ARTIFICIAL INTELLIGENCE CUTTING THROUGH THE HYPE

AHMADEUS

Technology Boutique

BACKGROUND


Goldman
Sachs


Morgan Stanley


CREDIT SUISSE


BARCLAYS


citi


HEARST corporation


NOVARTIS


verizon


WEISS
MULTI-STRATEGY
ADVISERS LLC


JUICE
PHARMA WORLDWIDE

WHAT IS THIS TALK ABOUT

- Exposing the internals of how an AI engine conceptually works, no technical jargon.
- What we see through our projects in AI.
- Not an exhaustive list of all AI algorithms used.
- AI algorithms tend to be pretty sophisticated. So rather than wading into the mechanics of how they work, were going to focus on what the algorithms do conceptually.

TALK OVERVIEW

- AI DEFINITION.
- MACHINE LEARNING DISCUSSION.
- NEURAL NETWORKS DISCUSSION.
- AI IN THE NEWS, THE FUTURE.

ARTIFICIAL INTELLIGENCE

- AI is a broad term for computer algorithms that helps us make decisions or predictions about data.
- Pattern recognition(sometimes hidden) in the data.
- Examples : detecting spam emails, brain tumors (CT scans), Floor cleaning robot.
- Useful, but not “intelligent”.

MACHINE LEARNING

- Martian geology.
- Deciding whether Martian rocks are volcanic (above surface) or plutonic (below surface).
- Classification.
- Classifier algorithm.
- Algorithms reducing complexity of real world objects and phenomena into features.

MACHINE LEARNING

- Features : Values that usefully characterize things we want to identify and classify.
- Features (granularity, density).
- Geology team on mars, training (labeled)data.
- Unlabeled data.

SCATTER PLOT

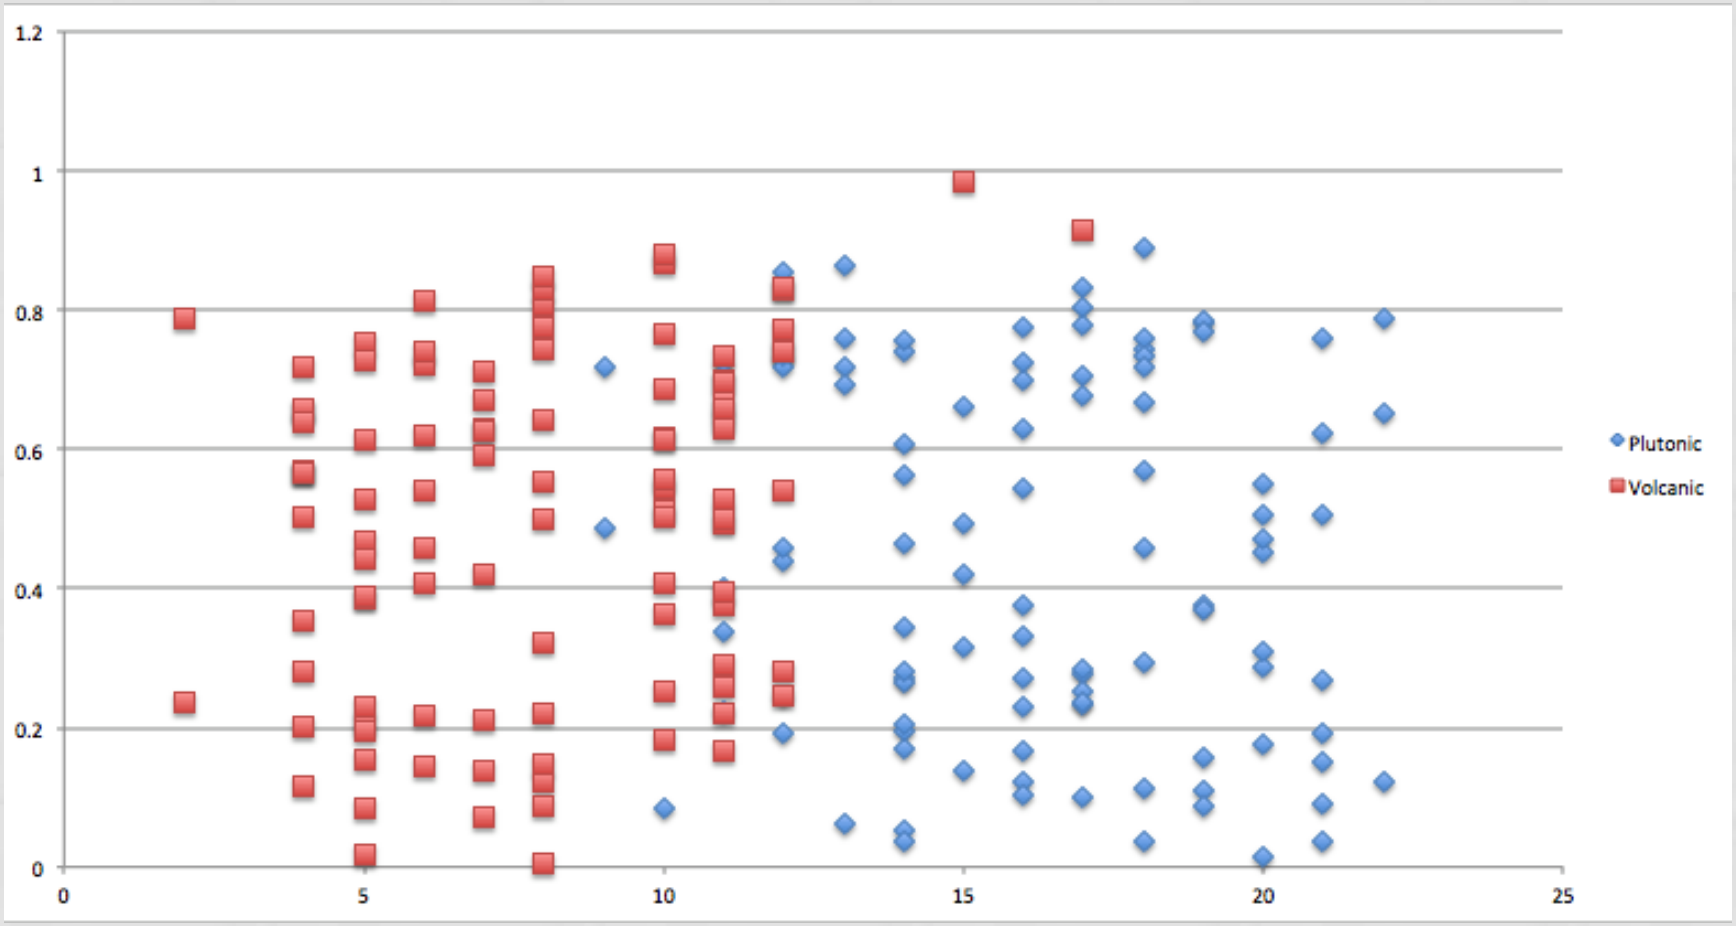

SCATTER PLOT

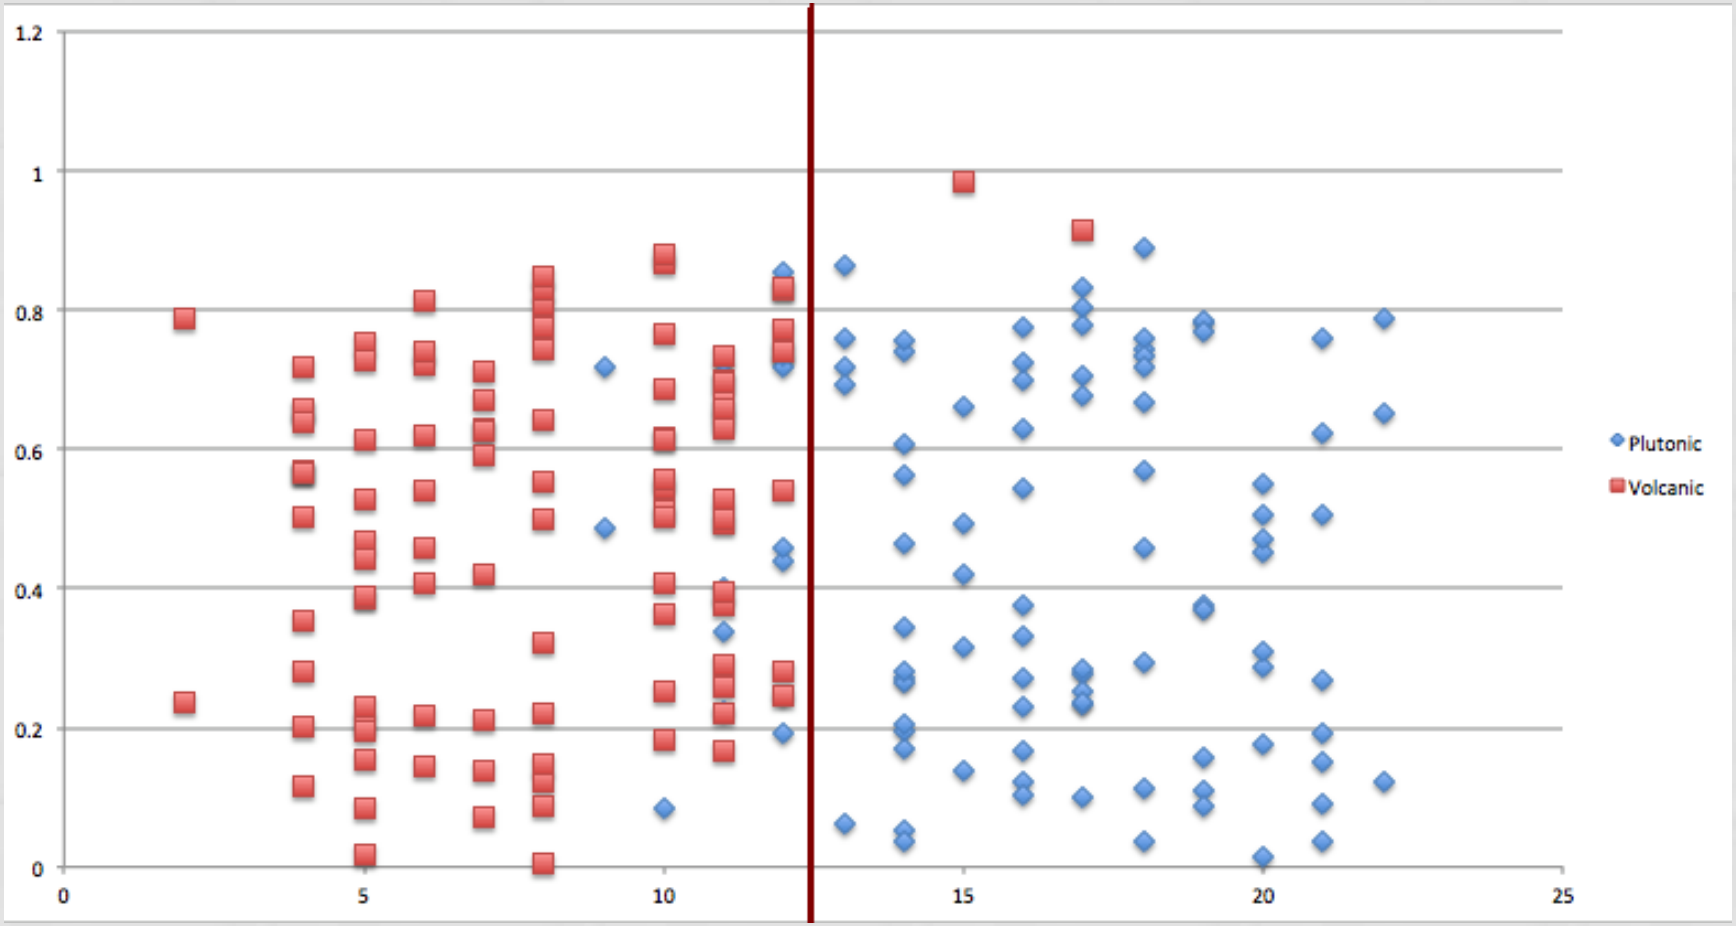

SCATTER PLOT

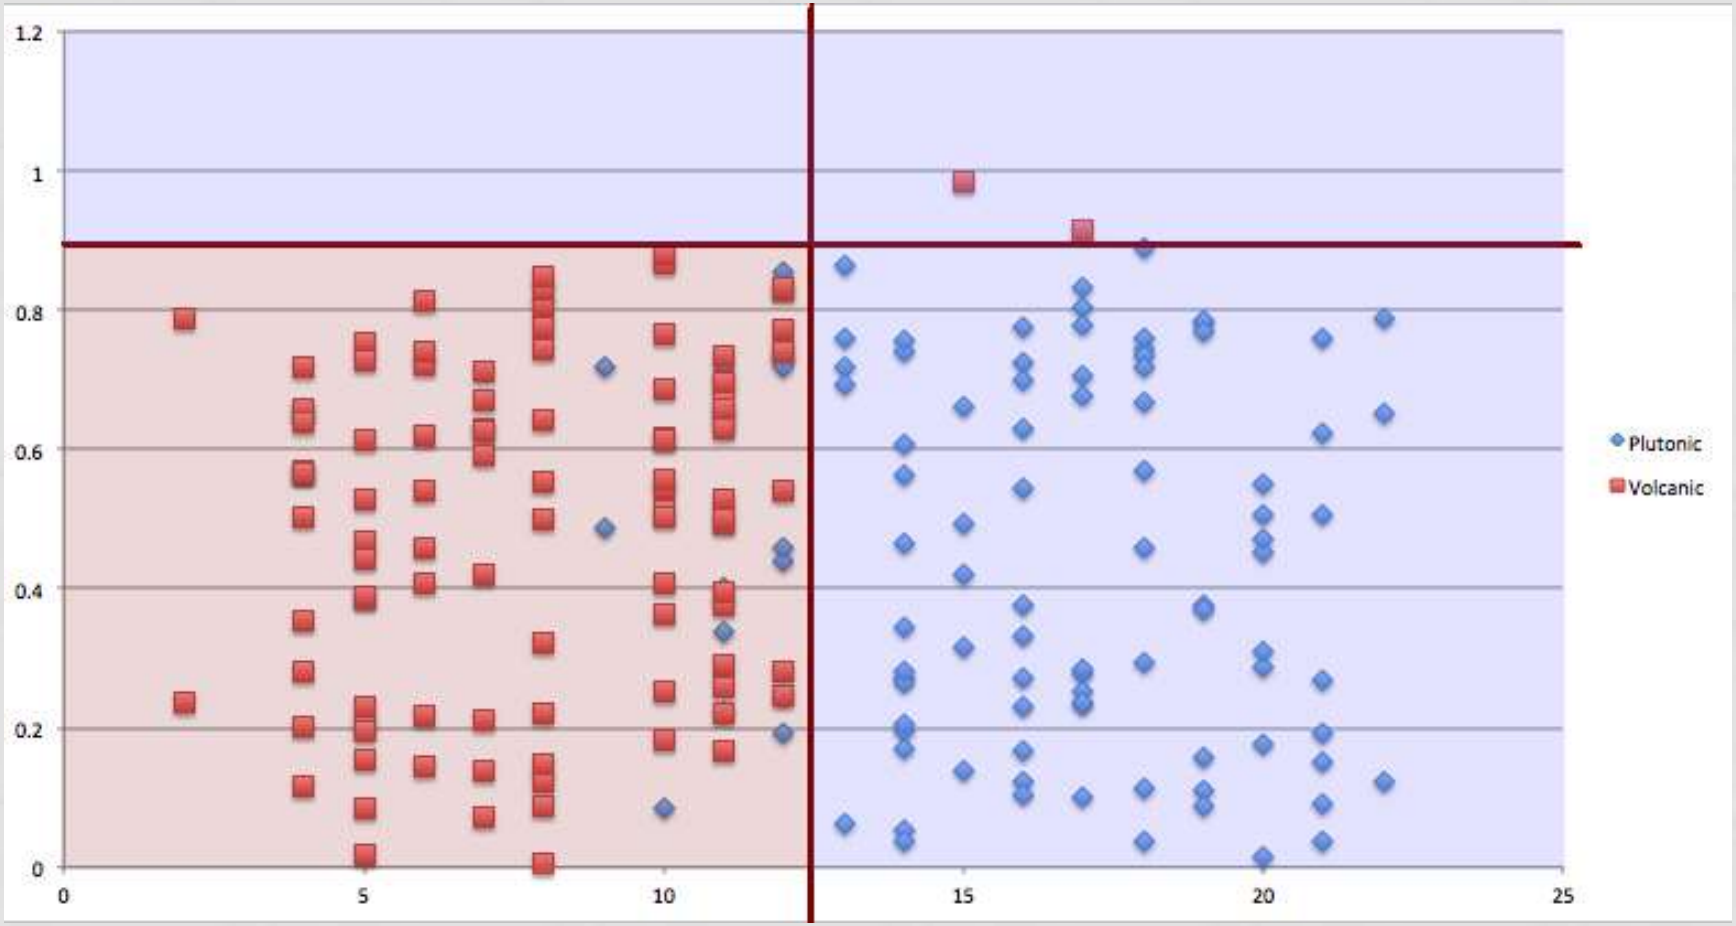

SCATTER PLOT

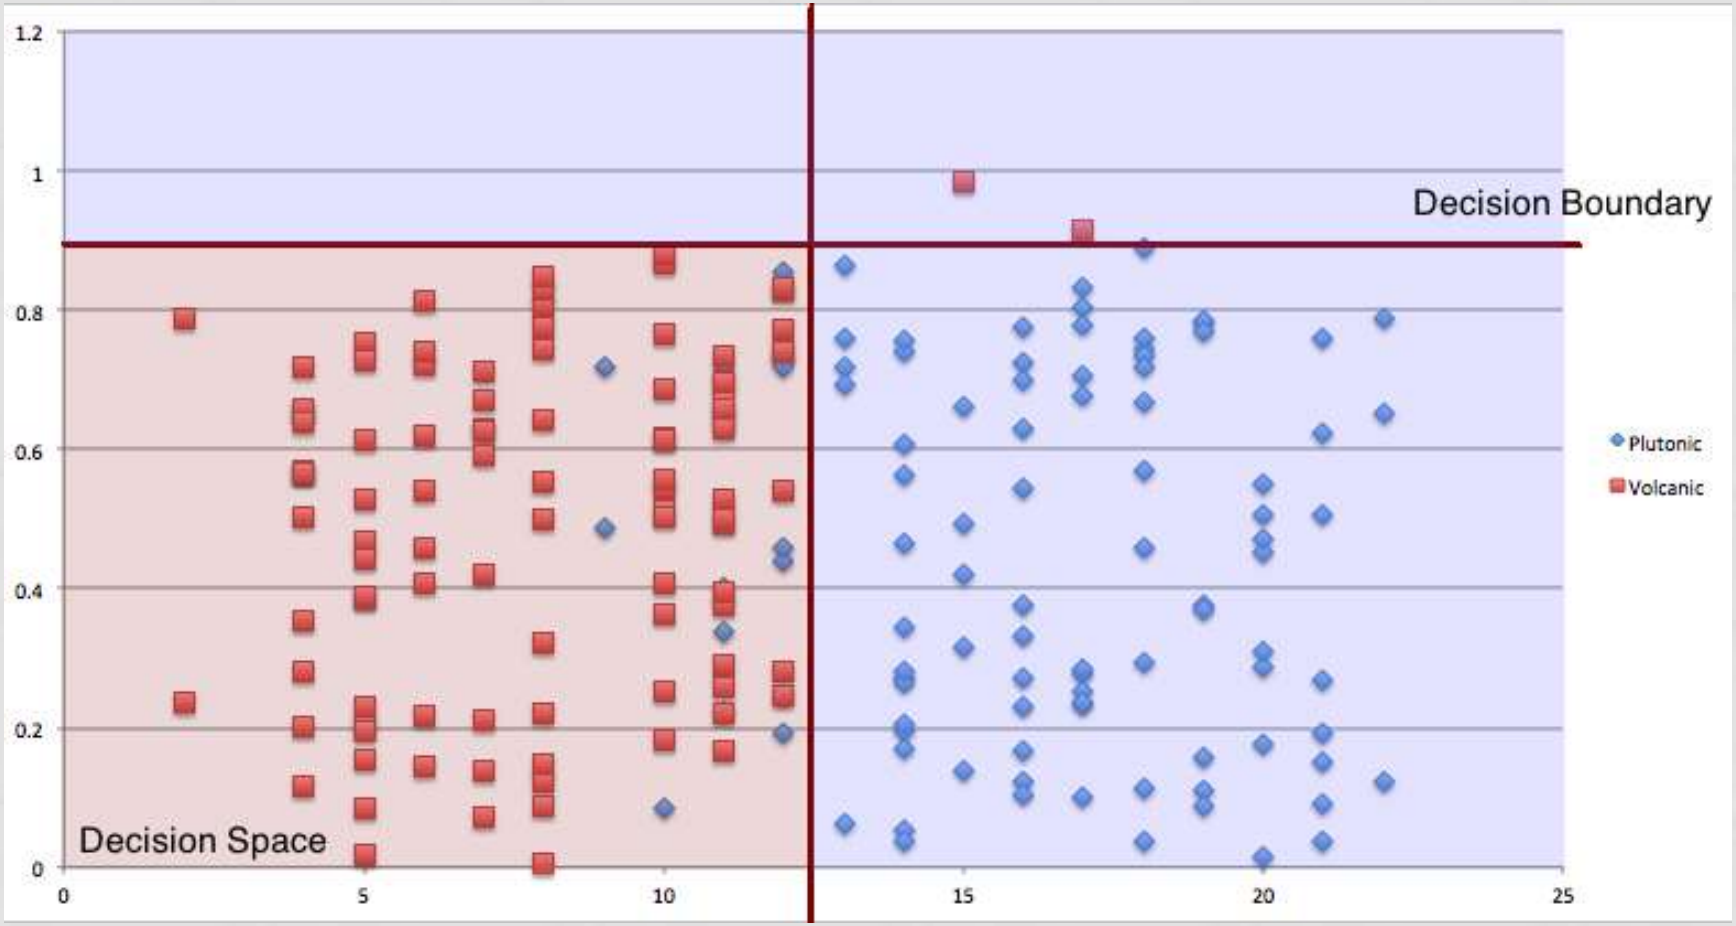

MACHINE LEARNING

- Machine learning algorithms : Max classification with min errors.
- 170 rocks correct, 30 rocks wrong, average classification accuracy 85%.
- Robots on Mars, unlabeled data.
- Accuracy Varies depending on industry and the project.
- AI is extremely context specific, no one size fits all.

DECISION TREE

```
1  if Granularity <= 12.5 then
2      if Density <= 0.9 then
3          Output ("Volcanic")
4      else
5          Output ("Plutonic")
6      end if
7  else
8       Output ("Plutonic")
9  end if
```

MACHINE LEARNING

- Decision tree (If else) .
- Divide the decision space using arbitrary lines, non linear.
- Multiple decision trees = Forests.
- Hundreds of ML Algorithms.

MACHINE LEARNING

- Non tree approaches, dividing decision space using curvy fancy mathematical notations (polynomials, other math).
- ML Algorithm job to figure out the best lines to provide most accurate decision boundaries.

MACHINE LEARNING

- 3 features.
- Viscosity.
- 2D Lines -> 3D Planes.
- Useful classifier, handling multiple rock types.
- Equation for a Hyperplane rippling through thousand dimensional decision space??
- Real world classifiers, 100s – 1000s features (google, amazon, facebook)

STATISTICS

- Decision trees, ML strongly rooted in statistics.
- Used before computers.
- Making confident decisions , using data.
- Other approaches to AI with no origin in statistics
- Most popular : Artificial Neural Networks (ANN's).

NEURAL NETWORKS

- Based on neurons in our brains.
- Biological neurons : Cells that transmit messages using chemical and electrical signals.
- Multi input from other cells, process signals, emit own signal.
- Huge interconnected networks in brain to process complex information.
- Artificial neurons : Similar approach.

NEURAL NETWORKS

- Artificial neurons : instead of chemical / electrical signals, take numbers in, spit numbers out.
- Organized into layers, forming network of neurons.
- Martian Rocks and Classification.
- Found rock (Unlabeled Data).

NEURAL NETWORKS

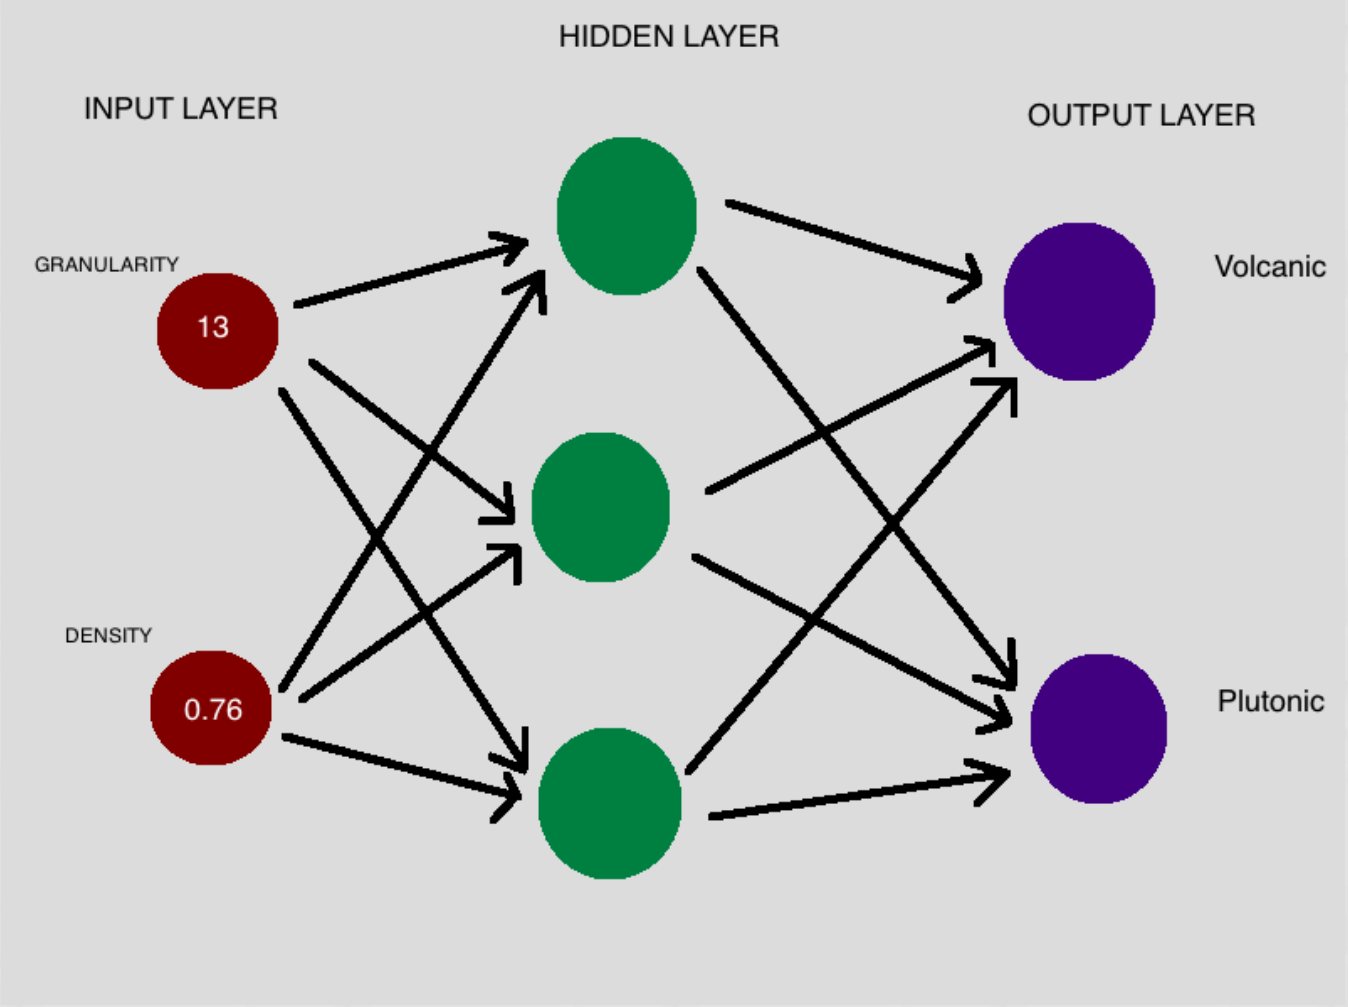

NEURAL NETWORKS

- Each neuron would have arbitrary values set.
- An algorithm tweaks the values.
- Each layer of neurons produces new values propagating forward to next layer.
- Labeled data, gradual improvements.
- Training and testing.
- Mimicking human learning.
- Deep learning, multiple hidden layers.

NEURAL NETWORKS

- Training complex neural networks requires intensive computation and data.
- Neural networks : 50 years old !!
- Deep learning recently practical.
- Cheap computers, cloud, etc.

NEURAL NETWORKS

- Google and Facebook (Facial detection in pictures)2015.
- Deep neural networks (autonomous cars, medical diagnosis, translating human speech).
- Sophisticated algorithms.
- Intelligent?
- Weak AI, Narrow AI : single task intelligent (Identifying rocks, driving cars).
- Doesn't mean it's not useful

ARTIFICIAL INTELLIGENCE

- Composing music? Cooking recipes?
- Not needed, is cool.
- Strong AI : General purpose, human like AI.
- No demo yet, maybe never.
- Explosion of digitized content (Wikipedia, twitter).

ARTIFICIAL INTELLIGENCE

- IBM WATSON.
- 200 Million pages of content (full text of wikipedia).
- Not a Strong AI.
- AI Platforms from large tech companies can absorb large amounts of data.
- Faster learning than humans in narrow tasks.

ARTIFICIAL INTELLIGENCE

- Google Alphago.
- Narrow AI playing GO.
- Played against millions of its clones.
- Learning new strategies, discovering completely new one.
- REINFORCEMENT LEARNING.
- Close to how humans learn (babies walking).

ARTIFICIAL INTELLIGENCE

- The future?
- Learning by trial and error.
- Reinforcement learning works for Narrow AI.
- Potential for Strong AI using reinforcement learning?
- Strong AI learning like Kids do?

AI PROJECTS

- WATERFALL AND AGILE.
- ITERATIVE WATERFALL.
- GATHERING LABEL DATA, BUILDING CLASSIFIER ALGORITHM.
- INVOLVING CLIENT.

THANK YOU

AHMADEUS

Technology Boutique