

# Building Better Minds

By: Andrei Şuiu



[meetup.com/IASI-AI/](https://meetup.com/IASI-AI/)



[facebook.com/AI.Iasi/](https://facebook.com/AI.Iasi/)



[iasiai.net](https://iasiai.net)

## Agenda:

1. AI applications in real world. A short history of my personal AI experience in business.
2. Modern AI projects. A list of currently most significant public AI projects and their aims.
3. AI terms dictionary.



## AI applications in real world. Using Machine Learning in virus detection

2005-2011 spent in BitDefender malware analysis lab. Thousands of viruses analysed and signed.

### Automatic detection

Traditional antivirus detection methods:

- Static signatures:
  - specific: known viruses (almost deterministic det.)
  - generic: families, variants (detect new unknown samples)
- Heuristics: unknown new viruses, behavior detection
  - custom hand-made algorithms
  - Supervised binary/multiclass classification

### Semi-automatic detection

- Clustering file samples in families, or prioritizing for analysis most likely viruses/false positives

## AI applications in real world. Using NLP & ML in paper grading

2011 working as NLP research at paperrater.com

(a free resource that utilizes Artificial Intelligence to help students write better, produces the most powerful automated proofreading tool available on the Internet today)

Mission:

- develop a service that grades(1 to 10) papers/essays submitted by students similarly how a human professor would do.
- use a training corpus (~100K papers), extract various features using NLP, train a regression/classification function.

Result:

- **FAIL**

## AI applications in real world. Make money using ML for robots trading on CME stock exchange

2012 working as team lead/data scientist in austrian company Mauve AG

### Targets:

- create a platform for stock exchange millisecond data collection - success
- platform for running independent High Frequency Trading robots - success
- robot/algorithm backtesting platform - success
- Hft price prediction - success
- Momentum day-trading algorithms - fail
- make real money - **FAIL**

## AI applications in real world. Metrics-driven communication assessment tool

2013 NLP researcher/team lead working on ExecutiveVoice automatic assessment module.

ExecutiveVoice - an online service that provides individual or group benchmarks to gauge progress in any training program. Communication assessment quickly measures more than 275 competencies, from articulation to cognition to strategy for native and non-native English speakers.

Target:

- Receiving a wav file - convert speech to annotated text, and generate a complete report on various speech indicators, mostly of these using NLP techniques, like sentiment analysis, sophisticated wording, speech speed, vague wording, verb accuracy, verbal junk, view point analysis, vocabulary analysis, meaning, etc...

## AI applications in real world. Natural language text rewriting service

Spinnerzilla is an English natural language text rewriting online service

Target:

- Rewrite a text article so it won't be detected by plagiarism detection tools like Copyscape.
- Main challenge - automatically measure the quality of generated text, development progress and QA



## AI applications in real world. ML in Robot Detection @ Amazon.com Inc.

2013-2016 Protect Amazon.com websites from automated web robots

Mission:

services availability, reduce load on HW to save money, clean business metrics, prevent financial data stealing

- Use supervised learning to train robot detection models
- clusterization techniques for data analysis
- Markov Chain models for human/robot actions prediction



## What is Intelligence?

### 71 definitions of Artificial Intelligence

If extracting commonly occurring features we find that intelligence:

- Is a property that an individual agent has as it interacts with its environment or environments.
- Is related to the agent's ability to succeed or profit with respect to some goal or objective.
- Depends on how able an agent is to adapt to different objectives and environments.

Putting these key attributes together produces the *informal definition* of intelligence that S. Legg and M. Hutter have adopted:

***“Intelligence measures an agent's ability to achieve goals in a wide range of environments.”***

So **Intelligence**, looks more like a metric, either uni or multi-dimensional, rather than being a binary attribute.

## Taxonomy

The field of AI research defines itself as the study of "Intelligent agents".

"Intelligent agent" is an autonomous entity which observes through sensors and acts upon an environment using actuators and directs its activity towards achieving goals. A simple example of an intelligent agent could be a thermostat.

Central problems of AI research:

- Knowledge Representation and Reasoning - expert systems, semantic nets, ontology engineering, fuzzy logic, Bayesian inference, expert system, automated theorem proving
- Automated planning and scheduling - related to decision theory
- Machine Learning - unsupervised, supervised, reinforcement
- Communication (with humans)/NLP - NLG, NLU, text data mining, IR, machine translation
- Machine perception (Computer Vision, Speech Recognition, etc.)
- ability to manipulate objects

## General AI vs Weak AI

- Informal definition: **GeneralAI**, or **StrongAI**, or **FullAI**, is the intelligence of a machine/agent that can perform any intellectual task as well or better than a human being.  
There is no definition that satisfies everyone
- **Weak AI**, in contrast to Strong AI, does not attempt to perform the full range of human cognitive abilities.
- Simulation of a human brain would require roughly 10 petaFLOPS (achieved in 2011 in one supercomputer)
- **AI-complete**, or AI-hard is an informal class of problems with the solving difficulty equivalent to that of solving the central artificial intelligence problem - making computers as intelligent as people, or StrongAI. Hypothesised examples: computer vision, NLU, machine translation, word sense disambiguation.

## Mathematical theory of "universal intelligence"

Marcus Hutter. *Universal Artificial Intelligence: Sequential Decisions based on Algorithmic Probability*. Springer, 2005.

[Theory overview slides](#)

The theory reduces all conceptual AI problems to pure computational questions. The necessary and sufficient ingredients are Bayesian probability theory; algorithmic information theory; universal Turing machines; the agent framework; sequential decision theory; and reinforcement learning.

It gives useful conclusions about general intelligences with infinite or infeasibly massive computational resources. In principle, a certain system could solve any problem given enough resources and time. But real-world systems are inevitably limited in resources, including finite amounts of time.



## Organizations&Projects: Partnership on AI

In October 2016 Amazon, Google, Facebook, IBM, Microsoft and Apple have established a non-profit partnership to formulate best practices on artificial intelligence technologies, advance the public's understanding, and to serve as a platform about artificial intelligence.

They stated: *"This partnership on AI will conduct research, organize discussions, provide thought leadership, consult with relevant third parties, respond to questions from the public and media, and create educational material that advance the understanding of AI technologies including machine perception, learning, and automated reasoning."*

- Other important partners that joined organization include OpenAI, Ebay, Intel, Sony, CogitAI, AllenAI, AAI.
- Goals: develop and share best practices, provide an open platform for discussion and engagement, advance public understanding, identify and foster aspirational efforts in AI for socially beneficial purposes

## Organizations&Projects: Machine Learning

ML enables software to improve its performance over time by learning from input or interaction.

Some opensource Machine Learning (there are hundreds of these):

- [TensorFlow](#) -dev by GoogleBrain, numerical computation on dataflow graphs, distributed on CPUs, GPUs, mobile devices (reinforcement, supervised, unsupervised)
- [OpenAI platforms](#) - reinforcement learning, use from Python
- [Weka](#) - classification, regression, clustering, feature selection, visualization
- [Apache Spark MLlib](#) - classification, regression, clustering, dim reduction, feature extraction, collaborative filtering
- [Scikit-learn](#) - most popular Python lib, supervised, unsupervised, dimensionality reduction, model selection, feature extraction&normalization
- [Apache Mahout](#) - running over Hadoop/Spark
- [Deeplearning4j](#) - opensource, distributed on CPUs/GPUs, some NLP tools
- [Torch](#) - ML, computer vision, signal processing, image, video, audio. Extensive support of CUDA.
- [MOSES](#) - an "evolutionary program learner" that is capable of learning short Python programs that capture patterns in input datasets. ([wiki](#))



## Organizations&Projects: General AI

- [OpenCOG](#) - is open-source software initiative aimed at directly confronting challenge of creating beneficial artificial general intelligence with broad capabilities at the human level and ultimately beyond. The most advanced open-source general cognition/reasoning system. Includes an NLP subsystem, reasoning, learning, 3d virtual avatar interfaces, robotics interfaces. But it's just an experimental research platform.
- [Numenta](#) - develops a cohesive theory, core software, technology and applications based on the principles of the neocortex. Their machine intelligence technology is called hierarchical temporal memory (HTM), and is a computational theory of the neocortex. When applied to computers, HTM is well suited for prediction, anomaly detection, classification and ultimately sensorimotor applications.
- [MIRI](#) -non-profit organization for research safety issues related to the development of Strong AI
- [IBM Watson](#) - General Question Answering system, this approach is currently applied to real-world domains like health care and retail.

## Organizations&Projects: Neuromorphic Hardware & NN simulations

### Neuromorphic hardware

- [DARPA SyNAPSE](#) - attempt to build a computer with similar function to the mammalian brain
- [SpiNNaker](#) -brain simulator machine, massively parallel, low power, neuromorphic supercomputer

### Neural network simulation

- [Blue Brain](#) -reverse engineer the human brain and recreate it at the cellular level inside a computer simulation
- [Spaun brain simulation](#) - a simulated brain that recognizes lists of numbers, does simple arithmetic, and solves reasoning problems. It's biologically realistic in its simulation of spiking neurons and neurotransmitters
- [Synthetic Cognition at Los Alamos](#) - simulated the primate visual system including the retina, LGN, and visual cortex areas V1, V2, and V4. The models are trained to visually recognise objects such as enemy vehicles on a battlefield - a possible application being their use in military UAVs

# Thank you

[andrei.suiu@gmail.com](mailto:andrei.suiu@gmail.com)



[meetup.com/IASI-AI/](https://meetup.com/IASI-AI/)



[facebook.com/AI.Iasi/](https://facebook.com/AI.Iasi/)



[iasiai.net](https://iasiai.net)

---