


# **SUBSTRATE AI'S TECHNOLOGY**

substrate**AI**

## ► What makes our AI technology better than the alternatives?

If you ever asked Siri, Alexa or Google Assistant things like "find me the toll-free phone number of this or that institution", it is likely the result was frustrating. This is because current AI technology is only capable of dealing with one task at a time. In our example, we are asking it for two sequential tasks: the first one to find the phone number of institution X, and the second one that it is also to be toll-free. Substrate AI's technology can do it.

What makes Substrate AI's technology different, and a breakthrough in the AI industry, is that it is able to learn faster and be more adaptable than other approaches. Importantly, it is able to be trained for real-world applications that you cannot simulate by using a teacher to guide our agent learning to make decisions about tasks divided between numerous objectives. This places Substrate AI technology in a middle ground between what is known as Weak AI and Strong AI.



Our Machine Learning (ML) is a **biologically inspired technology**. It uses advanced patent pending Reinforcement Learning (RL) technologies that enables RL to be used in a diverse set of real-world applications. From sample efficiency, transfer learning, risk control and more, our use of biologically based algorithms has resulted in an agent that has

demonstrated flexible control across multiple testing paradigms (see our white paper, [Integrated Multi-Task Agent Architecture with Affect-Like Guided Behavior](#), J. Worth and M. Sei, 2001 Annual International Conference on Biologically Inspired Cognitive Architectures).

Designed for decision support applications, our technology adjusts its behavior to reduce risky behavior by using emotions, learning to handle multiple objectives and to continually learn based on experience. Part of our architecture contains an **Integrated Multi-Task Model (IMT)** which uses the environment context and changes in reward over time to learn goals, attentional features, and action affordances to regulate behavior and build enhanced reference frames. Experiments performed on the Go game shows that our agent successfully demonstrated improved gaming performance and cumulative Q-value results over game episodes, compared to the vanilla Q agent. In addition to learning, our patent pending RL training algorithm supports agent pre-training using expert examples and allows our agents to learn to interact with non-simulatable environments while minimizing training bias. This significantly opens the doors to many real-world applications where it is not possible to have an agent act naively such as in medical applications.

For the past year we have been testing and refining a cloud native multi-agent and task agent that can combine multiple models and their respective modalities to perform ever more complex functions. This system employs advanced skill extraction technology that can learn transferable skills from experts in a semi-supervised way. Collectively, **these technologies and their future enhancements provide Substrate AI a competitive edge in building cloud scalable generalizable agents capable of continual and multi-task learning.**

## ► **Current deployment:**

### **1. Financial market investments**

We have tried our RL technologies in the real-world, especially in financial markets, as part of our fintech vertical. They have been under test for three years in real market conditions and optimized for professional investors. Financial markets are a perfect field test for our technology since they are non-stationary environments by nature. Our algorithm has

successfully adapted to changing market conditions and managed to balance risk/reward to achieve highest gains over time while minimizing drawdown.

The trading workflow includes:

1. Selecting securities to trade in for Long positions.
2. Computing Long/Flat/Cash portfolio allocation.
3. Rebalance portfolio on weekly or longer basis depending on market conditions. The adaptive allocation management allows the portfolio to distribute resources into long, flat (fixed income ETFs) and cash.

Our testing has shown that the adaptive allocation model significantly improves long term performance of the portfolio while controlling execution costs and complexity.

## RESULTS

- Our models have averaged in excess of **40% annual returns in live environments over the past three years.**
- They have resulted in multiple models being **in the top 5% in our reference platform Collective2 over the same period of time** (see below chart).
- Our algorithm has demonstrated a learning performance that is in excess of 700% over standard Q Learning based agent [1].

[1] What is q-learning? Q-learning is an off policy reinforcement learning algorithm that seeks to find the best action to take given the current state. It's considered off-policy because the q-learning function learns from actions that are outside the current policy, like taking random actions, and therefore a policy isn't needed. More specifically, q-learning seeks to learn a policy that maximizes the total reward.

What's 'Q'? The 'q' in q-learning stands for quality. Quality in this case represents how useful a given action is in gaining some future reward.

<b>38.9%</b> Annual Return (Compounded)	<b>(24.6%)</b> Max Drawdown	<b>281</b> Num Trades	<b>69.8%</b> Win Trades	<b>2.5</b> :1 Profit Factor	<b>70.4%</b> Win Months
--	--------------------------------	--------------------------	----------------------------	--------------------------------	----------------------------

Hypothetical Monthly Returns (includes system fee and Typical Broker commissions and fees)

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	YTD
2019												+3.3%	+3.3%
2020	+2.2%	(12.9%)	+7.2%	+16.9%	+11.0%	+2.1%	+2.2%	+3.1%	(1.6%)	(1.9%)	+9.2%	+5.3%	+47.9%
2021	(3.6%)	+14.7%	+6.3%	+7.8%	+5.5%	+2.2%	+0.1%	+0.6%	(3.4%)	+4.6%	(2.8%)	+3.5%	+40.1%
2022	(2.5%)	(0.2%)											(2.7%)

The above chart tracks the US ETF model managed by Substrate AI's financial model. This model has ranked consistently in the top 5% on Collective2 platform and yielded trading performance greater than the benchmark S&P 500.

## Statistics

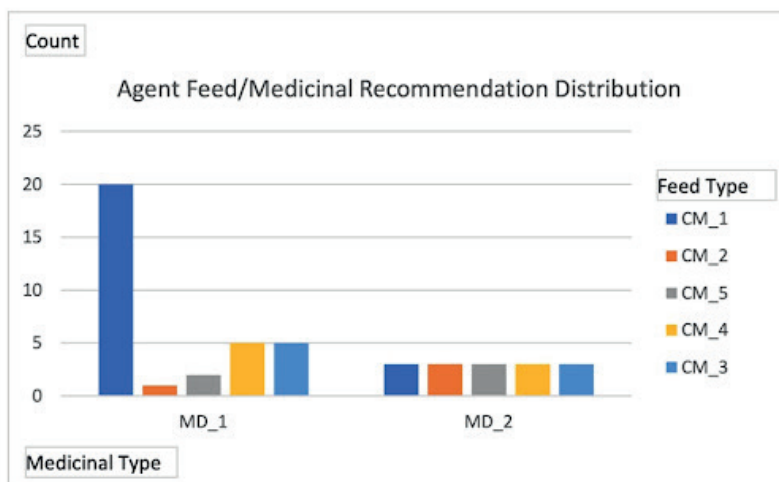
Strategy began	12/2/2019
Suggested Minimum Cap	\$15,000
Strategy Age (days)	811.53
Age	27 months ago
What it trades	Stocks
# Trades	281
# Profitable	196
% Profitable	69.80%
Avg trade duration	26.9 days
Max peak-to-valley drawdown	24.61%
drawdown period	Feb 20, 2020 - March 07, 2020
Annual Return (Compounded)	38.9%
Avg win	\$473.85
Avg loss	\$463.38

## 2. Milk producing livestock

We have tested our technology in milk producing livestock too. Milk producing animals come in numerous species. Challenges from data collection and execution of actions that optimize milk production and quality over time. An example of this is seen in the management of goat and sheep milk vs cow production management. In cow production each animal is highly instrumented and considered their own factory with significant data automation. In

smaller milk producing animals this is not possible. It presented a number of challenges which included collection, availability and quality of data. Additional issues included the lack of standardized approaches to measure impact and performance of feed supplements on animals.

**Our approach included creation of data collection automation through web scraping of laboratory results and mobile app to collect farm only available data.** Additionally, the mobile application notifies farmers and veterinarians of herd issues and management recommendations. In a recent trial, with six months of data collected twice a week, the agent was designed and pre-trained to behave similarly to our veterinarian expert. In the below chart the model had to select from 5 different feeds and 2 medicinal blends to be given to a milk producing herd. The agent correctly selected the CM\_1 and MD\_1 which was the veterinarian's choice based on the farm's milk lab data. This test was performed over 40 runs to verify the robustness of the result. The chart displays the resultant counts from these tests demonstrating that the agent was able to consistently reproduce the veterinarian's recommendation very robustly. In addition, our solution improved milk production and performance per animal. **In our dairy farm project, average milk production per goat was at 338 liters. With our solution, it increased to 372 liters (+10%). The average income per goat also increased from 145€ to 178€ (+22.8%).**



AI model trial results showing robust pre-training of the agent to select veterinarian's recommended CM\_1 and MD\_1 feed and medicinal product

Our milk producing livestock project was initially financed by Spain's ICEX - a publicly owned business-oriented entity of the Spanish Ministry of Industry and Trade. **Our commercial solution PAM (Prediction Animal Management)** is now marketed via Substrate AI's subsidiary Boalvet and will expand in European markets this year and next year.

## **Business strategy & Partnership program**

Our growth model and business strategy relies on collaboration and development agreements. Either we participate in the shareholding of our partners or we develop a project together to solve their most pressing business issues. Our business strategy allows us to have solid fundamentals to support the interim between data collection and the development of agents for specific solutions. Most companies face problems that have to do with data which is not only unstructured but must be generated in the future so agents must be trained with data harvested on the fly. This sometimes requires more than one or two years for the data sample to be significant. This means that many AI companies cannot withstand the passage of time without having revenues to support data collection and agent training, so the projects wind up not being developed.

Our partnership program provides us with experts in top universities (Rensselear Polytechnic Institute in New York and Aivancity School for Technology, Business and Society in Paris) and an anchor client to develop solutions that enable Artificial Intelligence to have a predefined growth model. We have adopted a **'lighthouse' strategy**. We work with companies leaders in their field such as Canadian Solar, Ornavera (cannabis for medical purposes), Delaval (milk production), Ibanfirst (payment industry), Mapfre (insurance), Exante (trading and financial investment) and many more. Our solutions are developed and implemented in the anchor client's business environment first. We will then use those customers to scale the solutions globally, as many of them will sell the solutions to third parties as part of the partnership agreement. This strategy avoids us having to use our own resources to generate sales and, on top of that, we are certain that the solution is being used successfully.

## ► Appendix

### Patents and Patent Applications

Substrate AI is in the process of patenting part of its scientific findings in the U.S helped by the law firm Cooley, as described below: above-mentioned patents are based on thorough research led by the CTO, Bren Worth, whose basis to develop these solutions is the ability of a machine learning agent to use the environment context to learn goals, attentional features, and action affordances to regulate behavior and build enhanced reference frames, via the reinforcement learning algorithm. The model utilizes functional and conceptual features shared by humans and animals to support reasoning and continual learning.

The patents pending apply for:

Method to automate the management of intensively managed milk producing livestock to produce customized product depending on end-use using machine learning

Method to adaptively optimize feed blend and medicinal selection using machine learning to optimize animal reproduction rate

Method To Adaptively Optimize Feed Blend And Medicinal Selection Using Machine Learning To Optimize Animal Milk Production And Health

Method To Automatically Perform Temporal Abstraction In Reinforcement Learning Options

Method To Automatically Tune Reinforcement Learning Hyperparameters Using

Hyperparameter Models That Use Sharpe Ratio Reward Signal To Optimize For Risk Adjusted Returns By The Agent Over Time



Method To Learn Repertoire Of Behavior For Reinforcement Learning Agent Using Options

Method To Create Cognitively Inspired Hierarchical Agent That Includes Models That Select Subgoals And Shape Agent Attention And Action To Be Used By An Enhanced Experiential Model For Agent Action Execution

Method To Pretrain Reinforcement Learning Agent From Tabular Data And Imperfect Expert Action Examples When Simulated Environment Is Not Available

Method To Create Reinforcement Learning Imagination System Through Synthetic State-Action Transitions And Their Associated Reward Signals And Facilitate Agent Planning And Creation Of Option Candidates

Method To Detect And Automatically Adjust Reinforcement Learning Agent Behavior Based On Multiple Objective Signal That Includes Bias Signal Value

Method To Extract Options From Demonstration Experience And Initialize Agents With Learned Options To Support Transfer Learning From Demonstrator

Method To Reduce Model Data For Inclusion Into Dqn Approximator By Building Empty Value Statistics Estimation

substrate **AI**

SPAIN (CENTRAL)

C/ Colón, 4 - 5ª A  
46004 Valencia

EEUU

706 Gunsmoke Dr. Bailey  
Colorado 80421

PORTUGAL

Rua Pedro Nunes, 11 4DT  
1050-169 Lisbon

[info@substrate.ai](mailto:info@substrate.ai)

[www.substrate.ai](http://www.substrate.ai)