Invention and Artificial Intelligence

Gary Bradshaw

Psychology Department, P.O. Box 6161,
Mississippi State University, MS 39762, USA
glb2@ra.msstate.edu

Abstract. Invention, like scientific discovery, sometimes occurs through a heuristic search process where an inventor seeks a successful invention by searching through a space of inventions. For complex inventions, such as the airplane or model rockets, the process of invention can be expedited by an appropriate strategy of invention. Two case studies will be used to illustrate these general principles: the invention of the airplane (1799-1909) and the invention of a model rocket by a group of high school students in rural West Virginia in the late 1950's. Especially during the invention of the airplane, inventors were forced to make scientific discoveries to complete the invention. Then we consider the enterprise of artificial intelligence and argue that general principles of invention may be applied to expedite the development of AI systems.

1 Heuristic Search and Invention

Humans live in a world that has been shaped by invention: the clothing we wear, the food we eat, our houses, our transportation, our entertainment – all depend on a vast aggregation of technology that has been developed over the millennia. Some invented artifacts, including stone knives and hammers, even predate homo sapiens. Given the importance of invention in the contemporary world, it is worth some effort to understand how new inventions are developed.

Even a superficial review of the history of technology and invention shows that many different paths can lead to an invention. Basalla [1] and Petroski [2] have discussed the similarities between biological evolution and technological invention. In reviewing several case studies, Basalla provides strong evidence that some inventions arise when inventors produce random mutations of existing inventions, and society determines which inventions are "fit" for reproduction. One example is the paper clip, which appeared about the same time as steel wire and wire-bending jigs became available. Figure 1 illustrates different clip shapes from three different American patents. The familiar double-oval Gem clip was never patented in the U.S., but other patents were filed that described wire loops of various shapes in an effort to create a clip that was easy to slip over a set of papers, did not tear the papers as it was used, and held the collection tightly.

A system of invention based on random mutation (by inventors) and natural selection (by society) does not appear to require any intelligent activity on the part of inventors themselves: how difficult can it be to bend wires, after all?

A. Hoffmann, H. Motoda, and T. Scheffer (Eds.): DS 2005, LNAI 3735, pp. 1–13, 2005.
© Springer-Verlag Berlin Heidelberg 2005

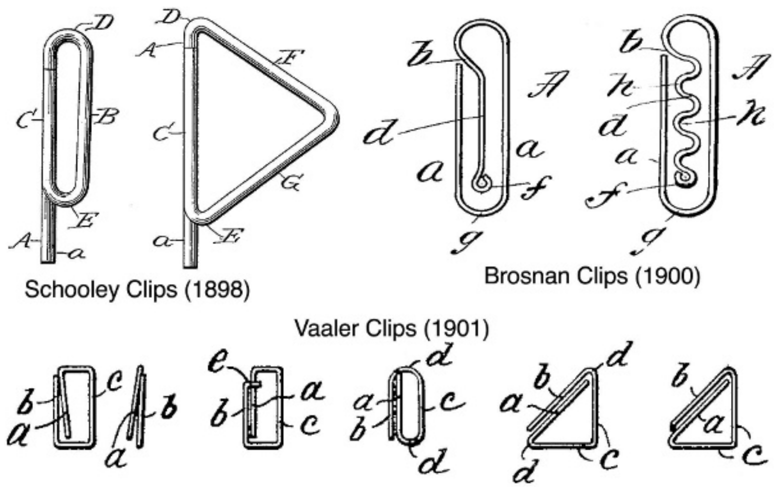

Fig. 1. Various forms of paper clips patented in the U.S. around the turn of the century

But in spite of the variability of all of the patented paper clips, they all shared some important features in common. First, all clips are relatively flat, having a structure that is mostly two-dimensional rather than 3-dimensional. Next, clips tend to loop over themselves in a way that allows them to pinch together a stack of papers. If inventors were simply producing random bends in wire, they would most commonly produce non-planar bends that created a 3-dimensional structure and would most commonly produce forms that did not have the necessary loops to hold together a paper stack. In spite of the diversity of paper clip forms, it does seem clear that the various alternatives were not produced by blind or random mutation, but rather by strategic alteration, perhaps akin to genetic algorithms in use today.

It should also be evident that the "inventive play" described by Basalla, where inventors produce different forms through manipulation of an existing artifact, is not sufficient to account for inventions where a number of parts are necessary to the performance of the whole. Consider, for example, the television. Although it might be possible for an infinite monkey team to wire together tubes, resistors, capacitors, and transformers to produce a television set, it seems unlikely this would have happened so quickly after the invention of the electron tube. Similarly, random mutation does not seem to be sufficient to produce inventions like the airplane or the telephone in a reasonable amount of time. These more complicated inventions appear to call for a more sophisticated method of invention.

1.1 Invention via Heuristic Search

Several researchers, following in the footsteps of Newell and Simon [3], appear to have independently developed the idea that inventions could be realized through heuristic search. Weber and Perkins [4] adapted contemporary problem solving

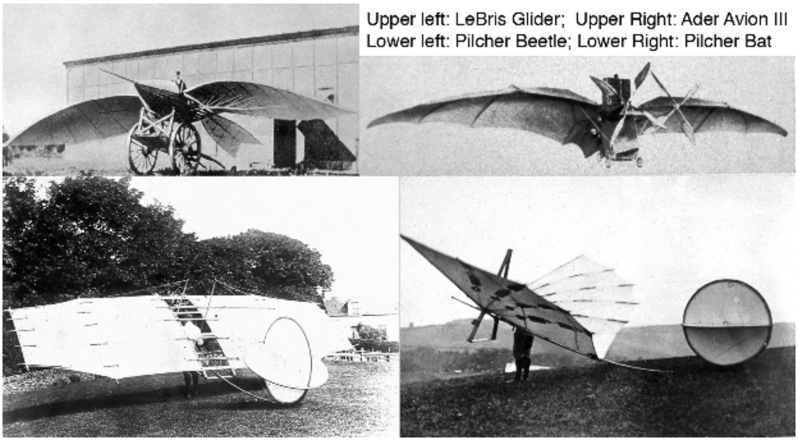

Fig. 2. Designs created by *analogy to nature*. These inventors deliberately copied the structure of flying creatures in their designs. Although birds were a popular source of analogy, other flying creatures like bats and beetles were used as well.

to account for the process of invention. Following Simon's ([5], [6]) model of scientific discovery as heuristic search, Weber and Perkins described a set of heuristics to produce new inventions by performing a goal-directed search through a space of possible inventions based on a series of working-forward heuristics that enable a new set of inventions to be developed from existing ones. For example, the *join* heuristic creates a new invention by combining separate inventions together. The awl and the scraper can be combined using the join heuristic to produce a pointed knife. The knife, in turn, can be combined with a screwdriver, a pair of scissors, a corkscrew, and a saw into the contemporary Swiss army knife, a lightweight and versatile tool.

The join heuristic is a weak method, but is still far more powerful than random recombinations. Weber and Perkins restrict the join heuristic to combine functionally related objects. The heuristic would not be invoked to join a can opener with a computer monitor. Although such an implement would have greater functionality than the original inventions, the purposes of can openers and monitors have little to do with one another, so there is no reason to construct this awkward marriage.

In accounting for the invention of the airplane, Bradshaw & Lienert [7] and Bradshaw [8] also adopted a problem-solving perspective. Four design heuristics were identified from historical records of numerous different attempts to create a workable airplane from 1799 to 1909. Two of these heuristics, *analogy to nature* and *copycat* were the source of full designs, while two other heuristics, more is better and make small changes, were used to revise an existing design. Figure 2 illustrates different examples of the use of analogy, where various flying creatures (songbird, bat, and beetle) were used to inspire an airplane design. Figure 3 illustrates copycat, where a design produced by one inventor is adopted by another.

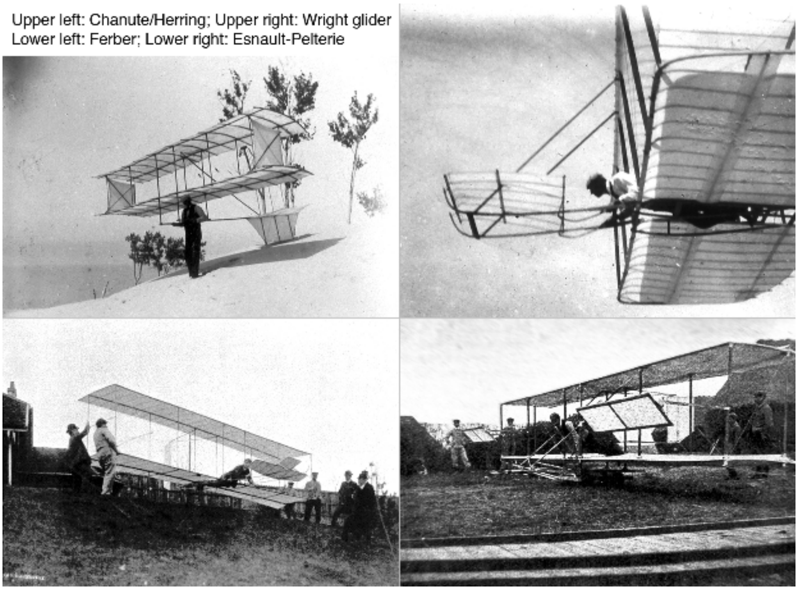

Fig. 3. Airplanes designed by *copycat*. The Wrights copied the sturdy biplane design introduced by Octave Chanute and Augustus Herring, including the Pratt system of trussing the wings for strength. Ferber and Esnault-Pelterie both copied the Wright design. The failure of Esnault-Pelteries "perfect copy" of the Wright glider convinced Europeans that the Wrights were "bluffing" in their claims to have built airplanes.

Once an airplane had been designed, two additional heuristics could be used to revise the design. The first such heuristic is called more is better. Inventors added additional wings, propellers, tail surfaces, and other components to see if their design could be improved. Phillips, an English inventor, produced a design that had at least 196 different airfoils in four racks, not realizing that the turbulence produced by the forward wings would spoil the airflow over rearward wings. The second design-revision heuristic is known as *make small changes*. This covers a multitude of minor modifications made to an airplane, usually for ad hoc reasons.

Strategies for Searching Through the Design Space. These four heuristics alone can generate billions of different airplane designs given the parameters shown in Table 1. We will refer to the set of possible designs as the *design space* of the invention. Effective solutions (airplanes that are airworthy) are rare in this design space, so inventors need to find efficient means of searching through such a large space for rare solutions.

Most inventors relied upon a simple *design and test* strategy: They would design and build a craft, then take it out to the field to test it. Some craft were launched from a hill, others were launched from a rail system, while others attempted to fly from grassy meadows. Figure 6 illustrates the best performance

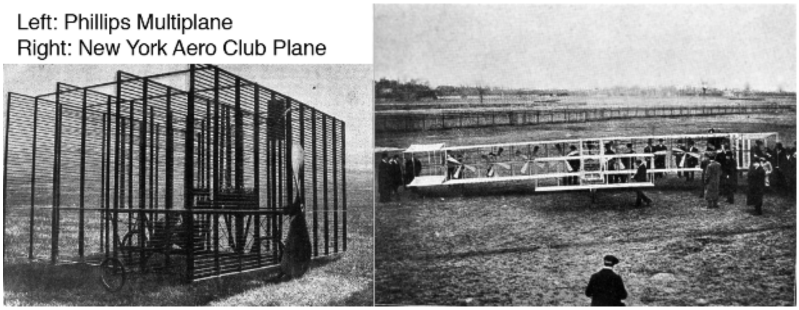

Fig. 4. Airplanes modified through the *more is better* heuristic. The Phillips Multiplane had approximately 196 different airfoils, while the New York Aero Club plane featured 8 propellers.

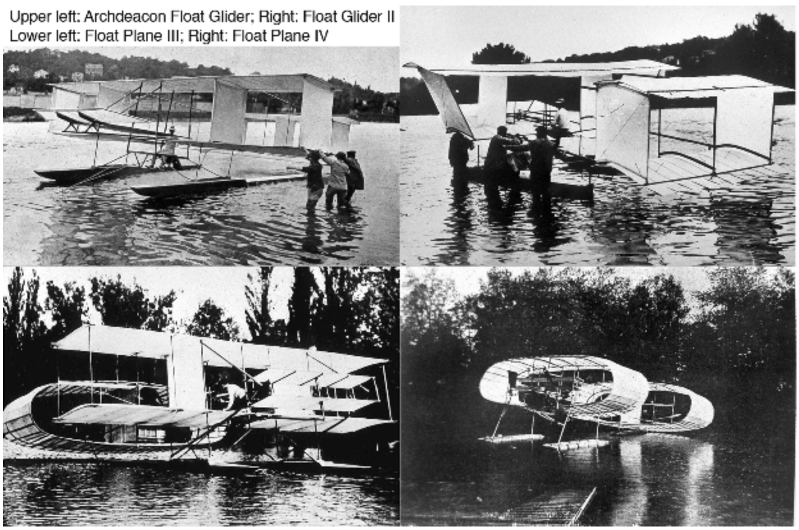

Fig. 5. Archdeacons airplane was modified several times *using make small changes*. The original design (a crude copy of the Wright craft) was changed by reducing the length of the lower wing, which led to the side curtains being placed at an angle. Then the rear wing was replaced with an ellipse (which suffers from particularly bad aerodynamic characteristics) and finally both wings were replaced by an ellipse.

of a number of craft between 1799 and 1909. The dotted line is the regression line for non-Wright craft in the period 1799 until December, 1905. In 1906 the Wrights patent first appeared, and so after that time inventors had access to information about an airworthy craft. The 1799-1905 regression line actually has a slightly negative slope, indicating that more capable craft were built near the beginning of this period than could be constructed at the end of the era: There was no substantial improvement in performance of airplanes over a 110-year

Table 1. Design Features of Para-Planes

Design Parameter	Possible Values
Number of wings.	1–196
Wing position	1–3 (monoplane, tandem, etc.)
Placement	stacked, tandem, staggered
Lateral arrangement	anhedral, flat, dihedral
Camber of wings	1–12, 1–6, etc.
Wingspan	6'–104'
Wing Chord	3'–10'
Shape of wings	bird-like, rectangular, bat-like, insect-like
Tail placement	forward (canard), rear, mid
Lateral control	none, wing warping, ailerons
Number of Propellers	0–8

span. The Wrights were a clear exception to this trend: their initial gliders were among the best-performing aircraft from the beginning, and they made steady progress. The 1903 craft (which was their first powered design) represents the transition between gliders and powered craft, and fits nicely on both regression functions.

The data shown in Figure 6 represents a considerable mystery: How were the Wrights able to develop an effective initial glider and sustain progress, while

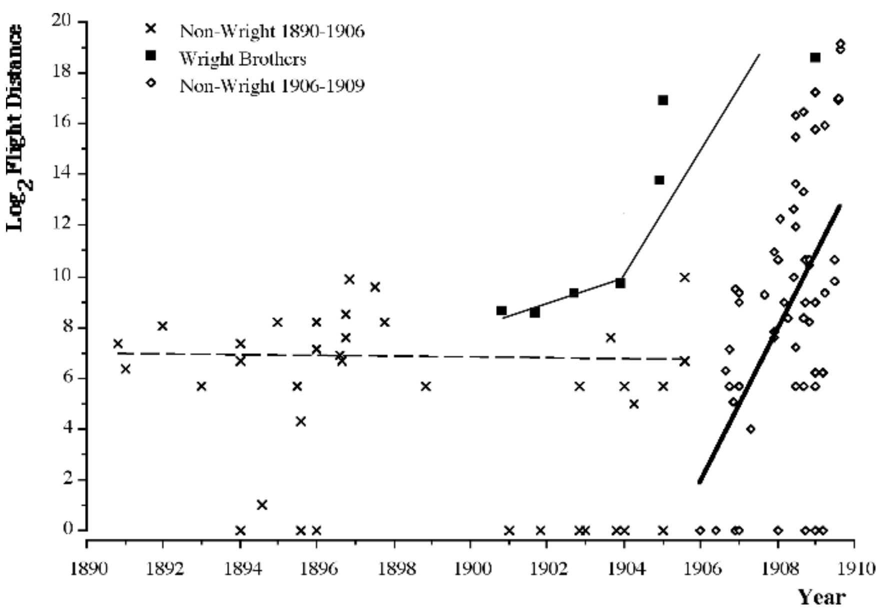

Fig. 6. Attempts to master the airplane. Each point represents the best flight made by a particular craft on the date the flight was made.

other inventors did not show similar improvements? To explain this mystery, we first need to examine the logic behind the design-and-test strategy of invention. Suppose that an inventor has designed an airplane, taken it out to the field, and flown it. Between 1799 and 1909, about the greatest distance achieved by any craft was 100 meters. Now the inventor would like to improve upon the design. As long as an inventor can create a craft that is even slightly better than the last, he can continue improvements and produce a positive performance slope. This represents a familiar hill-climbing strategy of solving problems, and hill-climbing is effective for solving many problems.

Why did this strategy fail in the invention of the airplane? Bradshaw & Lienert [7] argued that the failure arose because making a test flight does not provide diagnostic information about the strengths and weaknesses of a design: The wings may not have produced sufficient lift to keep the craft in flight, the airframe might be causing too much drag, the center-of-lift might not coincide with the center-of-balance, or the pilot may have made a mistake in flying the airplane. Under these circumstances, inventors had no reliable information about what was wrong with their craft, or what specifically to change to improve it.

So how did the Wrights escape this trap? Although they did build and test gliders and airplanes, their approach towards invention differed in substantial ways from their contemporaries. In particular they followed a *functional decomposition* strategy: They isolated different functional subsystems (wings, power plant, elevator), identified specific performance requirements for each of the systems (i.e., the wings must produce 110 kilos of lift), and employed what little was known about aerodynamics to produce a subsystem that met their design requirements. This led to the first glider, which they tested in 1900. When the glider did not perform as designed, the Wrights realized there was something wrong with their computations of lift. This led to their construction of a series of wind tunnels and the development of instruments to measure lift and drag. Somewhere between 80 and 200 different wing shapes were tested in the wind tunnel. Their results demonstrated that the current value of the *coefficient of lift* was incorrect, revealed a much better approximation for that coefficient, and showed that long and thin wings had better characteristics than short and broad ones. Wind tunnel tests also revealed that an airplane wing with its highest point near the leading edge of the wing had better changes in the travel of the center of lift than did a wing with the highest point in the center of the wing. Testing a functional subsystem in isolation produced diagnostic information about which wing designs were good, and which ones were bad. Their information was so precise that the Wrights were able to build a glider in 1902 that had excellent flight characteristics.

Between 1902 and 1903 the Wrights performed more tests in their wind tunnel – to determine the best shape for an airplane propeller. Once again their approach was quantitative: the Wrights knew how much horsepower their engine could produce, so they needed a propeller with a specific level of performance to produce the thrust their craft needed. The Wrights were so certain of the success of their 1903 flyer that the pair, who were known for their modesty and caution, wrote a press release announcing their success before leaving Dayton for Kitty Hawk. Their father issued the press release following the receipt of a telegram from the Wrights told of their actual success.

As an invention strategy, *functional decomposition* has several advantages over *design-and-test*. We have already discussed the lack of diagnosticity in design and test. Because of the decomposition inherent in functional decomposition, testing of a part is not confounded by the performance of other parts of the system. It was also possible for the Wrights to develop far more precise indices of performance of subsystems than it was to evaluate the performance of the part in a complete craft. In 1900 and 1901, for example, the Wrights suspected their wings were not producing as much lift as calculated. They attempted to measure the lift of their glider by weighting down the wings with some chain, flying their glider as a kite, and measuring the angle of attack and wind speed simultaneously. But the winds were not perfectly steady and their glider reacted by changing its elevation and angle of attack in reaction to changes in wind speed, so the Wrights could not determine with any degree of accuracy how much lift their wings were actually producing. They suspected that the value for the coefficient of lift was incorrect, but did not have proof. Once they built their wind tunnel, they were able to precisely measure lift, drag, and the coefficient of lift. A third advantage of functional decomposition arises from the combinatorics of *divide and conquer*: When one wing is shown to have poor lift and drag characteristics, the Wrights were able to exclude from consideration tens- or hundreds-of-thousands of airplanes: any design that included that wing was a bad choice.

Another advantage of the Wright approach arose from their use of precise performance specifications. The Wrights knew their gliders had to support the weight of the pilot along with the weight of the craft. For an 80-kilo pilot and a 45-kilo glider, the wings must produce 125 kilos of lift. Having these performance requirements allowed the Wrights to *satisfice* in their designs: once they found a wing that could produce the necessary lift at the target weight, they did not have to look further to find an optimal wing design. Also, having performance specifications allowed the Wrights to determine when their design had failed. The Wright's second glider, built in 1901, performed nearly as well as any other craft, powered or unpowered, had done to date. Rather than being satisfied with their near-world-record performance, the Wrights were discouraged that the craft had not performed as designed: this dissatisfaction led directly to their development of a wind tunnel to determine why the glider was not generating the computed lift.

One more aspect of the Wright's approach deserves mention: the use of theory and math to substitute for search. Previous research had uncovered a *lift function* that enabled the Wrights to test small models of wings just a couple of inches long, then predict the performance of a large-scale wing with considerable accuracy. If this function were not know, the Wrights might have been forced to test full-scale wings in a large wind tunnel. Such research would most likely have been prohibitive given their modest means.

Through all of these efficiencies, the Wrights were able to develop an airworthy glider in just three years, then take only three more to produce a practical airplane capable of extended flight. These accomplishments were made while they maintained a successful small business, and with quite modest financial resources. Only when others were able to study the Wright craft and the advances they made were they able to produce competitive airplanes.

1.2 Principles of Effective Invention

This review of the invention of the airplane suggests that there are ways to achieve considerable efficiency in the process of invention. These efficiencies will be most evident for complex inventions where various elements contribute to the success of the whole. Under such circumstances, search can be reduced by:

1. Identifying the functions to be performed by the invention;
2. Specifying functional requirements that the system must meet;
3. Developing subsystems that meet these functional requirements;
4. Testing the subsystems in isolation from the whole system;
5. Focusing attention on subsystems that fail to perform as designed; and
6. Utilizing theory to generalize results.

Whenever these strategies can be practically employed, they will reduce the complexity of the invention.

2 The Invention of AI Systems

In discussing invention, it should be clear that many AI methods, particularly those in learning and discovery systems, have application as a way to produce new inventions, just as they can learn and make new discoveries. Perhaps in the near future a discovery system will build a better mousetrap or a learning system will produce a better user interface. By exploring such applications we can increase the utility of our systems and methods: we can apply them to invention problems as well as discovery and learning problems. But there is another reason for discussing the process of invention: AI systems are not discovered they are invented. As such, they are governed by the same principles of invention that have just been described above.

Why are AI systems best considered as inventions and not discoveries? Clearly artificial intelligence falls within the"Sciences of the Artificial" as described by Herbert Simon [6]. AI may draw upon research findings in psychology, but clearly AI methods and systems are the product of human enterprise, and so can be understood as an invention. Given this status, we may consider how the lessons of invention can be applied to AI systems as a special case. We begin by considering the design space of AI systems. Table 2 illustrates some of the choices that investigators face as they are putting together a new AI system:

Many of the entries in the table refer to a family of related methods. For example, using parametric statistics to handle noise in the data might include

Table 2. Design Features of AI Systems

Design Parameter	Possible Values
Knowledge Representation	Symbols; Schemas; Propositions; Productions; Distributed Sub-Symbolic Nodes ...
Thought Processes	Productions; Bayesian Probabilities; Spreading Activation; Predicate Calculus; Schema Inference; Markov Transitions ...
Learning	Proceduralization; Composition; Backpropagation; Genetic Algorithms; ...
Noise	Parametric Statistics; Non-parametric statistics; Bayesian Probabilities; Signal Detection; ...
Test Database	Iris; Solar Flare; Credit Card; ...
Competing System	C4.5; Soar; ACT-R; Harmony; Neural Network; ...

something simple, like computing the mean, to finding a regression line, or even using the standard deviation to find outliers that are treated as a special case. Clearly we have produced a rich set of alternatives from which researchers can choose in developing a new AI system.

Let us suppose, for a moment, that a researcher decides to build a new AI system drawing upon the alternatives shown in Table 2. Choosing a system based upon schemas for knowledge representation, spreading activation and productions for thought processes, and proceduralization for learning, the researcher then adds a new method of dealing with noise based on non-parametric statistics and Bayesian probabilities, known as F.A.K.I.R., further adding to the pool of methods available in AI. The new system, including the F.A.K.I.R. algorithm, is called HOLIER.THAN.THOU.[1] The researcher decides to compare HOLIER.THAN.THOU. against C4.5 [9] using a database of credit card transactions.

On the first comparative test, HOLIER.THAN.THOU. does not perform well classifying database transactions. The researcher examines the errors made by the F.A.K.I.R. and identifies some problems with the new procedure for accommodating noise in the data. By adjusting several parameters and making other tweaks to the system, HOLIER.THAN.THOU. now outperforms C4.5 on the database by 2%.

We might ask, "What value is the new F.A.K.I.R. noise reduction technique for the AI community?" Clearly the researcher has demonstrated that, under some conditions, HOLIER.THAN.THOU. can out-perform C4.5. But grave questions remain about the generality of this claim, about the significance of the 2% difference, and about the source of the performance difference. Let us consider each of those issues in turn.

[1] Any resemblance between F.A.K.I.R/HOLIER.THAN.THOU. and an actual AI systems or methods is purely coincidental, and we do not suggest that any existing AI system has been developed in this way.

Generality of the Result. AI researchers often begin with an understanding of the problems that certain databases represent for learning and discovery algorithms. But we lack a deep conceptual understanding of the fundamentals: "How much adaptation does an adaptive system have to perform?" or "What kinds of learning does a learning system need to do?" We can answer these questions with respect to certain well-known databases, but not with respect to an entire class of problems. For this reason, we cannot readily determine how many different databases are needed to demonstrate the generality of a new system or algorithm. Should we test each new system on 10 different databases? Which ones? At what point are we certain that we have tested a new system against all of the interesting problems an intelligent system might face? The Wrights were lucky enough to have mathematical equations that allowed them to predict the performance of a full-size wing from a small model wing. Would anyone care to predict how HOLIER.THAN.THOU will do on a database of credit card transactions *vis a vis* a neural network system?

We may never enjoy the situation the Wrights found themselves in, where a simple mathematical function can predict how a system will behave under different situations. One way researchers have responded to questions about generalization is to test their system against multiple databases – a method that does help to establish the generality of a new method or system. Yet even still another serious issue remains: how good does our system or method need to be in order to be useful across an interesting range of problems? Remember that the Wrights could specify in advance how much lift their system needed to generate in order to fly. That allowed them to find a satisfactory solution to the problem, without the necessity of finding an optimal one. Are we now looking at solutions that are sub-satisfactory, satisfactory, approaching optimal, or optimal?

Significance of a Performance Difference. The researcher found a 2% performance advantage for HOLIER.THAN.THOU when compared to C4.5. But this advantage only occurred after careful fine tuning of HOLIER.THAN.THOU. Was the same care taken to ensure that C4.5 was performing at its best? Perhaps, perhaps not. Even if C4.5 was adjusted to perform at its best, we are still uncertain about what the best possible performance is on the credit card database. It might seem possible to attain a 100% accuracy on classifying all items in the database simply by memorizing each item. However, some databases could suffer from an impoverished description of items where two items with the same description belong to different classes, or God could be playing dice with the Universe, and no description would be adequate to classify every item correctly.

A more difficult question arises when we try to determine whether the 2% improvement in accuracy rate is *better* or *worse.* At first glance the question seems foolish: HOLIER.THAN.THOU performed more accurately on the database. How could that not be better? Yet experience has taught us that adaptive systems can learn the noise in the database as well as the signal. These systems do not generalize as well to new data as ones that only learn the true generalizations present in the database.

A final awkward question arises when we consider the statistical and practical significance of a 2% improvement in classification performance. If C4.5 correctly classifies 97% of the transactions and HOLIER.THAN.THOU correctly classifies 99%, the difference could be significant and important. But if C4.5 correctly classifies 20% of the transactions and HOLIER.THAN.THOU classifies 22%, neither system seems very impressive. Researchers in AI commonly use some sort of split-part reliability computation, which helps to determine the reliability and statistical significance of the results.

Credit Assignment for F.A.K.I.R and HOLIER.THAN.THOU. Once we convince ourselves that the 2% difference is significant and important, we are left with one more awkward question: where did this difference come from? Was it due to some advantage of the F.A.K.I.R. algorithm in isolating noise and concept drift from the signal? Or was it due to a difference in the initial representation of the data between HOLIER.THAN.THOU and C4.5? Or was there some other difference between the two systems? Answering the question has greater importance than at first appears. Perhaps the F.A.K.I.R. algorithm is improving the performance of HOLIER.THAN.THOU by 25%, but other limitations of HOLIER.THAN.THOU reduce the advantage by 23%. We might then combine F.A.K.I.R. with C4.5 and achieve a better result yet. By knowing how well each element of the system is doing its job, we can produce the best possible combination of elements.

The analogy between the invention of the airplane and the invention of AI systems can be pushed too far: Computer programs are typically designed through a series of function calls, and it is possible to determine if the calls are operating as designed. This lends a transparency to computer programs that airplane inventors did not enjoy. Function calls often map roughly onto the functional specifications for a system, although there are many internal function calls that do not have an obvious connection to the larger functional subsystems of the program. Yet by considering AI as an invention, it raises two important questions: "What are we trying to invent?" and "Are we working efficiently toward that goal?"

2.1 Reflections on Invention and AI

When one examines recent developments in AI, it is clear that AI researchers are inventive and are pouring tremendous creative energy into developing new heuristic and algorithmic methods to address difficult problems. Evidence of this inventiveness is present in the the two volumes published last year for ALT 04 [10] and Discovery Sciences 04 [11], each of which presented a number of important papers in their respective fields. As a result of this worldwide enterprise, researchers are now faced with an embarrassment of riches in the number of different methods they have available for the construction of new AI systems.

But there still seems to be an important gap in our knowledge – we don't fully understand the relationship between where we are and where we want to be. Are we building AI systems like the pre-Wright airplanes that struggled to 'fly' 100 meters? Or are we improving capable airplanes to extend their range from trans-continental flight to inter-continental flight? With aviation, everyone knew that airplanes needed to fly long distances quickly and to carry as much weight as possible. AI has no such simple goals: it may be quite valuable to build an expert system like XCON [12] to design the backplane of Vax computers, even if the system has only a limited expertise and lifetime. Or we may set our goals higher to develop more versatile and capable AI systems.

We can, of course, continue to develop even more new learning and discovery methods. But hopefully the energy being spent to develop new methods can be balanced with an effort to better document *what our systems need to do* and *how well they need to do those things.* As the baseball player Yogi Berra said, "If you don't know where you are going, you will wind up somewhere else." Through a better understanding of the fundamental problems of learning, discovery, and AI, we can work towards functional specifications that tell us how well our systems need to perform, then choose methods that will get us where we want to be, instead of somewhere else.

References

[1] Basalla, G.: The Evolution of Technology. Cambridge University Press (1988)
[2] Petroski, H.: The Evolution of Useful Things. Alfred A. Knopf. (1993).
[3] Newell, A., & Simon, H.A.: Human Problem Solving Prentice-Hall (1972)
[4] Weber, R.J. & Perkins, D.N.: How to invent artifacts and ideas. New Ideas in Psychology, **7** (1989) 49–72.
[5] Simon, H.A.: Scientific discovery and the psychology of problem solving. In R. Colodny (Ed.), Mind and Cosmos. (1966) Pittsburgh: University of Pittsburgh Press.
[6] Simon, H.A.: The Sciences of the Artificial, Second Edition. The MIT Press. (1981)
[7] Bradshaw, G.L., & Lienert, M.: The Invention of the Airplane Proceedings of the Thirteenth Annual Cognitive Science Conference (1991), pp. 605-610.
[8] Bradshaw, G.L.: The Airplane and the Logic of Invention. In R.N. Giere (Ed.), Cognitive Models of Science. The University of Minnesota Press. (1992)
[9] Quinlan, J.R.: C4.5: Programs for Machine Learning. Morgan Kauffman. (1993)
[10] Ben David, S, Case, J., & Maruoka, A.: Algorithmic Learning Theory 15th International Conference, ALT 2004, Padova, Italy. Springer-Verlag (2004)
[11] Suzuki, E., & Arikawa, S.: Discovery Science 7th International Conference, DS 2004, Padova, Italy. Springer-Verlag (2004)
[12] McDermott, J.: R1: A Rule-based Configurer of Computer Systems. Artificial Intelligence, **19** (1982), 39–88