

Conflict Resolution as Discovery in Particle Physics

SAKIR KOCABAS

Department of Electronic and Electrical Engineering, King's College London, London WC2R 2LS, U.K.

Editor: Pat Langley

Abstract. In this paper we focus on the general problem of identifying and resolving conflicts such as inconsistency and incompleteness in scientific discovery. The underlying idea is that conflicts in empirical knowledge can lead to new discoveries, provided that they are clearly identified. Our observations are based on the behavior and the results of an incremental discovery program, BR-3. The system models the discoveries of a series of quantum laws by physicists in this century. BR-3's discoveries are directed by a small set of consistency and completeness constraints. We evaluate the system in terms of its knowledge representation, discovery operators, and their behavior, and we describe its relation to other work in scientific discovery.

Keywords. Theory revision, conflict resolution, contradiction, incompleteness

1. Introduction

The basic requirements for a scientific theory include explanatory power, freedom from contradictions, and simplicity. The first of these requirements is concerned with completeness, whereas the second one is related to consistency. A scientific theory is expected to explain any phenomenon describable by its terms in a systematic way that is free from contradictions.

We use the terms *completeness* and *consistency* in an informal sense, in a way different from their formal counterparts. The formal concept of completeness implies that all well-formed formulae are provable within the formal system in question, whereas consistency in such systems means freedom from contradicting statements. However, in theoretical systems we are concerned with explainability rather than provability, and freedom from contradictions with facts rather than logical contradictions.¹

We can now define conflicts in theoretical systems in terms of these informal notions of incompleteness and inconsistency: A knowledge system is in a state of *conflict* when it is faced with facts in its domain that it cannot explain, or when any of its propositions contradicts those facts.

In theoretical systems, incompleteness and inconsistency are usually local problem states as long as they do not involve fundamental theoretical concepts. Resolution of conflicts in knowledge systems depends on their identification in a clear way. In general, incompleteness in such systems can be resolved by introducing independent beliefs or hypotheses that can explain previously unexplainable facts. On the other hand, contradictions can be resolved by deleting those beliefs or hypotheses that contradict the facts or by reducing their generality.

The BR-3 program described below is capable of identifying and resolving such conflicts within the general strategies just stated. The system models the discoveries of a series of quantum properties and the related conservation laws. However, before we describe the system and its behavior in detail, it is appropriate to present some background information about its task domain, particle physics, making use of the similarities and differences between this field and a more familiar one, chemistry.

1.1. *The domain of particle physics*

Just as many chemical elements react to produce chemical compounds, the elementary particles interact to produce other such particles. As some chemical compounds decompose into their simpler constituents, most heavier elementary particles decay and produce lighter particles (i.e., particles with smaller rest masses). In addition, chemical reactions and particle interactions can be represented in the same general formula

Reaction materials \rightarrow Reaction products.


From the standpoint of computational representation, this means that particle reactions can be represented in the same data structures as used in STAHL (Zytkow & Simon, 1986), STAHLp (Rose & Langley, 1986), and GLAUBER (Langley, Simon, Bradshaw & Zytkow, 1987) for the representation of chemical reactions. Moreover, some of the inference rules (e.g., substitution and reduction) used by these systems are also applicable to particle reactions.

There are also parallelisms between chemical reactivity and the conditions that control particle interactions. Just as the occurrence of chemical reactions are constrained by a number of chemical laws, so particle reactions are constrained by a set of quantum laws (e.g., see Davies, 1985). These laws dictate the conservation of certain quantum properties in particle reactions. These quantum properties include energy, electrical charge, lepton number, baryon number, spin, and strangeness. Table 1 shows most of the well-known particles with their quantum properties and values.

Despite the similarities, there are also some basic differences between particle reactions and chemical reactions. We shall address these later, but for now let us continue with our survey. The earliest known conservation laws in particle physics were the conservation of energy and of electrical charge. The latter can be stated as follows:

The sum of the electrical charges of the particles entering a reaction is equal to the sum of the charges of the resultant particles.

Two reactions that conserve electrical charge and that have been "observed" by physicists are



where p , n , π , π_0 , and γ designate the proton, neutron, pion, pion-zero, and gamma particles respectively. It has been known since early this century that the proton and electron

Table 1. Elementary particles and their quantum properties.

Particle	Mass (MeV)	Charge State	Spin	Lepton Number	Baryon Number	Strangeness
g	0	0	2	0	0	0
γ	0	0	1	0	0	0
ν	0	0	0.5	1	0	0
e	0.511	-1	0.5	1	0	0
μ	105.66	-1	0.5	1	0	0
τ	1784.2	-1	0.5	1	0	0
π	139.57	1	0	0	0	0
π_0	134.96	0	0	0	0	0
K	493.67	1	0	0	0	1
K_0	497.67	0	0	0	0	1
η	548.8	0	0	0	0	0
p	938.28	1	0.5	0	1	0
n	939.57	0	0.5	0	1	0
Λ	1115.60	0	0.5	0	1	-1
Σ^+	1189.36	1	0.5	0	1	-1
Σ^0	1192.46	0	0.5	0	1	-1
Σ^-	1197.34	-1	0.5	0	1	-1
Ξ^0	1314.9	0	0.5	0	1	-2
Ξ^-	1321.3	-1	0.5	0	1	-2
ω	1672.5	-1	1.5	0	1	-3

Symbols: g : graviton, ν : neutrino, e : electron, μ : muon, τ : tauon, π : pion, π_0 : pion-zero, K : kaon, K_0 : kaon-zero, η : eta, p : proton, n : neutron, Λ : lambda, Σ^+ : sigma-plus, Σ^0 : sigma-zero, Σ^- : sigma-minus, Ξ^0 : xi-zero, Ξ^- : xi-minus, ω : omega. The antiparticles—not shown in the table—are designed by “-” on the original particle symbol, and their quantum values are the opposites of the original particles. Some elementary particles, such as g , γ , and η , do not have distinguishable antiparticles.

have opposite and unit electrical charges. The neutron has been known to be unstable, decaying into a proton, an electron, and an antineutrino in what is called “beta decay,” or

$$n \rightarrow p + e + \bar{\nu}.$$


However, a proton decay has never been observed, and the stability of this particle had puzzled the physicists. Why does it not decay into lighter particles? Reactions such as

$$\begin{aligned} p &\rightarrow \pi + \pi_0 \\ p &\rightarrow \bar{e} + \gamma \end{aligned}$$

never happen despite the fact that they apparently obey the charge conservation law. A theoretical framework based only on the charge conservation law would not be capable of explaining the absence of these reactions. In other words, such a theory would be incomplete concerning particle reactions. The discrepancy between the theoretically valid and physically observable reactions² was a conflict that had to be resolved if particle physics were to become a consistent and complete theory.


1.2 Theory development by conflict resolution

Physicists resolved such conflicts by postulating new quantum properties and laws so that the physically unobservable reactions were made theoretically invalid by these laws (Omnes, 1970). Their conclusion was that these reactions did not conserve a quantum property, a new kind of “charge,” possessed by some of the particles in the reaction formulae. The next question was: Which particles have this new charge? To illustrate how physicists resolved such conflicts, let us consider a reaction that conserves electrical charge but has consistently escaped observation:



Let us assume that this reaction violates the conservation of a new property (e.g., the “protonic charge”). Now, if we arbitrarily assign the new charge value to the proton as one and assume that the other particles, π and π_0 , do not have this charge (i.e., they both have zero protonic charge), then the reaction would be unbalanced by the new charge. In this way the reaction $p \rightarrow \pi + \pi_0$ would be invalid by the conservation of the protonic charge values as indicated by the inequality $1 \neq 0 + 0$, and this would explain why the reaction had never been observed. Nevertheless, the value set [1, 0, 0] is not the only one that makes the reaction unbalanced, as the values [0, 1, 1], [0, 1, 0], [0, 0, 1], and [1, 1, 1] produce the same effect.

On the other hand, although the new charge values render the other unobserved reaction ($p \rightarrow \bar{e} + \gamma$) unbalanced, and thereby can explain its absence, the values make some unobserved reactions unbalanced, as in



These reactions conserve electrical charge, but not the “known” values of the new charge. This can be seen by substituting the protonic charge values:

$$\begin{aligned} 1 + 1 &= 1 + n + 0 \\ \bar{\pi} + 1 &= n + 0. \end{aligned}$$

This suggests that some of the other particles taking place in these reactions must have nonzero protonic charge. Here, if we assign the protonic charge value of one to the neutron and zero to $\bar{\pi}$, the reactions would be balanced. This resolves the immediate problem, but some other valid and observed reactions may conflict with the assigned values, and we may have to revise some of the assumptions about the protonic charge values of particles accordingly.

This example shows that the initial values (or “cause values”) assigned to the particles [p , π , π_0] determine the values that can be given to the particles (“effect values”) in the reactions made unbalanced by the former. Thus, if we assign the cause values [0, 1, 0] instead of [1, 0, 0] to the particles [p , π , π_0], these would make some other valid reactions unbalanced, and as a result, the effect values would be different.

The BR-3 system described below simulates the discovery of quantum properties like those discussed above. It starts with a general rule of conservation of quantum values, the electrical charge values of the elementary particles, a set of physically observed particle reactions, and a set of reactions that are not observed but valid by the charge conservation law. The system postulates new quantum properties to explain the absence of certain particle reactions, and assigns the particles their quantum values for the postulated properties. BR-3 revises its “beliefs” about the quantum values of the particles when additional data about such reactions are not consistent with the quantum values that it knows. In this way it models the discoveries of the lepton, baryon, electron, and muon number properties by physicists in this century.

The next section presents an overview of the system. Section 3 describes the program’s behavior in modeling the discoveries of the quantum properties and the related quantum laws. Section 4 provides a discussion on the system’s research goals, knowledge representation, theory revision and search methods. The section also discusses the generality of the program and provides some ideas on how it can be improved. Section 5 presents a brief review of the related work. Finally, Section 6 concludes with a summary of the results.

2. An overview of BR-3

This section explains the representation of BR-3’s descriptive and definitive knowledge, describes its overall inputs and outputs, and discusses the system’s operators in terms of their inference rules, inputs and outputs, and their behavior.

2.1. Knowledge representation

BR-3 represents its descriptive and definitive knowledge in terms of predicate statements,³ which fall into three main categories: (1) logical knowledge, (2) formal knowledge, and (3) theoretical, hypothetical, or empirical knowledge. Logical knowledge contains the definitions of logical functions and relations such as “and,” “or,” “for all” and “if-then-else.” By definition, this knowledge is totally domain independent. Formal knowledge includes definitions, class memberships, and class-superclass relationships. Finally, theoretical knowledge includes statements expressing the physical properties of domain objects and certain relationships between them.

Examples of BR-3’s formal knowledge include the beliefs that:

- A proton is a particle.
- A neutron is a particle.
- An antiparticle is designated by \bar{x} where x is a particle.

Here, the first two statements express member-class relationships, while the last one gives a definition of “antiparticle.” Initial theoretical knowledge contains physical data about the elementary particles, namely their electrical charge values:

- The electron’s electrical charge is -1 .
- The muon’s electrical charge is -1 .
- The proton’s electrical charge is 1 .
- The neutron’s electrical charge is 0 .

Theoretical knowledge of BR-3 also includes statements in which particle reactions are represented as ordered pairs of lists as in GLAUBER (Langley, Simon, Bradshaw, & Zytkow, 1987). In this case the elements of the first list are the reactant particles and the second specifies the resultant particles. For example, the reaction $p + p \rightarrow p + p + \pi_0$ is represented as [p, p], [p, p, π_0]. Valid reactions are identified by the internal label “reaction,” while physically unobserved reactions by the label “unobserved reaction.”

2.2. *Inputs and outputs*

The system’s initial inputs are the basic formal and theoretical knowledge of its domain, a set of observed reactions, and a set of unobserved reactions like those given above. Its outputs are a set of new beliefs, categorized as theoretical statements, about postulated properties that can explain the system’s previous state of incompleteness and/or inconsistency. Some example outputs are

The qpl value of p is 0.
The qpl value of e is 1.

where qpl indicates the postulated quantum property, and 0 and 1 are the values of the property for p and e , respectively.

BR-3 is an incremental discovery system. Once it reaches a consistent state in terms of its formal and theoretical domain knowledge, new observed and/or unobserved reaction formulae can be given to it for evaluation. These may initiate a new search for a consistent and complete state, leading to further changes in theoretical knowledge. This cycle continues as long as new inputs arrive.

2.3. *BR-3’s discovery operators*

BR-3 revises its theoretical knowledge by using three operators—CHECK-CONSISTENCY, CHECK-COMPLETENESS, and REVISE BELIEFS—which control a series of inference rules or procedures. Table 2 summarizes these operators, and Figure 1 shows the interactions among them with the arrows showing the flow of control. Briefly, the CHECK-CONSISTENCY operator checks for consistency between the quantum values and observed reactions. If the test succeeds, then the operator passes the control to CHECK-COMPLETENESS; otherwise, control goes to the REVISE-BELIEFS operator.

CHECK-CONSISTENCY tests the existing valid reactions in the knowledge base against the quantum values of the elementary particles by using its *test-reaction* rule. Those particles that have no value for a quantum property are given the default value of zero, although default values are not added to the knowledge base. CHECK-CONSISTENCY decides that a reaction is consistent when, for each quantum property, the total of quantum values of the resultant particles is equal to that of the reactant particles. If some valid reactions violate a quantum law, this means that the system’s knowledge of quantum values are in contradiction with the reactions. CHECK-CONSISTENCY labels all such reactions as “unbalanced,” and passes control to REVISE-BELIEFS.

Table 2. High-level descriptions of BR-3's discovery operators.

CHECK-CONSISTENCY

Check consistency between quantum values and observed reactions.
 If the system is consistent,
 then check completeness;
 else revise beliefs.

CHECK-COMPLETENESS

Get the unobserved reactions.
 Test them against the conservation rules.
 If the test fails,
 then report completeness and stop;
 else postulate a new quantum property;
 generate cause hypotheses;
 add cause beliefs; and
 check consistency.

REVISE-BELIEFS

generate-effect-hypotheses:
 If there are no effect beliefs,
 then generate effect hypotheses;
 add effect beliefs; and
 check consistency.

revise-effect-beliefs:
 If there are effect beliefs,
 then revise effect beliefs; and
 check consistency.

revise-cause-beliefs:
 If there are no effect hypotheses,
 then revise cause beliefs; and
 check consistency.

When all valid reactions are consistent, BR-3 activates the CHECK-COMPLETENESS operator, which first checks if there are any unobserved reactions in the knowledge base. If there are none, then this means that, with respect to the quantum values and particle reactions, the system's knowledge is complete. If there are unobserved reactions, the operator checks each unobserved reaction against the quantum conservation laws in order to explain their absence by the violation of one such law. If all the unobserved reactions violate some quantum law, then BR-3's domain knowledge is complete. When the knowledge base is consistent and complete, the system comes to a halt.

On the other hand, if an unobserved reaction is found to be "valid" by all the quantum laws that BR-3 knows, then the CHECK-COMPLETENESS operator postulates a new quantum property. Then it transforms the unobserved reaction formula into an inequality to find the sets of quantum values that would make the reaction invalid in terms of the new property. These sets of values are called *cause hypotheses*. Every cause hypothesis designates a set of beliefs about the quantum values, which BR-3 adds to its theoretical knowledge

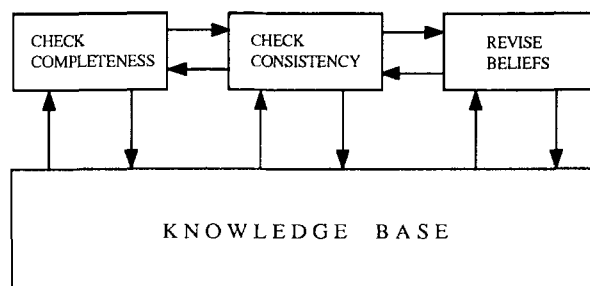


Figure 1. BR-3's operators and their interactions.

as “cause beliefs.” Once the system has added these beliefs, it must check for their consistency with the observed reactions in the knowledge base, and so passes the control to CHECK-CONSISTENCY.

The REVISE-BELIEFS operator has three main components, which are themselves complex inference rules. The first, which we call *generate-effect-hypotheses*, is activated when the cause beliefs contradict some valid reactions. As mentioned earlier, such reactions are labeled as “unbalanced” by the CHECK-CONSISTENCY operator. The *generate-effect-hypotheses* rule transforms the unbalanced reactions into linear equations, and the solutions to these equations are recorded as the *effect hypotheses*. Every effect hypothesis designates a set of beliefs about the quantum values for the particles in the unbalanced reactions. REVISE-BELIEFS arbitrarily selects one of these hypotheses, and adds the quantum values to its theoretical knowledge as the “effect beliefs.” Once BR-3 has added these beliefs, it must check them for consistency against the valid reactions, so control once again passes to CHECK-CONSISTENCY.

BR-3 uses a simple linear equation solver to generate its effect beliefs. The equation solver uses algebraic constraints to eliminate a large number of superficial solutions before passing the control to CHECK-CONSISTENCY. As a result, rather than producing several hundreds or thousands of effect hypotheses for each set of cause beliefs, the system generates only a few.

If there are effect beliefs and the current beliefs are still inconsistent, then BR-3 activates the *revise-effect-beliefs* component. This rule deletes the effect beliefs by using the current effect hypothesis as a pointer, deletes the effect hypothesis itself, and adds the new effect beliefs designated by the next effect hypothesis. Then control passes to CHECK-CONSISTENCY to check for the consistency of the new set of beliefs. If there are no other effect hypotheses to try, then the system activates the *revise-cause-beliefs* component. This rule deletes the cause beliefs by using the current cause hypothesis as a pointer, adds the new cause beliefs, and then passes the control to CHECK-CONSISTENCY. Control is transferred from one operator to another as described above, until the system reaches a consistent and complete knowledge state regarding the particle reactions and quantum properties, or until it is interrupted. Figure 2 shows how the program generates and tests its cause and effect hypotheses.

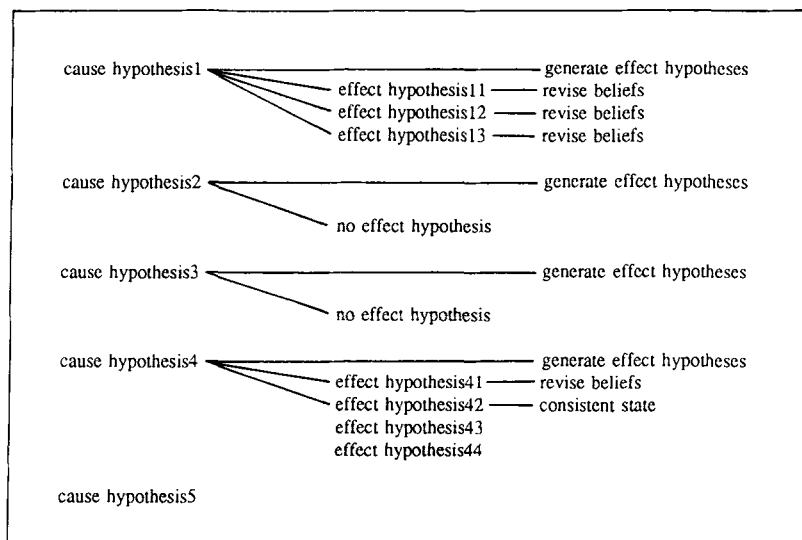


Figure 2. BR-3 generates cause hypotheses and checks each set of cause beliefs for consistency, generates effect beliefs when necessary, and checks for consistency and revises its beliefs until it reaches a consistent state. Some of the cause hypotheses do not yield effect hypotheses due to algebraic constraints.

In summary, BR-3 achieves consistency by revising its cause and effect beliefs, using a simple depth-first search with backtracking to look for complete and consistent hypotheses. Accordingly, the system tests all the effect hypotheses before it revises (and deletes) the cause hypothesis itself. The effect hypotheses are generated for each cause hypothesis separately.

3. The system's behavior in particle physics

Having described BR-3 in terms of its knowledge representation, inputs and outputs, and its operators, we can now examine how the program actually achieves and maintains completeness and consistency by postulating new quantum properties, assigning quantum values to the elementary particles, and revising these values as necessary.

3.1. Rediscovery of the lepton number property

In a test run, BR-3 was given the fourteen observed and four unobserved reactions in Table 3. Initially it knew only about the electrical charge values of 36 elementary particles, some of which are shown in Table 1, and a general conservation rule of quantum properties.

When provided with this knowledge, BR-3 checks its knowledge of valid reactions and quantum properties for consistency under the control of CHECK-CONSISTENCY. Since

Table 3. Some physically observed and unobserved reactions.

Observed reactions	Unobserved reactions
$p + p \rightarrow p + p + \pi_0$ (1)	$p \rightarrow \bar{e} + \gamma$
$p + p \rightarrow p + \pi + n$ (2)	$p \rightarrow \pi + \pi_0$
$p + \pi \rightarrow \pi + p$ (3)	$p \rightarrow \pi + \gamma$
$\bar{\pi} + p \rightarrow \bar{\pi} + p$ (4)	$p \rightarrow \pi + \pi + \bar{\pi} + \pi_0 + \pi_0$
$\bar{\pi} + p \rightarrow \pi_0 + n$ (5)	
$\bar{\pi} + p \rightarrow p + \pi + \bar{\pi} + \bar{\pi}$ (6)	
$\gamma + e \rightarrow \gamma + e$ (7)	
$e + p \rightarrow e + p$ (8)	
$\pi_0 \rightarrow \gamma + \gamma$ (9)	
$\bar{\pi} \rightarrow \mu + \bar{\nu}$ (10)	
$\pi \rightarrow \bar{\mu} + \nu$ (11)	
$\mu \rightarrow e + \nu + \bar{\nu}$ (12)	
$n \rightarrow p + e + \bar{\nu}$ (13)	
$\bar{\pi} + p \rightarrow \Lambda + K_0$ (14)	

it initially knows only about the electrical charge, the program tests the valid reactions one by one against this property. It finds that all satisfy the charge conservation law, and under the control of CHECK-COMPLETENESS, checks for completeness regarding the unobserved reactions. In order to explain the latter's absence, the program tests them for invalidity by finding out if they violate the charge conservation law. It tries the first reaction $p \rightarrow \bar{e} + \gamma$ first, but fails to explain its absence, because this reaction conserves electrical charge.

Faced with the incompleteness of its knowledge, BR-3 concludes that there must be another quantum property, and that this reaction violates the conservation of this property. Hence it postulates a new quantum property—qpl—for which it tries to find the quantum values of the particles in the unobserved reaction $p \rightarrow \bar{e} + \gamma$. To constrain search, it uses a list of three possible quantum values: 1, 0, -1. These values constitute a simple discrimination set, although one could consider a larger set of values at the cost of more search, as we will see later. The system also uses an antiparticle rule to assign only the values of 0 and -1 to antiparticles, and 0 and 1 to the original particles. In order to find the qpl values, BR-3 transforms the reaction formula into an inequality like $x \neq y + z$, where x , y and z represent the qpl values of p , \bar{e} , and γ respectively. The solution of this inequality yields five qpl value sets for p , \bar{e} , and γ .

$$\begin{aligned}
 0 &\neq 0 + 1 \\
 0 &\neq -1 + 0 \\
 1 &\neq 0 + 0 \\
 1 &\neq -1 + 0 \\
 1 &\neq -1 + 1
 \end{aligned}$$

The system records the combinations of qpl values that satisfy the inequality as cause hypotheses, which are

qpl values: $p = 0, \bar{e} = 0, \gamma = 1$.
 qpl values: $p = 0, \bar{e} = -1, \gamma = 0$.
 qpl values: $p = 1, \bar{e} = 0, \gamma = 0$.
 qpl values: $p = 1, \bar{e} = -1, \gamma = 0$.
 qpl values: $p = 1, \bar{e} = -1, \gamma = 1$.

BR-3 tries the first cause hypothesis, adding the cause beliefs about the qpl values to its theoretical knowledge as

The qpl value of p is 0.
 The qpl value of \bar{e} is 0.
 The qpl value of γ is 1.

and by using the antiparticle rule mentioned earlier, it finds the qpl values of \bar{p} and e , adding them to its theoretical knowledge as

The qpl value of \bar{p} is 0.
 The qpl value of e is 0.

Then, under the control of CHECK-CONSISTENCY, it checks for the consistency of its knowledge about the qpl values and valid reactions. This operator gives the default value zero to the particles with no recorded values of the same property in the system's theoretical knowledge. BR-3 finds that, among the valid reactions, the new beliefs render the valid reaction (9) in Table 3 unbalanced, and records this fact as

unbalanced reaction: $\pi_0 \rightarrow \gamma + \gamma$.

This means that the belief system is in an inconsistent state with respect to qpl, because the unbalanced reaction is physically valid. To achieve consistency, BR-3 must find the effect values that would turn the unbalanced reaction into a balanced one. To do this, it transforms the unbalanced reaction into a linear equation after substituting the qpl value of γ as


$$\pi_0 = 1 + 1.$$

However, this equation has no valid solution, because π_0 can only take the values 1 or 0. Hence, the equation can yield no effect hypothesis that would balance the reaction. Here the system backtracks, and REVISE-BELIEFS takes control, which forces the system to change its cause hypothesis. Thus, the current cause hypothesis is deleted from memory and the related cause beliefs from theoretical knowledge. The next cause hypothesis was

qpl values: $p = 0, \bar{e} = -1, \gamma = 0$.

Accordingly, REVISE-BELIEFS adds the following cause beliefs about the qpl values: $p : 0, \bar{e} : -1, \gamma : 0$. From the known values, it also finds the qpl values of the antiparticles

and adds them to the system's theoretical knowledge as $\bar{p} : 0, e : 1$. However, together with the default values, the new qpl values render the valid reactions (12) and (13) unbalanced:



In order to resolve the contradiction, the system transforms these reaction formulae into linear equations:

$$\begin{aligned}n &= 0 + 1 + \bar{\nu} \\ \mu &= 1 + \nu + \bar{\nu}.\end{aligned}$$

The solutions for the first equation are $[0, 1, 1, -1]$ and $[1, 0, 1, 0]$, and those for the second are $[0, 1, 0, -1]$, $[1, 1, 0, 0]$, and $[1, 1, 1, -1]$. BR-3 finds the effect hypotheses from the combinations of these sets of values, disregarding the inconsistent combinations to ensure that the unique quantum values are given to the particles that appear in both equations. In this way the effect hypotheses are found to be

$$\begin{aligned}\text{qpl values: } \mu &= 1, \nu = 1, \bar{\nu} = -1, n = 0. \\ \text{qpl values: } \mu &= 1, \nu = 0, \bar{\nu} = 0, n = 1. \\ \text{qpl values: } \mu &= 0, \nu = 1, \bar{\nu} = -1, n = 0.\end{aligned}$$

REVISE-BELIEFS adds effect beliefs to the system's theoretical knowledge in accordance with the first effect hypothesis: $\mu : 1, \nu : 1, \bar{\nu} : -1, n : 0$, along with two additional beliefs generated by the antiparticle rule: $\bar{\mu} : -1, \bar{n} : 0$.

BR-3 invokes CHECK-CONSISTENCY to test its new beliefs about the qpl values for consistency against the valid reactions. Upon finding that no valid reaction is made unbalanced, the system reports that its knowledge about qpl is consistent.⁴ The final state of its theoretical knowledge about the qpl values of the particles includes: $p : 0, \bar{e} : -1, \gamma : 0, \bar{p} : 0, e : 1, \mu : 1, \nu : 1, \bar{\nu} : -1, n : 0, \bar{\mu} : -1, \bar{n} : 0$. When compared with the values in Table 1, one can see that these qpl values are exactly the lepton number values of the elementary particles. The qpl values of those particles that have no such recorded values are assumed as zero by default.

If there were no unobserved reaction other than $p \rightarrow \bar{e} + \gamma$, BR-3 would report the completeness of its knowledge and would come to a halt. But its existing knowledge cannot explain why the next unobserved reaction $p \rightarrow \pi + \pi_0$ does not occur, for neither the electrical charge, nor the qpl values (the lepton numbers) of the particles, make the unobserved reaction unbalanced. Therefore, when the control passes to CHECK-COMPLETENESS the system is forced to postulate a new property to resolve the conflict. Let us now examine how the system proceeds from this point.

3.2. BR-3's rediscovery of the baryon number property

Seeing that its knowledge regarding the unobserved reaction $p \rightarrow \pi + \pi_0$ is incomplete, BR-3 postulates a new quantum property, qp2, and tries to find the values of this property


for the elementary particles that it knows. In pursuing this goal, the system proceeds the same way as it does in finding the qp1 values. Since the unobserved reaction must be unbalanced by the new quantum property, BR-3 transforms the reaction formula $p \rightarrow \pi + \pi_0$ into an inequality. This inequality yields five solutions for p , π and π_0 :

$$\begin{aligned} 0 &\neq 0 + 1 \\ 0 &\neq 1 + 0 \\ 0 &\neq 1 + 1 \\ 1 &\neq 0 + 0 \\ 1 &\neq 1 + 1. \end{aligned}$$

Corresponding cause hypotheses are recorded as

$$\begin{aligned} \text{qp2 values: } p &= 0, \pi = 0, \pi_0 = 1 \\ \text{qp2 values: } p &= 0, \pi = 1, \pi_0 = 0 \\ \text{qp2 values: } p &= 0, \pi = 1, \pi_0 = 1 \\ \text{qp2 values: } p &= 1, \pi = 0, \pi_0 = 0 \\ \text{qp2 values: } p &= 1, \pi = 1, \pi_0 = 1. \end{aligned}$$

BR-3 arbitrarily selects the first cause hypothesis and adds each cause belief about the qp2 values to its theoretical knowledge as $p : 0$, $\pi : 0$, $\pi_0 : 1$, $\bar{p} : 0$, and $\bar{\pi} : 0$. Since π_0 has no antiparticle, the rule disregards it. BR-3 now checks its knowledge of qp2 for consistency. This set of values render the reactions (1), (5), and (9) unbalanced,



where the particles n and γ are given the default qp2 value of zero. Again, this means an inconsistent state of knowledge for the property qp2. BR-3 passes control to REVISE-BELIEFS, which attempts to achieve consistency by finding the effect hypotheses that would make the unbalanced reactions balanced. For this reason, after substituting the qp2 values of p , π , $\bar{\pi}$, and π_0 , it transforms the unbalanced reaction formulae into linear equations to obtain

$$\begin{aligned} 0 + 0 &= 0 + 0 + 1 \\ 0 + 0 &= 1 + n \\ 1 &= \gamma + \gamma. \end{aligned}$$

BR-3 tries to solve the first equation, but notes that the equation has no solution. The program backtracks to try the next cause hypothesis, and adopts the new cause beliefs after deleting the earlier ones. However, the new cause beliefs render the valid reactions (2), (5), (10), (11), and (14) unbalanced. After substituting the new qp2 values, BR-3 turns these reaction formulae into linear equations, but the equations do not yield a consistent solution. The system backtracks again, and REVISE-BELIEFS forces it to change cause beliefs again.

The next set of cause beliefs render reactions (1), (2), (5), (9), (10), (11), and (14) unbalanced, and the corresponding linear equations lead to an arithmetical contradiction. BR-3 is forced to change its cause beliefs once again, so its new cause beliefs about qp2 become $p : 1$, $\pi : 0$, $\pi_0 : 0$, $\bar{p} : -1$, and $\bar{\pi} : 0$.

These make the valid reactions (2), (5), (13), and (14) unbalanced. After substituting the qp2 values of p , π , π_0 , and $\bar{\pi}$, it finds the following equations:


$$\begin{aligned} 1 + 1 &= 1 + 0 + n \\ 0 + 1 &= 0 + n \\ n &= 1 + e + \bar{\nu} \\ 0 + 1 &= \Lambda + K_0. \end{aligned}$$

By solving these equations, it finds four hypotheses:


$$\begin{aligned} \text{qp2 values: } \Lambda &= 1, K_0 = 0, n = 1, e = 1, \bar{\nu} = -1 \\ \text{qp2 values: } \Lambda &= 1, K_0 = 0, n = 1, e = 0, \bar{\nu} = 0 \\ \text{qp2 values: } \Lambda &= 0, K_0 = 1, n = 1, e = 1, \bar{\nu} = -1 \\ \text{qp2 values: } \Lambda &= 0, K_0 = 1, n = 1, e = 0, \bar{\nu} = 0. \end{aligned}$$


BR-3 tries the first effect hypothesis, adding the corresponding effect beliefs to its theoretical knowledge. Control passes to CHECK-CONSISTENCY, which in turn finds that these beliefs are inconsistent, because they render the valid reaction (10) unbalanced. Control passes to REVISE-BELIEFS, which forces the system to change its effect beliefs. After BR-3 deletes the current effect hypothesis and the related effect beliefs, it adds the new effect beliefs $\Lambda : 1$, $K_0 : 0$, $n : 1$, $e : 0$, and $\bar{\nu} : 0$. Also, by the antiparticle rule it generates $\bar{\Lambda} : -1$, $\bar{K}_0 : 0$, $\bar{n} : -1$, $\bar{e} : 0$, and $\nu : 0$.

At this point the system transfers control to CHECK-CONSISTENCY again, to see if its knowledge about qp2 is consistent. BR-3 notes that no valid reaction is made unbalanced by its latest set of beliefs, and concludes that it is consistent. These consistent beliefs about qp2 are in fact nothing but the baryon number values of the elementary particles. Having established consistency for qp1 and qp2 in this way, it checks the remaining two unobserved reactions



under the control of CHECK-COMPLETENESS, and finds that the absence of both are explainable by their violation of the conservation of qp2. By now, BR-3's knowledge is both consistent and complete regarding the charge states, the qp1 and qp2 properties, and the observed and unobserved reactions. Note that BR-3 would be equally satisfied with the qp2 values of 1 and 0, instead of 0 and 1, for K_0 and Λ . As far as its present knowledge of valid reactions is concerned, it would not make any distinction between the two sets of values. However, when the following set of valid reactions are added,






the system revises its beliefs about the qp2 values of K_0 and Λ . In this case, the REVISE-BELIEFS operator changes the effect beliefs about the qp2 values and adopts the correct beliefs.

3.3. Rediscovery of other quantum properties


In addition to discovering the lepton and baryon numbers, the current version of BR-3 can rediscover the spin (angular momentum) and strangeness properties. However, it cannot find the correct spin and strangeness values for some of the particles. This is because the spin conservation law is based on group theoretical calculations, rather than the arithmetical ones applicable to the other laws. 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.⁵ In addition, some particle reactions (e.g. weak interactions) do not conserve strangeness.

Nevertheless, the system was tested on a set of observed and unobserved reactions from Omnes (1970), to determine its behavior on the strangeness property with its existing quantum value set, 0 , 1 , and -1 . As expected, after postulating the new quantum property, BR-3 generates a series of cause and effect hypotheses. However, after trying all such hypotheses, it fails to achieve consistency, and in the end, it deletes all such beliefs about quantum values of the new property, and thus returns to its prior consistent (regarding all valid reactions that it knows) but incomplete (regarding the latest unobserved reaction) knowledge state.

BR-3 can also simulate the discovery of the quantum properties that distinguish between the three additional types of particles: the electron, muon, and tauon neutrinos. The history of this discovery goes back to early 1960s. Experiments conducted in Columbia University in 1962 revealed that these three neutrinos, together with the corresponding particles (i.e., the electron, muon, and tauon) have a series of new quantum properties (see Trefil, 1980). The purpose of the experiment was to look for the reactions



which would be initiated by the neutrinos and antineutrinos obtained from pion decays



The most striking feature of the experiments was that only the reactions given above were observed. No reactions of the type




Table 4. Observed and unobserved reactions used in modeling the discovery of electron and muon numbers.

Observed reactions	Unobserved reactions
$p + p \rightarrow p + p + \pi_0$	$\nu_\mu + n \rightarrow p + e$
$p + \pi \rightarrow \pi + p$	$\bar{\nu}_\mu + p \rightarrow n + \bar{e}$
$\pi_0 \rightarrow \gamma + \gamma$	$\nu_e + n \rightarrow p + \mu$
$\bar{\pi} \rightarrow \mu + \bar{\nu}_\mu$	$\nu_e + p \rightarrow n + \mu$
$\pi \rightarrow \bar{\mu} + \nu_\mu$	$p + \tau \rightarrow \bar{\nu}_\mu + n$
$\mu \rightarrow e + \nu_\mu + \bar{\nu}_e$	
$n \rightarrow p + e + \bar{\nu}_e$	
$\bar{\nu}_\mu + p \rightarrow n + \bar{\mu}$	
$\nu_\mu + n \rightarrow p + \mu$	

were seen at all, whereas with the neutrinos emitted from beta decay ($n \rightarrow p + e + \bar{\nu}$) such reactions (e.g., $\bar{\nu}_e + p \rightarrow n + \bar{e}$) had been observed. This indicated that there were at least two different kinds of neutrinos, one associated with electrons and the other with muons.

When BR-3 is given the observed and unobserved reaction formulae shown in Table 4, with its formal knowledge slightly modified to accommodate the distinction between the neutrinos by their origins, it postulates two new quantum properties (electron and muon numbers), and finds the corresponding quantum values shown in Table 5. (Those particles that do not appear in the table are given the quantum value of zero by default.) Note that the system does not postulate a corresponding tauon number for the tau particle and its neutrino, because the electron and muon numbers are sufficient to account for the unobserved reactions. It is interesting to note from BR-3's behavior that two quantum numbers were sufficient to distinguish between these three leptons, their neutrinos, and their antiparticles. This behavior of the system can be viewed as an instance of cognitive economy or Occam's Razor.

Table 5. Some of the quantum properties of e , μ , τ , ν_e , ν_μ , ν_τ , and their antiparticles.

Particle	Electrical Charge	Lepton number	Electron number	Muon number
e	-1	1	1	0
μ	-1	1	0	1
τ	-1	1	0	0
\bar{e}	1	-1	-1	0
$\bar{\mu}$	1	-1	0	-1
$\bar{\tau}$	1	-1	0	0
ν_e	0	1	1	0
ν_μ	0	1	0	1
ν_τ	0	1	0	0
$\bar{\nu}_e$	0	-1	-1	0
$\bar{\nu}_\mu$	0	-1	0	-1
$\bar{\nu}_\tau$	0	-1	0	0

BR-3 can even rediscover the electrical charge property, starting with only sets of observable and unobserved particle reaction formulae. In a test run with a slightly more generalized version that allows the original particles to have negative quantum values (by generalizing the antiparticle rule), the system was given the observed and unobserved reactions in Table 6 with no knowledge of any quantum property whatsoever. BR-3 begins by checking its knowledge of observed reactions for consistency. Since it has no quantum values, the consistency test succeeds by default. However, when it checks for completeness, it notices that the unobserved reaction ($n \rightarrow p$) is not explainable within the scope of its knowledge. The system postulates a quantum property, qp_0 , and determines the values for the particles in the observed and unobserved reactions. In this way, it postulates the correct electrical charge values for the elementary particles to be $p : 1, e : -1, n : 0, \mu : -1, \nu : 0, \pi : 1, \pi_0 : 0, \bar{p} : -1, \bar{e} : 1, \bar{\mu} : 1, \bar{\nu} : 0, \bar{\pi} : -1, \bar{n} : 0$, and $\pi_0 : 0$, which agree with the modern values.

We did not mention this earlier in the paper for historical reasons. Long before the discovery of many of the elementary particles, scientists knew that electrical charge was conserved in chemical reactions. The discovery of the charge conservation law in particle physics took place in a gradual process that went hand in hand with the discoveries of the earliest known particles such as the electron, proton, and neutron. Therefore the “simulation” cannot be said to have historical validity. However, it serves to show how a simple general concept of conservation, together with algebraic constraints, can lead to the discovery of a series of important laws in physical science.

4. Evaluation of BR-3

Having described the system’s simulation of the discoveries of the lepton, baryon, electron, and muon numbers, we can now discuss BR-3’s research goals, knowledge representation, predictive abilities, and methods for theory revision and search. This section concludes with an assessment of the generality of BR-3’s methods and the effects of data order and size on the system’s behavior.

Table 6. Data used in rediscovering the electrical charge property.

Observed reactions	Unobserved reactions
$p + p \rightarrow p + p + \pi_0$	$n \rightarrow p$
$p + \pi \rightarrow \pi + p$	
$\bar{\pi} + p \rightarrow \pi_0 + n$	
$\pi_0 \rightarrow \gamma + \gamma$	
$\bar{\pi} \rightarrow \mu + \bar{\nu}$	
$\pi \rightarrow \bar{\mu} + \nu$	
$\mu \rightarrow e + \nu + \bar{\nu}$	
$n \rightarrow p + e + \bar{\nu}$	

4.1. Research goals

BR-3 differs from STAHL (Zytkow & Simon, 1986), STAHLp (Rose & Langley, 1986), and REVOLVER (Rose & Langley, 1988) in the scope and generality of its goals. The main goal of STAHL and STAHLp is to infer valid reactions and components from their existing knowledge of chemical reactions and componential formulae; belief revision is secondary to their main function. REVOLVER, which finds the quark compositions of the elementary particles from known compositions and reactions, is closer in its research goals to the current program. In contrast to all three systems, BR-3's main goal is to attain consistency and completeness in its domain knowledge, and discovery is a side effect rather than the main concern.

Perhaps the most important difference between BR-3 and the earlier systems is that it tries to achieve completeness while maintaining consistency. By giving priority to consistency over completeness, it maintains its usefulness as a discovery system while it increases its knowledge. As is known by Gödel's theorem (Gödel, 1962), even in formal systems that include arithmetic, consistency and completeness cannot always be simultaneously satisfied. Hence consistency must have priority over completeness.

Accordingly, if the input data (e.g., some of the unobserved reactions) are incorrect and BR-3 fails to reach a consistent knowledge state concerning a quantum property after trying all cause and effect hypothesis combinations, it deletes the last set of cause and effect beliefs about that property. This brings the system to its prior incomplete but consistent knowledge state. In this way, BR-3 indirectly identifies inconsistencies in the input data.

Contradictions can appear in various forms, such as unbalanced reactions and compositions, multiple values of a property for the same object, and contradictions between general and particular knowledge. STAHL, STAHLp, and REVOLVER identify unbalanced reactions and componential formulae as contradictions. BR-3 also identifies unbalanced reactions as contradictions. In addition, it identifies algebraic equations with no solution and multiple quantum values as contradictions.

Similarly, there can be various forms of incompleteness in knowledge systems. BR-3 can identify two types: an unobserved reaction whose absence cannot be explained by existing knowledge, and the absence of quantum values for some known particles. The system resolves the former type by postulating new quantum properties, and the latter by assigning default values during consistency checks. Default values provide quasi-complete knowledge states, and are particularly useful in transitional states (i.e., until the real values are found). To facilitate belief revision, BR-3 does not add the default quantum values to its theoretical knowledge.

4.2. Knowledge representation

At the object level, BR-3's representation of particle reactions is similar to GLAUBER's (Langley et al., 1987) and STAHL's (Zytkow & Simon, 1986) representation of chemical reactions as ordered pairs of lists. However, the current system's knowledge is divided in three main categories as logical, formal, and theoretical knowledge. The division of descriptive and definitive knowledge into categories provides clarity and functional uniformity

in high-level reasoning.⁶ This division can be further detailed within each category; for example, theoretical knowledge can be further divided into empirical, hypothetical, and systemic knowledge. Moreover, BR-3's generated knowledge has exactly the same syntax as its original knowledge of the same category. This ensures that all the inference rules that are applicable to the original knowledge are also applicable to its generated knowledge.

As described earlier, BR-3's prescriptive knowledge is organized into three interacting system operators: CHECK-CONSISTENCY, CHECK-COMPLETENESS and REVISE-BELIEFS. The first of these operators checks if the particle reaction formulae are balanced by treating them as mathematical equations of quantum values. In these equations, the total quantum values must be balanced in order to satisfy the related conservation law. The concept of "balanced reaction" has been effectively used by STAHL and STAHLp in revising inconsistent chemical reaction formulae and componential models in their simulation of the chemical discoveries in the 18th century. However, the concept of "balanced reaction" of BR-3 is significantly different: As opposed to balancing the elements and/or compounds in chemical reactions, the system balances the total quantum values of the particles in a reaction. In this, it provides an important insight into a fundamental conceptual difference between chemistry and particle physics. In chemistry, the two main conservation laws involve energy and matter. The material conservation law dictates the conservation of the chemical elements in the reactions. In curious contrast, in particle physics, it is not the elementary particles themselves, but the elementary quantum properties that are conserved.

In the earlier systems mentioned above, there is no counterpart to BR-3's CHECK-COMPLETENESS operator, which postulates new quantum properties in order to resolve states of incompleteness. Discovery of a new property means that a large number of reactions previously thought theoretically possible are actually impossible. Initially, the system identifies only those reaction formulae that violate the charge conservation law as invalid and hence unobservable. However, its predictive ability improves as it rediscovers the other quantum properties, as explained in detail below.

4.3. The system's predictive abilities

BR-3 is an incremental system in the sense that it remembers its earlier discoveries and makes effective use of its accumulated knowledge in testing the validity of new reaction formulae by means of the *test-reaction* rule of its CHECK-CONSISTENCY operator. Thus, one can view it as an effective learning system as well as a discovery system, for its predictive performance improves with experience.

To illustrate the system's predictive abilities, consider the decay reactions of the type $n \rightarrow x + y$ and $n \rightarrow x + y + z$ for the neutron, where x , y and z are any of the 36 particles that BR-3 knows. The number of possible reaction formulae of the form $n \rightarrow x + y$ is $36^2/2! = 648$. These reactions are generated by a simple inference rule. When the system tests each of these reactions against the charge conservation law, it qualifies only 228 of them as "valid" reactions. When BR-3 rediscovers the lepton conservation law, the number of reactions valid against charge and lepton conservation laws is found to be 157. Finally, rediscovery of the baryon conservation law brings the figure down to 34 reactions. In other words, when the program "discovers" the charge, lepton, and baryon conservation laws,

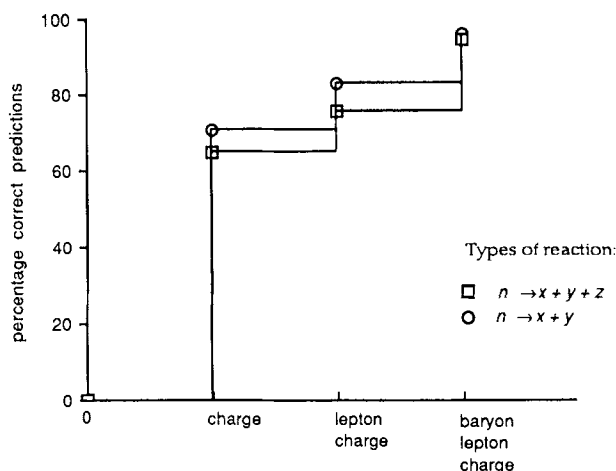


Figure 3. Increases in BR-3's predictive ability on the decay reactions of the neutron after the discovery of various quantum laws.

it makes correct predictions about the invalidity of 420, 491, and 614 such reactions after each stage. These correspond to 64.8%, 75.8% and 94.7% of all the reaction formulae of type $n \rightarrow x + y$. Figure 3 illustrates how the system's predictive ability improves by the discovery of these laws.

Similarly, of the possible $36^3/3! = 7776$ reaction formulae of the form $n \rightarrow x + y + z$, BR-3 qualifies 70.6%, 82.9%, and 96.0% as invalid after the rediscovery of the charge, lepton, and baryon number conservation laws, respectively. This shows that the program improves its ability to make correct theoretical predictions about the validity of a large number of reactions every time it "discovers" a new quantum property.

4.4. Belief revision and search methods

In its quest for consistent and complete knowledge states, BR-3 uses a combination of heuristic search and depth-first search with backtracking. The space of cause hypotheses is constrained by the rules of algebraic inequalities, while the space of effect hypotheses is constrained by simultaneous linear equations and the antiparticle rule. The system's search strategy effectively involves generating a search space that is as small as possible, and then applying an exhaustive search method.

The dimensions of the search space for cause hypotheses are dependent upon certain constraints. As mentioned earlier, particles can have a quantum value of 1, 0, or -1 for any of the quantum properties considered by the system. Hence, an unbalanced reaction formula with n particles opens a search space of 3^n cause hypotheses. The antiparticle rule, which dictates that the original particles can only have the values 1 or 0, and that antiparticles can only have the values 0 or -1 , reduces this search space to about 2^n . The

algebraic inequality constraint further reduces this to $2^n - n(n - 1)/2$ for unobserved collision reactions (e.g., $p + p \rightarrow \pi + \pi + \gamma$), and to $2^n - n$ for particle decays⁷ (e.g., $p \rightarrow \bar{e} + \gamma$). Table 7 presents data from a series of test runs showing the effects of these constraints in reducing the search space for cause hypotheses. The last column indicates the number of cause hypotheses generated in each search.

Similarly, in generating the effect hypotheses a search space of 3^n is reduced to about 2^n by the antiparticle rule, where n is the total number of particles in the unbalanced reaction formulae. The algebraic constraints sharply reduce this, such that the number of effect hypotheses generated (for any cause hypothesis) is no more than the number of unbalanced reactions for the set of allowable quantum values. The data from a series of test runs in Table 8 shows the effects of the antiparticle rule and the algebraic constraints in reducing the search space for effect hypotheses. The last column shows the number of effect hypotheses generated. These results provide convincing evidence for the importance of algebraic and domain knowledge in controlling search in scientific reasoning.

4.5. Generality of BR-3's methods

As we have described above, BR-3 is capable of rediscovering several quantum properties such as the lepton, baryon, spin, strangeness, electron, and muon number properties from sets of observed and unobserved reaction formulae. It can even rediscover the electrical charge property starting with no knowledge of quantum properties, except a general rule of conservation. The system's discovery strategy is quite general: Secure consistency first, and then try to achieve completeness. In a consistency check, first identify contradictions, and then revise beliefs. In a completeness check, postulate a new hypothesis, and then check for consistency. At this level, the approach is independent of any specific domain. In the physical sciences, the absence of certain effects or phenomena are explained by hypothesizing the existence of some objects or causes that prevent the occurrence of the former. For example, the absence of superconductivity in certain compounds is explained by the presence of magnetic impurities. In chemical research, when an expected reaction does not take place, the researcher often looks for some cause or causes of the effect.

Table 7. Reduction of the search space in generating cause hypotheses.

Unobserved reaction	Initial search space	Reduced by antiparticle rule to	Reduced by algebraic constraints to
$p \rightarrow \pi + \pi_0$	27	8	5
$\bar{\nu} + n \rightarrow p + \mu$	81	16	12
$p + p \rightarrow p + \pi$	81	16	2
$\bar{\pi} + p \rightarrow e + p$	81	16	6
$\nu_\mu + n \rightarrow p + e$	81	16	10

Table 8. Reduction of the search space in generating effect hypotheses.


Number of unbalanced reactions (i.e., linear equations)	Total number of particles with unknown quantum values	Number of different particles with unknown values	Initial search space	Reduced by antiparticle rule to	Reduced by algebraic constraints to
7	8	7	3^8	2^8	0
4	5	4	3^5	2^5	4
8	23	13	3^{23}	$2^{23}(8.3 \times 10^6)$	8
9	17	9	3^{17}	$2^{17}(1.3 \times 10^5)$	2
10	23	12	3^{23}	$2^{23}(8.3 \times 10^6)$	8

BR-3's representation of descriptive and definitive knowledge is also general and applicable to other fields in the physical sciences. Its categorization of such domain knowledge into formal and theoretical statements constrains the search activity of inference rules that are applicable only to knowledge in a certain category. For example, in reasoning about the elementary particle reactions (e.g., during consistency checks), BR-3 does not search for reactions in its logical or formal knowledge, but only in its theoretical knowledge.


We cannot say that BR-3's system operators, which represent its prescriptive knowledge, are general. Their application area is limited to particle reactions. However, some of the inference rules or procedures of these operators are general. For example, the *test-reaction* rule is applicable to simple chemical reactions and nuclear reactions (especially to fusion reactions). Similarly, BR-3's linear equation solver is a general mathematical inference procedure.

4.6. Effects of data order and size


As described in the previous pages, BR-3 has been tried with several different sets of valid and unobserved reactions with successful results. However, because there can be many valid combinations of quantum values, and because the system carries out a depth-first search through this space, its behavior is dependent on the order of the data given to it. For example, a change in the order of the unobserved reactions given in Table 3 causes BR-3 to discover the baryon number principle first, though the final results are the same. In most cases, the final results are unaffected by the order in which valid reactions introduced. However, in one experiment with the valid reactions



and the unobserved reaction



in addition to the valid and unobserved reactions given in Table 3, BR-3 combines the two properties (the baryon and lepton numbers) into a single property. Consequently, it finds the values of this combined property for the particles Λ , p , n , e , μ , and ν as 1, the values for their antiparticles to be -1 , and the values for the rest as 0. This was initially surprising, but it makes sense, because the nonzero lepton and baryon numbers of the elementary particles constitute mutually exclusive sets and do not cause any contradiction. However, the introduction of another unobserved reaction



which was valid by BR-3's latest knowledge (for it conserves the electrical charge and the new combined baryon-lepton property) causes the system to postulate a separate lepton property. In this case BR-3 retains the combined properties as the baryon number property, as it was consistent with all the given valid particle reactions.

An interesting order effect also occurs in the simulation of the discoveries of the electron and muon numbers. For certain orders of the reactions and particles in Table 4, BR-3 assigns nonzero quantum values to muon, tauon, their neutrinos, and their antiparticles. Similarly, in other cases, electron, tauon, their neutrinos, and their antiparticles are assigned nonzero quantum numbers.

Another surprising order effect takes place in the rediscovery of the electrical charge values. For some orders of the observed reactions given in Table 6, BR-3 finds the electrical charge values of the particles to be $p : 1$, $e : 0$, $n : 0$, $\mu : 0$, $\pi : 1$, $\pi_0 : 0$, $\bar{p} : -1$, $\nu : 1$, $\bar{n} : 0$, $\bar{\mu} : 0$, $\bar{\pi} : -1$, $\pi_0 : 0$, and $\bar{\nu} : -1$. These values are consistent with all the valid reactions tested by the system. This means that, as far as BR-3 is concerned, the electrical charge can be exchanged between the electron and its antineutrino, and between the positron and its neutrino, without the violation of any quantum law. Had this been physically true, it would have had very important implications in explaining a number of physical phenomena, perhaps including superconductivity. However, neutrinos are believed to be massless, and to our knowledge no charged particles with zero rest mass have yet been observed in isolation.

Although BR-3 can process input incrementally, batch processing can be more efficient if all data are available at the outset. Table 9 shows the amount of processing that resulted when we presented the system with data sets of three different sizes, both one at a time and in batch mode.⁸ The figures (which are averaged over 12 random orders of the data) suggests that, as the total amount of data increases, batch processing becomes more economical in terms of the number of beliefs added and revised. However, when the total amount of data is small, the effect of size is negligible.

5. Relation to other work

Our approach to machine discovery is closely related to earlier work in this area. This section provides a brief survey of recent research on theory revision, comparing each of the related systems to BR-3 in terms of their primary task domain, knowledge representation, inputs and outputs, and methods of discovery and theory revision.

Table 9. Effects of data size.

Tests	A1	A2	A3	B1	B2	B3
Number of reactions	7	14	21	7	14	21
Number of cause hypotheses generated	10	10	10	10	9	10
Number of effect hypotheses generated	2	8	15	1	4	17
Number of unbalanced reactions	4	17	27	4	17	30
Cause hypotheses revised	2	4	4	2	4	4
Effect hypotheses revised	1	6	13	0	2	5
Number of beliefs added	22	93	186	21	65	116
Number of beliefs revised	8	68	156	8	42	84

In tests A1–A3, all observed reactions were entered one by one, while in tests B1–B3 the same reactions were entered all at once. The tests were carried out over 12 random orders of data. (The figures were rounded to the nearest integer.) Comparison of the results of A1 with B1, A2 with B2, and A3 with B3 shows that as the total amount of data increases, batch processing becomes distinctly more economical.

5.1. Zytow and Simon's STAHL

As mentioned earlier, STAHL (Zytow & Simon, 1986) was the first system to apply belief revision methodology (Doyle, 1979) to modeling scientific discovery. The task domain of this system is the empirical discoveries of the chemists in the 17th century. STAHL's theories consist of chemical reactions and componential models of chemical substances, which are represented in list structures with labeled arguments. Its inputs are chemical reaction formulae, from which the componential models and some other reactions are deduced using several inference rules.

One of these rules infers, from reactions of the form $A \rightarrow B + C$ or $B + C \rightarrow A$, the componential model $A = B + C$. A second rule reduces reactions of the form $A + B + C \rightarrow A + D$ into $B + C \rightarrow D$ by deleting A from both sides, from which the componential model $D = A + C$ can be inferred by the first rule. A third rule substitutes the known components of substances that appear in a reaction formula. This formula is then simplified by the second rule, allowing the system to arrive at a simpler formula from which componential models are directly inferred.

STAHL recognizes two types of inconsistency: different componential models for the same substance, and reduced reactions with inputs but no output, or vice versa. When faced with such an inconsistency, the system reconsiders all the reactions that contain substances involved and rules out the inconsistent componential models. For example, from the reactions $A + B \rightarrow C + D$ and $C \rightarrow B + E$, STAHL infers $C = B + E$ and $A = D + E$. Suppose that later the componential model $A = F + E$ was found based on some additional data. This would contradict the previously inferred componential model of A . This causes the program to rule out the componential models of A in other inferences. Suppose that still later $D = F + B$ is found. By substituting the components of D and C in the first formula, the system finds $A = F + E$. This verifies the last componential model of A , which it accepts as valid.

Although its design borrows ideas from STAHL and STAHLp (described below), BR-3 differs from the former in its categorized knowledge representation, research goals and

methods of discovery. BR-3 has a much larger initial search space than that of STAHL, and rather than relying on domain-specific heuristics, it uses a depth-first search in a space constrained mainly by algebraic rules. Also, as mentioned earlier, BR-3 has a more detailed classification of conflicts in its knowledge states. Finally, STAHL's main goal is to find consistent models of its domain objects, whereas the current system's goal is to achieve a consistent and complete knowledge state. BR-3's search for completeness constitutes a clear advance over its predecessor.

5.2. Rose and Langley's STAHLp and REVOLVER

Another system that operates in the same task domain is STAHLp (Rose & Langley, 1986). This system uses a similar data structure to represent its domain knowledge and it can accept the same inputs, but there is a significant difference in its representation. STAHLp uses source tags in its reaction formulae and componential models to indicate the premises from which each substance in the reactions and models was derived. These source tags play an important role in the system's theory revision process. STAHLp uses the inference rules of its predecessor with an important modification on the reduction rule to eliminate certain types of errors in STAHL's reasoning about balancing chemical reactions.

Unlike its predecessor, when STAHLp infers an unbalanced null reaction such as $nil = A + B$ from its data, it tries to balance this reaction by a series of operations such as adding A and B to the left, or adding A to the left and deleting B from the right. Each of these changes yield new statements that constitute STAHLp's *effect hypotheses* concerning the null reaction. The system next determines what changes in the premises must be made so that they lead to the desired effect hypothesis. These possible modifications are called *cause hypotheses*. In its quest for cause hypotheses, the system can be viewed as using a form of abductive inference.

STAHLp uses an evaluation function in selecting the best effect hypothesis among the alternatives. Once the best hypothesis is chosen, the system generates a new set of beliefs and models, based on the revised premises. When a contradiction arises, the system traces the statements causing the contradiction and revises only those hypotheses that are responsible for the contradiction. In this way, STAHLp can even change its original beliefs as necessary. In resolving one such contradiction involving the oxidation of mercury, the system provides a partial explanation of the transition from the phlogiston theory to the oxygen theory.

There are some similarities between STAHLp and BR-3 in the way they process data. Both are incremental discovery systems, in that they can revise their domain theory in response to new data. Also, both systems use similar data structures as inputs. However, BR-3 has new capabilities over its predecessor, such as using algebraic rules to constrain search in its theory revision, and the ability to generate independent beliefs to explain the absence of certain domain events.

In contrast to BR-3's simpler approach to belief revision, STAHLp relies on an assumption-based (de Kleer, 1986) approach. The system utilizes source tags to indicate the origins of the inferred statements, such as reactions and componential models. This assumption-based method certainly provides an effective search control, and supports independence

and parallelism in truth maintenance. However, source tags complicate representation, especially when inferences are based on a large number of antecedents.

REVOLVER (Rose & Langley, 1988) is a variant of STAHLp that takes smaller steps through revision space and uses more evaluation criteria. The system was constructed to find models of objects consistent with initial beliefs about them, and to revise these models in response to contradictions with the new data. Its task domains are particle physics and chemistry. In particle physics, its task is to find the quark models of the elementary particles from their observed reactions and a set of initial models.

REVOLVER uses STAHLp's representation and inference rules with some domain related modifications in its search methods for consistent models. The system employs an additional domain-dependent rule in reducing the reaction formulae or quark models, which lets it move antiparticles from one side of the formula to the other, transformed into the original particle and vice versa. For example, in the neutron's decay reaction $n \rightarrow p + e + \bar{\nu}$, the antineutrino can be moved to the left side as the neutrino, resulting in $\nu + n \rightarrow p + e$ without the violation of any essential quantum conservation law.

The system employs three heuristics in its theory revision. One involves preferring premises that support a fewer number of beliefs. Another one, which is based on group-theoretic constraints (as opposed to BR-3's linear algebraic constraints), minimizes the complexity of models. The third heuristic increases the chances that each premise leading to an inconsistency will eventually be revised if the search begins to cycle among the same beliefs.

REVOLVER uses a hill-climbing search method in its quest for consistent domain theories. Hill-climbing methods rely on an evaluation function to guide search. An earlier version of BR-3 used a hill-climbing search, but this was abandoned because the local peaks did not always lead to the correct set of quantum values, and in some cases the system ran into loops. Rose (1989) describes the use of a "minimum revision criterion" to prevent a similar effect in REVOLVER's behavior in finding the quark models of elementary particles. BR-3's belief-revision method guarantees the correct solution provided that the input data are correct.

5.3. *Kulkarni and Simon's KEKADA*

Another system that can revise theories is KEKADA (Kulkarni & Simon, 1990). This system constitutes an important step in viewing scientific discovery as the product of a series of interrelated research activities such as formulating research goals, choosing among known strategies, proposing experiments, predicting the experiment results by forming expectations, generating hypotheses, and deleting them or modifying their confidences. KEKADA simulates the discovery of the urea cycle in biochemistry in the early 1930s. It has also been tested in several other domains.

The system's descriptive domain knowledge is represented in frame-like data structures. Experiments are represented in terms of several parameters, such as the purpose of the experiment, independent terms, methods, experimental conditions, measurements, and expectations. The program represents some of its hypotheses as a sequence of processes. Thus, its representation of descriptive domain knowledge is more structured than that of the earlier systems and BR-3.

KEKADA's prescriptive knowledge is represented as a set of condition-action rules that are grouped together in the form of several system operators. The program's research tasks are carried out by these operators in a sequential mode. In contrast, BR-3's operators are activated by the current state of the system's knowledge base.

The inputs to Kulkarni and Simon's system include domain knowledge about substances, methods, and processes, along with a set of initial hypotheses. The program also accepts as inputs the results of the experiments it has proposed. Its outputs are a set of research goals, proposed experiments, and generated hypotheses.

The system uses the two-space search model of Simon and Lea (1974) in its search activities. It proposes experiments on the basis of the current state of its "hypothesis space" and interprets the outcomes in the "instance space." An important feature of KEKADA is that it can formulate its own research goals by focusing its attention on surprises, that is, on facts that contradict a current hypothesis or expectation. As a result, its goals are more varied than those of the current system; KEKADA cannot identify an incomplete knowledge state as BR-3 does.

KEKADA's overall strategy can be regarded as best-first search, because it chooses between initial alternative research goals by using a set of qualitative criteria, and when it proposes experiments after choosing its research goal and strategy, the results of the experiment can produce a new research goal with higher priority.

The system employs several strategies to constrain search in its problem spaces. One such strategy is magnification, which means studying a research phenomenon in detail by controlled changes in its variables. Another strategy is a form of divide and conquer, which involves studying a phenomenon in its subprocesses. A third strategy is to study whether a surprising phenomenon belongs to a general class, and a fourth involves eliminating the irrelevant variables in the current research goal. Finally, a fifth strategy involves looking for similarities between the subject of the current research goal and other related phenomena. These strategies have no counterparts in the earlier systems and BR-3.

On the other hand, KEKADA is an active empirical discovery system in the sense that its discoveries are mostly the results of the experiments it proposes, rather than theoretical analysis as we observe in STAHL, STAHLp, REVOLVER, and BR-3. Therefore, it relies heavily on experimentation in its theory revision task and discoveries.

5.4. O'Rorke, Morris and Schulenburg's AbE

O'Rorke, Morris and Schulenburg (1990) propose abduction as a general method for theory formation and revision. They argue that an observation contradicting a prediction of a given theory can be explained in terms of the basic principles of the theory, claiming that this process can lead to new hypotheses by abductive inference. They also describe AbE, a system that uses this strategy for theory formation and revision.

The primary task domain of AbE is the 18th century chemistry, in which the system simulates the stages of the revolutionary transition from the phlogiston theory to the oxygen theory. These theories were developed to "explain" a series of chemical processes such as combustion, oxidation and calcination.

Recent discovery systems are increasingly using methods that allow more structured and fine grained representations of domain knowledge. AbE represents chemical processes in

qualitative process schemas (Forbus, 1984). Each schema represents a qualitative process description in terms of the substances involved in the process, the preconditions for the process to take place, the process conditions, and a set of relationships that indicate the qualitative changes in the process.

In this way the assumptions of a theory are also captured in the representation. For example, the phlogiston theorists' belief that all combustibles are compounds containing the element phlogiston is captured in the "process conditions" slot. The basic laws of AbE's domain theory are represented in simple and complex predicate statements. One such law is: "Combustion decreases the amount of phlogiston in charcoal." Compared to AbE's structured representation of objects and processes, BR-3's representation is simplistic.

The inputs to AbE are its domain knowledge, which consists of process descriptions and the laws of the domain theory (including the general laws of qualitative physics), and which it represents as schemas and predicate statements. The system's output is a set of explanations and a revised theory.

AbE's basic method of discovery is abductive inference on the qualitative descriptions of domain processes. The system employs a best-first search method to construct explanation trees to account for an observation. When an observation contradicts its domain theory, AbE deletes those hypotheses that are responsible for the contradiction, and attempts to construct the explanation tree for the observation. When it fails to do so, it uses its basic knowledge of domain theory (physics) to complete the explanation by introducing new hypotheses. In this way, it arrives at a partially modified theory. As these modifications take place in an incremental fashion, one domain theory (e.g., phlogiston theory) gradually transforms into another (e.g., oxygen theory) that is consistent with new observations as well as the old ones.

There are basic differences between AbE's methods of discovery and those of BR-3. The latter not only revises hypotheses that contradict observations, but also postulates new and independent hypotheses to explain a hitherto unexplained phenomenon. In BR-3's task domain, the reasoning applied in the discovery of the baryon number property involves a slightly more complex inference than abduction. This can be explained as follows: Before the discovery of the baryon and lepton conservation laws, physicists knew that electrical charge was conserved in particle reactions. Hence, they could explain the absence of certain reactions by the violation of the charge conservation law. When they realized that some other reactions did not occur even though they conserved electrical charge, they made an abstraction from one quantum property to the existence of a class of such properties, and then applied abductive inference. In symbolic terms, in which C stands for the charge conservation law, and u and v represent unobserved reactions,

$$\begin{array}{l} \text{explains}(C, u) \\ \sim \text{explains}(C, v) \end{array}$$

therefore,

$$\begin{array}{ll} \exists Q, C \in Q & \text{(abstraction)} \\ \exists B, B \in Q, B \neq C, & \\ \text{explains}(B, v). & \text{(abduction)} \end{array}$$

This analysis shows that BR-3's rediscovery of the first quantum property after the electrical charge property (which it knows at the outset) cannot be explained only by an abductive inference but involves a combination of abstraction and abduction. The new hypotheses that AbE generates by abductive inference are derived from its basic knowledge of domain theory, whereas an abstraction expands a theory beyond a system's basic knowledge. This suggests that BR-3's theory development processes also involve some activities beyond theory revision. On the other hand, AbE carries abduction over more structured and detailed representations than BR-3, suggesting that the former gives a more detailed account of theory revision.

5.5. Rajamoney's COAST

Rajamoney (1990) describes COAST, a system that employs a knowledge-intensive theory revision method. This system has been tested on problems from physics that involve a variety of processes such as evaporation and heat transfer.

COAST's theoretical knowledge is represented in qualitative process schemas, as in the AbE system. Domain objects are described using similar data structures. Physical systems (or "scenarios") are represented in frames with layout and behavior components. The former specifies the objects of the system, their arrangement, and the initial physical relationships between them. The latter represents the qualitative changes in the physical state of the objects.

COAST's inputs are the observed changes in the physical system under study, for which it tries to construct explanations. When the program fails to explain an observation, it revises its domain theory. The first step in this process involves identifying three different types of failures: incomplete explanations, contradictions, and mutually inconsistent multiple explanations. In general terms, COAST's identification of failures has similarities to STAHLp's and BR-3's identification of conflicts. However, as described earlier, BR-3 has a more varied notion of conflicts.

The theory revision methods of COAST focus on making partial changes in the description of the physical system behavior (or scenario), such that the new scenario is consistent with the current set of observations. The program generates a number of alternative theories during its revision process, whereas BR-3 generates only one complete theory at a time, extending and/or revising it when conflicts arise upon the introduction of new data.

Rajamoney (1990) defines theory development as a complex and continuous process that involves three major activities: theory formation, theory revision, and paradigm shifts. Like the AbE system, his program combines theory revision and paradigm shifts, whereas BR-3 integrates theory revision with theory formation. This constitutes a new capability on the part of BR-3, but unlike AbE and COAST, the current system cannot model paradigm shifts. Another difference is that, in contrast to AbE and COAST, which are knowledge-intensive systems, BR-3 behaves more like a data-driven system initially, but as its theoretical domain knowledge increases through its discoveries over the input data, it gradually acquires the characteristics of a theory-driven system.

COAST designs experiments to test and evaluate its alternative theories to select the best one. Unlike BR-3's reliance on algebraic constraints, this system uses general heuristic methods to constrain search in theory revision. The program generates representative partial models (or "exemplars") of its domain knowledge, and tests its proposed theories against

them for validity. The remaining theories are examined for their ratings in terms of their structural and explanatory simplicity and predictive power. The system's use of partial models to guide search constitutes a clear improvement on BR-3 and the earlier systems. Also, its ability to design experiments to constrain search for alternative theories is a feature that BR-3 lacks.

However, COAST has been tested only on limited physical systems. The number of its domain individuals and their possible physical states seem to be small compared with the number of the elementary particles, the possible quantum states, and the numerous reactions that BR-3 has to consider. COAST's current methods would not be capable of dealing with the type of problems that BR-3 faces, for the former cannot explain the existence or absence of objects, events, and processes in a physical system, but only the effects of such processes. On the other hand, COAST's theory revision methods are more general than those of BR-3.

5.6. Karp's GENSIM/HYPGENE

Karp (1990) views hypothesis formation as a design problem, and describes GENSIM/HYPGENE, a composite system that models theory revision in a specialized area of biochemistry. The first component, GENSIM, represents theories and experiments as process schemas, and domain objects in a taxonomic hierarchy of frames. The system maintains its domain knowledge in three different knowledge bases: a "class knowledge base," a "simulation knowledge base," and a "process knowledge base." The first represents GENSIM's formal knowledge and object descriptions, the second contains the descriptions of experiments, and the last includes the system's knowledge about qualitative processes of its domain. Its object and class descriptions are much more complex than BR-3's representation. GENSIM predicts the outcomes of experiments using a rule-based interpreter that applies processes to the objects (e.g., chemical compounds) present in the experiments.

The other component, HYPGENE, uses a set of design operators and GENSIM's knowledge base to formulate hypotheses and to revise its theories. As with BR-3's operators, the inputs of HYPGENE's operators are arbitrary predicate calculus assertions and definitions. HYPGENE's inputs include a description of the initial conditions of an experiment, GENSIM's prediction of the outcome of the experiment, a description of the error in this prediction, and the domain theory used in the prediction. The system's output is a set of hypotheses that are used in revising GENSIM's theory. Theory revision takes effect when GENSIM's predictions contradict the outcomes of an experiment. In revising hypotheses, HYPGENE uses a combination of best-first search and assumption-based belief revision in contrast to BR-3's algebraic and depth-first search methods.

HYPGENE's theory revision basically consists of making changes in the process descriptions. In its theory revision the system uses backward chaining from the experiment results to the process conditions, whereas BR-3 uses both forward chaining (in its theory revision) and backward chaining (in its theory formation). BR-3's activities involve both theory formation and theory revision, whereas GENSIM/HYPGENE is essentially a knowledge-intensive theory revision system. However, the latter has more detailed knowledge representation and complex theory revision methods, and operates over a larger amount of knowledge.

Also, the program's inclusion of a planner constitutes an improvement over the earlier systems and BR-3, which enables it to reason about a series of interlinked processes with complex propositions.

6. Conclusions

In recent years, computational modeling of scientific discovery and theory revision has become an increasingly active field, and various researchers have described methods for eliminating contradictions between the predictions of a theory and observations. In this paper we examined the more general problem of identifying and resolving conflicts (incomplete and inconsistent knowledge states) in an area of particle physics.

We described BR-3, an incremental theory revision system that rediscovers a series of quantum properties and the associated quantum laws. The system's domain knowledge is organized into three basic categories—logical, formal, and theoretical knowledge. Its knowledge organization highlights the significant differences between the formal and theoretical concepts of consistency and completeness.

BR-3's descriptive and definitive knowledge is represented in labeled predicate statements, in which the labels indicate the category of the statement. Recent systems such as AbE (O'Rorke, Morris & Schulenburg, 1990), COAST (Rajamoney, 1990), and GENSIM/HYPGENE (Karp, 1990) have more structured representations for certain data types, such as events and processes, an improvement that enables these systems to provide more detailed accounts of theory revision.

The current system identifies two types of incomplete knowledge states: the unexplained absence of an event, and the absence of qualitative or quantitative values of a property for some domain objects. The program resolves the first type of conflict by postulating new domain properties, which in turn are used in new domain hypotheses about the values of these properties for the domain objects. The second type of conflict is resolved by default value assignments, which are revised when they contradict new data.

BR-3 identifies three types of contradictions: an unbalanced particle reaction, multiple values of a property for a domain object, and algebraic contradictions (i.e., a linear equation with no solution). The system employs a combination of heuristic search and depth-first search in its theory revision. When BR-3 achieves a consistent knowledge state, it checks for the completeness of its theoretical knowledge concerning domain events, and generates new domain properties and relationships as necessary. New data can be given to the system when it is in a consistent knowledge state. BR-3's predictive ability concerning the validity of particle reactions increases as the system develops its theory by formulating new quantum laws upon examining experimental data. The program successfully reproduces the discovery of the lepton, baryon, electron, and muon number properties and the corresponding quantum values for the elementary particles.

Although BR-3 is an incremental system, the order and size of input data affects its behavior. When changes are made in data order, the program generally finds consistent and complete discrimination sets of quantum properties and values, though in some cases these do not correspond to the correct physical values, and in others the program produces surprising results, which nevertheless make sense. The tests also indicate significant effects

of data size on BR-3's behavior in its search for consistent solutions. As the total amount of data increases, batch processing becomes more economical than incremental data in terms of the number of beliefs added and revised.

BR-3 differs from previous discovery systems in that it integrates theory revision with theory formation. In its quest for completeness the system increases its domain knowledge, and in its quest for consistency it revises its theories. This approach introduces the identification of incompleteness as an important issue, and opens a new perspective for computational modeling of scientific discovery. Incomplete theories have been the motivation of a number of important discoveries in the history of science. In particle physics, the advancement of the quark theory was motivated by the lack of explanations for the absence of elementary particles with certain combinations of quantum properties. Similarly, in astronomy, perturbations in the orbits of the outer planets led the scientists to postulate the existence of new planets such as Neptune and Pluto before they had been observed.

BR-3 could be further developed to provide a more comprehensive model of theory revision in its task domain. Although the current version can rediscover the other quantum properties such as spin and strangeness, its belief revision methods cannot find the correct strangeness values without high computational cost, and the spin values which must be calculated by group axioms rather than the algebraic constraints used by the system. The program's current search methods could be supplemented with other methods for a complete simulation of the discovery of the strangeness values. Also, its current representation could be further developed in order to accommodate more structured data types, so that the system can simulate paradigm shifts such as the transition from the earlier theory of particle physics to the quark theory. Finally, the system's categorized predicate logic representation could be integrated with schemas to represent events and processes more effectively. In the longer term, we hope to incorporate these ideas into an intelligent aid for research in particle physics.

Acknowledgments

The author wishes to express his gratitude to Pat Langley, whose detailed comments on the earlier drafts were extremely valuable in developing the project as well as improving the organization, style, and understandability of the paper. Thanks are also due to the anonymous reviewers for their comments and suggestions.

Notes

1. For more detailed definitions of consistency and completeness in causal theoretical systems, see Dean and Boddy (1987).
2. Naturally, the statistical effects had been taken into account in these observations.
3. The system has been implemented in Prolog, and the runs described in this paper were carried out on an IBM AT. A program listing of an earlier version of BR-3 can be found in Kocabas (1989).
4. The system retains the remaining cause and effect hypotheses, because some new valid reactions to be added in the future may force it to revise its beliefs.

5. The initial search space for the set of actual strangeness values is 7^n . With the set of the seven quantum values, BR-3 generates 306 cause hypotheses for an unbalanced reaction with three particles (e.g., $p \rightarrow \pi + \pi_0$). Compare this with the figures in Table 7.
6. For a detailed account of the categorization of descriptive and definitive knowledge, see Kocabas (1989).
7. The negative components $n(n-1)/2$ and n correspond to the equalities in the combinations.
8. The same set of unobserved reactions were used in all the tests.

References

- Davies, P.C.W. (1985). *The forces of nature* (2nd ed). Cambridge: Cambridge University Press.
- Dean, T., & Boddy, M. (1987). Incremental causal reasoning. *Proceedings of the Tenth International Conference on Artificial Intelligence* (pp. 196–201). Milan, Italy: Morgan Kaufmann.
- de Kleer, J.S. (1986). An assumption-based TMS. *Artificial Intelligence*, 28, 127–162.
- Doyle, J. (1979). Choices without backtracking. *Proceedings of the Fourth National Conference on Artificial Intelligence*, (pp. 79–85). Austin, TX: Morgan Kaufmann.
- Forbus, K.D. (1984). Qualitative process theory. *Artificial Intelligence*, 24, 85–168.
- Gödel, K. (1962). *On formally undecidable propositions of Principia Mathematica and related systems*. Tr. by B. Meltzer. New York: Basic Books.
- Karp, P.D. (1990). Hypothesis formation as design. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Kocabas, S. (1989). *Functional categorization of knowledge: Applications in modeling scientific research and discovery*. Doctoral Dissertation, Department of Electronic and Electrical Engineering, King's College London, University of London.
- Kulkarni, D., & Simon, H.A. (1990). Experimentation in machine discovery. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Langley, P., Simon, H.A., Bradshaw G.L., & Zytkow, J.M. (1987). *Scientific discovery: Computational explorations of the creative processes*. Cambridge, MA: The MIT Press.
- Omnes, R. (1970). *Introduction to particle physics*. Tr. by G. Barton. London, Wiley Interscience.
- O'Rourke, P., Morris, S., & Schulenburg, D. (1990). Theory formation by abstraction. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Rajamoney, S.A. (1990). A computational approach to theory revision. In J. Shrager & P. Langley (Eds.), *Computational models of scientific discovery and theory formation*. San Mateo, CA: Morgan Kaufmann.
- Rose, D. (1989). Using domain knowledge to aid scientific theory revision. *Proceedings of the Sixth International Workshop on Machine Learning* (pp. 272–277). Ithaca, NY: Morgan Kaufmann.
- Rose, D., & Langley, P. (1986). Chemical discovery as belief revision. *Machine Learning*, 1, 423–452.
- Rose, D., & Langley, P. (1988). A hill-climbing approach to machine discovery. *Proceedings of the Fifth International Conference on Machine Learning* (pp. 367–373). Ann Arbor, MI: Morgan Kaufmann.
- Simon, H.A., & Lea, G. (1974). Problem solving and rule induction: A unified view. In L. Gregg (Ed.), *Knowledge and cognition*. Lawrence Erlbaum, Hillsdale, N.J.
- Trefil, J.S. (1980). *From atoms to quarks*. London: The Athlone Press.
- Zytkow, J.M., & Simon, H.A. (1986). A theory of historical discovery: The construction of componential models. *Machine Learning*, 1, 107–137.

Table of Contents: Volume 6 (1991)

Number 1

Editorial	<i>Jaime Carbonell</i>	5
Rigel: An Inductive Learning System	<i>Roberto Gemello, Franco Mana,</i> <i>..... and Lorenza Saitta</i>	7
Instance-Based Learning Algorithms	<i>David W. Aha, Dennis Kibler</i> <i>..... and Marc K. Albert</i>	37
Information-Based Evaluation Criterion for Classifiers' Performance		67
.....	<i>Igor Kononenko and Ivan Bratko</i>	
Technical Note: A Distance-Based Attribute Selection Measure for Decision Tree Induction	<i>R. López de Mántaras</i>	81
Technical Note: Improved Estimates for the Accuracy of Small Disjuncts ...		93
.....	<i>J.R. Quinlan</i>	
Book Review: <i>Exemplar-Based Knowledge Acquisition</i>	<i>Yoram Reich</i>	99
A Reply to Reich's Book Review of <i>Exemplar-Based Knowledge Acquisition</i>	<i>Ray Bareiss</i>	105

Number 2

Symbolic and Neural Learning Algorithms: An Experimental Comparison ...		111
.....	<i>Jude W. Shavlik, Raymond J. Mooney, and Geoffrey G. Towell</i>	
A "PAC" Algorithm for Making Feature Maps	<i>Philip Laird</i> <i>..... and Evan Gamble</i>	145
Letter Recognition Using Holland-Style Adaptive Classifiers		161
.....	<i>Peter W. Frey and David J. Slate</i>	
Technical Note: A Critical Look at Experimental Evaluations of EBL.....		183
.....	<i>Alberto Segre, Charles Elkan, and Alexander Russell</i>	
Book Review: <i>Learning Search Control Knowledge: An Explanation-Based Approach</i>	<i>Roland J. Zito-Wolf</i>	197
A Reply to Zito-Wolf's Book Review of <i>Learning Search Control Knowledge: An Explanation-Based Approach</i>	<i>Steven Minton</i>	205

Number 3

Complexity Results on Learning by Neural Nets	<i>Jyh-Han Lin</i> <i>..... and Jeffrey Scott Vitter</i>	211
The Use of Background Knowledge in Decision Tree Induction		231
.....	<i>Marlon Núñez</i>	
A Nearest Hyperrectangle Learning Method	<i>Steven Salzberg</i>	251
Conflict Resolution as Discovery in Particle Physics	<i>Sakir Kocabas</i>	277