Reinventing the Wheel: An Experiment in Evolutionary Geometry

Josh C. Bongard
Computational Synthesis Laboratory
Sibley School of Mechanical and Aerospace Engineering
Cornell University
Ithaca, NY 14850 USA
josh.bongard@cornell.edu

Hod Lipson
Computational Synthesis Laboratory
Sibley School of Mechanical and Aerospace Engineering
Cornell University
Ithaca, NY 14850 USA
hod.lipson@cornell.edu

ABSTRACT

In the domain of design, there are two ways of viewing the competitiveness of evolved structures: they either improve in some manner on previous solutions; they produce alternative designs that were not previously considered; or they achieve both. In this paper we show that the way in which designs are genetically encoded influences which alternative structures are discovered, for problems in which a set of more than one optimal solution exists. The problem considered is one of the most ancient known to humanity: design a two-dimensional shape that, when rolled across flat ground, maintains a constant height. It was not until the late 19th century—roughly 7000 years after the discovery of the wheel—that Franz Reuleaux showed that a circle is not the only optimal solution. Here we demonstrate that artificial evolution repeats this discovery in under one hour.

Categories and Subject Descriptors

J.2 [**Computer Applications**]: Physical Sciences and Engineering

1. INTRODUCTION

It is widely believed that the wheel—a circular object that allows much less friction during transport than dragging—was discovered in ancient Mesopotamia in the 5th millennium BC. It is possible that there was an independent discovery of the wheel in China in 2800 BC, but there is less historical evidence supporting this discovery. Despite the overwhelming utility of this structure, many major civilizations throughout history failed to discover it, including those in Sub-saharan Africa, Australia and the Americas[1] [11].

The circle is an optimal shape for a wheel because a circle's geometric center maintains a constant height when it is rolled over flat ground. This is desirable for vehicles, in which the axle passes through the wheel's geometric center. However objects can also be rolled over wheels or cylinders because circles also maintain a constant height: the height of a circle remains constant as it rolls. The circle is only one of an infinite set of optimal shapes that maintain a constant height during rolling. These shapes are known as curves of constant width [5].

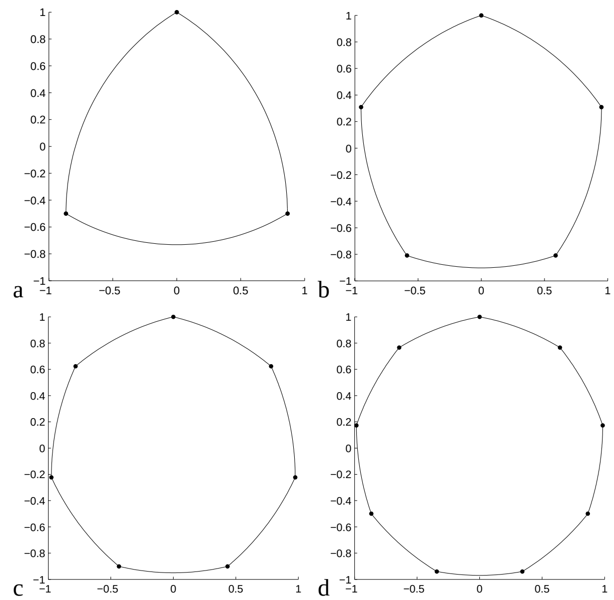

Figure 1: a: The Reuleaux triangle. b-d: Reuleaux 5-, 7- and 9-polygons, respectively.

The time at which it was realized that the circle is not the only known curve of constant width is not well known. However the Reuleaux triangle, named after Franz Reuleaux (1829-1905), a German professor of engineering and machine design, is a non-circular curve of constant width. The Reuleaux triangle was first mentioned in 1876 [9] (and reprinted in 1963 [10]), so taking this as an approximate date of the discovery of non-circular curves of constant width, and the discovery of the first curve of constant width—the wheel—in the 5th millennium BC, it took humanity roughly 7000 years to discover that there is more than one curve of constant width. The Reuleaux triangle, along with three other non-circular curves of constant width, are shown in Figure 1. Pairs of Reuleaux triangles connected together to serve as smooth rollers are shown in

[1]Children's toys from the Incan civilization suggest that that society was at least aware of wheel-like shapes, even if they were not used for utilitarian purposes.

GECCO'05, June 25–29, 2005, Washington, DC, USA.

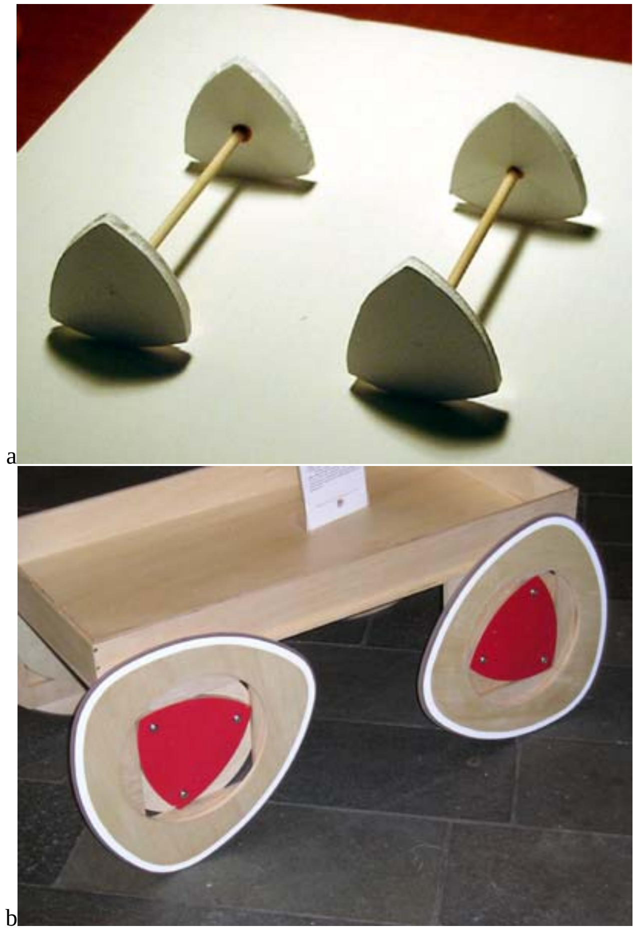

Figure 2: a: A pair of rollers constructed from Reuleaux triangles. b: Cart with Reuleaux triangles as wheels. Source: UNESCO exhibit “Experiencing Mathematics”. Photo: David W. Henderson.

Figure 2a. Reuleaux polygons can also serve as wheels if their axle is also made to have the same profile with opposite, compensating motion. Figure 2b shows a cart with Reuleaux triangles as wheels and as axles [1].

In this paper we present three genetic algorithms that differ only in their genetic representation: all of them consistently design curves of constant width. We will show that several different curves are discovered, including the circle, but that the method of genetic encoding biases which solution is discovered, a well-documented process in evolutionary computation (eg. [2], [4]). In the next section we describe the set of curves of constant width; in section 3 we describe a methodology for, and results from the evolution of curves of constant width. In section 4 we provide some concluding remarks.

2. CURVES OF CONSTANT WIDTH

For closed convex planar bodies whose boundary is a smooth curve, there are exactly two parallel tangent lines to the boundary curve in any given direction. In a physical setting, these lines can be considered to be flat ground, and a line parallel to the ground that the body intersects at some point when it rolls through 2π radians. The width of the curve in a given direction is taken to be the perpendicular distance between the tangents perpendicular to

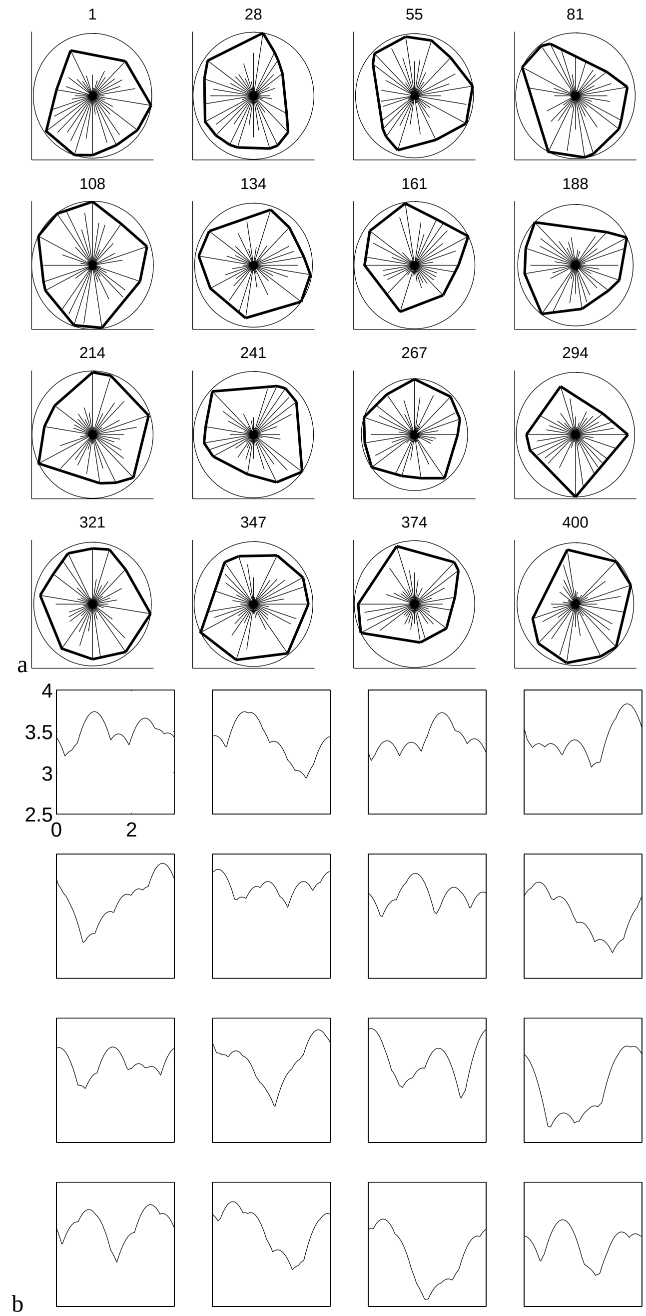

Figure 3: a: A sample of solutions from an initial random population (thick line), with no bias toward symmetric or asymmetric shapes. The rays indicate the 60 evolved radii; the polygons indicate the resulting convex hulls. The circles indicate the bounding circle, which is for visual comparison only; it does not influence fitness. Figures indicate the index of the solution within the population. b: The tracks left by the height of these shapes, when they are rolled along a flat surface from 0 to π.

that direction. For physical shapes, this definition is simpler: at any point of rotation, when the shape touches the ground, its width is taken as the distance from the ground to the shape’s maximum height.

According to this definition, the circle qualifies as a curve of constant height. However there are also an infinite set of other shapes

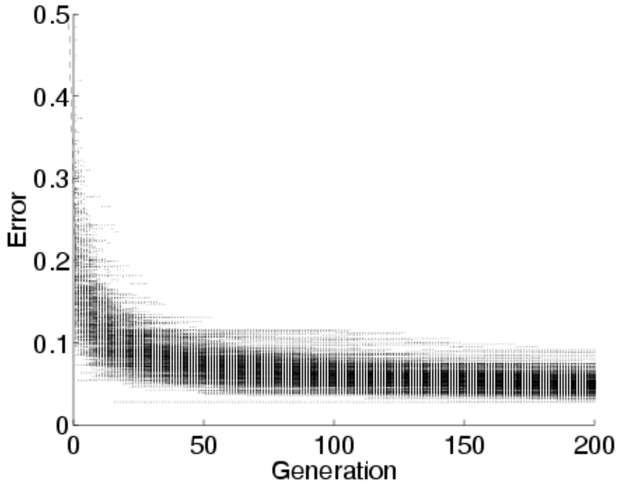

Figure 4: The errors of each shape from a single evolutionary run.

which satisfy these conditions. Rather than giving a formal definition here, we instead provide a constructive definition. Given some n selected from the infinite set of odd integers $\{3, 5, 7, \ldots\}$, construct an n-polygon with equal length sides. For each point of the polygon A, draw an arc connecting the two opposing points B and C on the far side of the polygon, where the radius of the arc is $|A - C| = |A - B|$. Once all n arcs have been drawn, the shape composed of these connecting arcs is a curve of constant width. The curve of constant width with $n = 3$ is the Reuleaux triangle (Figure 1a); the curve with $n = \infty$ is the circle. The British 20p and 50p coins are curves of constant width with $n = 7$, which—for a vending machine—makes them indistinguishable from a round coin.

The members of the set of curves of constant width have several distinguishing features. For example the circle has maximum area, while the Blaschke-Lebesgue theorem proves that the Reuleaux triangle has the least area [6]. Also, the Reuleaux triangle can be used to dig a square hole [5]; as n increases, the shape of the hole produced by rotating the corresponding curve of constant width has increasingly rounded corners until when $n = \infty$, the hole dug by rotating a circle is a circle.

3. EVOLVING CURVES

Evolving shapes is a popular application of evolutionary computation. Rechenberg, in one of the first evolutionary computation experiments ever carried out, evolved the shape of a curved pipe in order to maximize fluid flow from a vertical input nozzle to a horizontal output nozzle: surprisingly, he found that a pipe describing a perfect quarter-circle arc is not the optimal solution to this problem [8]. Other methods for evolving shapes can be found in [3].

Three genetic algorithm variants were devised to evolve two-dimensional shapes, which are distinguished only by their genetic encoding. In all three variants, each linear genome encodes 60 floating-point values in $[0.0, 1.0]$. The values are used to construct 60 radii; there is an angle of $\frac{60}{2\pi}$ radians between each pair of radii. Radii have a length within $[0.1, 2.0]$. The fitness of a constructed shape is determined as follows. For each genome, a convex hull[2] is constructed using the tips of the 60 radii as the points to be considered. Figure 3a shows the radii and resulting convex hulls produced by a sampling of random genomes.

[2] A convex hull produces a shape that contains no concavities

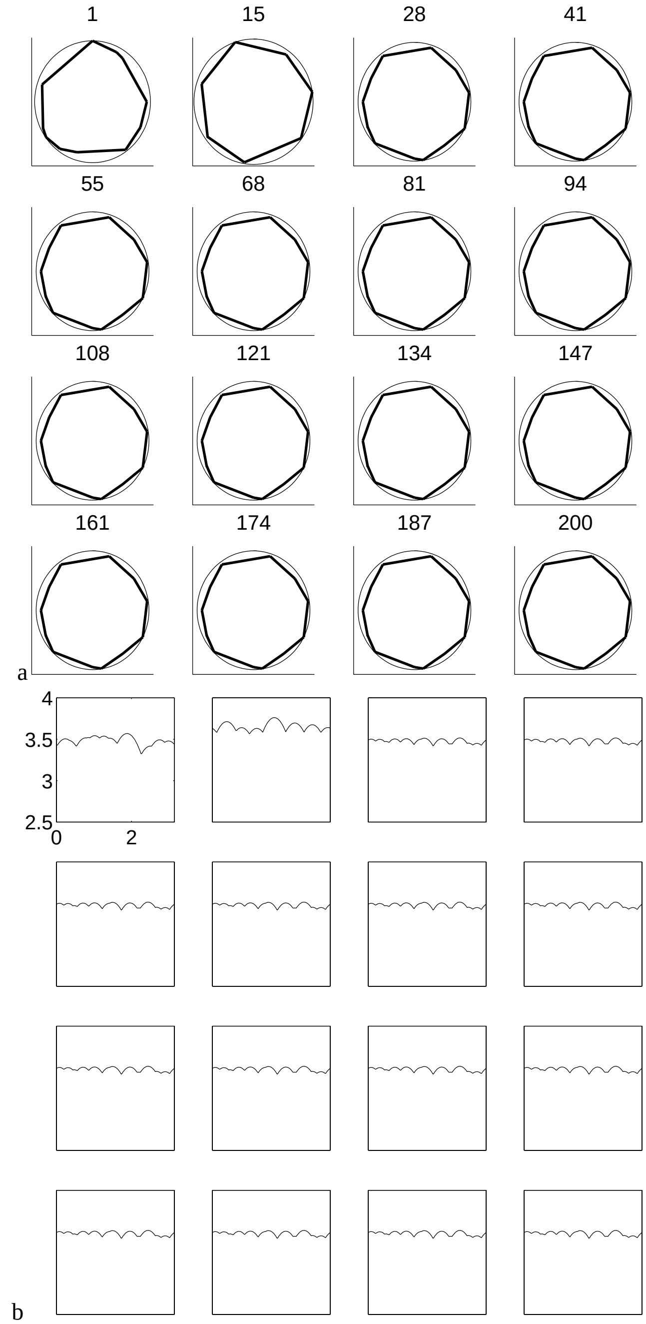

Figure 5: a: The best solutions from a sampling of generations in a single run. Figures indicate the generation. The best solution from the first initial population is shown top-left; the best solution in the final population is shown bottom-right. b: The tracks left by the heights of these shapes.

The set of widths W are calculated as the set of heights of the bounding boxes at the 100 orientations $[p_1, p_2 \ldots p_{100}]$ equally spaced between 0 and π. The error of a genome is then given as

$$e = \frac{\max(W) - \min(W)}{\min(W)}.$$

Intuitively, e represents the amount of 'bumpiness' experienced by an object when rolled π radians over the shape in question. The

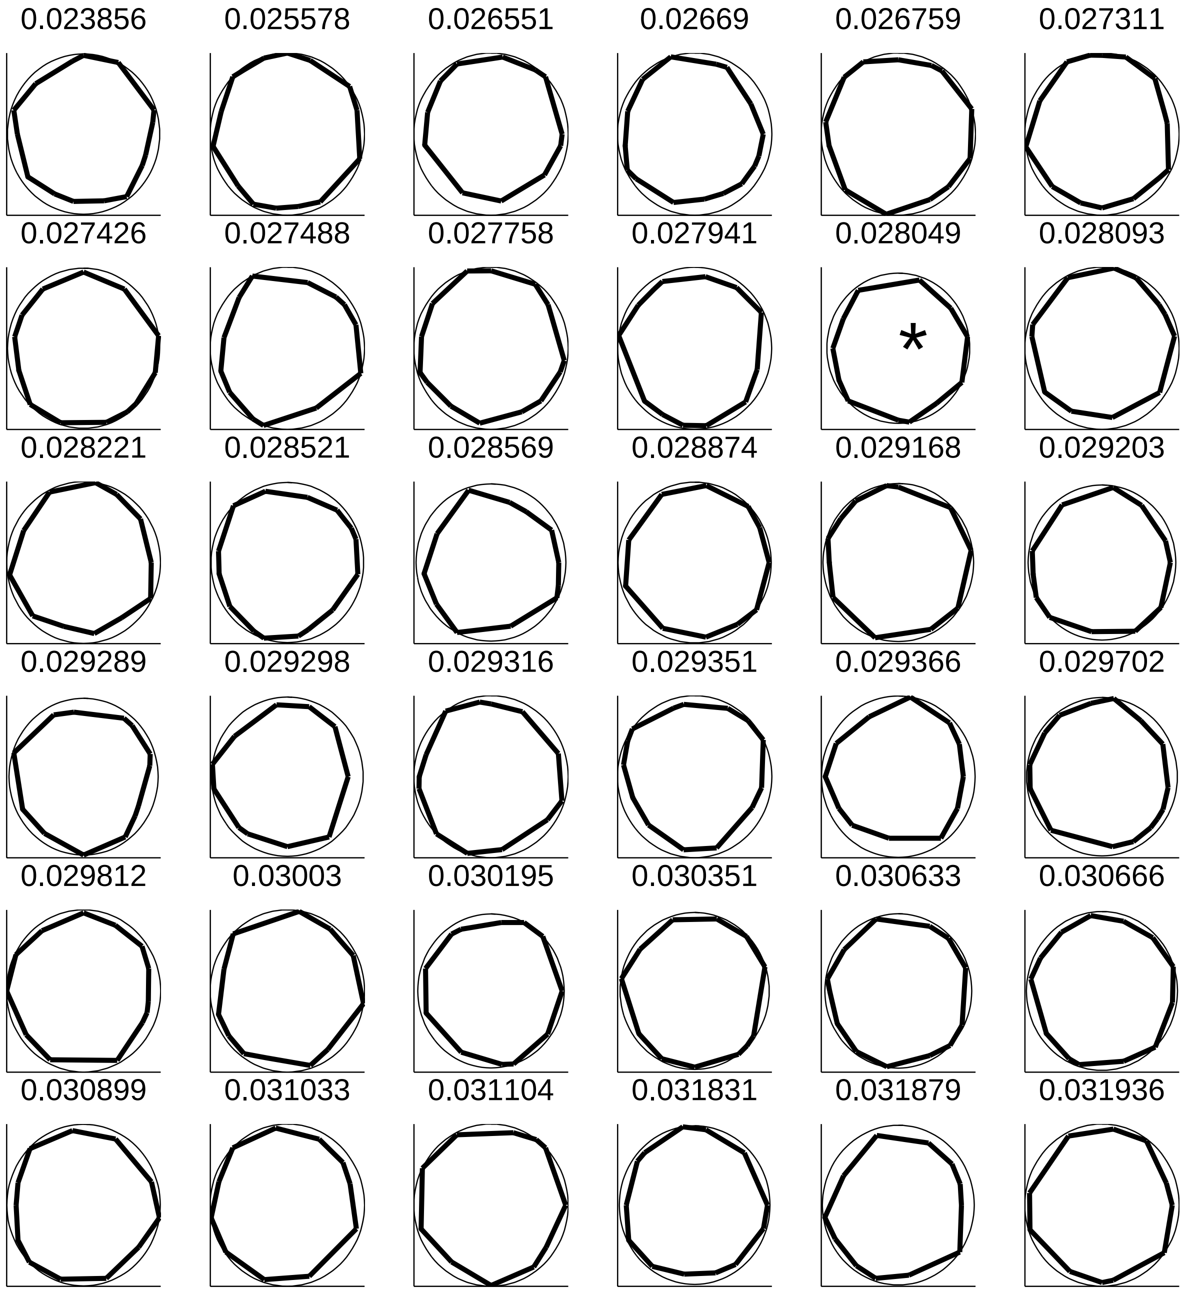

Figure 6: The best solutions from a sampling of generations in a single run, with no bias toward symmetric or asymmetric shapes. Figures indicate the generation. The best solution from the first initial population is shown top-left; the best solution in the final population after the last generation is shown bottom-right. The asterisk indicates the run reported in figures 4, 3 and 5.

genetic algorithm attempts to minimize e.

For all three algorithm variants, an initial population of 400 genomes was constructed, and evolved for 200 generations. Deterministic crowding [7] was used to maintain diversity. Each new child genome underwent mutation. Mutation rate was set to $\frac{1}{60}$, so that on average one gene was mutated per genome duplication. A Gaussian-type mutation operator was used: an original gene value was increased by e^{α} with probability 0.5, and decreased by e^{α} otherwise, where α is a random value selected uniformly from $[-10, 0]$. If the new value was less than zero it was set to zero; if it was greater than one it was set to one. This operator ensures that small mutations happen more often than large mutations. Each child pair produced by deterministic crowding underwent one-point crossover.

For all three algorithm variants, 36 independent runs were conducted. In the first algorithm variant the 60 values encoded in the genome are simply scaled from $[0, 1]$ to $[0.1, 2.0]$, producing the 60 radii. Figure 3 reports some initial random shapes from a sample run, along with their corresponding tracks. A track is simply the set of line segments that connect the 100 heights stored in W; it represents the change in height that an object would experience if rolled over the shape along flat ground for one-half of the shape's rotation.

Figure 4 reports the errors of the genomes for the same run, while figure 5 reports the best solutions from a sampling of generations, along with their corresponding tracks. Clearly, in this run, the genetic algorithm converged on the curve of constant width with $n = 7$ early in the run, around generation 28. Before this point, a straight-edged 7-sided polygon with sides of differing length dominated the population, as seen in generation 15.

However, this algorithm variant did not consistently converge on this curve of constant width. Figure 6 reports the best shapes produced by all 36 runs of this variant. Henceforth only the convex hulls and bounding circles are shown, and not the underlying radii, for visual clarity. As can be seen in that figure, some runs converged on an approximation of a circle, as indicated by how fully they occupy their bounding circle (a circle drawn using the largest radius from among the 60 radii of the shape). For example the second best solution found (second from the left in the top row of Figure 6) nearly completely fills its bounding circle. However several of the other best solutions approximate 3-, 5- or 7-sided curved polygons of constant width. In this respect this algorithm is competitive with human ingenuity, in that it indicates, after less than 1 hour, that this problem has several alternative solutions, a fact that took humanity about 7000 years to realize.

The consistent evolution of odd-sided curved polygons with more or less equal sides and roughly equal curvature (within each polygon type) suggests that there is a common symmetry among the set of alternative solutions. Because the set of optimal solutions is known, it is confirmed that symmetry does indeed exist in all solutions: for each curve of constant width, each side is of equal length, and each arc is of equal curvature. Based on the evolved solutions, we chose to create two additional algorithm variants that enforce symmetry. This is a common technique in evolutionary computation: observation of evolved solutions motivates the programmer to revise their fitness function or genetic encoding.

In the second algorithm variant, genomes still encode 60 floating-point values. However, the first value is rounded to an integer in $[2, 6]$, s. The 2nd to $(\lfloor \frac{60}{s} \rfloor + 1)$th gene values are then extracted from the genome, and rounded to $[0.1, 2.0]$. (The remaining values in the genome are not used in this algorithm variant.) This subset of radii are then copied and concatenated s times. If i less than 60 radii have been created, they are filled in using the first i radii from the subset are copied. This encoding ensures that between 2 and 6 sets of radii will be repeated around the circumference of the shape, increasing the symmetry of the shapes. Another 36 independent runs were conducted using this encoding scheme, but with the same population size, number of generations, and selection, mutation and crossover operations as in the first variant. Figure 7 depicts some of the initial random shapes produced by this variant. As can be seen, this encoding tends to produce polygons with curved sides. With the help of the radii within the repeated subset, polygons with sides greater than six and with sides of differing lengths can be produced, such as the 10-sided polygon produced by genome 241.

Figure 8 reports the best shapes produced by each of the 36 runs for this second algorithm variant. Clearly, the algorithm favors polygons with five sides, although this variant does produce near-perfect circles in several runs. This variant could then cause us to conclude that this problem has two equally good solutions: circles and curved five-sided polygons. However, as is known, there are actually an infinite number of equal and optimal solutions. Figure 8 therefore indicates that this new, symmetry-enforcing encoding has introduced a bias that favors some solutions over others.

This result is supported by Figure 9a, which reports how many sides each genome encodes, over all 36 runs, and grouped according to which generation it appears in. Clearly, genomes that encode five sides quickly come to dominate each population early during the evolutionary process. However, there is a minority group of three-sided polygons that exists for some time before being completely replaced after about the 30th generation. The three even-sided possibilities are replaced even more rapidly, even though they appear in equal numbers in the initial random population.

It is hypothesized that the reason that three-sided polygons are replaced that five-sided polygons is that because there are fewer radii to optimize in this latter case, evolution can better optimize the curvature of these polygons to approach the corresponding curve of

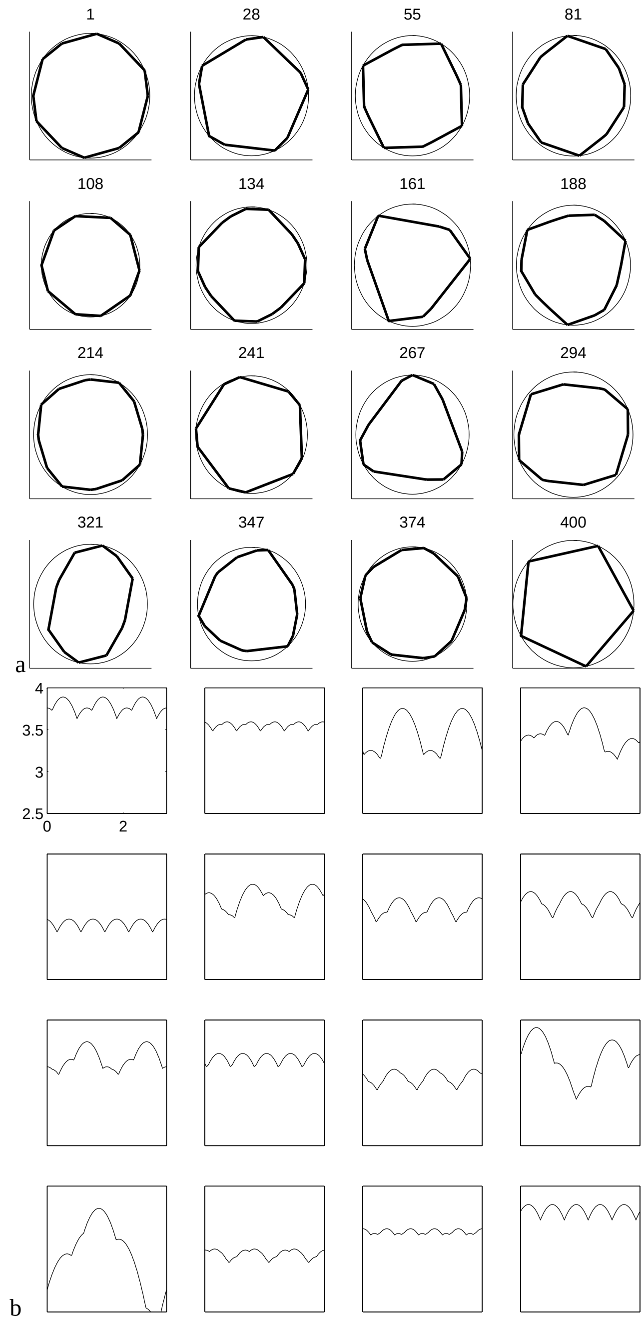

Figure 7: a: A sample of solutions from an initial random population (thick line), with no bias toward symmetric or asymmetric shapes. b: The tracks left by the heights of these shapes.

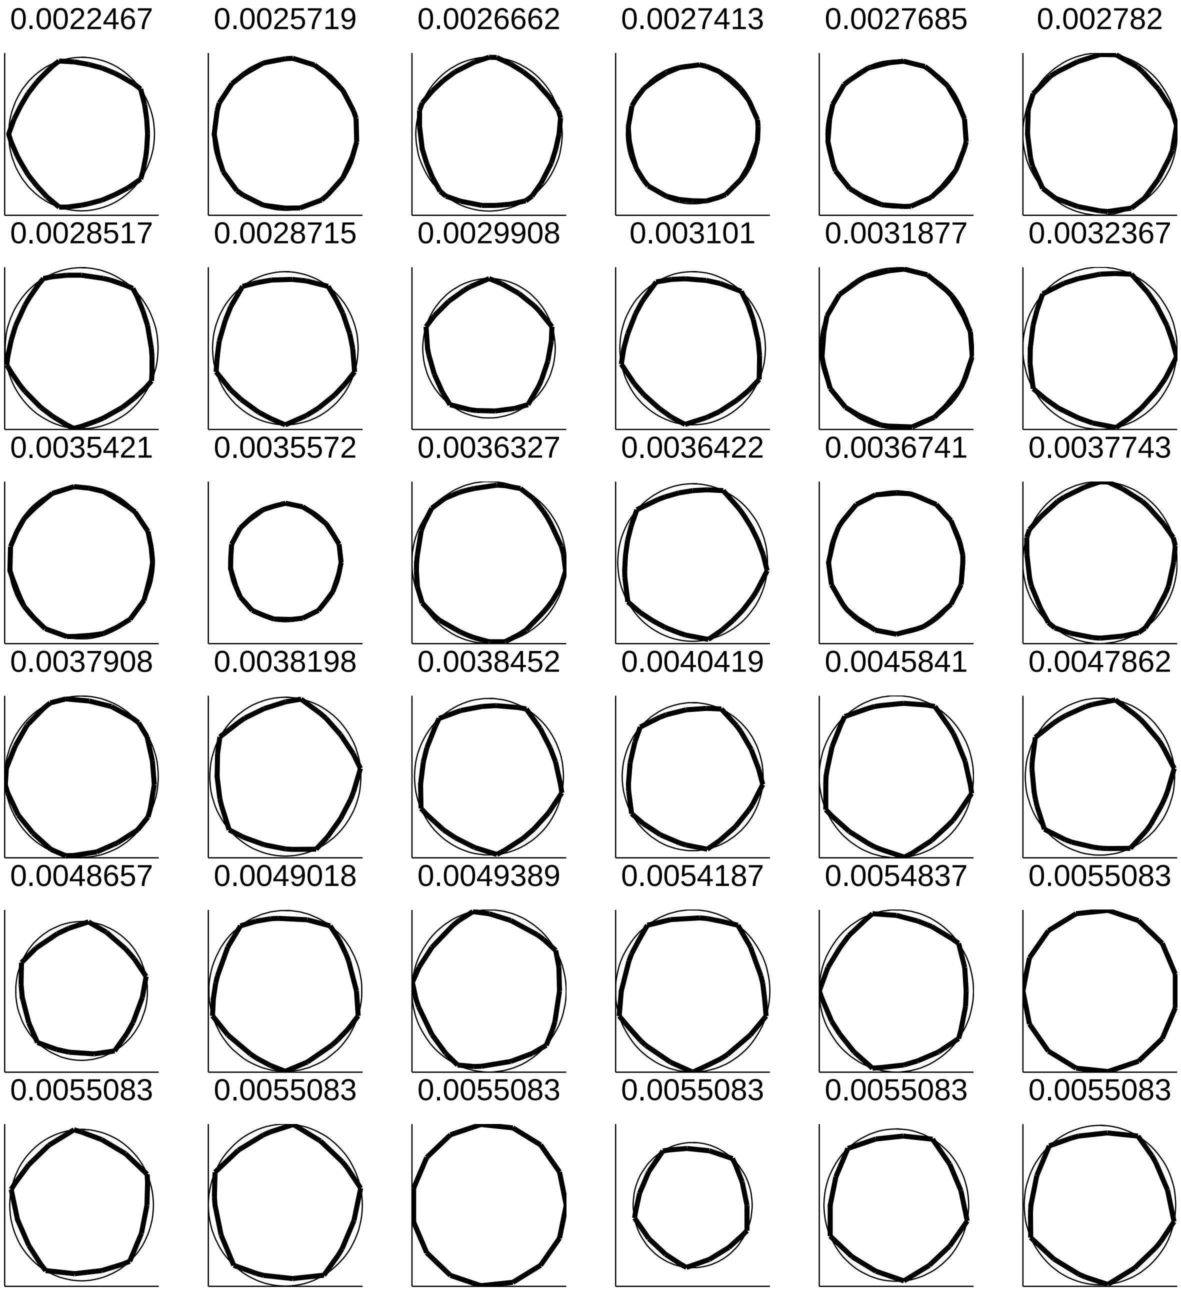

Figure 8: The best solutions from a sampling of generations in a single run, with enforced symmetry.

constant width. In order to test this hypothesis, a third variant was devised that is identical to the second variant, except that the first value is scaled to $[2, 8]$ instead of $[2, 6]$. Once again, 36 runs were conducted using the same evolutionary parameters as the first two variants.

Figure 9b reports the number of genomes that encode which number of sides. In this case seven-sided polygons come to dominate the population. This is somewhat surprising, as $s = 7$ genomes do not produce symmetric shapes: the subset containing $\lfloor 60/7 \rfloor = 8$ radii is repeated seven times, giving 56 radii; the remaining four radii are filled in from the first four radii in the repeated subset, leaving a shorter, eighth side. Also, five-sided polygons clearly dominate all populations early on in evolution, but are gradually replaced by the seven-sided polygons starting around the 50th generation. This observation seems to fit with our prior hypothesis that polygons with a greater number of sides can be better optimized toward a curve of constant width. However the reason why the five-sided polygons clearly dominate early on is not yet well understood. The surviving minority of five-sided polygons near the end of evolution is explained by the fact that three of the 36 runs produced clear five-sided curved polygons of constant width, with the remaining runs all converged on near-perfect circles (data not shown).

Thus, this encoding change between the second and third algo-

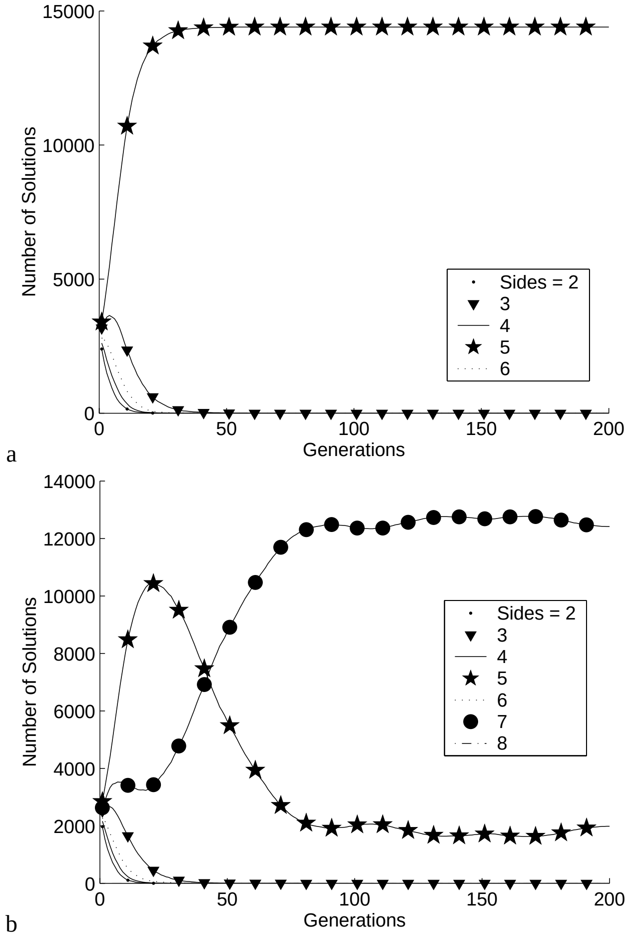

Figure 9: a: The number of sides of each shape found in each population, when symmetry is enforced, and the maximum number of sides is fixed at 6. Solutions from all 36 runs are considered, and are grouped according to which generation they appeared in. a: The number of sides of each shape found in each population, when symmetry is enforced, and the maximum number of sides is fixed at 8.

rithm variant produces results that suggest that the circle is a much better solution than the five-sided curved polygon, a fact that is known to be false.

4. DISCUSSION AND CONCLUSIONS

In this paper we have investigated the evolution of curves of constant width, of which there are known to be an infinite number of equal and optimal solutions. The most obvious of these is the circle, which has been known to Western civilization since the fifth millennium BC. However it was not until the end of the 19th century that it was discovered that the other solutions exist. Here we have documented the use of three evolutionary algorithm variants to explore this design space.

We found that when less constraints were placed on the problem—in this case, symmetry was not enforced—more of the optimal solutions were approximated. These included clear examples of approximations to the Realeaux triangle, as well as the five- and seven-sided curved polygons that are three instances of curves of constant width. Based on observation of solutions from several independent runs, a more restrictive encoding that enforced symmetry was used, which produced solutions with one order of magnitude better fitness (compare the errors of the best solutions reported in Figure 8 to those reported in Figure 6). By restricting the encoding even further—increasing the number of sides allowed in evolved polygons—the algorithm's bias is changed: in the second encoding it tends to converge on an equal number of five-sided curved polygons and circles; in the third encoding it converges almost exclusively on circles.

This work indicates that changing an encoding to increase the fitness of resulting solutions not only biases the algorithm toward solutions that fall within this restriction (which is to be expected), but also steers it away from optimal solutions that are still valid within the restricted encoding. For example in this problem, the Reauleaux triangle is an optimal solution and can be described by the symmetric encoding, but it is never output as the best solution by any of the 2×36 independent runs that enforce symmetry.

The findings of this paper are therefore two-fold. First, we have demonstrated that using a simple encoding, a genetic algorithm can discover in under one hour of real time (using a single processor) what took humankind over 7000 years to discover: that there is more than one curve of constant width. So although artificial evolution does not exceed human ingenuity in this particular design domain by producing a better design that those discovered previously, it has clearly been demonstrated that it is more adept at discovering that this problem has more than one optimal solution.

Second, we have shown that although restrictive encodings produce better evolved designs at the expense of excluding equally good designs that are not describable by the encoding, they can also, in some cases, exclude solutions that are describable by the encoding. This latter finding helps to better define how genetic encodings bias evolutionary search.

5. REFERENCES

[1] Traveling exhibit "Experiencing Mathematics" sponsored by UNESCO, at The International Congress of Mathematics Education (ICME-10), July 2004, Lyngby, Denmark.

[2] K. Balakrishnan and V. Honavar. Properties of genetic representations of neural architectures. In *Proceedings of the World Congress on Neural Networks*, pages 807–813, 1995.

[3] P. Bentley and D. Corne. *Creative Evolutionary Systems*. Morgan Kaufmann, San Francisco, CA, 2001.

[4] L. Davis. Bit-climbing, representational bias, and test suite design. In R. K. Belew and L. B. Booker, editors, *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 18–23, 1991.

[5] M. Gardner. Mathematical Games: Curves of Constant Width, one of which makes it possible to drill square holes. *Scientific American*, 208:148–156, 1963.

[6] E. M. Harrell. A direct proof of a theorem of blaschke and lebesgue. *Journal of Geometric Analysis*, 12:81–88, 2002.

[7] S.W. Mahfoud. Crowding and preselection revisited. In *Parallel Problem Solving from Nature – PPSN-1992*, pages 27–36, 1992.

[8] I. Rechenberg. *Evolutionsstrategie: Optimierung technischer System nach Prinzipien der biologischen Evolution*. Fromann-Holzboog, Stuttgart, DE, 1994.

[9] F. Reuleaux. *The Kinematics of Machinery: Outlines of a Theory of Machines*. MacMillan, London, 1876.

[10] F. Reuleaux. *The Kinematics of Machinery*. Dover Publications, New York, 1963.

[11] E. Tunis. *Wheels: A Pictorial History*. Johns Hopkins University Press, Baltimore, MD, 2002.