An Integrated Framework for Extended Discovery in Particle Physics

Sakir Kocabas[1] and **Pat Langley**[2]

[1] Space Engineering Department, ITU
80626 Maslak, Istanbul, Turkey
`ukoca@itu.edu.tr`
[2] Institute for the Study of Learning and Expertise
2164 Staunton Court, Palo Alto, CA 94306 USA
`langley@isle.org`

Abstract. In this paper we describe BR-4, a computational model of scientific discovery in particle physics. The system incorporates operators for determining quantum values of known particles, formulating new quantum properties, positing new particles, and predicting reactions among particles. BR-4 carries out heuristic search guided by constraints that its theory be consistent and complete with respect to observed reactions. We show that this control scheme is sufficient to model, with some manual intervention, an extended period in the history of particle physics, including the discovery of the neutrino and the postulation of baryon, lepton, and electron numbers. In closing, we compare BR-4 to other discovery systems and suggest directions for future research.

1 Introduction and Motivation

Computational research on scientific discovery falls into two broad categories. The first, typified by the work of Langley, Simon, Bradshaw, and Żytkow (1987), focuses on modeling the processes responsible for discoveries from the history of science. The second approach, exemplified by the work of Valdés-Pérez (1995) and Mitchell, Sleeman, Duffy, Ingram, and Young (1997), uses computational methods to discover new scientific knowledge. These two approaches share many ideas, and both have made valuable contributions to discovery science, but they have distinct goals and criteria for evaluation.

In this paper we describe results within the first, historical, approach to scientific discovery. Like Nordhausen and Langley (1993), we believe that there has been important progress in this area, but that most previous models have focused on one aspect of the scientific process to the exclusion of others. Like them, our goal has been to extend earlier models to account for a broader range of scientific enquiry during an extended period in science. We have not tried to model the processes in detail or to craft a precise theory of human cognition, but rather to provide an abstract but unified account of major activities and their order of occurrence. This has required us to develop an integrated framework that combines discovery mechanisms in a coherent way.

Nordhausen and Langley's work addressed empirical discovery in physics and chemistry, which led their IDS system to integrate mechanisms for forming taxonomies, finding qualitative laws, and detecting numeric relations. We have focused instead on the more theory-laden domain of particle physics, so that our BR-4 system integrates processes for constructing and revising structural theories, detecting and formulating problems, generating new theoretical terms, and predicting new events.

In the next section we present our integrated framework for scientific discovery and its implementation in BR-4. After this, we consider four examples from the history of particle physics, showing for each how the system simulates discoveries made during the period. These case studies include the postulation of the neutrino, the prediction of various reactions, the proposal of baryon and lepton numbers, and the discovery of electron and muon numbers. In closing, we review related computational work on discovery and consider directions for extending our framework.

2 A Framework for Discovery in Particle Physics

In this section we present a computational framework for explaining the processes that support scientific discovery in particle physics, starting with an analysis of the task. We then turn to the representational assumptions that underlie our framework, the heuristics that drive the discovery process, and the search algorithm that our model, BR-4, uses to explore the space of theories.

2.1 The Discovery Task

Particle physics studies the nature of elementary particles – the building blocks of matter – and interactions among these entities. The basic phenomena in this field take the form of reactions, similar in many ways to those found in chemistry. For instance, one such 'observed' reaction (typically inferred from tracks in cloud chambers) is $p + p \rightarrow p + n + \pi$, where the symbols p, n, and π represent the proton, neutron, and pion particles, respectively.

As in chemistry, physicists require that reactions among elementary particles obey certain conservation laws. One of the main tasks in particle physics concerns the assignment of values for *quantum properties* such that observed reactions conserve those properties. For example, the above reaction conserves the quantum property of electric charge, provided we assign the accepted charges 1 to p, 0 to n, and 1 to π. Other assignments are possible for this reaction, but they would not work for other particles and their observed interactions.

The notion of conservation also explains why some particle reactions are never observed. For example, proton decay, as in the reaction $p \rightarrow \bar{e} + \gamma$, has never been seen, despite its conservation of electric charge. However, one can explain its absence by positing that it fails to conserve another quantum property, the baryon number. Thus, another central task in particle physics involves explaining missing reactions by postulating new quantum properties.

Other activities include the inference of new particles, either on theoretical or empirical grounds, and the prediction of reactions that involve these particles in ways that satisfy known conservation laws. Testing such predictions leads into the realm of experimental particle physics, which we will not address here. But the above pursuits cover a wide range of the behaviors that occur in this scientific field.

2.2 Discovery Operators and Internal Representation

The above analysis of the discovery task suggests that four basic operators play a central role in particle physics. First, for a given set of particles, quantum numbers, and observed reactions, we must be able to determine a set of quantum values that satisfy conservation for those reactions. Second, we must be able to posit new quantum properties that account for the absence of unobserved reactions. Third, we require an operator that posits new particles and their role in known reactions. Finally, we need some mechanism for predicting reactions that have not yet been observed, but that follow from the current theory. We have incorporated these operators into the BR-4 model, where they support the process of theory formation and revision.

Operators of this sort must alter some internal representation that contains hypotheses about the particles, properties, and reactions that exist, and that also indicates specific quantum values for each pair of property and particle. This representation can take many forms, but, following Valdés-Pérez, Żytkow, and Simon (1993), we can view it as two related matrices. One matrix lists particles against quantum properties, with each matrix entry specifying the value for a specific particle on a specific property. The other matrix lists particles against reactions, with an entry containing the total number of times the particle occurs in the reaction. Our operator for determining quantum values alters entries in the first matrix, whereas the other operators each extend one or both matrices along one of their dimensions. In our examples, we will use the matrix notation to specify the properties of particles but not the reactions in which they occur, since the latter matrix would be largely empty.

2.3 Heuristics for Consistency and Completeness

Naturally, simply formulating the problem in this manner does not solve it. Given P particles and Q quantum properties with V values each, there are $V^{Q \cdot P}$ possible assignments of values to particle-property pairs. For small values of P, Q, and V, one could search this space exhaustively, but recall that one must also consider different numbers for these parameters themselves (i.e., different size matrices). In general, constrained search is preferable to blind search, and we have incorporated a number of heuristics into the BR-4 system that focus its attention in useful directions.

First, the system considers simpler theories first, starting with one that contains only directly 'observable' particles, quantum properties for which there exists separate evidence (such as electric charge), and a few observed reactions.

Second, BR-4 alters this theory only when it encounters evidence of some deficiency, and then it considers only those operators that promise to repair the problem. Finally, the model uses constraints on the problem domain, such as conservation, to limit the search within the space of repairs.

More specifically, BR-4's approach to discovery in particle physics relies on the notions of *consistency* and *completeness* to constrain the reasoning process. For example, the operator for determining quantum values applies only when the system detects that an observed reaction is inconsistent with some conservation law. In this case, it carries out a depth-first search through the space of values, continuing until it encounters a value combination that violates conservation, in which case it backtracks. When this process is complete, the resulting quantum values are guaranteed to be consistent with all reactions observed so far. To keep the process tractable, BR-4 considers only the values 0, 1, and -1 during its search.[1]

In some cases, the above revision process cannot eliminate the inconsistency, either because no combination of property values leads to conservation across all observed reactions or because the quantum values are determined experimentally (as for the spin number). This condition leads BR-4 to revise the unbalanced reaction by adding a 'hidden' particle in either the input or output, positing that it actually takes part in the reaction but for some reason is not directly observable. The system then computes the property values that would balance the reaction and associates them with the new particle.

The incompleteness constraint leads to complementary behavior. When BR-4 finds that its current theory fails to rule out a reaction that does not occur, it introduces a new quantum property that is *not* conserved by this reaction but that is conserved by those it has observed. Determining the values of this property requires search, first for the values of particles in the missing reaction (constrained to satisfy an inequality), and then an embedded search for the values of other particles (constrained to satisfy equalities corresponding to observed reactions). As before, if the system arrives at a partial combination of values that rules out an observed reaction or fails to eliminate the unobserved one, it considers alternative paths until it finds an acceptable set. In both searches, BR-4 considers smaller absolute values before turning to larger ones.

We can extend the notion of incompleteness to include theories that do not explicitly specify all reactions that follow from them, as occurs when one postulates a new particle. In this situation, BR-4 systematically generates all possible reactions of the new particle involving one, two, or three other known particles. Some of these reactions take the form of decays, whereas others involve collisions among particles. For each such tentative reaction R, the system predicts that R will occur if it conserves all known properties, and predicts that the reaction will not occur otherwise.

---

[1] Physicists assign to the spin property not only integers like 0 and 1, but also values like $\frac{1}{2}$ and $\frac{3}{2}$. BR-4 also considers these values for this property and, like physicists, calculates the spin number using group theory.

Table 1. The quantum values for elementary particles known (a) in 1930, prior to experimental detection of the neutron, and (b) after postulation of the neutrino.

	Particle	mass	charge	spin
(a)	γ	0.00	0	1
	e	0.51	-1	$\frac{1}{2}$
	p	938.26	1	$\frac{1}{2}$
	$\bar{e}$	0.51	1	$\frac{1}{2}$
(b)	n	939.55	0	$\frac{1}{2}$
	ν	0.00	0	$\frac{1}{2}$

3 Illustrative Examples from Particle Physics

In this section we consider four examples of discovery from the history of particle physics, involving the neutrino, baryon and lepton numbers, and electron and muon numbers. In each case, we recount the main historical events, and then examine BR-4's behavior when presented with similar observations. Our historical treatment is based upon a number of sources on particle physics, including Griffiths (1987), Ne'eman and Kirsh (1986), Omnes (1970), Pais (1986), and Trefil (1980).

3.1 Discovery of the Neutrino

Until the early 1930's, scientists had observed only a few elementary particles, shown in Table 1 (a) along with their mass and their values on three conserved quantum properties – energy, charge, and spin. The known reactions were also limited to a small set: $p+p \rightarrow p+p$, $e+\bar{e} \rightarrow \gamma$, and $\gamma \rightarrow e+\bar{e}$. This situation changed after Chadwick's experimental detection of the neutron in 1932, which also clarified another outstanding issue (Giancoli, 1995).

Much earlier, physicists had observed beta decay, a process in which an element emits an electron and is transformed into another element with a higher atomic number. This transformation appeared to violate conservation of both energy and spin, leading Bohr to suggest that these properties are truly not conserved within the nucleus. However, in 1930, Pauli proposed another explanation – that beta decay also emitted another particle that was difficult to detect.

Chadwick's experiments also revealed neutron decay, $n \rightarrow p + e$, which occurs in about 800 seconds on free neutrons. Like beta decay, this reaction appeared to violate energy and spin conservation, but in simplified form. Again, Pauli's account avoided this problem by postulating a new particle, also generated during the decay reaction, that would balance out the missing energy and spin. In 1934, Fermi formalized this proposal for the *neutrino*, which he posited as having zero rest mass, no electrical charge, and a spin of one half.

Table 2. Particle reactions that were (a) observed and (b) not observed in experiments after the introduction of the particles in Table 1 (b).

(a) Observed reactions	(b) Unobserved reactions
$p + p \rightarrow p + p$	$p \rightarrow \bar{e} + \gamma$
$e + \bar{e} \rightarrow \gamma$	$p \rightarrow \bar{e} + e + \bar{e}$
$\gamma \rightarrow e + \bar{e}$	$p \rightarrow \bar{e} + \gamma + \gamma$
$\gamma + p \rightarrow e + \bar{e} + p$	
$n \rightarrow p + e + \nu$	

Given the four reactions above and the quantum numbers in Table 1 (a), BR-4 responds in a similar manner. The system immediately detects an inconsistency concerning the spin values for neutron decay and attempts to correct it. (The current program does not address the issue of energy conservation.) BR-4 cannot modify the spin counts of the particles in the reaction, as these values are marked as having been established by observation. This leaves revision of the unbalanced reaction as the only solution.

One such revision adds an extra particle to the output side of the reaction, giving $n \rightarrow p + e + \nu$. Using the conservation laws as constraints, the system computes the mass, charge, and spin of the new particle, ν, as 0.0, 0, and $\frac{1}{2}$, respectively. Another possible revision would have added a new particle with the opposite spin to the input side of the reaction. However, we believe physicists favored the former solution because they were thinking in terms of a decay process, so we have biased BR-4 in this direction as well.

Our treatment of this episode ignores many details, including the role that conservation of energy, in addition to spin, played in driving proposals for the neutrino. But the general line of reasoning, that a new particle with certain quantum values was needed to preserve conservation, appears historically accurate, and BR-4's heuristics arrive at the same description for this particle as did Fermi and his colleagues.

3.2 Proposing the Baryon Number

The inference of the neutrino left physicists with six elementary particles, having the properties and values shown in Table 1 (a) and (b). Scientists realized that the existence of these particles, combined with the existing conservation laws, implied a variety of reactions. Subsequent experiments revealed evidence for some of these reactions, shown in Table 2 (a), but not for some others, shown in Table 2 (b). For some reason, the three predicted decays of the proton did not occur in nature; to remedy this problem, physicists proposed a new quantum property, known now as the *baryon number*.

Given the six particles in Table 1, our model follows a similar line of reasoning. BR-4 realizes that its current theory is incomplete, so it predicts all decay and collision reactions involving these entities (up to length three) that conserve

Table 3. The quantum values of particles known in 1953, after discovery of baryon and lepton numbers.

Particle	mass	charge	spin	baryon	lepton
γ	0.00	0	1	0	0
e	0.51	-1	$\frac{1}{2}$	0	1
p	938.26	1	$\frac{1}{2}$	1	0
n	939.55	0	$\frac{1}{2}$	1	0
$\bar{e}$	0.51	1	$\frac{1}{2}$	0	-1
ν	0.00	0	$\frac{1}{2}$	0	1
μ	105.60	-1	$\frac{1}{2}$	0	-1
$\bar{\mu}$	105.60	1	$\frac{1}{2}$	0	1
π	139.60	1	0	0	0
$\bar{\pi}$	139.60	-1	0	0	0
π_0	135.00	0	0	0	0

charge and spin, giving the seven reactions[2] in Table 2. These correspond to proposed experiments with the particles, or at least to suggestions for what to look for in such experiments. When informed that the reactions in Table 2 (a) occur but those in (b) do not, BR-4 infers that its theory is incomplete in a deeper sense and proposes a new property to correct the situation.

To determine the values of this new property, BR-4 selects one of the missing reactions, say $p \to \bar{e} + \gamma$, and turns it into a set of inequalities, each based on a different combination of values for the particles involved. In this case, it generates the four inequalities $1 \neq 0+0$, $1 \neq 1+1$, $0 \neq 1+0$, and $0 \neq 0+1$. The system then selects one of these value sets, say the first, $\{p = 1,\ \bar{e} = 0,\ \gamma = 0\}$, and inserts them into one of the observed reactions, say $n \to p + e + \nu$, this time treating it as an equality.

In this case, BR-4 obtains the expression $n = 1 + 0 + \nu$, which leaves the property values for n and ν unspecified. Two consistent value sets are possible for this pair, $\{n = 1,\ \nu = 0\}$ and $\{n = 0,\ \nu = -1\}$. BR-4 selects the first and uses it to check the observed reactions, introducing values for the remaining unassigned particles as necessary. Detection of an unbalanced reaction that violates conservation of the new property causes backtracking to one of the alternative value sets. If the search exhausts all such sets produced from the observed reactions, the system backtracks further and considers other value sets generated from the unobserved reactions.

---

[2] BR-4 also generates two other reactions, besides $n \to p+e+\nu$, that involve neutrinos: $\nu+p \to n+\bar{e}$ and $\nu+n \to p+e$. However, physicists showed little concern when they did not immediately detect these reactions, presumably because theory predicted that neutrinos interacted very rarely. Thus, we told the system to ignore them at this stage of our simulation.

Table 4. Some particle reactions that were (a) observed and (b) not observed in experiments after the discovery of mesons.

(a) Observed reactions	(b) Unobserved reactions
$\mu \rightarrow e + \nu + \nu$	$\mu \rightarrow e + \gamma$
$\pi \rightarrow \bar{e} + \nu$	$\mu \rightarrow e + \bar{e} + e$
$\pi \rightarrow \bar{\mu} + \nu$	$\pi_0 \rightarrow e + \bar{\mu}$
$\pi \rightarrow \pi_0 + \bar{e} + \nu$	$\pi_0 \rightarrow \mu + \bar{e}$
$\pi_0 \rightarrow \bar{e} + e$	
$\pi_0 \rightarrow \nu + \nu$	
$\pi_0 \rightarrow \gamma + \gamma$	

Given the experimental results in Table 2, BR-4 arrives at the value zero for all particles except the proton and neutron, to which it assigns the value one, as shown in the first six rows of Table 3. These settings correspond to those obtained by physicists for the baryon number, which successfully explain the absence of the reactions in Table 2 (b), since they fail to conserve this property. As new particles become known, BR-4 assigns baryon values to them as well, using the same search mechanism.

3.3 Mesons and the Lepton Number

In 1935, Yukawa proposed the existence of additional particles in the nucleus, with a mass of about 100 MeV. The reasoning behind Yukawa's proposal, which we have not attempted to model, involved energy calculations on atomic nuclei. Later, in the 1940s, observations of cosmic rays revealed five such particles: the muon (μ) and anti-muon ($\bar{\mu}$), the pion (π) and anti-pion ($\bar{\pi}$), and the pion zero (π_0). These suggested a variety of reactions, some that were observed by scientists and others that were not.

Konopinski and Mahmoud (1953) attempted to explain the mismatch between theory and data, focusing on the five detected reactions $\mu \rightarrow e + \nu + \nu$, $\mu + \nu \rightarrow e + \nu$, $p + \mu \rightarrow n + \nu$, $\nu + n \rightarrow p + \mu$, and $\nu + n \rightarrow p + e$ and on the single unobserved reaction $\mu \rightarrow e + \gamma$. In order to explain the absence of this decay, they proposed a new quantum property, the *lepton* number, with nonzero values for the muon, the electron, the neutrino, and their antiparticles.[3] However, Konopinski and Mahmoud assumed that the muon in the reactions was an antiparticle, which led them to assign it the lepton value −1. With the introduction of the lepton number, physicists had produced a theory, equivalent to that depicted in Table 3, that appeared consistent and complete. Many scientists had reservations about Konopinski and Mahmoud's theory, but it was the best account available at the time.

---

[3] Pais (1986) claims that he suggested the lepton number, including its name, earlier, in 1947, based on an analogy with the baryon number for heavier particles.

Table 5. Particle reactions that were (a) observed and (b) not observed in experiments after distinguishing between electron neutrinos (ν_e) and muon neutrinos (ν_μ).

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

BR-4 responds to the introduction of mesons in a similar manner. Given the five new particles, it predicts a variety of reactions, including four muon decays, five pion decays, and ten reactions that involve the pion-zero. Table 4 shows a sample of these predictions, some (a) that were observed and others (b) that were not. These differ somewhat from the ones addressed by Konopinski and Mahmoud, who presumably did not mention the observed decays that had been known since 1947 (Griffiths, 1987, p. 19, p. 25) and may have ignored some unobserved ones because the values for the lepton number forbid them.

Upon finding that the predicted reaction $\mu \rightarrow e + \gamma$ has not been observed, BR-4 attempts to introduce a new property with values that rule out this interaction. However, the system cannot find a consistent set of values for this property if, as usual, it considers only zero and positive values. For BR-4 to follow Konopinski and Mahmoud's reasoning, we must tell it (as the physicists concluded) that μ is an anti-particle, which lets the system consider negative quantum values. Table 3 shows the values generated by the system when given this assistance; they correspond to those inferred by Konopinski and Mahmoud, with the exception that μ and $\bar{\mu}$ are reversed.

3.4 Electron and Muon Numbers

In the year 1953, another important development took place. Additional experiments revealed indirect evidence for the predicted reaction $\nu + p \rightarrow n + \bar{e}$, which obeyed all known conservation laws and thus was required for the theory to be complete. Yet this reaction occurred when the neutrino (ν) had been generated through beta decay ($n \rightarrow p + e + \nu$), but not when produced through muon decay ($\bar{\mu} \rightarrow \bar{e} + \nu + \nu$).

To resolve this dilemma, scientists postulated that the two reactions actually generated two distinct types of neutrinos, calling the former an electron neutrino (ν_e) and the latter a muon neutrino (ν_μ). This distinction (and the analogous one for anti-neutrinos) introduced two additional rows in the table of particles. However, it also produced the unobserved reactions shown in Table 5 (b), which physicists sought to explain by introducing yet another property and which they named the *electron number*.

Our model cannot directly explain the historical distinction into two classes of neutrinos, but we believe it constitutes a variation on the heuristic for postulating new particles that originally led to inference of the neutrino. The situation also bears some similarity to the distinction inferred by Mendel 1865 to explain the different offspring of apparently identical peas, which Shen and Simon (1989) have modeled using a related mechanism. Langley et al. (1987) have used a similar technique to explain distinctions that occurred in the history of chemistry.

Once this difference has been introduced manually, BR-4 realizes that its current theory is incomplete, in that it cannot explain the unobserved reactions. Postulating a new property, it searches the space of values using the same process as it used for the baryon and lepton numbers. The resulting values agree with those proposed by physicists for the electron number, and they are sufficient to rule out the two unobserved muon reactions shown in Table 5 (b). Physicists also postulated yet another quantum property, called the *muon number*, on grounds of symmetry between electrons and muons. However, lacking any heuristics of this sort, BR-4 cannot reproduce this step in the human scientists' reasoning.

3.5 BR-4 as a Historical Model

We have implemented BR-4's operators and heuristics in PROLOG, and we have verified the system's ability to reproduce the historical discoveries reported earlier. In each case, we gave the system a set of particles, a set of known quantum properties, the hypothesized values for those properties, and a set of observed and unobserved reactions; in response, BR-4 generated the revised values, new particles and properties, and predicted reactions we have described. These formed a partial basis for the next inputs to the system, giving historical continuity to the model's behavior.

The resulting chain of reasoning carries BR-4 through more than two decades of major discoveries in particle physics. Moreover, the system relies on mechanisms that are consistent with our knowledge about the nature of human cognition. In particular, it carries out a limited heuristic search through a space of models that is guided both by knowledge about the domain and by observations. Moreover, this process occurs in an incremental fashion, with the system revising previous models as new phenomena become available and with new results becoming background knowledge for the next round of discovery.

As we have noted, BR-4 does not explain all of the major events in particle physics, even during the period we have attempted to simulate. In a number of cases, we had to intervene manually at selected points beyond the insertion of information about the outcomes of predictions. These steps included telling the system to ignore some unobserved reactions involving neutrinos, to assume that the muon is an antiparticle with nonpositive quantum numbers, and introducing the distinction between electron and muon neutrinos. Also, the system explains the historical sequence of events at a quite abstract level that ignores many details which occupied particle physicists' time and energy.

Thus, although BR-4 has let us model an extended period in the history of science, it remains an incomplete account. Each situation that required intervention suggests the need for additional mechanisms that should let its successor better match the historical record. These should include heuristics for ignoring predictions that are too difficult to observe, for considering wider ranges of quantum values, and for discriminating particles that appear the same but behave differently. Each such extension seem as general, at least in principal, as the existing operators and heuristics on which BR-4 relies.

4 Related Work on Computational Scientific Discovery

Our computational model of discovery draws many of its ideas from earlier work in this area. BR-4 is a direct descendant of Żytkow and Simon's (1986) STAHL, which modeled a variety of qualitative discoveries in the history of chemistry. The detection of inconsistencies in reactions played an important role in this system, with one of its responses being the introduction of new elements like phlogiston, which served much the same role in early chemistry as the neutrino did in particle physics.

Rose and Langley (1986) described STAHLp, a rational reconstruction of the earlier system that showed all of its discoveries could be explained in terms of inconsistencies and their resolution. In addition, they used STAHLp and REVOLVER (Rose & Langley, 1988), a similar system, to model a number of other reaction-oriented discoveries from the history of science, including some from particle physics. Moreover, their approach showed that dependency-directed reasoning simplified the theory-revision process, letting their systems handle problems with a search-control scheme that relied on incremental hill climbing rather than more systematic search.

Kocabas' (1991) BR-3 system extended this framework to include the detection of incomplete theories and the postulation of new properties to explain the absence of reactions. He applied this idea to the history of particle physics, using it to explain the origin of several quantum numbers and the particular values assigned to them. In related work, Kocabas (1992) adapted similar methods to discovery in the area of superconductivity. BR-3 was the immediate precursor of BR-4, with the former differing mainly in that it lacked the ability to postulate new particles and to predict new reactions.

Valdés-Pérez (1994) has described an alternative approach to discovery in particle physics, which he implemented in his PAULI system. This scheme uses a variation on linear programming to search the space of property values, subject to constraints that reflect observed and unobserved reactions. In addition, Fischer and Żytkow (1992) have reported on GELL-MANN, a system designed to explain the formation of the quark theory, which also carries out a form of constraint-satisfaction search to determine parameter values. Both systems have generated interesting models that differ from those found by human scientists, but these results, combined with their more powerful and nonincremental search methods, make them less plausible as historical accounts than the STAHL, STAHLp, BR-3, and BR-4 systems.

Despite their differences, each of these systems fits nicely within the framework proposed by Valdés-Pérez, Simon, and Żytkow (1993), which characterizes the discovery process in terms of operations on two related matrices. The various programs differ in their operators for altering the matrices, with BR-3 and BR-4 adding steps for introducing a property, predicting reactions, and positing a particle. PAULI and GELL-MANN also explore a matrix space but invoke different search regimens for selecting operators.

Other research on scientific theory revision, such as Rajamoney's (1990) work on theory-guided experiment generation in physics, seems less closely related. However, Kulkarni and Simon's (1990) KEKADA integrates theory revision, experiment design, and problem formulation to model Krebs' discovery of the urea cycle. The system includes heuristics for making predictions, redirecting attention when they are violated, and designing experiments to determine the underlying cause. The KEKADA work comes the closest to our own in spirit, in that both involve modeling an extended period in the history of science, rather than isolated events. However, Kulkarni and Simon's model operates at a finer granularity and better matches the historical details than does BR-4.

5 Directions for Future Research

Although BR-4 provides an abstract account for some important developments in particle physics, there remains considerable room for improvement. One problem is that the model's coverage of the historical process remains far from continuous. A more complete account would incorporate knowledge about the difficulty of detecting some reactions to explain why scientists chose to ignore some unobserved interactions (e.g., those involving neutrinos) while focusing their attention on others (e.g., those concerning proton decay). We should futher reduce reliance on human intervention by adding an operator like the one described by Shen and Simon (1989) that introduces a distinction between particles (e.g., electron and muon neutrinos) based on behavioral differences observed over time. Heuristics for proposing new particles and quantum properties on theoretical grounds would further strengthen the model.

We also hope to extend the system to introduce componential models, which describe particles at one level as combinations of more primitive ones. Langley et al.'s (1987) DALTON took some steps along these lines to explain relations between chemical molecules and elements, but we can incorporate similar methods into BR-4 to explain the origins of the quark theory and its alternatives. The basic task involves explaining why elementary particles with some quantum properties exist while others do not. BR-4's constraints of consistency and completeness seem well suited for this problem, which involves postulating new component particles (quarks), then searching the space of quantum values and their compositions that satisfy certain constraints (such as symmetry) for known particles and that violate these constraints for nonexistent ones.

Finally, although BR-4 implicitly models social aspects of the discovery process by addressing extended periods to which multiple scientists contributed, it accomplishes this at a very abstract level. A more detailed account of social interactions would include explicit communication among particle physicists, with theorists passing on predictions to experimentalists, who in turn report their observations to theorists. An extended model would also support competition in the development of theories to explain new findings and in finding evidence for predicted events. The history of particle physics is rich in examples of such interactions, and we believe that appropriate revisions to BR-4 would let us model at least some of them. To this end, we should assign different facets of the system's domain knowledge to different agents, which would communicate through a common representation; in addition, separate agents would explore different branches when the search process suggests alternative solutions.

6 Concluding Remarks

In this paper we presented BR-4, an integrated model of historical scientific discovery. We examined the system's behavior on four major problems that arose in particle physics, showing that it can replicate important steps in the historical development of this field, some of which were considered major discoveries when first introduced. In particular, BR-4 proposes the existence of the neutrino to avoid violating conservation of spin, it introduces baryon and lepton numbers to explain the absence of reactions involving proton decay, and it postulates electron numbers to rule out unobserved neutrino reactions. The system also finds appropriate quantum values for each particle and predicts the reactions implied by a set of particles and properties.

The BR-4 model achieves these results using simple processes that appear to have considerable generality. The system employs four basic operators for determining the values of a quantum property, creating new properties, positing new particles, and predicting reactions among known particles. Moreover, it uses consistency and completeness constraints to selectively apply these operators, and it incorporates a depth-first control scheme to carry out search when necessary. These activities operate in a continual loop, with incorrect predictions leading to revised models, which then become the starting point for new discoveries. Together, they let BR-4 explain, with occasional aid from its developers, an extended period in the history of particle physics. The simplicity and generality of these mechanisms suggest that we can explain other aspects of scientific discovery in similar terms, and we hope to test that hypothesis in future work.

Acknowledgements

Portions of this paper appeared in the *Working Notes of the AAAI Spring Symposium on Systematic Methods of Scientific Discovery*. The research described herein was supported by Grant No. N00014-94-1-0505 from the Office of Naval Research and by the Nippon Telegraph and Telephone Corporation.

References

Fischer, P., & Żytkow, J. M. (1992). Incremental generation and exploration of hidden structure. *Proceedings of the ML92 Workshop on Machine Discovery* (pp. 103–110). Aberdeen, Scotland.

Giancoli, D. C. (1995). *Physics: Principles with applications* (4th ed.). Englewood Cliffs, NJ: Prentice Hall.

Griffiths, D. (1987). *Introduction to elementary particles*. New York: John Wiley.

Kocabas, S. (1991). Conflict resolution as discovery in particle physics. *Machine Learning*, *6*, 277–309.

Kocabas, S. (1992). Elements of scientific research: Modeling discoveries in oxide superconductivity. *Proceedings of the ML92 Workshop on Machine Discovery* (pp. 63–70).

Konopinski, E. J., & Mahmoud, H. M. (1953). The universal Fermi interaction. *Physical Review*, *92*, 1045–1049.

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., & Żytkow, J. M. (1987). *Scientific discovery: Computational explorations of the creative processes*. Cambridge, MA: MIT Press.

Mitchell, F., Sleeman, D., Duffy, J. A., Ingram, M. D., & Young, R. W. (1997) Optical basicity of metallurgical slags: A new computer-based system for data visualisation and analysis. *Ironmaking and Steelmaking*, *24*, 306–320.

Ne'eman, Y., & Kirsh, Y. (1986). *The particle hunters*. Cambridge: Cambridge University Press.

Nordhausen, B., & Langley, P. (1993). An integrated framework for empirical discovery. *Machine Learning*, *12*, 17–47.

Omnes, R. (1970). *Introduction to particle physics* (Tr. by G. Barton). Wiley Interscience.

Pais, A. (1986). *Inward bound*. Oxford: Clarendon Press.

Rajamoney, S. (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., & Langley, P. (1986). Chemical discovery as belief revision. *Machine Learning*, *1*, 423–451.

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.

Shen, W. M., & Simon, H. A. (1989). Rule creation and rule learning through environmental exploration. *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence* (pp. 675–680). Detroit, MI: Morgan Kaufmann.

Trefil, J. S. (1980). *From atoms to quarks*. London: The Athlone Press.

Valdés-Pérez, R. E. (1994). Discovery of conserved properties in particle physics: A comparison of two models. *Machine Learning*, *17*, 47–67.

Valdés-Pérez, R. E. (1995). Machine discovery in chemistry: New results. *Artificial Intelligence*, *74*, 191–201.

Valdés-Pérez, R. E., Żytkow, J. M., & Simon, H. A. (1993). Scientific model building as search in matrix spaces. *Proceedings of the Eleventh National Conference on Artificial Intelligence* (pp. 472–478). Washington, DC: AAAI Press.