Zurich University
of Applied Sciences

Artificial Intelligence
V01: The field of Artificial Intelligence

Logistics of this module
What is AI?
A brief history
The state of the Art

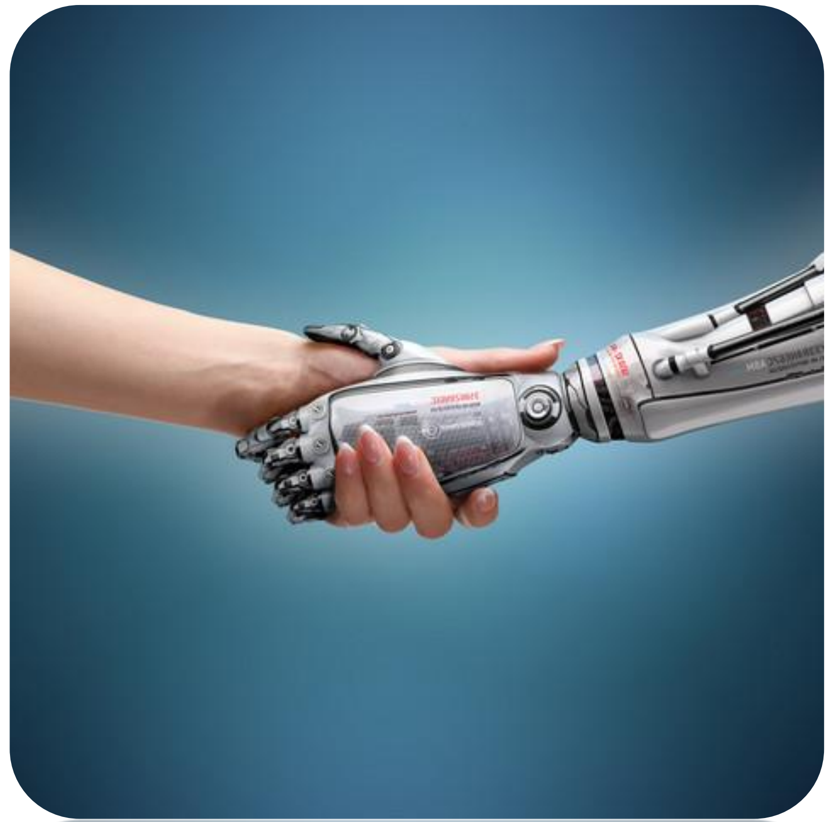

Based on material by Stuart Russell, UC Berkeley

0. LOGISTICS OF THIS MODULE

About me

Thilo Stadelmann
- Born 1980, married
- Studied computer science in Giessen & Marburg, then doctorate (AI, voice recognition)
- Passion for programming & artificial intelligence (20 years experience)

At ZHAW
- Email: stdm@zhaw.ch, office: TD 03.16 (Obere Kirchgasse 2)
- Tel.: 058 934 72 08, web: https://stdm.github.io
- Professor for AI/ML at InIT's Information Engineering group, head of ZHAW Data Science Laboratory

Interests




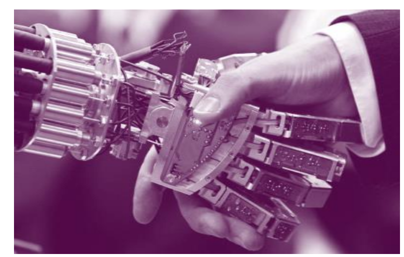

About You

About Information Engineering research group
at InIT Institute of Applied Information Technology

Your career at InIT:
we are looking for
excellent master
students and
assistants

Information Retrieval

- Search on unstructured and semi-structured textual and multimedia data
- Natural language processing
- Multilingual information retrieval
- Benchmarking of IR applications, advanced topics

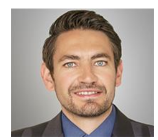

Artificial Intelligence

- Machine Learning algorithms and applications
- Intelligent systems and agents
- Data analytics and mining
- Multimedia analysis, reinforcement learning, deep learning

Databases and Information Systems

- Data Warehousing and Big Data processing
- Efficient storage, management & querying of data
- Complex data integration and enrichment processes
- Data product development in an enterprise context

About my research

See also https://youtu.be/efCyLhSACoU

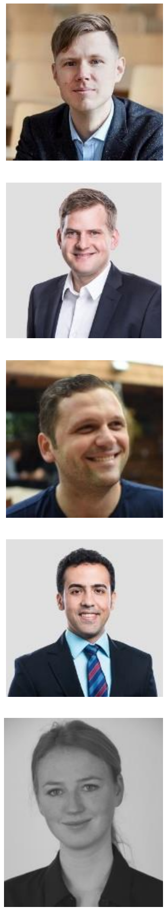

Machine learning-based Pattern Recognition

- Robust deep learning
- Voice recognition
- Document analysis
- Learning to learn & control

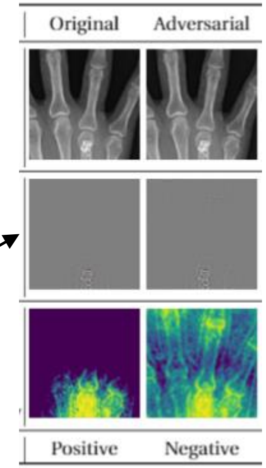

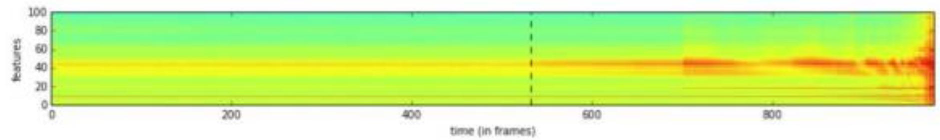

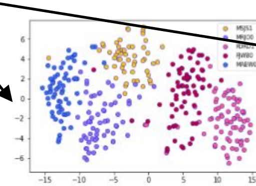

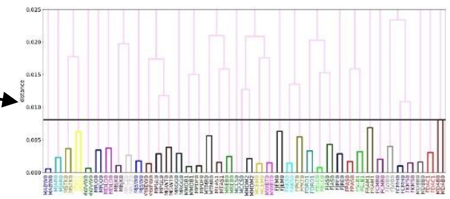

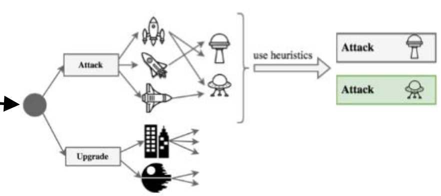

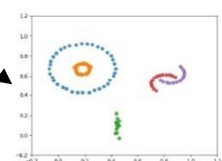

Zurich University of Applied Sciences

About ZHAW Datalab
Est. 2013

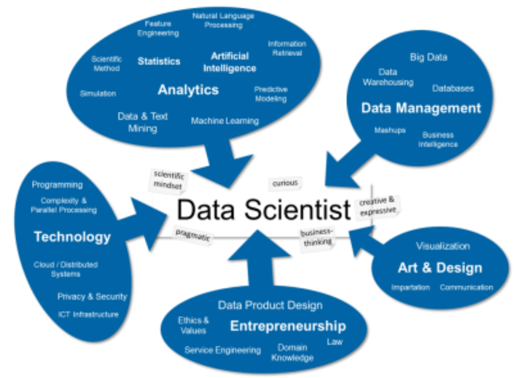

Forerunner
- **One of the first** interdisciplinary data science initiatives in Europe
- One of the first interdisciplinary centers at ZHAW

Foundation
- **People**: ca. **80 researchers** from **5 institutes** / **3 departments** opted in
- Vision: **Nationally leading** and **internationally recognized** center of excellence
- Mission: **Generate projects** through critical mass and mutual relationships
- Competency: **Data product design** with structured and unstructured data

Success factors
- **Lean** organization and operation → geared towards projects
- Years of successful **pre-Datalab collaboration**

Logistics

Lecture

- Starts on time, ends on time
- Gadgets (notebooks, tablets) are used for class purposes only

Self-study

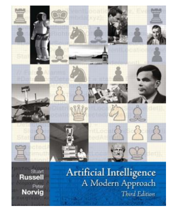

- Read & experiment as much as possible at home (ca. **3h / week**)
- **Reading** corresponding **AIMA** chapters is **highly** (!!) **recommended** (→ see later)

Material

- Everything on OLAT: https://olat.zhaw.ch/auth/RepositoryEntry/219152410

Grading

- Final exam (SEP): max. 80 points
  (90 minutes written test, **closed book**, **1 A4** sheet of **handwritten notes** allowed)
- Labs: max. 20 points
  (**2 graded** labs of **your choice**, **demonstrated** until one week after the official end of this lab)

➔ See terms & conditions on OLAT

Literature

Russell, Norvig, *«Artificial Intelligence – A Modern Approach»*,
3rd edition 2010

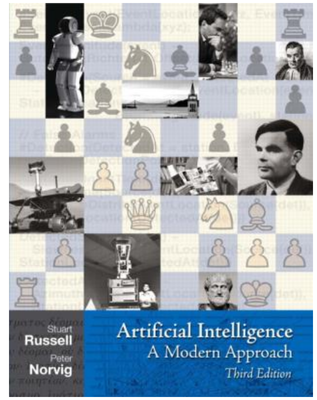

- Commonly known as *AIMA* or the *intelligent agent book*
- *The* textbook on AI: concise, complete, comprehensible
- Among the 25 most cited scientific reference on CiteSeer

➔ **Lectures build on certain AIMA chapters**

https://www.swissbib.ch/Search/Results?lookfor=russell+norvig+artificial+intelligence+modern+approach&type=AllFields
https://www.amazon.de/Artificial-Intelligence-Stuart-Russell/dp/0132071487
https://www.amazon.de/Pearson-Studium-Intelligenz-Stuart-Russell/dp/3868940987

Other good & informative reads

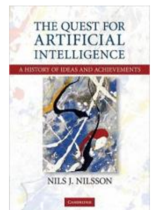

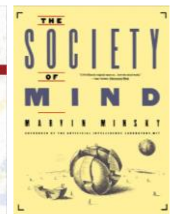

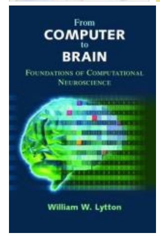

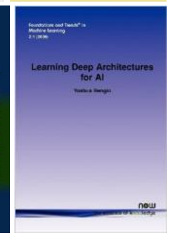

- Nilsson, *«The Quest for Artificial Intelligence», 2010*
  → *A History of Ideas and Achievements*
- Minski, *«The Society of Mind»*, 1985
  → A collection of essays to explain how intelligence emerges from unintelligent parts
- Lytton, *«From Computer to Brain», 2013*
  → *Foundations of Computational Neuroscience*
- Bengio, *«Learning Deep Architectures for AI»*, 2009
  → Deep machine learning from an AI perspective

Semester schedule HS 2018

Read this on time!

SW	Topic	Lecture	Lab	AIMA	Winti	ZH
1	**Introduction**	Welcome & P01: Esoteric AI	Will be discussed during lecture time	-	18.09.	21.09.
2		V01: The field of artificial intelligence	P01: Esoteric AI	ch. 1	25.09.	28.09.
3		V02: Intelligent agents	P02: 2048 (P02.1: Learn Python)	ch. 2	02.10.	05.10.
4	**Search**	V03: Problem solving through search	-"- (P02.2: Heuristic agent)	ch. 3	09.10.	12.10.
5		V04: Local and adversarial search	-"- (P02.3: Expectimax agent)	ch. 5 (+4)	16.10.	19.10.
6		V05: Constraint satisfaction problems	P03: CSP & Logic (P03.1: Sudoku)	ch. 6	23.10.	26.10.
7	**Planning**	V06a: Knowledge, reasoning & logic	-"- (P03.2: Datalog)	ch. 7	30.10.	02.11.
8		V06b: Datalog	-"-	ch. 8 (+9)	06.11.	09.11.
9		V07: Planning	AIMA catch-up week	ch. 10 (+11)	13.11.	16.11.
10	**Learning**	V08: Learning agents	P04: Decisions trees for data mining	ch. 18.1-18.6	20.11.	23.11.
11		V09: Ensemble learning	-“-	ch. 18.10-18.12	27.11.	30.11.
12		V10: Probabilistic learning	P05: Multimedia data analysis	ch. 20	04.12.	07.12.
13		V11: Generative modeling with neural nets	-”-	ch. 18.7	11.12.	14.12.
14	**Selected chapters**	V12: AI & society Exam preparation, FAQ	P01b: P01 revisited	ch. 26	18.12.	21.12.

All materials (V/P), up-to-date schedule, terms & conditions etc. → OLAT

Superior educational objectives

- You **Know** the breadth of AI problem solving strategies
- …thus **identify** such challenges in practice
- ...and **develop** corresponding solutions on your own.

- You can **explain** the discussed algorithms and methodologies
- ...and are able to **transfer** it to the real world.

➔ This course is concerned with **methodology** and a good understanding of **what** may work **where**

➔ It thus **sacrifices depth** in how to implement concrete solutions **or proofs** why things work

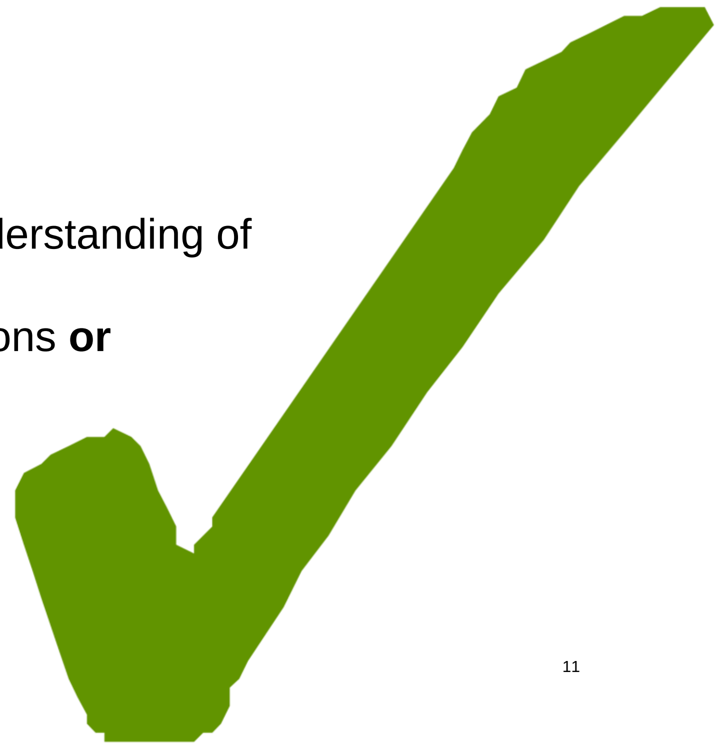

Zurich University
of Applied Sciences

Educational objectives for today

- **Know** the **history** and **breadth** of the discipline **of Artificial Intelligence**
- **Define** what is meant by **the term AI** (and what is not)
- Be able to **engage in an educated discussion** on the state of the art and future of AI (→ see P01)

"In which we try to explain why we consider artificial intelligence to be a subject most worthy of study, and in which we try to decide what exactly it is, this being a good thing to decide before embarking."

➔ Reading: AIMA, ch. 1

Another thing…
Further information: Anna-Flurina Kälin (kaeliann@students...)

Instruments of Creation
Talk by Mario Klingemann

In his talk, Mario Klingemann will give insights into his process of experimenting with the possibilities of creating visual art with artificial neural networks.

The ability to wield tools is one of the core qualities that define humanity. The evolution of human culture has always run parallel to the evolution of the tools and instruments at our disposal. Machine learning and what is commonly known as artificial intelligence is a very recent instrument that we have created, and we are starting to apply it to all possible areas, including the creation of art. As with any instrument, it takes time to learn how to use it skillfully or how to play it masterfully. Mario Klingemann has been experimenting with the possibilities of creating visual art with artificial neural networks for several years and is beginning to understand the potential and limitations of these instruments. In his talk, he will give insights into his process and show some of the latest developments in this fast-moving field.

Thursday, 4th October 2018, 5:30 - 6:30 pm
ETH Zurich, Main Campus CAB G11, Universitätsstrasse 6, 8006 Zürich

1. WHAT IS AI?

AI definitions

thinking

"The exciting new effort to **make computers think**... machines with minds, in the full and literal sense." (Haugeland, 1985)

"[The automation of] **activities that we associate with human thinking**, activities such as decision-making, problem solving, learning..." (Bellman, 1978)

"**The study of mental faculties through** the use of **computational models**." (Charniak and McDermott, 1985)

"The study of the **computations that** make it possible to **perceive, reason, and act**." (Winston, 1992)

← concerned with →

humanly ← standard measured by → rationally

"The art **of creating machines that perform** functions that require **intelligence** when performed by people." (Kurzweil, 1990)

"The study of **how to make computers do things** at which, at the moment, **people are better**." (Rich and Knight, 1991)

"Computational Intelligence is **the study of the design of intelligent agents**." (Poole et al., 1998)

"AI... is concerned with **intelligent behaviour in artefacts**." (Nilsson, 1998)

acting

Acting humanly: The Turing test

Zurich University of Applied Sciences

Turing, “Computing machinery and intelligence“, 1950

- *“**Can machines think**?” → “Can machines behave intelligently?”*
- Operational test for intelligent behaviour: the **Imitation Game**

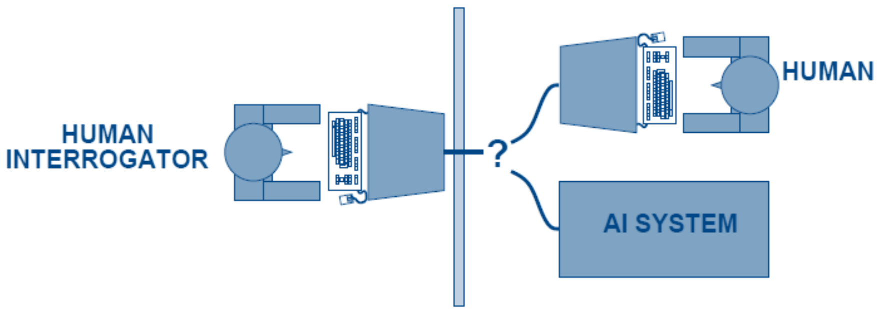

- Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- Anticipated all major arguments against AI in following 50 years
- Suggested all major components of AI: natural language processing, knowledge representation, automated reasoning, machine learning (+computer vision, robotics for full Turing test incl. video)

Problem: **Not reproducible**, **constructive**, or amenable to **mathematical analysis**

Thinking humanly: Cognitive science
Understanding the human mind by computer modelling

1960s “**cognitive revolution**”
- information-processing psychology replaced prevailing orthodoxy of **behaviourism**

Requires scientific **theories** of internal activities **of the brain**
- What level of abstraction? “Knowledge” or “circuits”?
- How to validate? Requires…
  1. …predicting and testing behaviour of human subjects (top-down) or
  2. …direct identification from neurological data (bottom-up)

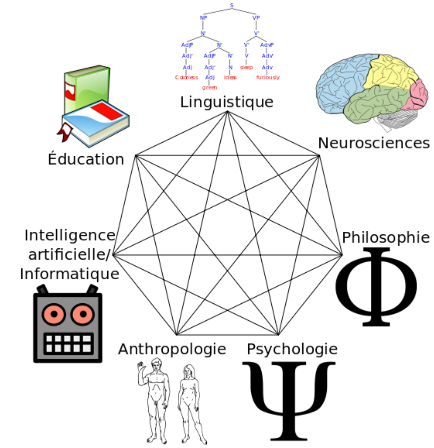

Today
- Both approaches (roughly, **cognitive science** and **cognitive neuroscience**) are now distinct from AI
- Both share with AI the following characteristic:
  - The **available theories do not explain** (or engender) anything resembling human-level general **intelligence**
  - Hence, all three fields share one principal direction!

Thinking rationally: Laws of thought
How to make provably correct inference?

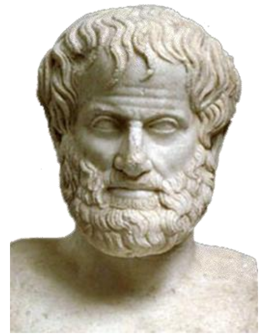

Several Greek schools developed various forms of **logic**
- Aristotle: What are correct arguments/thought processes?
- **Notation** and **rules of derivation** for thoughts
- May or may not have proceeded to the idea of mechanization

➔ **Normative** (or **prescriptive**) rather than descriptive
➔ Direct line through mathematics and philosophy to modern AI

Problems
1. Not all intelligent behaviour is mediated by logical deliberation (e.g., uncertainty exists)
2. What is the **purpose of thinking**? What thoughts **should** I have out of all the thoughts (logical or otherwise) that I **could** have?

Acting rationally

Rational behaviour: doing the right thing

that which is expected to maximize goal achievement,
given the available information

➔ Doesn't necessarily involve thinking (e.g.: blinking reflex) but thinking should be in the service of rational action

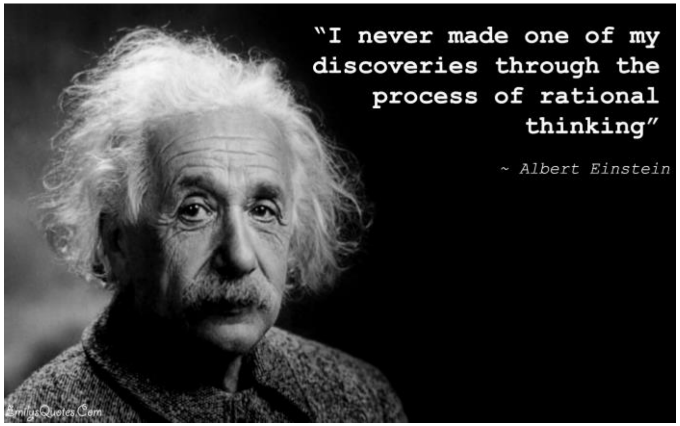

Aristotle (Nicomachean ethics)

- **Every art** and every **inquiry**, and similarly every **action** and **pursuit**, **is thought to aim at some good**

Rational agents
A practical way and goal of this course

Agents
- an **entity that perceives and acts**
- a **function from percept** histories **to actions** $f: P^* \rightarrow A$

Rational agents
- **For any** given class of **environment**s and **task**s, we **seek** the agent (or class of agents) with the **best performance**

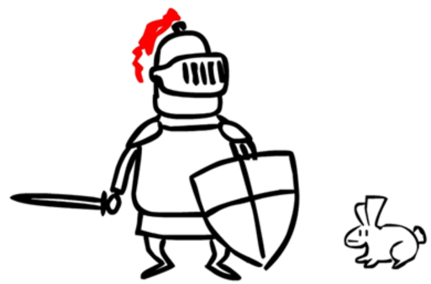

Caveat
- Computational limitations make perfect rationality unachievable
  ➔ **Design best program for given machine resources**

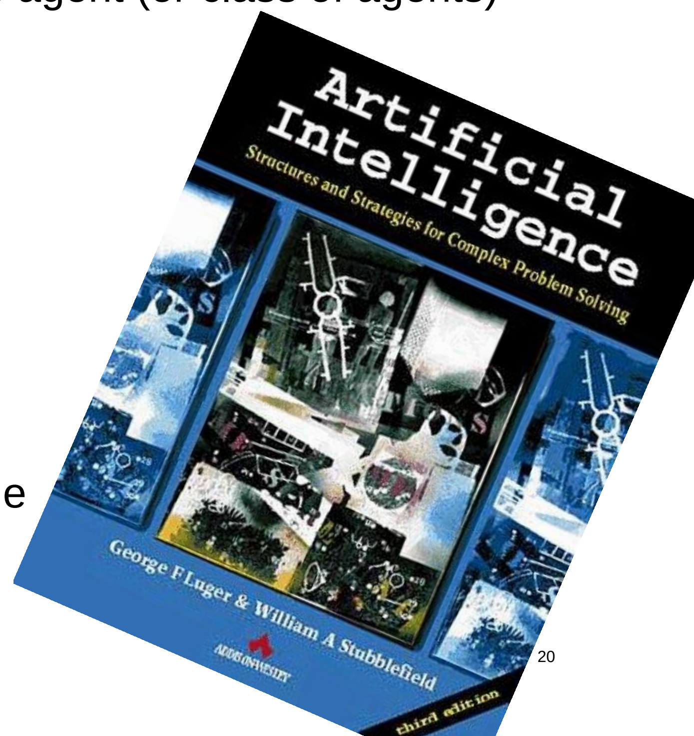

2. A BRIEF HISTORY

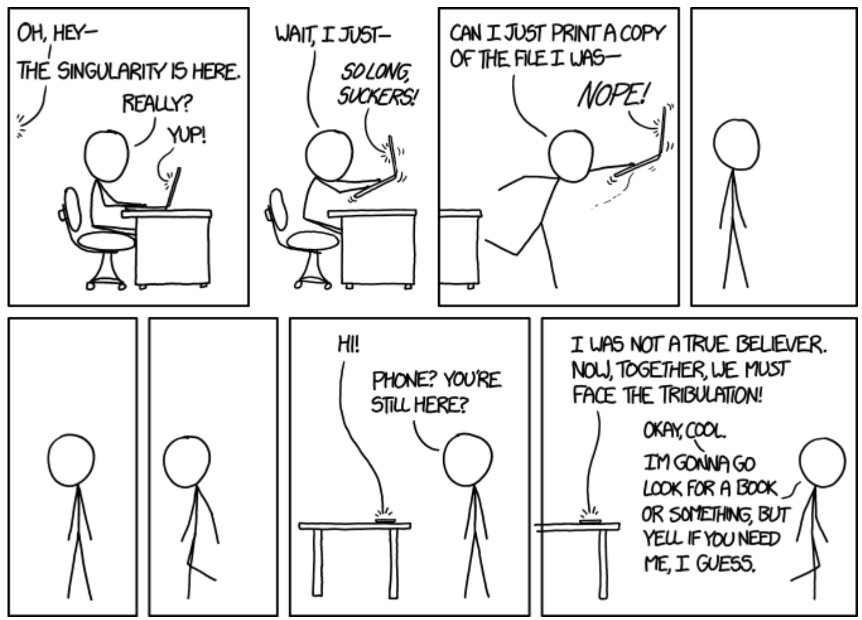

AI prehistory

AI paybacks to **computer science**: e.g. time sharing, interactive interpreters, GUI & mice, linked list, symbolic/functional/declarative/OO programming

Philosophy	**logic**, methods of reasoning mind as **physical system** (roots of calculation machines) foundations of learning, language, rationality (induction, empiricism, rationalism, utilitarianism)
Mathematics	**formal** representation and proof algorithms, **computation** (complexity, (un)decidability, (in)tractability) **Probability**
Psychology	adaptation phenomena of perception and motor control **experimental techniques** (controlled experiments, behaviourism)
Economics	formal **theory of** rational **decisions**, **game theory**
Linguistics	knowledge representation, **natural language processing** Grammar
Neuroscience	plastic physical substrate for mental activity
Control theory	Cybernetics (**homeostatic** systems: stability via feedback loops) simple optimal agent designs (**objective function** optimization)

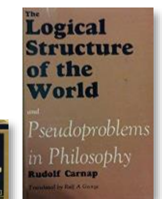

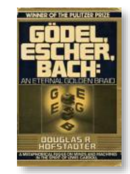

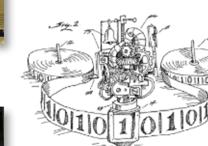

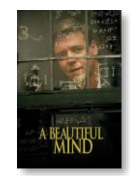

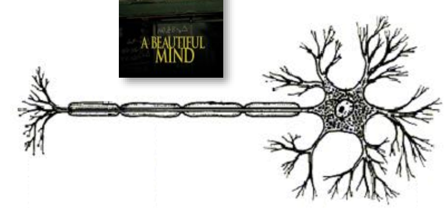

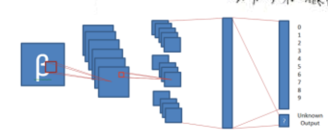

Potted history of AI

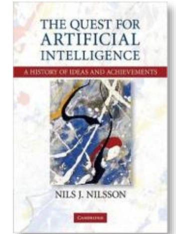

1943	McCulloch & Pitts: Boolean circuit model of brain
1950	Turing's *"Computing Machinery and Intelligence"*
1952-69	*"Look, Ma, no hands!"*
1950s	Early AI programs, including Samuel's **checkers program**, Newell & Simon's **Logic Theorist**, Gelernter's Geometry Engine
1956	**Dartmouth meeting**: *"Artificial Intelligence"* adopted
1965	Robinson's complete algorithm for logical reasoning
1966-74	AI discovers computational complexity **Neural network** (NN) research almost disappears: 1st "**AI Winter**"
1969-79	Early development of knowledge-based systems
1980-88	**Expert systems** industry **booms**
1988-93	Expert systems industry **busts**: 2nd "AI Winter"
1985-95	Neural networks return to popularity
1988	Resurgence of **probability**; general increase in technical depth "Nouvelle AI": Artificial Life, Genetic Algorithms, soft computing
1995	**Agents**, agents, everywhere…
2003	Human-level AI back on the agenda
Since 2010	**Machine learning** widely applied in industry (trends of big data, data science) Superhuman performance in pattern recognition via NN under guise of "**deep learning**"
Since 2016	**Buzzword** again, used to sell everything "digital"
2018	CLAIRE initiative for AI research in Europe

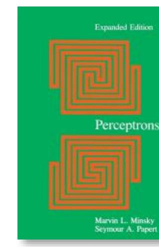

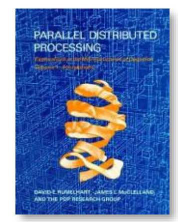

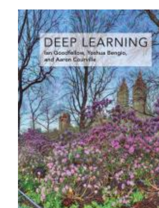

3. THE STATE OF THE ART

Which of the following can be done at present?

1. Play a decent game of **table tennis**	ok
2. Drive safely along a curving **mountain road**	ok
3. Drive safely along **Technikumstrasse** Winterthur	ok (only since recently)
4. Buy a week's worth of **groceries on the web**	ok
5. Buy a week's worth of groceries **at Migros**	no
6. Play a decent game of **bridge**	ok
7. Discover and prove a new mathematical **theorem**	not completely
8. Design and execute a **research program** in molecular biology	not completely
9. Write an **intentionally funny** story	no
10. Give competent **legal advice** in a specialized area of law	ok
11. Translate spoken English **into spoken** Swedish in real time	ok
12. Converse successfully with another person for an hour	no (but some minutes)
13. Perform a complex **surgical operation**	not completely
14. Unload any **dishwasher** and put everything away	no
15. Compete in the game show **Jeopardy!**	ok
16. Write clickbait articles fully automatized	ok
17. Write mathematical articles fully automatized	not completely

Learning to produce text char by char

See http://karpathy.github.io/2015/05/21/rnn-effectiveness

```
\begin{proof}
We may assume that $\mathcal{I}$ is an abelian sheaf on $\mathcal{C}$.
\item Given a morphism $\Delta : \mathcal{F} \to \mathcal{I}$
is an injective and let $\mathfrak q$ be an abelian sheaf on $X$.
Let $\mathcal{F}$ be a fibered complex. Let $\mathcal{F}$ be a category.
\begin{enumerate}
\item \hyperref[setain-construction-phantom]{Lemma}
\label{lemma-characterize-quasi-fi
Let $\mathcal{F}$ be an abelian qua
Let $\mathcal{F}$ be a coherent $\m
$\mathcal{F}$ is an abelian catena
\item The following are equivalent
\begin{enumerate}
\item $\mathcal{F}$ is an $\mathca
\end{lemma}
```

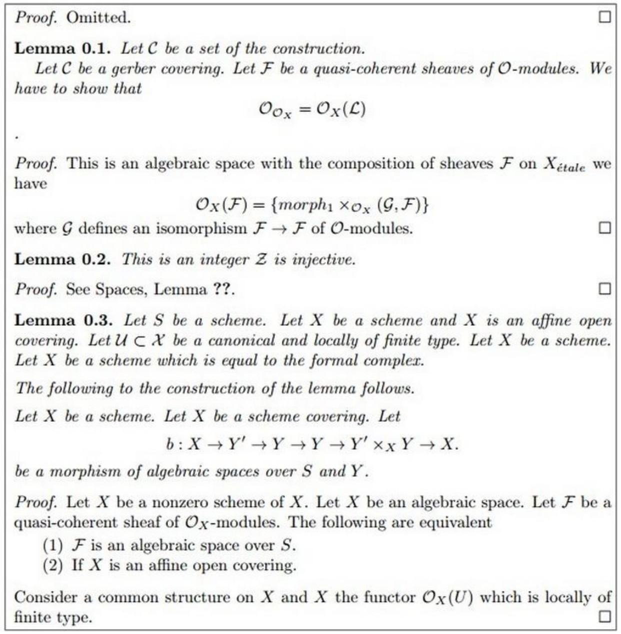

Proof. Omitted. □

Lemma 0.1. *Let $\mathcal{C}$ be a set of the construction.*

Let $\mathcal{C}$ be a gerber covering. Let $\mathcal{F}$ be a quasi-coherent sheaves of $\mathcal{O}$-modules. We have to show that

$$\mathcal{O}_{\mathcal{O}_X} = \mathcal{O}_X(\mathcal{L})$$

.

Proof. This is an algebraic space with the composition of sheaves $\mathcal{F}$ on $X_{\acute{e}tale}$ we have

$$\mathcal{O}_X(\mathcal{F}) = \{morph_1 \times_{\mathcal{O}_X} (\mathcal{G}, \mathcal{F})\}$$

where $\mathcal{G}$ defines an isomorphism $\mathcal{F} \to \mathcal{F}$ of $\mathcal{O}$-modules. □

Lemma 0.2. *This is an integer $\mathcal{Z}$ is injective.*

Proof. See Spaces, Lemma ??. □

Lemma 0.3. *Let S be a scheme. Let X be a scheme and X is an affine open covering. Let $\mathcal{U} \subset \mathcal{X}$ be a canonical and locally of finite type. Let X be a scheme. Let X be a scheme which is equal to the formal complex.*

The following to the construction of the lemma follows.

Let X be a scheme. Let X be a scheme covering. Let

$$b : X \to Y' \to Y \to Y \to Y' \times_X Y \to X.$$

be a morphism of algebraic spaces over S and Y.

Proof. Let X be a nonzero scheme of X. Let X be an algebraic space. Let $\mathcal{F}$ be a quasi-coherent sheaf of $\mathcal{O}_X$-modules. The following are equivalent

(1) $\mathcal{F}$ is an algebraic space over S.
(2) If X is an affine open covering.

Consider a common structure on X and X the functor $\mathcal{O}_X(U)$ which is locally of finite type. □

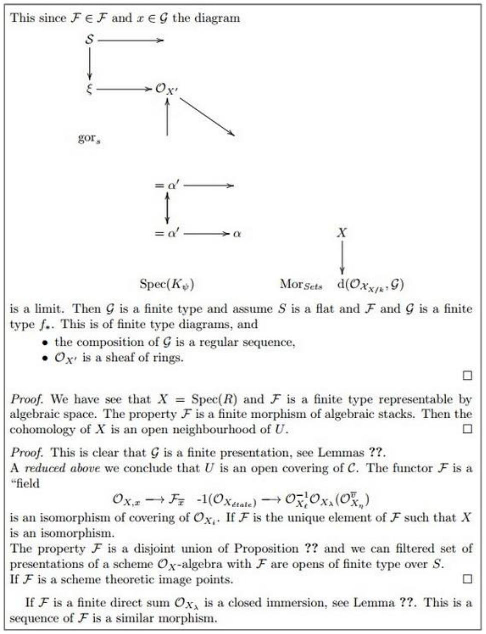

This since $\mathcal{F} \in \mathcal{F}$ and $x \in \mathcal{G}$ the diagram

$\mathcal{S}$ ξ $\mathcal{O}_{X'}$ gor_s $= \alpha'$ $= \alpha'$ α X $\mathrm{Spec}(K_\psi)$ Mor_{Sets} $\mathrm{d}(\mathcal{O}_{\mathcal{X}_{X/k}}, \mathcal{G})$

is a limit. Then $\mathcal{G}$ is a finite type and assume S is a flat and $\mathcal{F}$ and $\mathcal{G}$ is a finite type f_*. This is of finite type diagrams, and

- the composition of $\mathcal{G}$ is a regular sequence,
- $\mathcal{O}_{X'}$ is a sheaf of rings.

□

Proof. We have see that $X = \mathrm{Spec}(R)$ and $\mathcal{F}$ is a finite type representable by algebraic space. The property $\mathcal{F}$ is a finite morphism of algebraic stacks. Then the cohomology of X is an open neighbourhood of U. □

Proof. This is clear that $\mathcal{G}$ is a finite presentation, see Lemmas ??.
A *reduced above* we conclude that U is an open covering of $\mathcal{C}$. The functor $\mathcal{F}$ is a "field

$$\mathcal{O}_{X,x} \longrightarrow \mathcal{F}_{\overline{x}} \ \text{-}1(\mathcal{O}_{X_{\acute{e}tale}}) \longrightarrow \mathcal{O}_{X_\ell}^{-1}\mathcal{O}_{X_\lambda}(\mathcal{O}_{X_\eta}^{\overline{v}})$$

is an isomorphism of covering of $\mathcal{O}_{X_i}$. If $\mathcal{F}$ is the unique element of $\mathcal{F}$ such that X is an isomorphism.
The property $\mathcal{F}$ is a disjoint union of Proposition ?? and we can filtered set of presentations of a scheme $\mathcal{O}_X$-algebra with $\mathcal{F}$ are opens of finite type over S.
If $\mathcal{F}$ is a scheme theoretic image points. □

If $\mathcal{F}$ is a finite direct sum $\mathcal{O}_{X_\lambda}$ is a closed immersion, see Lemma ??. This is a sequence of $\mathcal{F}$ is a similar morphism.

Computer vision and control
From Yann LeCun's NIPS'2016 keynote

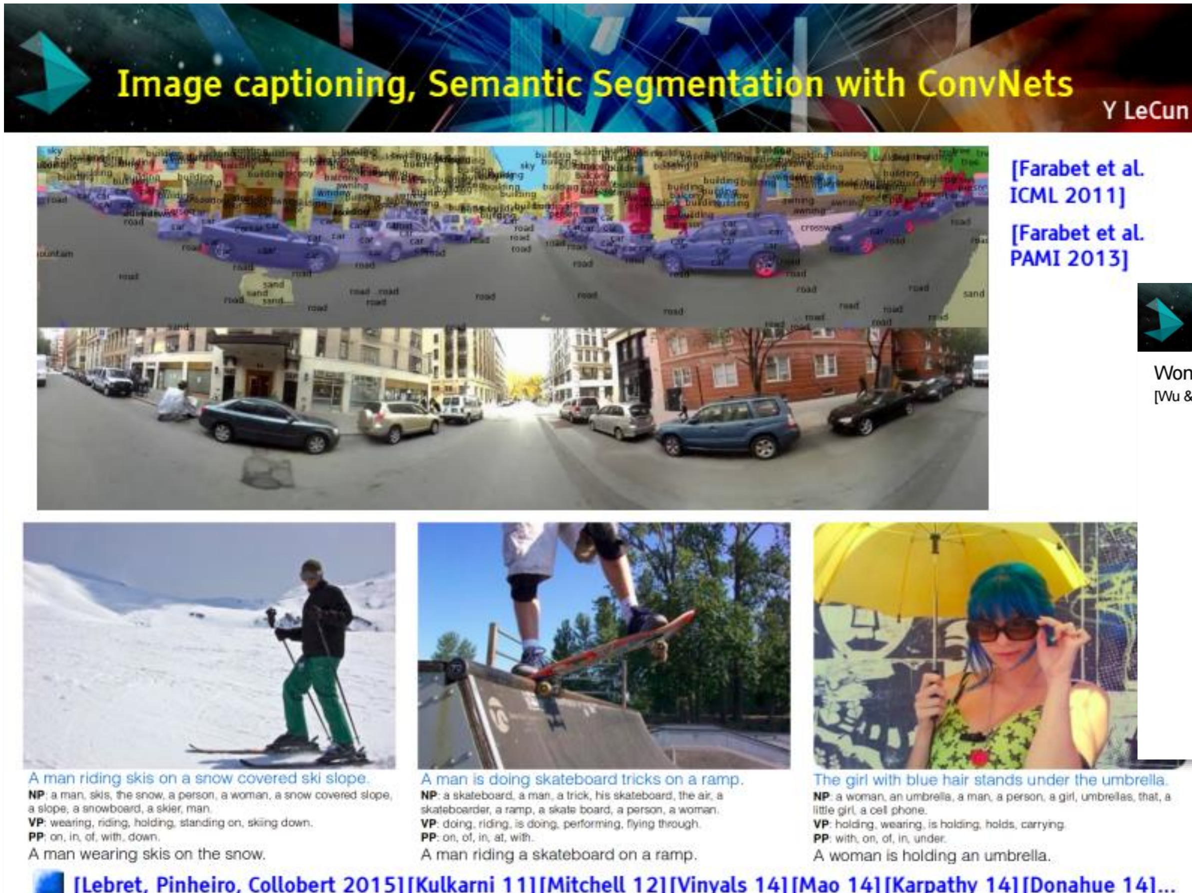

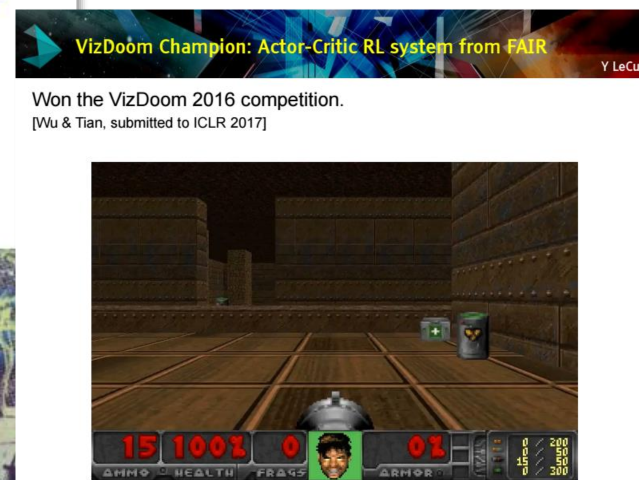

The story of Rocket AI @ NIPS'2016
Or: The danger of hype

Quoting from the blog post (https://medium.com/the-mission/rocket-ai-2016s-most-notorious-ai-launch-and-the-problem-with-ai-hype-d7908013f8c9#.9gjgyxre5):

Turns out anyone can make a multi-million dollar company in 30 minutes

…with a website editor whilst in a Spanish mansion found on Airbnb. *'Temporally Recurrent Optimal Learning'* is a combination of buzzwords we put together to spell out TROL(L) that were conjured up over breakfast. If we hadn't put significant effort into making sure people realized it was a joke, Rocket AI would be in the press right now.

Metrics for the Rocket AI launch party:


Email RSVPs to party: 316
People who emailed in their resume: 46
Large name brand funds who contacted us about investing: 5
Media: Twitter, Facebook, HackerNews, Reddit, Quora, Medium etc
Time Planning: < 8 hours
Money Spent: $79 on the domain, $417 on alcohol and snacks + (police fine)
For reference, NIPS sponsorship starts at $10k.

Estimated value of Rocket AI: *in the tens of millions.*

Review

- AI is a traditional **sub discipline of computer science** with strong **interdisciplinary roots**

- Among several definitions of the field, to “**act rationally**” lends itself best to **practical exploitation**

- The state of the art comprises **numerous human-level systems** for narrow tasks (“**Artificial Narrow Intelligence**”, as opposed to “Artificial General Intelligence”)

- **AI people** tend to be visionaries motivated by **solving certain applications**, *by the way* discovering new methodologies and principles (like programming paradigms, multi-threading, etc.): **“Just do it”**

APPENDIX

Zurich University of Applied Sciences

Secrets of success

Do **program**, do **take notes** (yourself)!

Formulate goals, cross-connect knowledge

Know your learning style (Kolb): do you need additional material to our deductive (general theory → example) approach to AI?

«Most of the things you *need* will *be brought* to you; most of the things you *want* you have to *go get*.»
(Bill Johnson)

Use self study possibilities

Unintentionally funny stories

One day Joe Bear was hungry. He asked his friend Irving Bird where some honey was. Irving **told him there was a beehive** in the oak tree. Joe threatened to **hit Irving if he didn't tell him where some honey** was. The End.

Joe Bear was hungry. He asked Irving Bird where some honey was. Irving refused to tell him, so **Joe offered to bring him a worm** if he'd tell him where some honey was. Irving agreed. But Joe didn't know where any worms were, so he asked Irving, who refused to say. So **Joe offered to bring him a worm** if he'd tell him where a worm was. Irving agreed. But Joe didn't know where any worms were, so he asked Irving, who refused to say. So **Joe offered to bring him a worm** if he'd tell him where a worm was**…**

Henry Squirrel was thirsty. He walked over to the river bank where his good friend Bill Bird was sitting. Henry slipped and fell in the river. **Gravity drowned**. The End.

Once upon a time there was a dishonest fox and a vain crow. One day the crow was sitting in his tree, holding a piece of cheese in his mouth. He **noticed that he was holding the piece of cheese. He became hungry, and swallowed the cheese.** The fox walked over to the crow. The End.

Google Acquires Artificial Intelligence Startup DeepMind For More Than $500M

Posted Jan 26, 2014 by Catherine Shu (@catherineshu)

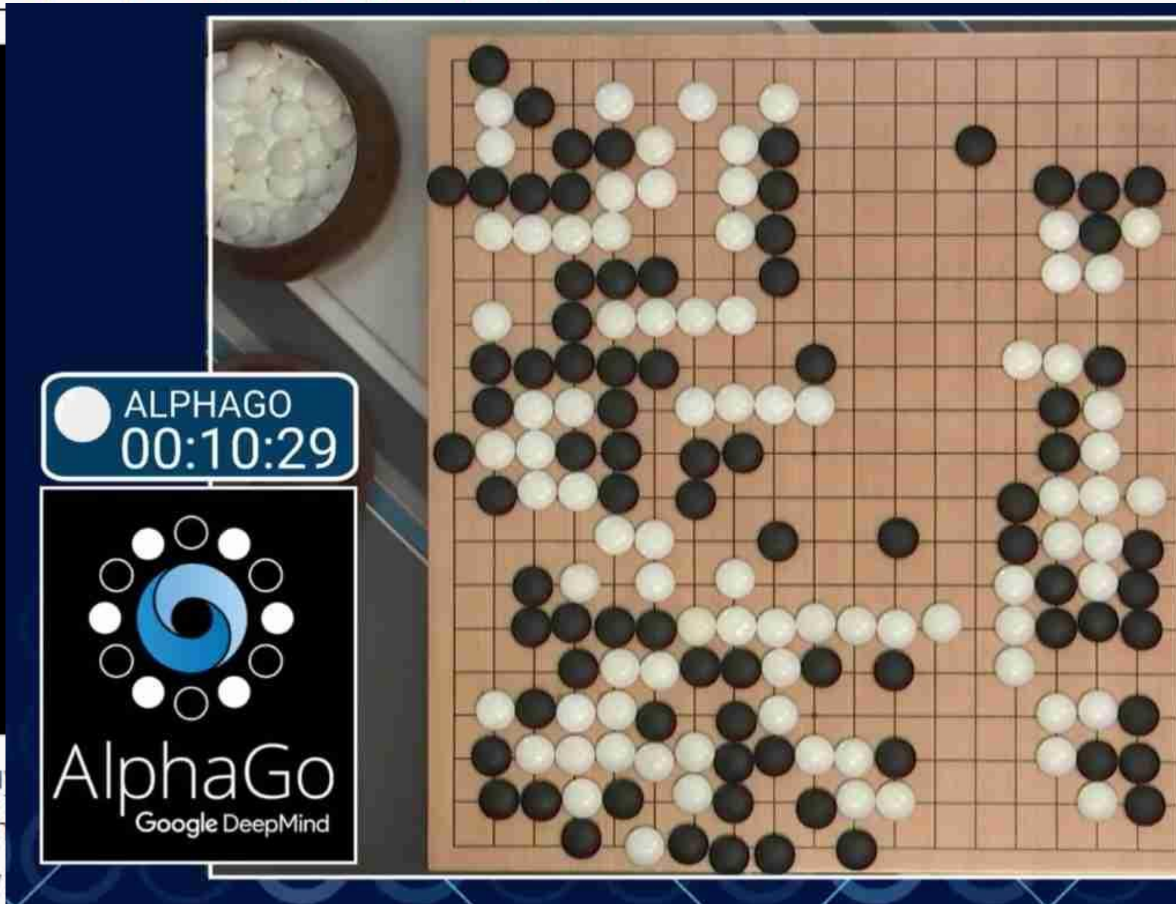

Google will bu
reports that th
in talks to buy
couldn't disclose deal terms.

The acquisition was originally confirmed by Google to Re/code.

Google's WaveNet uses neural nets to generate eerily convincing speech and music

Posted Sep 9, 2016 by Devin Coldewey

Generating speech from a piece of text is a common and important task undertaken by computers, but it's pretty rare that the result could be mistaken for ordinary speech. A new technique from researchers at Alphabet's DeepMind takes a completely different approach, producing speech and even music that sounds eerily like the real thing.

Early systems used a large library of the parts of speech (phonemes and morphemes) and a large ruleset that described all the ways letters combined to produce those sounds. The pieces were joined, or concatenated, creating functional speech synthesis that can handle most words, albeit with unconvincing cadence and tone. Later systems parameterized the generation of sound, making a library of speech fragments unnecessary. More compact — but often less effective.

WaveNet, as the system is called, takes things deeper. It simulates the sound of speech at as low a level as possible: one sample at a time. That means building the waveform from scratch — 16,000 samples per second.

Crunchbase

Google –

Zurich University of Applied Sciences

zh aw

Generated speech from text

Generated music out of creativity

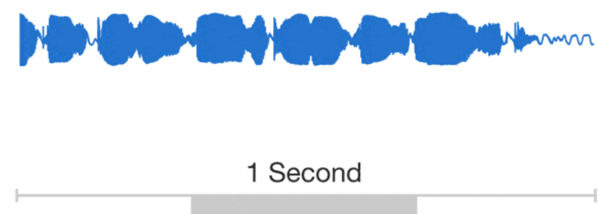

MIT
Techn
Revie

 +  = 

Computing

Algorithm Clones Van Gogh's Artistic Style and Pastes It onto Other Images, Movies

A deep neural network has learned to transfer artistic styles to other images.

by Emerging Technology from the arXiv May 10, 2016

 +  = 

The nature of artistic style is something of a mystery to most people. Think
of Vincent Van Gogh's *Starr* … , or
Edvard Munch's *The Scream* … e th
humans recognize easily.

 +  = 

Zurich University
of Applied Sciences

…and the list could be continued


Brandon Amos About Blog

Image Completion with Deep Learning in TensorFlow

August 9, 2016

- Introduction
- Step 1: Interpreting images as samples from a probability distribution
  - How would you fill in the missing information?
  - But where does statistics fit in? These are images.
  - So how can we complete images?
- Step 2: Quickly generating fake images
  - Learning to generate new samples from an unknown probability distribution
  - [ML-Heavy] Generative Adversarial Net (GAN) building blocks
  - Using $G(z)$ to produce fake images
  - [ML-Heavy] Training DCGANs
  - Existing GAN
  - [ML-Heavy]
  - Running DCG
- Step 3: Finding the
  - Image compl
  - [ML-Heavy]
  - [ML-Heavy]
  - Completing y
- Conclusion
- Partial bibliography
- Bonus: Incomplete

Introduction

Content-aware fill is a po completion and inpaintin do content-aware fill, im "Semantic Image Inpaint shows how to use deep l some deeper portions fo section can be skipped i from images of faces. I h completion.tensorflow.

We'll approach image co

1. We'll first interpret
2. This interpretation
3. Then we'll find the

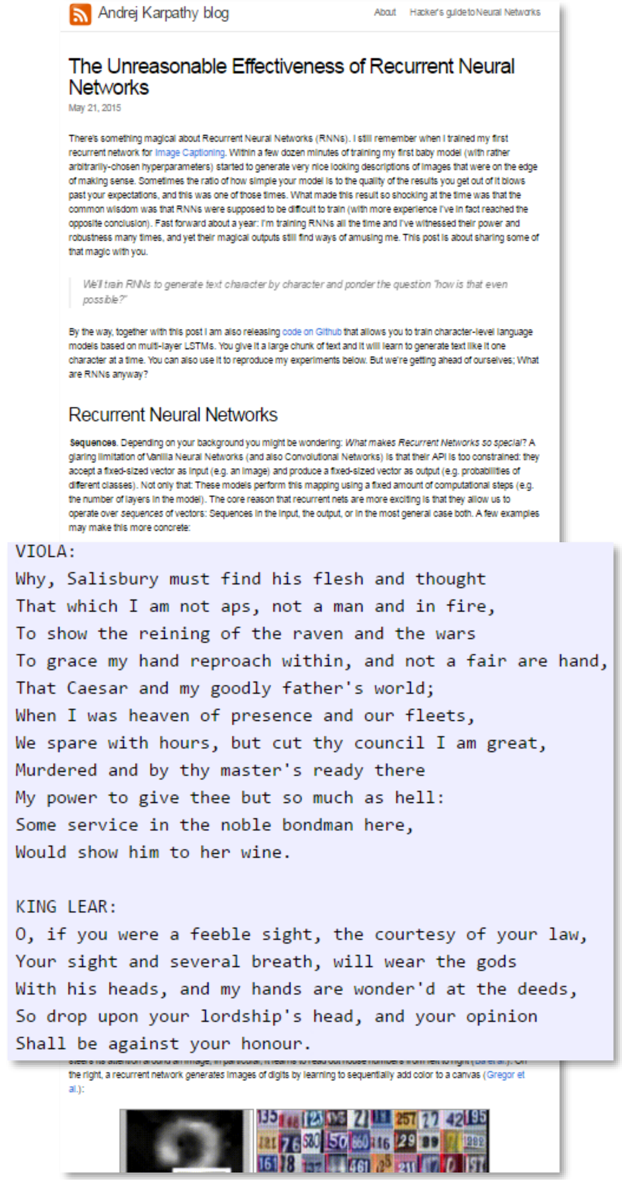
Andrej Karpathy blog About Hacker's guide to Neural Networks

The Unreasonable Effectiveness of Recurrent Neural Networks

May 21, 2015

There's something magical about Recurrent Neural Networks (RNNs). I still remember when I trained my first recurrent network for Image Captioning. Within a few dozen minutes of training my first baby model (with rather arbitrarily-chosen hyperparameters) started to generate very nice looking descriptions of images that were on the edge of making sense. Sometimes the ratio of how simple your model is to the quality of the results you get out of it blows past your expectations, and this was one of those times. What made this result so shocking at the time was that the common wisdom was that RNNs were supposed to be difficult to train (with more experience I've in fact reached the opposite conclusion). Fast forward about a year: I'm training RNNs all the time and I've witnessed their power and robustness many times, and yet their magical outputs still find ways of amusing me. This post is about sharing some of that magic with you.

> *We'll train RNNs to generate text character by character and ponder the question "how is that even possible?"*

By the way, together with this post I am also releasing code on Github that allows you to train character-level language models based on multi-layer LSTMs. You give it a large chunk of text and it will learn to generate text like it one character at a time. You can also use it to reproduce my experiments below. But we're getting ahead of ourselves; What are RNNs anyway?

Recurrent Neural Networks

Sequences. Depending on your background you might be wondering: *What makes Recurrent Networks so special?* A glaring limitation of Vanilla Neural Networks (and also Convolutional Networks) is that their API is too constrained: they accept a fixed-sized vector as input (e.g. an image) and produce a fixed-sized vector as output (e.g. probabilities of different classes). Not only that: These models perform this mapping using a fixed amount of computational steps (e.g. the number of layers in the model). The core reason that recurrent nets are more exciting is that they allow us to operate over sequences of vectors: Sequences in the input, the output, or in the most general case both. A few examples may make this more concrete:

```
VIOLA:
Why, Salisbury must find his flesh and thought
That which I am not aps, not a man and in fire,
To show the reining of the raven and the wars
To grace my hand reproach within, and not a fair are hand,
That Caesar and my goodly father's world;
When I was heaven of presence and our fleets,
We spare with hours, but cut thy council I am great,
Murdered and by thy master's ready there
My power to give thee but so much as hell:
Some service in the noble bondman here,
Would show him to her wine.

KING LEAR:
O, if you were a feeble sight, the courtesy of your law,
Your sight and several breath, will wear the gods
With his heads, and my hands are wonder'd at the deeds,
So drop upon your lordship's head, and your opinion
Shall be against your honour.
```

the right, a recurrent network generates images of digits by learning to sequentially add color to a canvas (Gregor et al.):

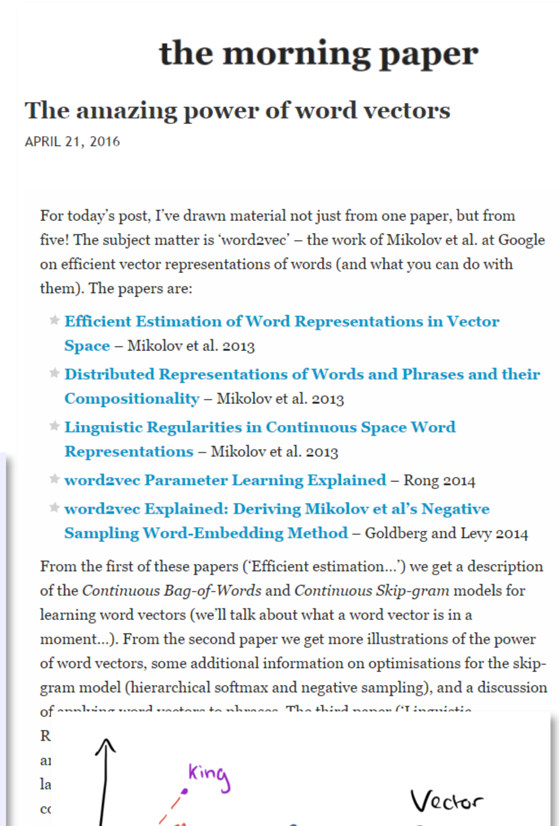
the morning paper

The amazing power of word vectors

APRIL 21, 2016

For today's post, I've drawn material not just from one paper, but from five! The subject matter is 'word2vec' – the work of Mikolov et al. at Google on efficient vector representations of words (and what you can do with them). The papers are:

- **Efficient Estimation of Word Representations in Vector Space** – Mikolov et al. 2013
- **Distributed Representations of Words and Phrases and their Compositionality** – Mikolov et al. 2013
- **Linguistic Regularities in Continuous Space Word Representations** – Mikolov et al. 2013
- **word2vec Parameter Learning Explained** – Rong 2014
- **word2vec Explained: Deriving Mikolov et al's Negative Sampling Word-Embedding Method** – Goldberg and Levy 2014

From the first of these papers ('Efficient estimation…') we get a description of the *Continuous Bag-of-Words* and *Continuous Skip-gram* models for learning word vectors (we'll talk about what a word vector is in a moment…). From the second paper we get more illustrations of the power of word vectors, some additional information on optimisations for the skip-gram model (hierarchical softmax and negative sampling), and a discussion