

Algebraic Reasoning about Reactions: Discovery of Conserved Properties in Particle Physics

RAÚL E. VALDÉS-PÉREZ

valdes@cs.cmu.edu

Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213 - USA

Editor: Paul Utgoff

Abstract. Kocabas (1991) describes a situation from particle physics in which quantum properties and conservation laws are postulated from lists of observed and unobserved reactions. Kocabas also presents a program named BR-3 that can rediscover some accepted quantum properties from textbook data, although it fails on a more difficult example from the same source. This paper describes PAULI, a program that solves the same task as BR-3 but uses a different problem-solving model. PAULI produces different, simpler solutions than does BR-3, and it can also handle the problematic example. After comparing the two programs, we conclude that PAULI offers distinct advantages over its predecessor, which we attribute to an *algebraic* approach to reasoning about sets of reactions.

Keywords: Machine discovery, reactions, particle physics, algebraic reasoning, model building

1. Introduction

Kocabas (1991) described a situation from elementary particle physics in which new particle properties are postulated from data on observed reactions among these particles, and from lists of reactions that have never been observed despite considerable effort (Omnes, 1971). Whenever the existing quantum properties and their conservation laws cannot account for the non-existence of the unobserved reactions, new quantum properties are sought that will provide *selection rules*: the observed reactions conserve all of the new properties, whereas each unobserved reaction violates at least one property. Finding such new properties provides a satisfactory phenomenological explanation for why certain reactions occur and others do not.

Kocabas's introduction of the problem into machine discovery is valuable because it involves a simple, powerful heuristic of physics that leads to the discovery of new conserved properties. Moreover, the problem differs from more familiar situations involving conservation-law induction from numeric data (Langley, Bradshaw, & Simon, 1981).

The purpose of this paper is to introduce a computational model of the task that offers certain advantages over the one given by Kocabas. The behavior of the present model, implemented in a computer program PAULI, is compared with the model reflected in the program BR-3.

The approach embodied in PAULI relies on a key idea that it shares with the MECHEM system (Valdes-Perez, 1994; Valdes-Perez, 1992a). Reasoning about sets of reactions, which is a common task throughout the natural sciences, is better accomplished by an algebraic representation of reactions, which allows the application of powerful mathematical and algorithmic tools to answer questions of scientific interest. Many processes in

science are similar to reactions, hence automated reasoning about reactions advances the field of machine scientific discovery by making yet one more class of tasks intelligible in terms of computation.

We begin by reviewing the scientific task and the design of BR-3. We then show how one can express the task constraints algebraically, and describe a program that can search for solutions that satisfy these constraints. The solutions found by the new program are compared with those found by its predecessor, and the former are shown to be simpler. Finally, we contrast the two task models and extract a lesson from this work, which we anticipate as follows: automated reasoning about reaction-like processes should make use of algebraic representations and draw on the powerful methods developed within mathematics and computer science for dealing with such representations.

2. Discovery of Quantum Properties

In their book *The Particle Hunters*, Ne'eman and Kirsh (1986) describe how modern physics found it necessary to expand on the conservation laws of classical physics, such as those governing energy, charge, and momentum. Certain reactions among elementary particles were never observed to occur, despite much experimental effort, even though they obeyed classical laws. For background, we quote at some length from Ne'eman and Kirsh (p. 132):

In classical physics, conservation principles are used to simplify difficult calculations. If one is to deal with a complicated multi-step process, it is helpful to know that certain quantities remain constant throughout the whole process. The relative importance of conservation laws in classical physics was, at any rate, secondary. [...]

In modern physics, not only were several conservation laws added, but their relative importance increased as well. And this is no surprise: in a world ruled by probability, where laws such as those of Newton are unable to predict just what will happen, laws which determine *what cannot happen* become increasingly important. If we know that no process which contradicts any conservation law can take place, we may infer that for any process which *does not* contradict any conservation law, a certain probability exists that it will occur. [...]

Let us assume that a given particle can decay into other particles without breaking a single conservation law. Will it actually do so? The answer is yes! All of the experimental data at hand indicate that it will do so as soon as possible! [...]

Let us now summarize this important change in outlook: instead of the set of deterministic laws predicting exactly what will occur, which ruled in classical physics, modern physics is ruled by conservation laws, which determine what *cannot occur*. And according to the present approach, any process which does not contradict these laws does indeed take place.

Physicists concluded that particle interactions must be governed by new conservation laws that are not revealed in the behavior of macroscopic objects, hence were not discov-

ered within the framework of classical physics. Some of these laws, and the quantities they conserve, were postulated quite *ad hoc*; the conservation of baryon number and strangeness are two such cases. These ad-hoc postulates addressed the specific problem of ruling out experimentally unobserved reactions; they were not derived theoretically. Neither the textbook by Omnes (1971) that Kocabas used as a source, nor a more recent popular account by Ne’eman and Kirsh (1986), describe any systematic method to discover these properties. Describing such a method is an important contribution of the Kocabas paper.

There is nothing sacrosanct about using conservation laws as explanations for why certain reactions occur and others do not. In chemistry, analogous explanations of chemical reactions are made using the concept of free energy, which is a defined quantity that tends to decrease, rather than increase or remain constant. In thermodynamics, the related concept of entropy is invoked to explain why certain processes that would *decrease* entropy cannot happen. Further comparisons between chemical and nuclear reactions are made by Godfrey, McLachlan, and Atwood (1991).

Interestingly, discovery can also proceed in the reverse direction. A well-established conservation law may surprisingly rule out a newly observed reaction. Rather than discard the conservation law, or narrow the scope of its application, the impasse may be resolved by postulating an unnoticed participant in the reaction whose quantum numbers are such that all established conservation laws are upheld. This was the reasoning followed by Wolfgang Pauli in 1930 in the context of beta decay, in which a nucleus apparently decayed into a lighter nucleus and an electron in a way that violated conservation of energy (Sutton, 1990). To uphold energy conservation, Pauli postulated a novel particle (later named ‘neutrino’ by Enrico Fermi) as a third product of the decay. We have named the present program PAULI after Wolfgang Pauli, even though he was not involved in postulating new quantum properties.¹

For brevity of exposition, we refer in this paper to the observed reactions as “good” reactions, and to the unobserved reactions as “bad.” The following simple problem illustrates the case of a single good and a single bad reaction:

$$\begin{array}{l} \text{good: } \bar{\pi} + p \rightarrow \pi^0 + n \\ \text{bad: } \quad \quad p \not\rightarrow \pi + \pi^0 \end{array}$$

The symbol ‘ \rightarrow ’ indicates that the particles on the left undergo a reaction that forms the particles on the right. Typically, reactions occur upon collision between particles, and new products are formed due to interactions among the reacting particles or their fragments. The symbol ‘ $\not\rightarrow$ ’ indicates an unobserved (bad) reaction.

To exclude the bad reaction above, while admitting the good one, it suffices to postulate a new quantum property having the value of unity for p and n and zero for the other particles. This property is conserved by the good reaction and violated by the bad one. Section 4.5 shows how the PAULI program is able to postulate this property.

3. Summary of BR-3

Kocabas describes the design of his Prolog program BR-3 in terms of three operators **check-consistency**, **check-completeness**, and **revise-beliefs**. The **check-consistency** operator tests whether the current assignment of quantum numbers to the particles results in conservation in the case of the good reactions. The **check-completeness** operator tests whether each bad reaction violates conservation of at least one quantum property. Finally, the **revise-beliefs** operator is charged with carrying out the search for quantum numbers that satisfy the given constraints, which are discussed by Kocabas and are also discussed in Section 4.2 of this paper.

The author summarizes BR-3's behavior as a backtrack search for quantum numbers over a small set of several integer values. There is a degree of look-ahead in the backtrack search, in the sense that the system uses a "simple linear equation solver" to carry out some simple algebraic inference. The program's search regimen is consonant with the control structure of Prolog, its implementation language. For full details on BR-3, we refer the reader to the original paper.

One limitation of BR-3 is that it searches for quantum numbers over a small fixed integer set, such $\{-1,0,1\}$ in some examples discussed by Kocabas. Of course, one could enlarge this set, but the generate-and-test regimen employed by the program renders large integer sets rather problematic, as is readily acknowledged by the author.

In our opinion, there is a second limitation that strictly concerns not the program itself, but its manner of use on several cases of discovery. BR-3 was used on bad reactions one at a time, and new quantum properties were postulated whenever the current stock of conservation laws could not rule out the new bad reaction. It is not clear to this writer whether the historical record exhibits this one-at-a-time consideration of unobserved reactions. Whether it does or not, we assert that it is preferable to consider all the existing data before postulating hidden properties, especially if the number of data is more like a dozen, rather than a million. Sections 5.1 and 5.2 below show how, given the same data, PAULI finds solutions having fewer postulated quantum properties than those found by BR-3. We do not believe that saving computation via incremental processing is an important issue in such cases of scientific discovery, since novel fundamental properties are not postulated every day.

4. The PAULI Program

PAULI's design is quite different from that of BR-3. One difference is that PAULI carries out a highly constrained search for quantum numbers, because the constraints on acceptable solutions are incorporated within the generator and serve to guide the search very selectively. Due to the formal structure of the problem, which permits formulating most of the constraints as linear equalities and inequalities, one can exploit linear programming as an efficient search algorithm. The usual algorithms for solving linear-programming problems never generate candidate solutions that violate the given constraints; their search is concerned entirely with finding an optimum among the feasible solutions. It is in this sense that PAULI carries out a highly constrained search.

4.1. Overview of linear programming

A *linear programming* problem consists of a set of constraints, which can be linear equations $\sum a_{ij}x_i = c_j$ or linear inequalities $\sum a_{ij}x_i \geq c_j$, and a linear expression $\sum b_i x_i$ to be minimized, called the *objective function*. The x_i are variables, whereas the a_{ij} , c_j , and b_i are known constants. By convention, all variables appearing in the constraints assume only non-negative values.

A linear-programming problem is solved by reporting one of three results: the constraints are inconsistent; the quantity to be minimized is unbounded; or values for the variables that minimize the objective function, along with the numeric value of the objective function. The simplex algorithm is a highly efficient procedure for solving linear-programming problems. Papadimitriou and Steiglitz (1982) give a detailed description of linear programming and the simplex algorithm.

4.2. Expression of the task constraints

We use the notation $r_1 + \dots + r_m \rightarrow p_1 + \dots + p_n$ to refer to a good reaction, and a similar notation with ' \rightarrow ' instead of ' \rightarrow ' to refer to a bad reaction; the 'r' suggests 'reactant' and the 'p' suggests 'product' (not 'proton' nor 'particle'). The sum of the values for quantum property q within a single reaction i is denoted as

$$\sum_k p_{ikq} - \sum_l r_{ilq}$$

where the sums are taken over the particles appearing among the l reactants and among the k products. We use $|S|$ to denote the cardinality of a set S , and \bar{x} to denote the antiparticle corresponding to particle x ; for example, \bar{p} is an antiproton. We use the symbol ' x ' to refer to particles generically.

Generally in science, a conservation law states that a certain property keeps the same aggregate value after undergoing a certain class of processes. In our case, aggregation is summation, and the processes are reactions. The requirement that the i th good reaction obey all Q conservation laws (one for each quantum property) can be expressed as the following conjunction of Q equations:

$$\begin{aligned} \sum_k p_{ik1} - \sum_l r_{il1} = 0 \quad \wedge \quad \sum_k p_{ik2} - \sum_l r_{il2} = 0 \quad \wedge \quad \dots \\ \dots \quad \wedge \quad \sum_k p_{ikQ} - \sum_l r_{ilQ} = 0 \end{aligned} \quad (1)$$

Conversely, the requirement that the j th bad reaction violate at least one of Q conservation laws is expressed as a disjunction of *inequations* thus:

$$\begin{aligned} \sum_k p_{jk1} - \sum_l r_{jl1} \neq 0 \quad \vee \quad \sum_k p_{jk2} - \sum_l r_{jl2} \neq 0 \quad \vee \quad \dots \\ \dots \quad \vee \quad \sum_k p_{jkQ} - \sum_l r_{jlQ} \neq 0 \end{aligned} \quad (2)$$

(an inequation ‘ \neq ’ is the logical negation of an equation ‘ $=$ ’). In linear programming, one can express the Conjunction (1) straightforwardly as a set of Q algebraic equations. The expression of Condition (2) is less straightforward, but can be accomplished via a representational device to be described shortly. Before doing so, and for simplicity of exposition, we ignore the fact that variables of linear programming are restricted to non-negative values; later we show how to restore the restriction while keeping the ability to postulate negative quantum numbers, which is allowed in physics.

When Q new quantum properties need to be postulated to explain the reaction data, the Disjunction (2) for the j th bad reaction can be reformulated as a single, linear inequation:

$$b^0 \left(\sum_k p_{jk1} - \sum_l r_{jl1} \right) + b^1 \left(\sum_k p_{jk2} - \sum_l r_{jl2} \right) + \dots \quad (3)$$

$$\dots + b^{Q-1} \left(\sum_k p_{jkQ} - \sum_l r_{jlQ} \right) \neq 0$$

in which b is an arbitrary positive constant raised to the various powers shown. To prevent the summation from equaling zero due to nonzero terms that cancel, one imposes Q additional conditions on the reaction sums:

$$-(b-1) \leq \sum_k p_{jk1} - \sum_l r_{jl1} \leq (b-1)$$

$$-(b-1) \leq \sum_k p_{jk2} - \sum_l r_{jl2} \leq (b-1) \quad (4)$$

$$\vdots$$

$$-(b-1) \leq \sum_k p_{jkQ} - \sum_l r_{jlQ} \leq (b-1)$$

These ensure that no nonzero terms in Inequation (3) cancel out, so that if the summation does equal zero, it implies that all Q terms within it are zero. In effect, a “number” of base b is formulated, each “digit” of which is the sum of the bad reaction j for one of the Q properties. Since the quantity b is a free parameter, it can be set to any value deemed “safe.” We have never needed to change b ’s default value of 100.²

Linear programming cannot directly express inequation constraints of the form *expression* \neq *constant*, of which Constraint (3) is an instance. This obstacle is overcome by searching on a case-by-case basis: first the constraint *expression* \geq *constant* $+ \Delta$ can be tried, followed by the converse constraint *expression* \leq *constant* $- \Delta$, where Δ is a small positive constant (we have used $\frac{1}{100}$). Finding a solution, or verifying that there is none for the given Q , may involve a search of complexity $2^{|bad|}$ ($|bad|$ is the number of bad reactions) in the worst case, since a binary search tree of depth $|bad|$ is generated by the two choices for inequality direction per bad reaction.

Another set of constraints arises from any specified particle/anti-particle pairs, since their quantum numbers are set opposite thus: $\bar{x} = -x$ for all Q quantum properties. Following Kocabas, quantum numbers for any antiparticles are also required to be non-positive. These two conditions contribute respectively $Q \times |antipairs|$ and

$Q \times |\text{antiparticles}|$ constraints, where $|\text{antipairs}|$ is the number of particle/anti-particle pairs (possibly none), which can be smaller than the number $|\text{antiparticles}|$ if a particle does not appear in the reactions, but its antiparticle is present.

Above, we temporarily ignored the fact that variables in linear programming must be non-negative, and proceeded to express certain key task constraints in which variables could be negative. Now we overcome the restriction to the non-negative numbers by another representational device: each occurrence of a variable x is replaced by $(x' - x'')$. Both of the variables x' and x'' are necessarily non-negative, but the replacement allows converting the domain of the represented quantity from the non-negative numbers to the entire numeric domain. For example, a negative number -5 can thus be represented as $0 - 5$, where x' is 0 and x'' is 5, i.e., both positive numbers. Even though this representational device doubles the number of variables, the added computational cost should be small, since typically the number of simplex iterations depends on the number of constraints (see below).

Finally, a linear program includes an objective function, which serves to select one optimized solution from among the potentially many allowed by the constraints. The objective function is a linear expression which is to be minimized. We wish to minimize the sum of absolute values of all quantum numbers, which, however, is not a linear function. Nevertheless, this wish is met by adapting a familiar trick from linear-programming practice (Chvatal, 1983), as follows. First, one formulates another set of constraints:

$$-E_{ij} \leq x'_{ij} - x''_{ij} \leq E_{ij} \quad (5)$$

where i and j span respectively the number of particles P (including anti-particles) and the number of quantum properties Q being postulated, thus yielding a total of $2PQ$ inequalities. Then, one selects the following objective function:

$$\sum_{i=1}^P \sum_{j=1}^Q E_{ij} \quad (6)$$

whose minimization will, in effect, force the overall sum of the quantum-number magnitudes to be minimal. In the absence of any constraints, this bias forces all quantum numbers to zero, which is a sensible result in scientific practice. With constraints, the bias is for smaller quantum numbers. Unless the constraints are contradictory, the simplex algorithm will always find a solution, since the objective function is bounded from below. All linear-programming variables are non-negative, so a lower bound on the minimization is clearly zero.

4.3. Search complexity

Taken together, the total number of constraints is:

$$Q|\text{good}| + (2Q + 1)|\text{bad}| + Q(|\text{antipairs}| + |\text{antiparticles}|) + 2Q|\text{particles}| \quad (7)$$

and the number of variables is:

$$4Q|particles| + (2Q + 1)|bad| + Q|antiparticles| \quad (8)$$

where all but $2Q|particles|$ of these variables are due to the usual “slack variables” needed to put all the inequality constraints in the standard form of equalities (Papadimitriou and Steiglitz, 1982). A textbook on linear programming (Chvatal, 1983) states that “the typical number of iterations increases proportionally to m ,” where m is the number of constraints in the linear program. Assuming that our problem is not pathological, it follows that the expected number of iterations is proportional to Expression (7). We can combine this expression with the cost of the exhaustive case-by-case search mentioned previously (admitting the peccadillo of combining typical-case with worst-case analysis) to obtain an order of complexity of:

$$2^{|bad|} Q(|good| + |bad| + |particles|) \quad (9)$$

where Q is the number of postulated quantum properties. If the quantity $2^{|bad|}$ becomes impracticably large, then a divide-and-conquer heuristic could be applied by considering subsets of the set *bad* separately and incrementally. In the extreme of one bad reaction at a time, one obtains the approach taken in Kocabas’s BR-3 system.

4.4. Search modes

PAULI can search in either “satisficing” or “optimizing” mode. In satisficing mode, the program reports the first solution it finds, which will always involve the smallest number of new quantum properties, since the program starts with $Q = 1$ and increments from there as needed. In optimizing mode, even after finding a solution, PAULI continues searching (but keeping the number of quantum properties constant) by varying the inequality directions generated by the implementation of Constraint (3) above. Recall that this constraint induces a case-by-case search, since linear programming cannot directly express an inequation constraint of the form $\sum_i a_i x_i \neq 0$. After examining all of the remaining cases with Q fixed, PAULI reports the solution that minimizes the criterion discussed in the next paragraph.


Satisfaction of the task constraints by the simplex algorithm actually leads to fractional values of quantum numbers. These numbers are then converted to integers by multiplying by the least-common-multiple of all the denominators. Multiplication of all quantum numbers (corresponding to a single property) by a non-zero constant does not change whether a given reaction conserves or violates that property. In optimizing mode, the program will report the integer assignment that minimizes $\sum_i \sum_q |x'_{iq} - x''_{iq}|$, i.e., the sum of absolute values of all the quantum numbers. Hence, smaller assignments are favored over larger ones.

4.5. An illustrative example

Table 1. Constraints for one Bad Reaction, one Good Reaction, and one Quantum Property.

Constraint	Reference
$(\pi_1^0 + n_1) - (\bar{\pi}_1 + p_1) = 0$	[Constraint (1)]
$\pi_1 + \pi_1^0 - p_1 \geq 1/100$	[Constraint (3)]
$\pi_1 + \pi_1^0 - p_1 \leq 99$	[Constraint (4)]
$\pi_1 + \pi_1^0 - p_1 \geq -99$	[Constraint (4)]
$\bar{\pi}_1 + \pi_1 = 0$	[particle/antiparticle constraint]
$\bar{\pi}_1 \leq 0$	[antiparticles are non-positive]
$\pi_1^0 \leq E_{\pi_1^0}$	[remainder are Constraints (5)]
$-E_{\pi_1^0} \leq \pi_1^0$	
$n_1 \leq E_{n_1}$	
$-E_{n_1} \leq n_1$	
$\bar{\pi}_1 \leq E_{\bar{\pi}_1}$	
$-E_{\bar{\pi}_1} \leq \bar{\pi}_1$	
$p_1 \leq E_{p_1}$	
$-E_{p_1} \leq p_1$	
$\pi_1 \leq E_{\pi_1}$	
$-E_{\pi_1} \leq \pi_1$	

Let us reconsider the simple problem from Section 2 that involved a single good and a single bad reaction:



As predicted by Expression (7), PAULI sets up 16 linear constraints as it postulates a single quantum property; these constraints appear in Table 1. The system actually represents a property like π_1 as $\pi_1' - \pi_1''$; we have avoided these details for ease of exposition. The second constraint in the Table is one case of the case-by-case search needed to implement Constraint (3); the only other case here would involve the constraint $\pi_1 + \pi_1^0 - p_1 \leq -1/100$. The program indeed reports that a single quantum property is enough to include the good reaction and exclude the bad one: both p and n receive quantum numbers of unity and the rest are zero.

If two quantum properties were considered for this problem, then suitably subscripted duplicates of the constraints would be added to the linear program, since the quantum numbers are assigned independently. Even though two quantum properties are not needed to exclude a single bad reaction, we illustrate in Table 2 what one set of constraints would be in this case.

4.6. Comments on PAULI'S inductive bias

PAULI's inductive bias is to minimize, first, the number of postulated quantum properties, and second, the sum of absolute values of all quantum numbers. We believe that this

Table 2. Constraints for one Bad Reaction, one Good Reaction, and two Quantum Properties.

Constraint	Reference
$(\pi_1^0 + n_1) - (\bar{\pi}_1 + p_1) = 0$	[Constraint (1)]
$(\pi_2^0 + n_2) - (\bar{\pi}_2 + p_2) = 0$	[Constraint (1)]
$(\pi_1 + \pi_1^0 - p_1) + 100(\pi_2 + \pi_2^0 - p_2) \geq 1/100$	[Constraint (3)]
$\pi_1 + \pi_1^0 - p_1 \leq 99$	[Constraint (4)]
$\pi_1 + \pi_1^0 - p_1 \geq -99$	[Constraint (4)]
$\pi_2 + \pi_2^0 - p_2 \leq 99$	[Constraint (4)]
$\pi_2 + \pi_2^0 - p_2 \geq -99$	[Constraint (4)]
$\bar{\pi}_1 + \pi_1 = 0$	[particle/antiparticle constraint]
$\bar{\pi}_2 + \pi_2 = 0$	[particle/antiparticle constraint]
$\bar{\pi}_1 \leq 0$	[antiparticles are non-positive]
$\bar{\pi}_2 \leq 0$	[antiparticles are non-positive]
$\pi_1^0 \leq E_{\pi_1^0}$	[remainder are Constraints (5)]
$-E_{\pi_1^0} \leq \pi_1^0$	
$n_1 \leq E_{n_1}$	
$-E_{n_1} \leq n_1$	
$\bar{\pi}_1 \leq E_{\bar{\pi}_1}$	
$-E_{\bar{\pi}_1} \leq \bar{\pi}_1$	
$p_1 \leq E_{p_1}$	
$-E_{p_1} \leq p_1$	
$\pi_1 \leq E_{\pi_1}$	
$-E_{\pi_1} \leq \pi_1$	
$\pi_2^0 \leq E_{\pi_2^0}$	
$-E_{\pi_2^0} \leq \pi_2^0$	
$n_2 \leq E_{n_2}$	
$-E_{n_2} \leq n_2$	
$\bar{\pi}_2 \leq E_{\bar{\pi}_2}$	
$-E_{\bar{\pi}_2} \leq \bar{\pi}_2$	
$p_2 \leq E_{p_2}$	
$-E_{p_2} \leq p_2$	
$\pi_2 \leq E_{\pi_2}$	
$-E_{\pi_2} \leq \pi_2$	

is a generally accepted heuristic in science; in the absence of other factors, smaller numbers are preferable. This heuristic seems especially justified when the quantities that are involved potentially represent numbers of objects or processes. Whereas here we are concerned with abstract properties, physicists have given structural interpretations for these properties which indicate why small property values are preferred: the quantum value for a particle is the sum of quantum values for the particle's structural constituents, and one prefers to postulate fewer such constituents. Finding the structural constituents of particles is the very problem addressed by Fischer and Zytkow (1990) in their GELL-MANN program.

Instead of minimizing the aggregate sum of quantum-number magnitudes, one could instead use a minimax bias that minimizes the maximum quantum number (Chvatal, 1983). To adapt this new bias involves changing the Constraints (5) to the following:

$$-E \leq x'_{ij} - x''_{ij} \leq E \quad (10)$$

and changing the objective function (Expression (6)) to minimize, quite simply, the variable E , which serves to minimize the maximum quantum number. This bias is inferior to the bias that PAULI uses because, for example, the number of particles that share the maximum quantum number becomes irrelevant under this bias. We have implemented and discarded the minimax bias in PAULI after observing such inferior results.

5. Examples of Discovery using Historical Data

This section examines PAULI's performance on four examples taken from Kocabas (1991); each example reflects a historical instance of the discovery of a novel quantum property. In the first two cases, the present program finds solutions having one fewer quantum property than was found by its predecessor. In the third case, both programs agree that a single quantum property is enough, but they disagree on the exact quantum numbers. Finally, on the *strangeness* example that BR-3 was unable to solve, PAULI verifies that the accepted strangeness values are the simplest, given the constraints used in the actual discovery. Another example, created by collecting reactions from the first three cases, illustrates how PAULI can quickly detect insoluble problems.

5.1. Lepton and baryon numbers

In sections 3.1 and 3.2 of his paper, Kocabas reports that, given the reactions in his Table 3, BR-3 discovers the lepton and baryon numbers as a consequence of incrementally postulating two quantum properties to explain the reaction data. Given the same data all at once in optimizing mode, PAULI creates a linear program having 64 constraints that leads it to conclude after 40 seconds³ that a single quantum property suffices. The nonzero quantum numbers found are unity for each of the particles n (neutron), p (proton), and K^0 . The quantum numbers for the particles appearing in this solution correspond to the baryon numbers accepted in physics, except that K^0 instead of Λ receives a quantum value of unity. The lepton numbers were not needed, since only a single new property

was enough to account for the reaction data. In satisficing mode, the program finds after several seconds the same solution as in optimizing mode.

The fact that PAULI found different assignments to the quantum properties than both physicists and BR-3 is surprising. Moreover, PAULI finds different assignments for all of the four cases taken from the Kocabas paper. In each case, PAULI's values are simpler by our criterion; two of these cases involve one fewer quantum property, and two cases involve the same number of properties (one) but exhibit simpler numeric assignments. Section 5.5 tries to explain these surprises, but first we proceed by reporting the program's behavior on the remaining cases.

5.2. *Electron and muon numbers*

In his section 3.3, Kocabas reports that, given the reactions in his Table 4, BR-3 discovers the electron and muon numbers as a result of incrementally postulating two quantum properties to explain the reaction data. Given the same data in optimizing mode, PAULI creates a linear program having 62 constraints, and the program again reports after 20 seconds that a single new quantum property is enough; two properties are not needed. The nonzero values found in this case are $\mu = 1$, $\nu_\mu = 1$, $\bar{\mu} = -1$, $\bar{\nu}_\mu = -1$, and $\bar{\nu} = -1$. In satisficing mode, PAULI finds a single quantum-property solution having nine nonzero quantum numbers, rather than the previous five.

5.3. *Electric charge*

On page 293 of section 3.3, Kocabas states that, given the reactions in his Table 6, BR-3 discovers the accepted values of electric charge. Given the same data in optimizing mode, PAULI creates a linear program having 38 constraints. After one second the program confirms that a single new quantum property suffices, although the values it finds differ from BR-3's values. PAULI's nonzero values are $p = 1$, $\pi = 1$, $\nu = 1$, $\bar{\pi} = -1$, and $\bar{\nu} = -1$, which differ from the conventional values, since the electron is here assigned a zero charge. The same solution is found in satisficing mode. PAULI's solution is again simpler than BR-3's.

5.4. *Strangeness*

Kocabas also considers the strangeness property described by Omnes (1971) in chapter 2, section 5, who lists the twelve observed reactions and six unobserved reactions shown here in Table 3. Kocabas (p. 291) observes that for BR-3, "the strangeness property is problematic because it takes on integer values between -3 and 3 inclusively, while the current version of BR-3 assigns the values -1 , 0 , and 1 for reasons of search control."

We speculate that another reason that finding the strangeness property is problematic for BR-3's incremental approach is that by examining only one of the six bad reactions at a time, many more quantum properties will be postulated than by considering all the

Table 3. Reactions Giving Rise to the Strangeness Property.

Observed reactions		Unobserved reactions	
$p + p \rightarrow p + \Lambda + K^+$		$p + n \not\rightarrow p + \Lambda$	
$K^- + p \rightarrow \Lambda + \pi^0$		$p + p \not\rightarrow p + n + K^+$	
$\bar{p} + p \rightarrow \pi^+ + \pi^- + \pi^+ + \pi^-$		$p + n \not\rightarrow p + p + K^-$	
$\bar{p} + p \rightarrow K^+ + \bar{K}^0 + \pi^-$		$K^- + p \not\rightarrow \Omega^- + \pi^+$	
$\bar{p} + p \rightarrow K^- + K^0 + \pi^+$		$K^- \not\rightarrow \pi^- + \gamma$	
$K^- + p \rightarrow \Sigma^- + \pi^+$		$\Sigma^+ \not\rightarrow p + \gamma$	
$K^- + p \rightarrow \Sigma^0 + \pi^0$			
$K^- + p \rightarrow \Sigma^+ + \pi^-$			
$K^- + p \rightarrow \Xi^- + K^+$			
$K^- + p \rightarrow \Xi^0 + K^+ + \pi^-$			
$K^- + p \rightarrow \Omega^- + K^0 + K^+$			
$\Sigma^0 \rightarrow \Lambda^0 + \gamma$			

Table 4. Strangeness values given by Omnes and PAULI.

particle	Omnes's values	PAULI's values (run I)	PAULI's values (run II)	particle	Omnes's values	PAULI's values (run I)	PAULI's values (run II)
K^+	1	0	1	\bar{p}	0	0	0
Λ	-1	0	-1	K^0	1	1	1
p	0	0	0	\bar{K}^0	-1	-1	-1
n	0	2	0	Σ^-	-1	1	-1
K^-	-1	0	-1	Σ^0	-1	0	-1
π^+	0	-1	0	Σ^+	-1	-1	-1
π^-	0	1	0	Ξ^-	-2	0	-2
π^0	0	0	0	Ξ^0	-2	-1	-2
γ	0	1	0	Ω^-	-3	-1	-3
Λ^0	-1	-1	-1				

bad reactions at once. According to Ne'eman and Kirsh (1986), a number of bad reactions were known and considered when M. Gell-Mann and K. Nishijima independently postulated the strangeness property in 1953.

Given the 18 reactions from Omnes, PAULI creates a linear program having 72 constraints in optimizing mode. After three minutes, the program ascertains that a single new quantum property involving only the integer set $(-1,0,1,2)$ is enough to account for the observed and unobserved reactions. However, these quantum numbers (shown as columns labelled "run I" in Table 4) differ from the accepted strangeness values given by Omnes. In fact, PAULI's values are somewhat "better" according to the criterion discussed in Section 4.4. That is, the absolute-value sum of the discovered values is smaller in PAULI's case.

These discrepant values of "strangeness" were surprising to us. However, Ne'eman and Kirsh (1986) explain that a further constraint used to postulate strangeness is that

particles of the nucleon and pion families possess zero strangeness. When this constraint is added to the others, PAULI does find the accepted values of strangeness after four minutes of search, and further verifies (in optimizing mode) that they are the simplest solutions. These results are labelled “run II” in Table 4.

According to Kocabas (personal communication, February 12, 1993), BR-3 did not make use of the constraint on nucleons and pions because that program found the accepted strangeness values without the constraint. PAULI’s analysis indicates that the nucleon and pion constraint is strictly needed to justify the accepted strangeness values, since otherwise simpler solutions exist.

5.5. *Why doesn’t PAULI find the accepted quantum numbers?*

Using only Kocabas’s data, PAULI finds solutions that disagree with the accepted values in physics in every case. The program disagrees with BR-3 in every case except strangeness, which the latter program did not handle. In all cases, PAULI’s values were simpler. How does one make sense of these facts?


We can see three possible explanations:

1. Physicists erred by proposing unnecessarily complex assignments of quantum values.
2. Physicists used further constraints to postulate the new quantum properties.
3. PAULI’s inductive bias is different from (and inferior to) the biases used by physicists.

The second of these explanations is correct for the case of strangeness, since Ne’eman and Kirsh (1986) discuss a further constraint (the nucleon and pion families possess zero strangeness) that was not mentioned in Omnes (1971) nor used by Kocabas in BR-3. Yuval Ne’eman (personal communication, October 26, 1993) has pointed out that other constraints beyond the data on good and bad reactions should be incorporated in order to express other experimental facts. His initial suggestions can be expressed as *inequation* constraints between pairs of particles, but further work and experimentation are needed to determine whether the current framework should be extended to handle more complex constraints, if any. Since PAULI’s current framework is already a hybrid of search and linear optimization, adding more complex nonlinear constraints would imply a more complex search procedure.

5.6. *Detecting insoluble problems*


As a last example, let us consider PAULI’s behavior when we input the union of the reactions from Sections 5.1, 5.2, and 5.3 of this paper, while keeping the distinction between good and bad reactions. PAULI determines, before beginning the search for quantum properties, that no conservation laws can account for the reaction data. This is because there exists a bad reaction, $\nu_\mu + n \rightarrow p + e$, that is a linear combination of five good reactions:



The first reaction is taken from the problem in Section 5.2, whereas the remaining four come from Section 5.3; the full reaction data for all problems are given in the Kocabas paper.

Since the sum of any conserved quantum property across a good reaction is zero, any linear combination of good reactions must also possess a zero sum. Hence the problem data are insoluble, since the dependent bad reaction cannot be made to violate any quantum property that the good reactions conserve.

A reason for the above contradiction was offered by Yuval Ne'eman (personal communication, October 26, 1993): the reactions data in Sections 5.2 and 5.3 (taken from the Kocabas paper) do not reflect the modern distinction between "neutrinos associated with muons" and "neutrinos associated with electrons," rather, both types are simply called "neutrinos." For example, the reaction of beta decay in Kocabas (p. 279) is



whereas the modern interpretation is that the third product is instead the antineutrino electron $\bar{\nu}_e$. We quote again from Ne'eman and Kirsh (p. 72):

During the 1930s a particle named the muon ... was discovered. It was found that in some of the processes in which the muon participates, a neutrino, similar in its properties to the neutrino of the electron, is also produced. However, in 1962 the Americans Lederman, Schwartz, and Steinberger found that the muon neutrino is in fact different from the 'electron' neutrino - it 'remembers' its origin and in reactions which it initiates only muons, and not electrons, are produced.

Later sections of Kocabas's paper involve the modern distinction between neutrinos, although the earlier sections do not, in order to reflect the historical situation.

The contradiction found when considering all the reactions as a whole illustrates the benefits of an algebraic representation of reactions as matrices. In fact, detecting such contradictions can motivate postulating two types of a single particle (as was done with neutrinos) in a data-driven manner.

To detect insoluble reaction data, first one finds the rank of the good-reactions matrix. Then, for each bad reaction B, one finds the rank of the good-reactions matrix augmented with B. If the rank of the augmented matrix remains the same, then clearly B is linearly dependent on the good reactions, and the problem cannot be solved with the use of conservation laws. The problem is soluble only if the rank increases in every case.

By taking a matrix-algebraic viewpoint on the task, one also sees that the number of good reactions can be reduced at the start by finding a full-rank subset of the reactions, since any good reaction that is linearly dependent on this subset will conserve any property that the subset does. PAULI carries out this reduction of the good reactions.

6. Discussion

This section compares in some detail the approaches embodied in BR-3 and PAULI. The two main contrasts between the approaches lie in the use of algebraic methods, and also in whether to examine bad reactions incrementally or as a whole. In closing, we analyze the limitations of the newer program, and comment on possible relevance to current physics research.

6.1. Comparison on general approach

An important difference between the PAULI and BR-3 approaches is that PAULI reasons about reactions at a highly algebraic level, whereas BR-3 deals with reactions mostly at the arithmetic level. A key difference between arithmetic and algebra is their level of abstraction. Whereas arithmetic deals with specific numbers, algebra introduces the replacement of numbers by variables. The immediate consequence of this replacement is greater generality. Moreover, mathematics has contributed powerful reasoning tools at the algebraic level. BR-3 does make use of a simple linear equation solver, but it is unclear from the paper what sorts of algebraic manipulation this solver can carry out.

We illustrate what we mean by the arithmetic level as follows. If a program approached the current task completely by naive generate-and-test, in which a small set of integer values are assigned in all ways to the particle properties, and balance were tested by summing the integer assignments, then that program would be entirely arithmetic, and would make no use at all of algebraic reasoning. The DALTON program described in Langley, Simon, Bradshaw, and Zytkow (1987) operated largely in this manner.

PAULI represents sets of reactions as explicit algebraic equations by replacing the ‘ \rightarrow ’ by ‘ $=$ ’ and interpreting the particles as variables over the domain of numbers. The resulting set of homogeneous equations can be manipulated with the familiar methods of elementary algebra, matrix algebra, and linear programming. That is, algebraic constraints (by definition, containing variables) are thoroughly manipulated in order to carry out the intended inference, whether to solve for the variables, test linear dependence, or some other purpose.

The algebraic representation, which makes connection with powerful mathematical and algorithmic methods, is the source of PAULI’s power. It likewise is one source of the power of MECHEM (Valdes-Perez, 1994) in comparison with the programs in Langley et al. (1987) dealing with chemical reactions. The STAHLp program (Rose and Langley, 1988) had some reasoning at the algebraic level, in the sense of substituting variables in expressions, but without the full algebraic representation and manipulation of the type present in PAULI and MECHEM.

A second important difference concerns strictly how BR-3 and PAULI were used rather than the programs themselves. Kocabas used BR-3 incrementally to examine one bad reaction at a time, whereas in this paper PAULI considers all of the available data at once. We argue below that the latter strategy is preferable; our evidence is that PAULI consistently found simpler solutions given the same facts and constraints. In two cases,

PAULI's solutions were incontrovertibly simpler, since they involved postulating one quantum property rather than two.

6.2. Comparison on original criteria

In section 4 of his paper, Kocabas evaluates BR-3 along several dimensions. The present subsection compares BR-3 and PAULI along those same dimensions that we consider to reveal interesting differences.

Research goals. The goal of BR-3 is seemingly to postulate a conserved quantity when confronted with a single unexplained bad reaction. One never needs more than one such quantity to explain one bad reaction, so the program ordinarily postulates one quantity at a time. Kocabas used BR-3 mostly in this incremental way, although he mentions having tested the program in "batch" mode. PAULI, on the other hand, aims to postulate the minimal number of conserved quantities Q required to explain all the known good and bad reactions, and to find the simplest combination of quantum numbers for the given Q . Of course, nothing limits PAULI to running only once; new data on bad reactions could prompt calling the program with a new state of knowledge. We do not view incrementalism as a necessary virtue in such areas of scientific discovery, since new conserved quantum properties are not postulated very often.

Belief revision and search methods. Kocabas explains that BR-3 revises its beliefs as it backtracks upon finding that an assignment of quantum numbers has led to a contradiction, e.g., a good reaction is found to violate a property. Further, the program's search method is largely inherited from Prolog: integer assignments to particle properties are generated and tested for contradictions. In contrast, PAULI formulates the constraints explicitly at the start, and then uses an efficient search algorithm (linear programming with simplex) that involves hill climbing from one feasible solution to a better one, finally arriving at an optimum. The design of PAULI exploits the formal structure of the problem, and uses algebraic methods to carry out much of the inference.

Generality of methods. There are two forms of generality typically of interest in AI. One refers to the ability to handle many problems within a single task domain, and another refers to the range of distinct tasks that are susceptible to similar methods. We compare BR-3 and PAULI in both respects.

We have some evidence that PAULI can handle more problems within the same domain than can BR-3, since it was able to solve the strangeness problem of Section 5.4. We conjecture that BR-3 would output too many new properties for the strangeness case when used incrementally. In addition, PAULI searches a larger space of assignments than BR-3's discrete set $\{-1,0,1\}$. The latter program could of course execute with a larger set, but the run time would grow exponentially with this factor under BR-3's backtrack search regimen, whereas the linear programming in PAULI automatically considers the whole range of non-negative numbers.

The core of PAULI is an optimization carried out by linear programming. Experience has shown that the simplex algorithm of linear programming is applicable across a wide range of tasks from various disciplines, including scientific inference. For example, we have used it within MECHEM to determine an overall stoichiometry of a set of chemical

reactions (Valdes-Perez, 1991), and to infer whether a set of reactions is consistent with experimental evidence of catalysis (Valdes-Perez, 1992b).

Effects of data order and size. PAULI has little dependence on the order of data presentation. The minimal number of quantum numbers needed will never depend on the order, although in satisficing mode the exact assignment found may vary. BR-3's behavior does depend on the order, as Kocabas points out. The effect of problem size on PAULI was given as Expression (9) in Section 4.3. Kocabas does not give an analogous result, but does point out that for BR-3 batch processing becomes more economical than incremental processing as the number of reactions increases. We have also demonstrated above that batch processing gives rise to simpler solutions.

6.3. Comparison on simplicity

Comparing the relative simplicity of two computer models is not straightforward. It is generally agreed that a model M_1 is simpler than another model M_2 if M_1 has fewer free parameters, *ceteris paribus*. Additionally, in the case of AI programs, "free parameters" can include ad-hoc heuristics of meager generality, either across tasks or within the same task domain. PAULI contains a free parameter that acts to bound the sum of quantum numbers across any bad reaction.⁴ BR-3 does not have this free parameter, but for the reason that its search is limited by an ad-hoc heuristic to the small set of integer values $\{-1,0,1\}$. Hence, the two systems appear roughly equal in terms of simplicity.

6.4. Limitations of the current approach

There are two main limitations to PAULI. First is the expected exponential growth in computation with the number of bad reactions. This has not proved important in the historical examples addressed here. However, on a much larger set this factor could present problems. We have not studied this issue, but we conjecture that divide-and-conquer heuristics could be introduced, together with an incremental version of the simplex algorithm, to reduce computation time.⁵

The second limitation is that it remains unclear whether there are other physical constraints one would want to incorporate that cannot not be expressed within the current framework of linear constraints augmented with linear inequations. Certainly one can conceive of such constraints (e.g., arbitrary nonlinear ones), but whether they are needed to solve any problems of scientific interest in this domain is not known to us.

6.5. Relevance to current physics

As in the case of strangeness, physicists sometimes postulate phenomenological selection rules in order to explain why observed reactions happen and unobserved ones do not. Selection rules are still being invented today, as revealed by searches of databases of physics literature using keywords like 'selection rule' and 'novel.'

Nevertheless, the relevance of a program such as BR-3 or PAULI to current physics is still unclear. It might be interesting to run these programs, or their successors, on all the observed and unobserved reactions known to physics to determine a minimal set of selection rules or conserved quantum properties. As a first step, we have submitted a short article to a physics journal (Valdes-Perez, 1993) in order to acquaint physicists with PAULI (which is available by emailed request to valdes@cs.cmu.edu). The paper also discusses how the program can be modified to discover entropy-like selection rules, in which observed reactions *increase or maintain* a quantity, rather than exactly conserve it, and unobserved reactions diminish the quantity.

7. Conclusion

An important goal of computational scientific discovery is to impose order on the seemingly great variety of scientific inference by means of computational models. Reasoning about reaction-like processes is one category of scientific inference found throughout natural science; biology, chemistry, petrology, and physics all deal with reactions or with analogous processes. We have even argued elsewhere that the production rules of cognitive psychology bear potentially useful similarities to reactions (Valdes-Perez, Simon, and Murphy, 1992).

The present work, together with some of our other work (Valdes-Perez, 1994; Valdes-Perez, 1992a) imposes a semblance of order on reasoning about reactions by developing algorithms that answer various scientific questions, including determination of the simplest pathway (set of concurrent reactions) consistent with experimental data. The main argument in this paper is that reasoning about sets of reactions is better accomplished with the aid of mathematical and algorithmic tools that accompany an algebraic representation of reactions. Our description of an alternative, arguably superior model for the discovery of conserved quantum properties in particle physics illustrates the power of this idea. While we have also pointed out an advantage to using all of the available data, rather than proceeding incrementally, it is not our intention to formulate this specific point into a general lesson.

Finally, we have recently reported unsuspected similarities in the search spaces of several discovery tasks, including the task of BR-3 and PAULI (Valdes-Perez, Zytkow, and Simon 1993). We expect that these similarities, and the concept and notation introduced to express them, will lead to further applications of the central lesson of this paper.

Acknowledgments

The author thanks P. Pandurang Nayak, formerly of Stanford University, for making available his Common Lisp implementation of the simplex algorithm, and Sakir Kocabas for discussions about BR-3. Yuval Ne'eman pointed out the slight errors in the particle reactions. Jan Zytkow and the anonymous reviewers made extensive comments on drafts of this paper. This work was partially supported by the W.M. Keck Center for Advanced Training in Computational Biology, by a Science and Technology Center grant from the

National Science Foundation, #BIR-8920118, and by a High Performance Computing and Communications grant from the National Science Foundation, #ASC-9217091.

Notes

1. GELL-MANN would be a natural name for the present program, since Murray Gell-Mann was a co-discoverer of strangeness (Ne'eman and Kirsh, 1986). However, M. Gell-Mann also proposed quark models, hence the name of the program in Fischer and Zytow (1990).
2. Actually, the Constraints (4) are not an absolute guarantee that no solution will be missed. For example, if $\sum_k p_{jk1} - \sum_l r_{jl1}$ equals 10, $\sum_k p_{jk2} - \sum_l r_{jl2}$ equals $-1/10$, and there are only two quantum properties being considered, then the sum $10 + 100 \times (-1/10)$ equals zero, and this satisfactory solution will be rejected. However, it seems unlikely that solutions involving fractional quantum numbers exist while integral solutions do not, given the relatively small sizes of the reactions (in terms of number of reactants and products). The Constraints (4) constitute a reliable heuristic implementation of Disjunction (2).
3. All execution times in this paper refer to solitary use of a Silicon Graphics Indigo.
4. In any case, we have not needed to adjust this parameter on any problem the program has handled, since in all cases the program found that a single new quantum property was sufficient.
5. The term 'incremental' has a different meaning here. An incremental version of simplex allows adding and deleting one or a small number of constraints and re-doing the optimization with less work than formulating the constraints from scratch.

References

- Chvatal, V. (1983). *Linear programming*. New York, NY: W.H. Freeman.
- Fischer, P. and Zytow, J. M. (1990). Discovering quarks and hidden structure. *Proceedings of the Fifth International Symposium on Methodologies for Intelligent Systems* (pp. 362–370). Knoxville, TN: North Holland.
- Godfrey, J., McLachlan, R., and Atwood, C. H. (1991). Nuclear reactions versus inorganic reactions. *Journal of Chemical Education*, 68, 819–824.
- Kocabas, S. (1991). Conflict resolution as discovery in particle physics. *Machine Learning*, 6, 277–309.
- Langley, P., Bradshaw, G. L., and Simon, H. A. (1981). Bacon.5: The discovery of conservation laws. *Proceedings of the Seventh International Joint Conference on Artificial Intelligence* (pp. 121–126). Vancouver, Canada.
- Langley, P., Simon, H.A., Bradshaw, G.L., and Zytow, J.M. (1987). *Scientific discovery: Computational explorations of the creative processes*. Cambridge, MA: MIT Press.
- Ne'eman, Y. and Kirsh, Y. (1986). *The particle hunters*. New York: Cambridge University Press.
- Omnes, R. (1971). *Introduction to particle physics*. New York: Wiley Interscience.
- Papadimitriou, C., and Steiglitz, K. (1982). *Combinatorial optimization: Algorithms and complexity*. Englewood Cliffs, NJ: Prentice-Hall.
- Rose, D., and Langley, P. (1988). A hill-climbing approach to machine discovery. *Proceedings of the Fifth Machine Learning Conference* (pp. 367–373). Ann Arbor, MI: Morgan Kaufmann.
- Sutton, C. (1990). Happy birthday, little neutral one. *New Scientist*, 128, 34–36.
- Valdes-Perez, R. E. (1994). Conjecturing hidden entities via simplicity and conservation laws: Machine discovery in chemistry. *Artificial Intelligence*, 65, 247–280.
- Valdes-Perez, R. E. and Erdmann, M. (1993). Systematic induction and parsimony of phenomenological conservation laws. *Computer Physics Communications*, in press.
- Valdes-Perez, R. E. (1991). On the concept of stoichiometry of reaction mechanisms. *Journal of Physical Chemistry*, 95, 4918–4921.
- Valdes-Perez, R. E. (1992a). Theory-driven discovery of reaction pathways in the MECHEM system. *Proceedings of Tenth National Conference on Artificial Intelligence* (pp. 63–69). San Jose, CA: AAAI Press.

- Valdes-Perez, R. E. (1992b). A necessary condition for catalysis in reaction pathways. *Journal of Physical Chemistry*, 96, 2394–2396.
- Valdes-Perez, R. E., Simon, H. A., and Murphy, R. F. (1992). Discovery of pathways in science. *Proceedings of the ICML 1992 Workshop on Machine Discovery* (pp. 51–57). Aberdeen, Scotland.
- Valdes-Perez, R. E., Zytow, J.M., and Simon, H.A. (1993). Scientific model-building as search in matrix spaces. *Proceedings of Eleventh National Conference on Artificial Intelligence* (pp. 472-478). Washington, D.C.:AAAI Press.

Received October 1, 1992

Accepted November 15, 1993

Final Manuscript November 15, 1993