to find a relation between the two expressions. As with conservation of momentum, the result is a linear relation with an intercept of 0 and a slope of -1.0. The ratio $d_1m_2(f_2$ • $t_2)/d_2m_1(f_1$ • t,) is defined and stored as a constant. This equation is a simple restatement of the second formulation of Black's law shown above.

In discovering Joule's second relationship, BACONS is presented with the same container of liquid and a new objecj, an electric circuit. The liquid has the same associated terms as before, while the circuit is associated with two independent variables • its resistance R and its current I. In generating its experimental design, the system recognizes a familiar situation. Rather than varying the values of L_2, m_2, and t_2 systematically, it decides to compute values of $m_2(f_2 - t_2)/d_2$. This is done for a number of values of the current I, and the two sets of numbers are compared. Although $m_2 (f_2 \bullet t_2)/d_2$ and I_1, increase together, the ratio $m_2 (f_2 \bullet t_2)/d_2 I_1$ is not constant. However, the new term is linearly related to I, with an intercept of zero. The second ratio $m_2(f_2 - t_2)/d_2 I$ is defined and found to be constant. When the resistance R, is varied next, a linear relation with a zero intercept is discovered, and the ratio $m_2(f_2 \bullet t_2)/d_2 I^2 R_t$ is defined. This complex term has a constant value, and BACON.5 has arrived at Joule's second formulation of energy conservation.

Generality of the Heuristics

In addition to the laws described above, BACONS has rediscovered a number of other empirical laws. These include the ideal gas law, Coulomb's law, Kepler's third law, Ohm's law, Archimedes' law of displacement, and Galileo's pendulum law and law of constant acceleration. These discoveries differ from those just discussed in that reasoning by analogy could not be used to reduce search. However, most of the remaining heuristics played a role in each of these taws.

The program's success on these tasks suggests that the primary goal of our research has been met: BACON.5 is a general discovery system capable of operating in many different domains. However, an intelligent system is only as general as the heuristics from which it is constructed. As we have seen, BACONS incorporates five main heuristics: noting constancies; finding linear relations; detecting monotonic trends; postulating intrinsic properties; and reasoning by analogy. Table 7 presents those heuristics applied in each of the program's discoveries. The degree of use indicates that BACON.S'S heuristics are indeed general principles useful in discovering a wide range of laws.

	CONSTANCY	LINEAR	MONOTONIC	INTRINSIC	ANALOGY
IDEAL GAS	X	X	X		
COULOMB	X	X	X		
KEPLER	X	X	X	X	
OHM		X	X	X	
ARCHIMEDES	X	X		X	
GALILEO	X	X	X	X	
SNELL		X		X	X
MOMENTUM		X		X	X
BLACK	X	X	X	X	X
JOULE	X	X	X	X	X

Table 7. Heuristics used in BACON.5's discoveries.

Conclusions

In this paper we described BACON.5. a system that has discovered a number of laws from the history of science. The program incorporates heuristics for noting constancies, finding linear and monotonic relations, defining theoretical terms, and postulating intrinsic properties. We focused on a heuristic for reasoning by analogy, which significantly reduced BACON.S'S search for data, theoretical terms, and hypotheses. This heuristic proved useful in discovering Snail's law, conservation of momentum, Black's specific heat law, and two forms of conservation of energy.

Several issues present themselves lor attention in the future. The first of these is noise. Although BACON.5 s heuristics are capable of handling considerable noise, the control structure of the system is not suited for dealing with the increased search which can result. In its current form, BACON.S can consider only a single hypothesis at a time. Future versions should be able to entertain multiple hypotheses, and design critical experiments to distinguish between them. Another issue relates to the variables presented to BACON.S. In all of the discoveries we have discussed, the program was only presented with terms relevant to the final law. For example, in the discovery of Snell's law, we carefully avoided giving BACON.S an extra term which would have let it define ratios equivalent to cosine *0* and tangent *6.* The next instantiation of BACON should have facilities for dealing with many irrelevant variables, so that its input need not be so carefully crafted.

A final limitation involves the form of the laws BACON.S can discover. The current implementation finds product and power laws, linear relations, and complex combinations of these forms. However, we hope to attain much more complex functions by incorporating Gerwin's [6] method of residuals. In this strategy, one formulates an hypothesis about the relation between "two variables, much as BACON.S does at present. But then one considers the residual variation left unaccounted for by the law. If this variation is random, the program is satisfied. However, if significant regularity remains, one searches for a new hypothesis to explain the variation. Once found, the new law is incorporated into the original rule, and the system iterates by examining the remaining variation.

In conclusion, we feel that BACON.S'S heuristics are general discovery principles with a wide range of application, and that the program has responded well in the areas it has been tested. In its present form, the system has certain limitations: it cannot adequately deal with noise; it cannot handle irrelevant variables; and the class of laws it considers is overly constrained. However, we feel confident that generalizations and extensions of the data-driven, Baconian approach will suggest solutions to these problems, and that more interesting discoveries lie ahead.

References

[1] Lenat, D. B. Automated theory formation in mathematics. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence,* 1977,833 842.

[2] Buchanan, B. G., Feigenbaum, E. G , and Sridharan, N. S. Heuristic theory formation. In 0. Michie (ed.) *Machine Intelligence 7.* New York: American Elsevier Publishing Co., 1972, 267-290.

[3] Langley, P. Rediscovering physics with BACON.3. *Proceedings of the Sixth International Joint Conference on Artificial Intelligence,* 1979, 505-507.

[4] Bradshaw, G. L, Langley, P., and Simon, H. A. BACON.4: The discovery of intrinsic properties. *Proceedings of the Third National Conference of the Canadian Society for Computational Studies of Intelligence,* 1960,19-25.

[5] Langley, P., Bradshaw, G. L., and Simon, H. A. Rediscovering chemistry with BACON.4. To appear in R.Michalski, J. Carbonell, and T. Mitchell (eds.), *Machine Learning.*

[6] Gerwin, D. G. Information processing, data inferences, and scientific generalization. *Behavioral Science,* 1975, *19,* 314-325.