



Neurala Overview 6/2017

Neurala

- *The vendor of the Neurala Brain: 10 yrs. old software company*
- Team: 35 people, strong expertise in Neural Networks/Deep Learning
- Early and pioneering IP (GPU computing, 2006) considerable head-start in Neural Network deployment on GPUs
- Experience in *building and commercializing* AI
- \$25M of R&D and private investment
Confidential Information Redacted

The Neurala Brain



What NASA asked us for

- Give us a 'brain' for a Mars Rover
- No Wi-Fi connection, no GPS
- Fit on a 'brain' on a single GPU
- Use only a passive camera & IMU
- Learn 'as you go': learn what's in the environment, and update robot perceptions and maps

What we did

- Lifelong Deep Learning: L-DNN
- Emulates 'whole brains' in software
- No need of active sensors: inexpensive cameras are OK
- Beyond object recognition: collision avoidance, navigation, planning



Deploying NASA technology on everyday devices



\$ 10.0b+

CONSUMER ELECTRONICS

- Learn, recognize and react to new classes of objects on-the-fly after deployment
- Engaging and interacting devices
- Location awareness
- Satisfy privacy concerns (no cloud)



\$ 12.0b+

AUTONOMOUS VEHICLES

- Very fast on device intelligence
- Learn on-the-fly based on new use cases (e.g., misclassifications)
- Fast processing



\$ 11.0b+

DRONES

- Learn very specific object of interests (e.g., rust on a cell tower, safe areas to land, anomalies)
- Learn to follow safely
- Avoid (or seek) obstacles



TBD

SMART DEVICES

- Learn something new (e.g., a person to find, a new product to pick, an abandoned backpack) often without cloud support
- Inspect items in production
- Recognize, follow and track

Neural AI, LLC, Exhibit 2005

Page 2005 - 4

Our solution: Brains For Bots SDK™

Brains For Bots™ is a Software Development Kit for building smart machines AI architectures. Based on Neurala's Lifelong DNN (L-DNN™), the AI that learns after deployment, It sits between the Customer Application and the device OS to enable:

- Continuous learning at the Edge
- Learning, recognition, tracking, and finding objects in the video stream
- Fast Edge execution speed and optimal use of CPU and GPU (if available)
- Software modularity
- Wide compatibility with different processors

Customer Application

Code Example

Interface Layer



Algorithm Abstraction Layer

Core Abstraction Layer

OpenCV, NVIDIA VX, Caffe, etc.



NVIDIA, x64, ARMv




Unique differentiation (not available in other frameworks):

With Lifelong Deep Neural Networks (L-DNN) technology, Neurala's software is the 1st of its kind to learn in real-time, on the device, like humans do.

- *No need for long training & servers*
- *Does not require Cloud connection*



And we are *fast*: 51% faster than TensorFlow



- Neurala inference time is 51% faster than TensorFlow
- Using just 100MB more memory (basically free)

CALCULATION using GPU.



..... and 600% faster and lighter than Caffe



- 60x faster than Caffe2
- Using a little less memory (basically free)

CALCULATION using GPU.

.... and faster than Microsoft/CNTK



- 186x faster on GPU
- Using more memory (basically free)


CALCULATION using GPU. Microsoft/CNTK is so slow because it saves memory by allocating chunks and hence must ask for memory again and again. But, GPU processors like the NVIDIA Jetson TX1 have plenty of available memory and does not need to be managed tightly.

Growing & early IP portfolio

Neurala IP dates back to 2006, where we patented a method for running difference and differential equations on GPUs (US [8,648,867](#)) Since then, Neurala has grown a patent, trade secret, and trademark portfolio ranging from using GPUs for Deep Learning and Neural Networks, sensory (vision, audition, and multimodal) perception, spatial navigation, collision avoidance, and user interface and smart tagging to build and deploy brains for intelligent devices.

Patent Portfolio:

- 3 US Utility Patents
- 5 Foreign Design Patents
- 13 Pending Patent Applications:
 - 2 US provisional applications
 - 1 PCT application
 - 1 Canadian patent application
 - 3 European patent applications
 - 1 Korean patent application
 - 5 US utility patent applications.



Brains For Bots SDK™

Enables users to embed a custom Neurala Brain for deployment on a target device with C++ API

Neural AI, LLC, Exhibit 2005

Page 2005 - 11

Example: selfie dronie



- 1st DNN-powered AI for drones
- Early 2016
- App store and Google Play

Example: obstacle detection for automotive



- US and EU OEM and Tier 1 suppliers
- Use case: obstacle avoidance for self-driving vehicles, drive assist, ground truth for large datasets
- Requirements: speed, lifelong learning for new use cases (e.g., recognize driver's home, parking, etc)

Example: inspections with drones



- Customers: drone manufacturers, telecom & energy companies, logistics
- Use case: inspection of telecommunication towers, solar panels, roofs, power lines
- Requirements: onboard intelligence, speed, accuracy, refine learning over time

Example: low-cost smart toy



- Low-cost smart toy producers
- Other customers: Smart device manufacturers, camera manufacturers
- Use case: interactive toys, local intelligence without cloud support

In collaboration with Air Shepherd and Lindbergh Foundation




- Drone for anti-poaching
- Persistent surveillance
- Analyze in real-time RGB and IR data
- Requirements: speed, high FPS, refine learning over time

Neurala & nVidia


Core premise of entire company: *we could run AI on GPUS*



- Started programming GPUs in **2 BC**
(**BC** = **B**efore **C**UDA – 1st release on 6/2007)
- Solved a key problem: hardware limitation made AI deployment impractical
- U.S. Patents 8,648,867 and 9,189,828 (2006) while solving this problem



This chart made us believe Neurala could be a business



If our customers want high performance, we tell them 'get nVidia'



CONSUMER ELECTRONICS

- Advanced robots/drones that self-localize, perceive, and learn at the **EDGE**

\$ 10.0b+

Confidential Information Redacted



AUTONOMOUS VEHICLES

- Automotive use cases where learning at the Edge and fast processing unlock new use cases

\$ 12.0b+

Confidential Information Redacted




DRONES

- Professional drones for inspection, logistics, ...

\$ 11.0b+

Confidential Information Redacted



SMART DEVICES

- Learn something new (e.g., a person to find, a new product to pick, a bag that has been abandoned) often without cloud support
- Recognize, follow and track

TBD


Confidential Information Redacted

Neural AI, LLC, Exhibit 2005


Page 2005 - 19

Neurala: shifting customers from CPU $\Rightarrow$ GPU

- Remember: more processing power = more speed & accuracy (these are tradeoffs)
- GPU systems offer significant performance advantages over CPU systems (using about the same memory)
- By using GPU, Neurala unlocks high-valued applications where inference & learning speed and accuracy are mission critical.
- These include:
 - Automotive: high FPS & precision
 - Drone inspections: speed and accuracy
 - Industrial applications: very accurate smart cameras
 - Personal robots with Edge intelligence & learning



We can work together in at least 3 ways



Company info

Confidential Information Redacted



Confidential Information Redacted



Neurala Inc.
8 St. Mary's Street, Suite 613
Boston, MA 02215

info@neurala.com
sales@neurala.com
tel. +1.671.418.6161 x 701