Cognizant

Evolutionary AI

Go Beyond Prediction with LEAF

Today’s world presents us with extremely complex problems – and staggering amounts of data. Cognizant’s powerful new Learning Evolutionary AI Framework (LEAF) is creating opportunities for artificial intelligence to move beyond human limits to deal with this complexity. By taking advantage of massive datasets and distributed computing capacity, it generates complex models and discovers novel decision-making strategies automatically, thereby optimizing business outcomes.

August 2019

Cognizant

Create New Solutions

Harness machine creativity to design better products, enhance personalization, streamline and automate operations, improve cost containment and reduce risk.

LEAF makes it possible to discover new solutions that optimize several such objectives at once – thus augmenting human decision making. It allows businesses to rapidly apply Evolutionary AI™ to scale and optimize their current solutions.

> LEAF builds a model of the environment, and then evaluates millions of candidates with the model before they are actually deployed in the environment, vastly accelerating its arrival at optimal solutions.

Problem-Solving Beyond the Scope of Humans

As business problems become more complex, so do the processes necessary to solve them. Japan's newest Maglev trains, for instance, have tested at 375 mph – an astonishing speed for ground-based public transportation. Yet, over the years, the designs for the front of bullet trains suggest engineers' ongoing battle with speed's greatest enemy: air resistance.

When they accelerate, trains contend with rising wind pressure and dangerous resistance entering tunnels. The variables involved in modeling air dynamics present an incredibly complex equation. Speed. Temperature. The size of a tunnel's entrance and its length. The ability to expel volumes of air through vents. No engineer can work out all the possible permutations of problems such as these. Even teams can struggle, over years, and at enormous cost to come up with solutions.

As it happens, engineers drew from nature to solve these complex problems – using principles derived from evolution. Evolutionary AI™ creatively designed the shape of the bullet trains in Japan.

Variables, by their very nature, continue to change, and often new variables are introduced, making the systems more complex.

And it's not just engineering. Complexity arises in:

- Researching the most effective cancer treatments
- Developing strategies for financial quantitative trading to maximize returns while minimizing risk
- Analyzing cyber-agriculture techniques to optimize flavor and increase yield
- Optimizing website design and architecture for the highest engagement and sales
- Evaluating in-store product displays to determine the best arrangements to boost sales

Cognizant LEAF brings evolutionary computation to the most complex business problems, and thereby allows solving them with machine creativity.

LEAF maps billions of outcomes, using extremely efficient evolutionary computation methods to recommend optimal decisions.

AI's New Frontier is Evolutionary Learning

Cognizant LEAF is a unique evolutionary AI platform that uses the most advanced evolutionary algorithms, deep learning and distributed computation technology to tackle complex problems and derive meaningful results. Together, these technologies enable businesses to solve the most challenging issues.

Today's AI-based machine learning algorithms help us predict outcomes by modeling how a fixed set of variables affects performance. In contrast, LEAF is much more than a prediction engine. It moves AI towards machine creativity, utilizing a virtual model of the world to examine new opportunities and solutions that have never been seen before. It also provides a process for dealing with an ever-changing deployment environment, even exploring how its own actions affect that environment. It builds a model of the environment, and then evaluates millions of candidates with the model before they are actually deployed in the environment, vastly accelerating its arrival at optimal solutions.

Case in Point: Human Health Solutions

Imagine having to test a complex range of compositions of a promising drug. Now imagine having to test its effects on patients whose biological makeup is as varied as the human genome. Mapping genomic permutations to optimal treatments depends on examining thousands of outcomes from a range of potential treatments. Such trials entail risk and are costly.

In contrast, using its model of the environment, Evolutionary AI can evaluate numerous compounds on numerous genomic variations, to find good, more effective treatments before they are ever tested in a clinical trial. Using Evolutionary AI, LEAF provides not just an avenue to explore: it offers millions of avenues to discover novel and surprising solutions – towards personalized medicine and finding a cure.

How It Works: The DNA of Better Decision-Making

Inspired by biological evolution, Cognizant LEAF can perform evolutionary computation on many different representations, including neural networks, rule sets and parameter vectors to solve specific problems. LEAF begins by generating a diverse set of solutions to form the initial population (Figure 1). It assigns a fitness score to each member of this population based on how well it performs on the problem.

The superior candidates produce offspring, obtained by recombining their components and mutating them. The offspring then replace the inferior members of the population.

Even if the initial population of randomly generated candidates all perform poorly, some candidates will prove to be "less bad" than others and they will form the parents for the next generation. Over multiple generations, the good components will become more prevalent in the population, and eventually candidates will be discovered that solve the problem.

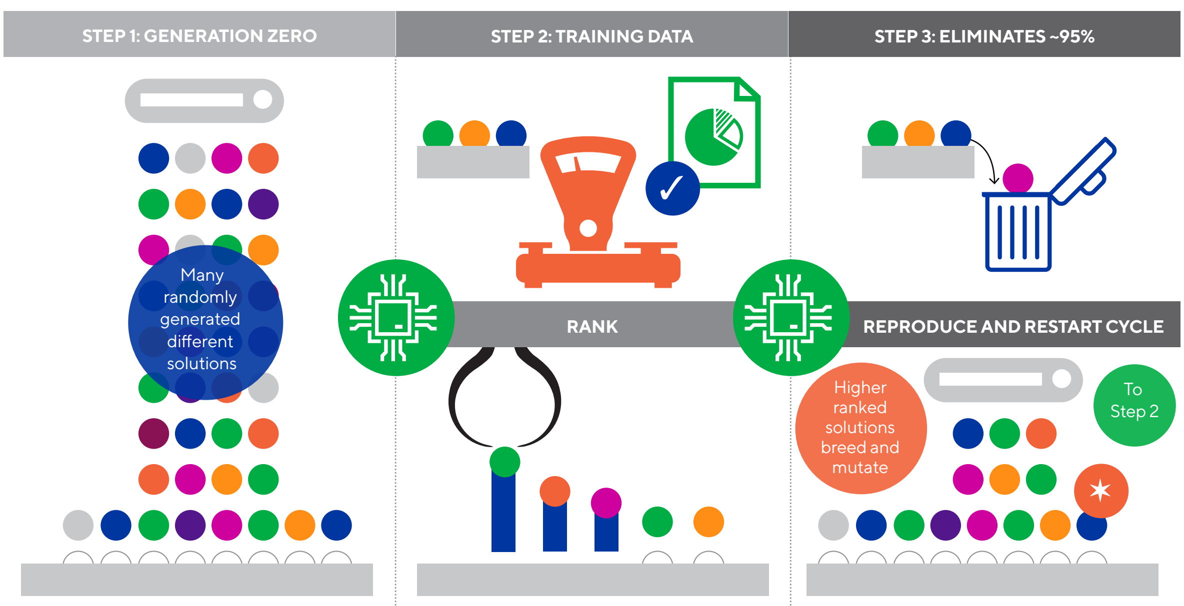

Figure 1. Evolutionary AI – A Revolution in Processing – In step 1, the system begins by randomly generating a population of candidate solutions. These candidates are evaluated and rank-ordered in step 2. In step 3, the least useful candidates are discarded , and new candidates generated from the most useful candidates through recombination and mutation. Steps 2 and 3 are then repeated rapidly over and over again until candidates that solve the problem sufficiently well are found.

What distinguishes LEAF from Predictive AI is its capacity to discover solutions that have never existed before and have never been evaluated before. In this sense, it is the next step in advancing AI capability and one that allows us to truly harness machine creativity.

Case in Point: Quant-Driven Investing

By 2017, less than 10% of stock trades were placed by discretionary traders. The other 90%? High-volume trading based on algorithms or index-driven and pattern-based strategies. "Quants" rule the field – and investors are always looking for an edge.

LEAF presents a next step for investors to develop novel trading strategies. Based on vast amounts of data on economic history, current performance, market indicators, aggregate financial statement data and policy options such as the actions of the U.S. Federal Reserve Bank, Evolutionary AI can be used to discover interactions between all of them and take full advantage of that information in the market.

Cognizant Evolutionary AI Offerings

LEAF currently consists of two offerings: the first one discovers optimal strategies for business decision making; the second one optimizes current machine learning-based business solutions.

Generating Optimized Prescriptions

Evolutionary AI Business Optimization is an engine that operates a continuous learning and optimization lifecycle loop. Most businesses are looking to build predictive models that human decision makers can use in order to come up with the best decisions themselves. Our AI optimization service goes well beyond predictive models, suggesting actual decisions that maximize the outcomes.

To build a prescriptive outcome-based system (Figure 2), the first step is to look at the customer data or decision-making process and identify what constitutes a **context**, what constitutes **actions** they take within various contexts (i.e., the degrees of freedom they can exercise in decision-making) and determine how they will measure the **outcome** of the decisions. Often, the customer has data capturing the context and action pairing that can be aligned with the outcomes of the past. Once that data is captured, a predictive model, for example a surrogate of the real world, can be built.

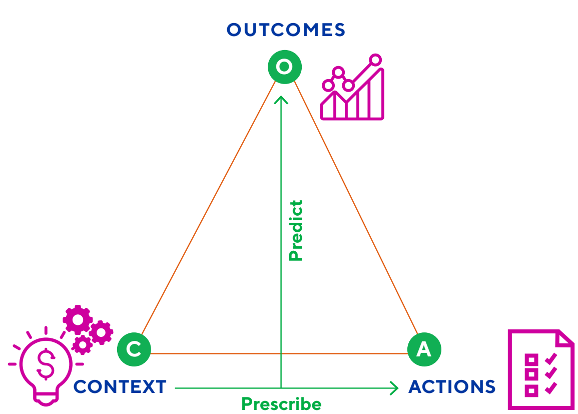

Figure 2. LEAF Optimization Service – The **predictor** is trained with historical data to estimate what outcomes result when given actions are executed in a given context. The **prescriptor** is evolved to suggest optimal actions in a given context. Together this service makes it possible to optimize decision making in many business problems.

Surrogate modeling plays a key role in the prescriptive process (Figure 3). Using real-world data, we can create virtual models of the problem, otherwise known as *digital twins*. These digital twins act as surrogates through which optimal solutions can be discovered. These solutions can then be applied back to the real world. More specifically, once a surrogate model is created, it is interrogated to find out how good the prescriptions are. LEAF thus uses the surrogate to build a **prescriptor** to map any given context to actions that will maximize outcomes.

Consider a retailer looking to place products on shelves in a way that maximizes revenue and minimizes inventory cost – all in the shortest amount of time with the least amount of resources. Using historic point of sale and shopper data, a surrogate model can be created to predict how well the different product-placement programs – *planagrams* – would work in different contexts. A **prescriptor** is evolved to generate candidate planagrams, to be tested against the surrogate, until the surrogate model predicts that they perform sufficiently well. These planagrams are then implemented in each store. The replacements create more data for training the prediction further, and the cycle repeats until the desired real-world outcomes are achieved.

While millions of planagrams can be quickly created and evaluated with the surrogate, only a small number of them need to be tested at the store.

Our Evolutionary AI Business Optimization service is most valuable when:

- Sufficient and consistent historical data of quality measurements on different actions performed in various contexts is available to bootstrap the surrogate model.
- Deployment allows for reasonable iterations of the latest prescriptors and resulting data collection for improving surrogate models.
- It is replacing decision loops currently operated by humans or rigid pre-programmed systems.

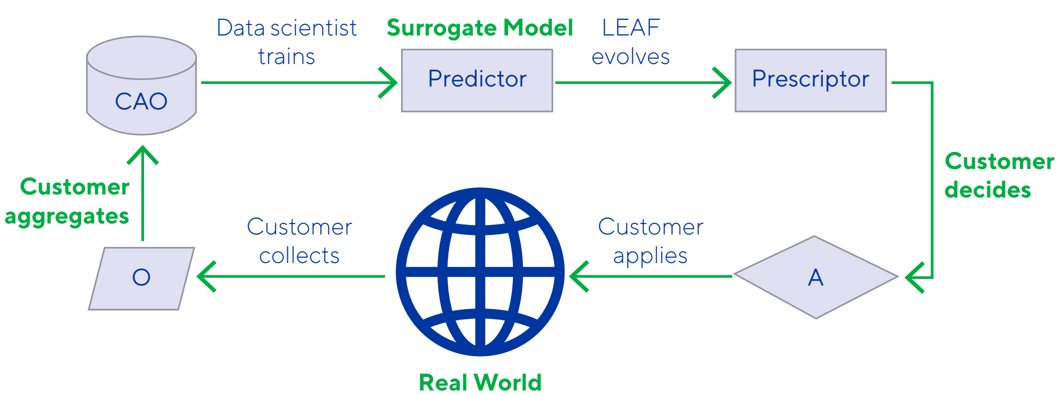

Figure 3. Prescriptive Continuous Learning and Optimization – In the above, the customer aggregates historical context, action and outcome (**CAO**) data from the real world. This data is then used to train the **predictor**, which acts as a surrogate model of the real world. The surrogate model in turn makes it possible to evolve a **prescriptor** that recommends optimal actions. A human decision-maker receives the suggested actions and may decide to implement them or modify them. When they are applied to the real world, more data is collected which makes it possible to improve the **predictor**, and consequently the **prescriptor**. This loop repeats until satisfactory outcomes are obtained. Note that most of the discovery is done with the **surrogate**, not the real world, thereby speeding up the process and making it safe as well as economical.

Generating Optimized Deep Learning Architectures

Cognizant Evolutionary AI Model Optimization service, or evolutionary AutoML, allows our clients to apply the power of Evolutionary AI to the optimization of deep learning architectures. This service is useful for four reasons: it develops models with high performance and accuracy, reduces the need for expert in-house talent, and it extends the model to a wider range of applications, including those where little data exists and when only limited computing and memory is available.

First, evolutionary AutoML makes it possible to customize deep learning for specific tasks, such as object recognition, language interfaces and sentiment analysis. It does so by discovering the principles of the domain and encoding them into the architecture, for instance by creating multiple paths, allowing several hypotheses to be evaluated at the same time. These architectures often perform better even than the best human designs.

The second goal of evolutionary AutoML is to democratize AI by making it possible for anyone to use it. Current AutoML systems, such as those of Google and Yelp, are limited in both scope and transparency. They focus on optimizing superficial features of the architecture, and they do not provide insight to the user. In contrast, Evolutionary AutoML optimizes the entire architecture, leading to insights into the domain. It makes cutting-edge architectures available to solve business problems even without extensive expertise in DNNs.

Third, evolutionary AutoML makes it possible to apply deep learning to solve tasks where millions of training examples and data are not available.

Case in Point: Revolutionizing Retail Sales

Through people's behavior online and in stores, retailers glean an enormous amount of data on their wants, needs, and activities. What designs work best to retain users, to increase revenue, to shift buyers to certain products? What pricing strategies are optimal and what makes users click and buy?

Historically, predictive machine learning has been used to test hypothetical answers to these questions, created by human experts. LEAF generates those hypotheses automatically, evaluates millions of them and often discovers solutions that the experts miss. For instance, it can be used to make recommendations for e-commerce site designs and in-store product displays that result in maximizing revenue.

Evolutionary AutoML can evolve an architecture that combines a small amount of data with datasets of other related tasks, such as other vision or other language tasks. By learning all these at once, each task is acquired more effectively than learning them in isolation. A customized architecture that evolves automatically is crucial in taking advantage of this process. This capability is highly useful in an environment where multiple related tasks exist, such as in recognizing handwritten characters in multiple different alphabets, recognizing spoken words in multiple different languages, or diagnosing multiple different diseases from X-ray images.

And fourth, many commercial applications of deep learning need to be run on smartphones or similar-sized processors in wearables, vehicles, appliances and potentially even toys. However, the hundreds of millions of parameters of modern deep neural networks (DNNs) cannot fit into the few gigabytes of RAM in most such devices. Evolutionary AutoML minimizes the complexity or size of the network, while simultaneously maximizing its performance.

Cognizant's Evolutionary AI Model Optimization service is an ideal approach for businesses when:

- Improvement to existing state of the art is desired.
- The cost of design and development of models needs to be reduced.
- There is insufficient domain-specific data.
- Other objectives beyond accuracy are involved, such as size, speed, or fit to a hardware architecture.

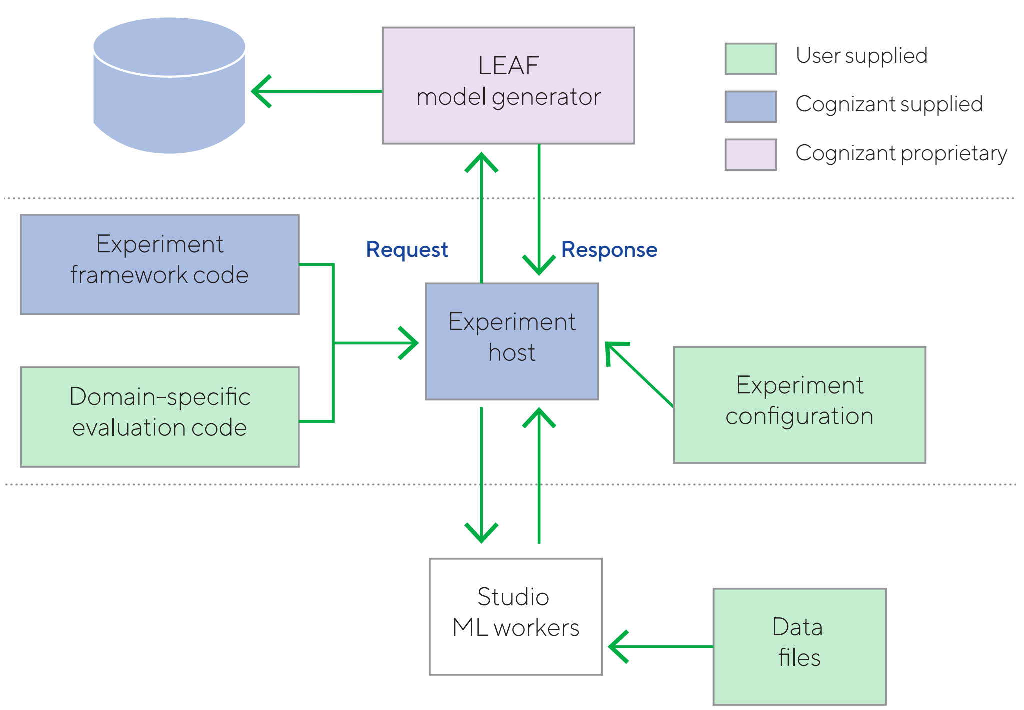

Figure 4. Evolutionary AutoML Workflow – The customer supplies a data set that can remain proprietary, as well as details of how the evolved networks should be evaluated on it and the goals of the experiment. The **experiment host** queries the **LEAF model generator** for candidate architectures and evaluates each one in parallel using studio.ml to distribute the evaluation jobs in the cloud resources. In this manner, the data, compute and AutoML system remain separate, safe and efficient.

Beyond Empirical Research and Predictive AI

Empirical research has historically been the most comprehensive and reliable method for investigating problems with multiple variables. Algorithmic analysis and AI currently focus on predictive modeling – asking what will happen if we change this variable or that one.

However, we are starting to reach the limits of these two approaches. It is difficult to design experiments that take into account the interactions of multiple variables when the subject is as complex as fluid dynamics, airflows at high speeds or the human genome. It is also difficult to model the domain when the number of variables is in the hundreds and possible permutations number in the millions or billions.

Cognizant LEAF provides the technology to make meaningful, measurable gains in such domains. Examples include:

- Recognizing and flagging toxic content on social media.
- Optimizing environmental conditions for growing crops, including demonstrable gains in increasing desirable characteristics of flavor and mass.
- Broadening testing regimens for complex bioengineering research, while lowering research and development costs as well as reducing patient risk.
- Analyzing a complex array of risk variables in insurance underwriting – moving beyond empirical actuarial models to identifying new and creative ways to design policies.
- Identifying lung disease in radiographic classifications.

LEAF optimizes complex business systems, moving beyond empirical research and predictive modeling, by allowing the exploration of millions of alternatives.

Cognizant LEAF moves beyond empirical research and predictive modeling, by allowing the exploration of millions of alternatives to optimize complex business systems. Our LEAF offerings discover entirely new designs and solutions and can be applied to problems across industries, including financial technology and investing, risk management, digital commerce and retail stores, cyber-agriculture, medical treatment, pharmaceuticals and bioengineering.

As this technology evolves, we expect to see much of the ad-hoc manual decision making and design be replaced by such automated systems that will empower managers in decision making, scientists in research and engineers in development.

LEAF is enabling this next step in advanced AI capabilities. Learn more by visiting www.Cognizant.com/EvolutionaryAI.

Cognizant

World Headquarters
500 Frank W. Burr Blvd.
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Fax: +1 201 801 0243
Toll Free: +1 888 937 3277

European Headquarters
1 Kingdom Street
Paddington Central
London W2 6BD England
Phone: +44 (0) 20 7297 7600
Fax: +44 (0) 20 7121 0102

India Operations Headquarters
#5/535 Old Mahabalipuram Road
Okkiyam Pettai, Thoraipakkam
Chennai, 600 096 India
Phone: +91 (0) 44 4209 6000
Fax: +91 (0) 44 4209 6060