



SOCIETY OF ACTUARIES

Article from:

# Forecasting & Futurism

July 2011 – Issue 3

# Are Genetic Algorithms Even Applicable to Actuaries?

By Ben Wadsley

**Editor's Note:** This article was first published in the February 2011 issue of *Risks & Rewards*, the Society of Actuaries' (SOA) Investment Section newsletter. Since publication in *Risk & Rewards*, some updates have been incorporated.

Several professional fields are currently using genetic algorithms for different applications. Genetic algorithms are being used to plan airplane routes,<sup>1</sup> develop equity market bidding strategies,<sup>2</sup> point antennae on military vehicles,<sup>3</sup> optimize an iterative prisoner's dilemma strategy,<sup>4</sup> and even work toward developing artificial intelligence.<sup>5</sup> While these applications are very useful to other professions—and quite interesting to study—they don't seem to have anything to do with actuaries. As I was being introduced to the idea of genetic algorithms through the Forecasting and Futurism Section of the SOA, my main question was, "If these people are so successful in using genetic algorithms, why can't actuaries?"

This essay intends to answer the question: "Are genetic algorithms even applicable to actuaries?" by first walking through the example of "Robby the Robot" as derived from the example in Melanie Mitchell's *Complexity, A Guided Tour*.<sup>6</sup> Also, I will look at what characteristics of this application are useful and then apply those characteristics to an example based on my use of this technique to solve a life insurance asset and liability management (ALM) problem. The goal is not only to describe one use of genetic algorithms, but also to help the reader explore this thought experiment and discover how genetic algorithms can be expanded to solve many other actuarial problems.

AS I WAS BEING INTRODUCED TO THE IDEA OF GENETIC ALGORITHMS THROUGH THE FORECASTING AND FUTURISM SECTION OF THE SOA, MY MAIN QUESTION WAS, "IF THESE PEOPLE ARE SO SUCCESSFUL IN USING GENETIC ALGORITHMS, WHY CAN'T ACTUARIES?"

## What is a Genetic Algorithm?

There are many different varieties of corn—some that are wind-resistant and some that produce many ears of corn. The objective of a seed corn company is to breed the two types of corn to hopefully develop a variety of corn that both produces a lot of corn and is wind-resistant. This is the exact idea that is being leveraged with the use of genetic algorithms—except instead of corn we are breeding computer programs and investment strategies.

## "Robby the Robot"

Robby the Robot is a great example through which the steps of implementing a genetic algorithm can be learned. Robby lives in a two-dimensional 10x10 matrix that is littered with empty soda cans. In this twist on Mitchell's example, Robby's job is to pick up the soda cans from the grid with increasing efficiency, while being blind and having no initial intelligence. Below is the process used to train Robby's brain through genetic algorithms (*See illustration on page 7*):

1. Generate an initial population of solutions. This is done by creating random 'individuals' from the universe of possible solutions. An important step here is the definition of individuals; in this case they are defined as different sequences of actions Robby can take. They are defined by a string of numbers that represent several actions {12315...} where 1=bend over to pick up can, 2=move north, 3=move east, etc.
2. Calculate the 'fitness' of each individual in the current population. The fitness is defined by how well the solution performed, defined here by how efficient Robby's actions are. He receives +10 points for picking up a can, -1 point for bending over to pick up a can when there isn't a can there, and -5 points for running into a wall.
3. Select some number of individuals to become parents of the next generation. These parents are selected by using a 'fitness function' that gives the individual a higher probability of being selected if it has a higher fitness as calculated in step 2.

- Pair up the selected parents through ‘recombining’ parts of the parents to make offspring. The offspring then mutate with a given probability. Recombining can be done in many ways, but is done here by taking a portion of the string from parent #1 and a portion from parent #2, creating offspring #1, and using the unused portion of the parent strings to form offspring #2. Mutation is done by randomly changing portions of the strings. Inspired by nature, mutation maintains diversity in the population and prevents the population from converging too quickly.
- Repeat steps 2–4 for a specified number of generations, or until a sufficient fitness is achieved.

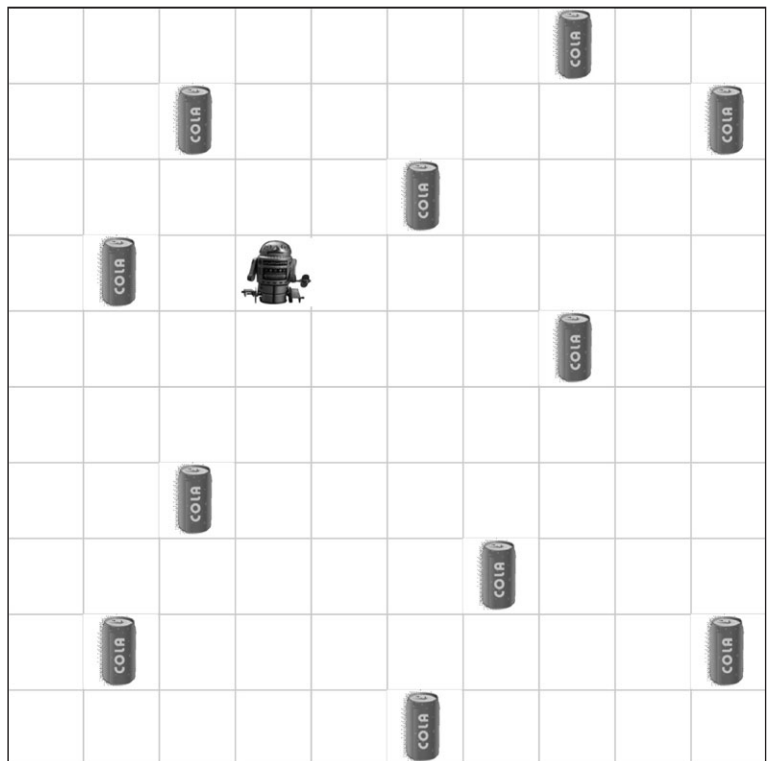
The result of this algorithm is a solution that, in Mitchell’s example, outperformed several solutions that were derived by computer scientists.

### Introduction to the Life Insurance ALM Problem

For our thought experiment, let’s consider a life insurance company that measures its economic capital requirement for interest rate risk for an in-force block using the Principal Component Analysis (PCA) as described in *Options, Futures, and Other Derivatives*.<sup>7</sup> PCA is an approach to measuring risk from groups of highly correlated variables, such as yield curve movements, into principal components that attempt to explain historical movements. Due to the orthogonal nature of the principal components, the principal components are uncorrelated, thus allowing us to measure our exposure to interest rates as:

$$f(x) = \sqrt{\sum_{n=1}^k (\text{Surplus Reduction from PC Shock } n)^2}$$

In short, the insurance company’s goal is to reduce variability in surplus for given shocks to the interest rate curve.




Since this is an in-force block, the main tool that we have to minimize variability in surplus is our choice in asset allocation. Here lies the problem—we have thousands of assets to choose from to create our portfolio. Which ones and how much of each shall we choose? In practice, we would probably develop several portfolios and test them against the capital function and implement the best one. We may use other simple optimizers. The question we need to answer here is: can we do better?

### Environments where Genetic Algorithms are Useful

There are several characteristics of problems for which genetic algorithms may be beneficial. Three of the

CONTINUED ON PAGE 8



characteristics and their applicability to our ALM problem are described below.

1. **The metric you are trying to optimize is not smooth or unimodal.** Many traditional search and optimization techniques will end up finding local minima. Consider the graph above:

If we used an optimization technique such as hill-climbing while trying to optimize the function given in the graph above, we may incorrectly identify a point as a global maximum. The basic principal of any variation of a hill-climbing algorithm is to set an initial point, test the fitness to either side of the point, move to the point with the highest fitness, and repeat until fitness cannot be improved.

In our ALM example, the fitness landscape is neither smooth nor well understood. A portion of this complexity comes from the way we measure fitness through the PCA approach and through the correlations of fixed income assets. If we were to compare two bonds with maturities one year apart, they would have similar market changes with a general move in rates, but a twist in the yield curve may cause them to act differently.

2. **The solution space is large.** If the number of solutions is finite and small, the best method is simply to try all of the options and choose the best one. Because we have thousands of assets to choose from and any dollar amount of each that can be purchased, there are infinite combinations of asset portfolios that we could try. The method that is often used is to narrow the universe of investable assets and limit the investment increments. However, there are still too many combinations to test, and if the universe is limited too far, we may have eliminated the best portfolio before beginning testing.
3. **It is a situation where good solutions tend to be made up of good building blocks.** If a portfolio of all short bonds does very well, the assumption is that short bonds are good building blocks of a great portfolio.

### Life Insurance ALM Application

In applying the genetic algorithm technique to solve this life insurance ALM problem, I used a fair number of variations from the standard procedures found in texts. It is important to remember that genetic algorithms are a tool; they should be modified to fit your needs and to develop new uses. I used the basic steps of genetic algorithms as described above and modified them to fit with this example.

As noted above, the universe of assets is immense. I limited the scope of my model to concentrate on the optimum maturity profile to manage interest rate risk. The asset choices were limited to an investment grade corporate portfolio with 30 bonds—one for each maturity year up to 30 years. Instead of choosing a random initial generation, I used a

ONCE THE BASE CODE IS TOGETHER (WHICH IS ACTUALLY QUITE EASY), THIS IS A POWERFUL TOOL THAT SHOULD BE A PART OF EVERY ACTUARY'S TOOLBOX!

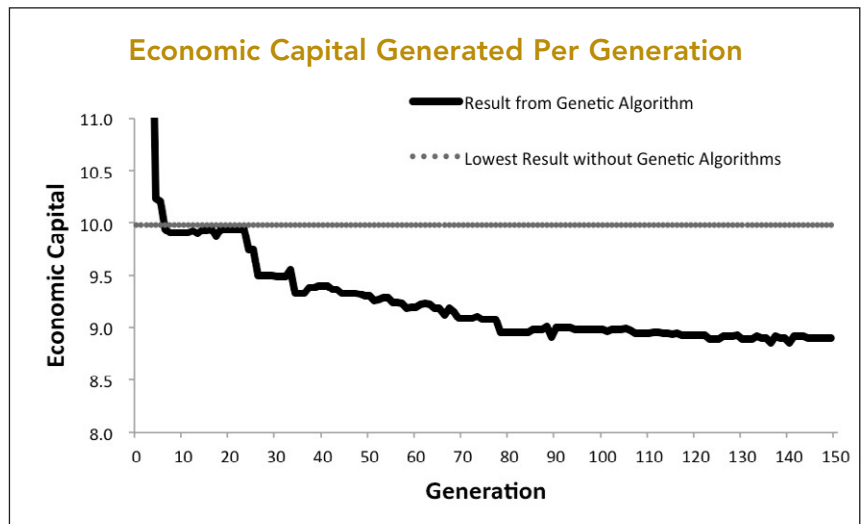
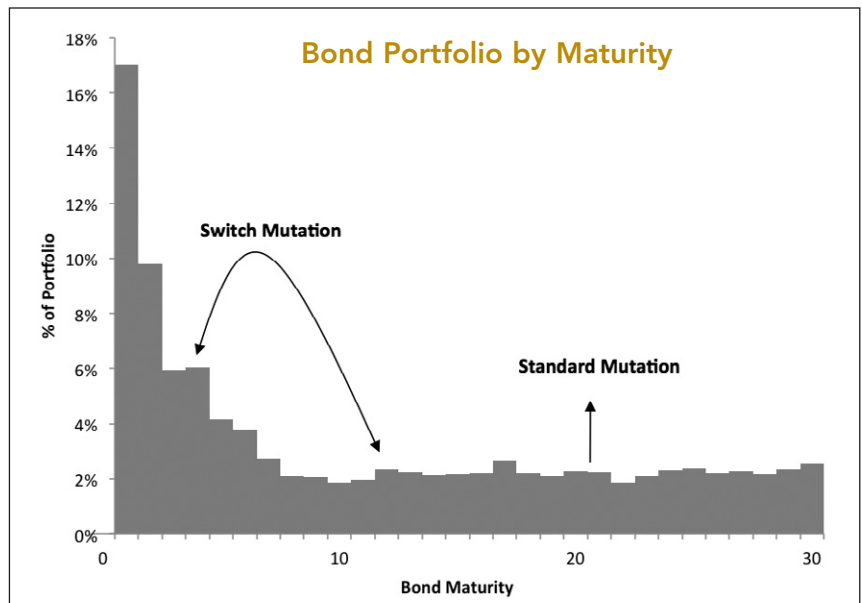
population size of 600, with each initial individual being a portfolio with the entire portfolio invested in a single bond. Rather than defining the individuals as a string, I defined the individuals as a 30-element array, with each element being the dollar amount invested in each of the 30 bonds. The fitness in my example is easily defined by the capital function described above.

Once the parent individuals were chosen, I recombined the strategies by weighted multiples of the two parents' strategies chosen with random weights. The mutation was done in two ways: first, a random maturity bucket could be set to a random weight; and, second, two maturity buckets could swap weights. This maturity bucket swapping was a great way to eliminate early convergence on local minima. After 150 generations, a suitable result was obtained. (See top graph to the right.)

The genetic algorithm solved for an investment strategy that reduced the capital by about 10 percent further than the other two methods attempted—hill climbing and trying large numbers of reasonable portfolios. Even though hill climbing was more structured, it wasn't robust enough to capture the global minimum.

The graph (bottom right) of the best investment strategy from each of three generations of the model. The model tended to learn in bursts—the best strategy was similar from generation to generation for a few iterations, and then a new portfolio that had a much better fitness emerged. For example, from generation four to generation five, the model learned to get the asset duration correct. In later generations, the model learned that a barbelled strategy worked better than a more bulletted one.

As you can see from the graph of economic capital (where less required economic capital is better), around generation 5 the genetic algorithm does about as well as our other



CONTINUED ON PAGE 10

methods, and then around generation 25 and beyond the algorithm discovers much better matched portfolios!

### Conclusion

Genetic algorithms have been used fruitfully in many other professions, and actuaries should be creative in finding ways to adapt this technique to make it a valuable tool for our profession. Not only did the genetic algorithm discover a better investment strategy, but it also gave me a structured way to solve for a result. We don't want to rely on luck to find a portfolio that does a good job of ALM matching. Many more uses for genetic algorithms are yet to be discovered. I recommend looking at examples in the resources listed in the endnotes and then programming some of the examples yourself. Once the base code is together (which is actually quite easy), this is a powerful tool that should be a part of every actuary's toolbox! ▼



Ben Wadsley

**Ben Wadsley, FSA, MAAA**, has worked for AEGON USA in Cedar Rapids, Iowa for eight years in a range of Asset/Liability Management, Investment, and Economic Capital roles. He is currently a risk manager for the Employer Solutions and Pensions division. He can be reached at [bwadsley@aegonusa.com](mailto:bwadsley@aegonusa.com).

### FOOTNOTES

- <sup>1</sup> Mitchell, M., *Introduction to Genetic Algorithms*, MIT Press, Cambridge, Mass. (1996).
- <sup>2</sup> Mitchell, M., *Introduction to Genetic Algorithms*, MIT Press, Cambridge, Mass. (1996).
- <sup>3</sup> Oh, C.K. & Hanley, B.K. *Self-Optimizing Adaptive Antenna*, 2006 NRL Review (2006). [www.nrl.navy.mil/content\\_images/06Information\(Oh\).pdf](http://www.nrl.navy.mil/content_images/06Information(Oh).pdf)
- <sup>4</sup> Mitchell, M., *Introduction to Genetic Algorithms*, MIT Press, Cambridge, Mass. (1996).
- <sup>5</sup> Mitchell, M., *Introduction to Genetic Algorithms*, MIT Press, Cambridge, Mass. (1996).
- <sup>6</sup> Mitchell, *Complexity: A Guided Tour*, Oxford University Press U.S. (2009).
- <sup>7</sup> Hull, John C., *Options, Futures, and Other Derivatives*, Prentice Hall (2006), 450–453.
- <sup>8</sup> Ho, Thomas S.Y. & Panning, William H., *Frontiers in Fixed Income Management*, Probus Pub. (1995).