# Commentary on Lectures By Hinton and LeCun

**John F. Sowa**

**Talks by Hinton and LeCun, 23 June 2019**

**Commentary: Revision of 15 February 2022**

# Lectures and Commentary

ACM named Yoshua Bengio, Geoffrey Hinton, and Yann LeCun recipients of the 2018 Turing Award for "conceptual and engineering breakthroughs that have made deep neural networks a critical component of computing."

On 23 June 2019, Hinton and LeCun presented versions of their 2018 talks. These slides are based on selected screen shots from those talks with commentary that relates them to issues in AI and cognitive science:


1. Lecture by Hinton . . . . . . . . . . . . . . . . . . . . . . . . . . Slide 3

2. Lecture by LeCun . . . . . . . . . . . . . . . . . . . . . . . . . . . . 21

3. Cognitive learning and reasoning . . . . . . . . . . . . . . . . . . . . 25


The 2019 video: https://www.youtube.com/watch?v=VsnQf7exv5I&feature=youtu.be

The complete video takes 92 minutes. Hinton's first slide begins at the 10 minute mark, LeCun's first slide begins at 42 minutes. Each screen shot is marked with the time it appears in the video.

To align the commentary to the video, move the red dot to the time point for any screen shot. Then listen to what Hinton or LeCun says about it.

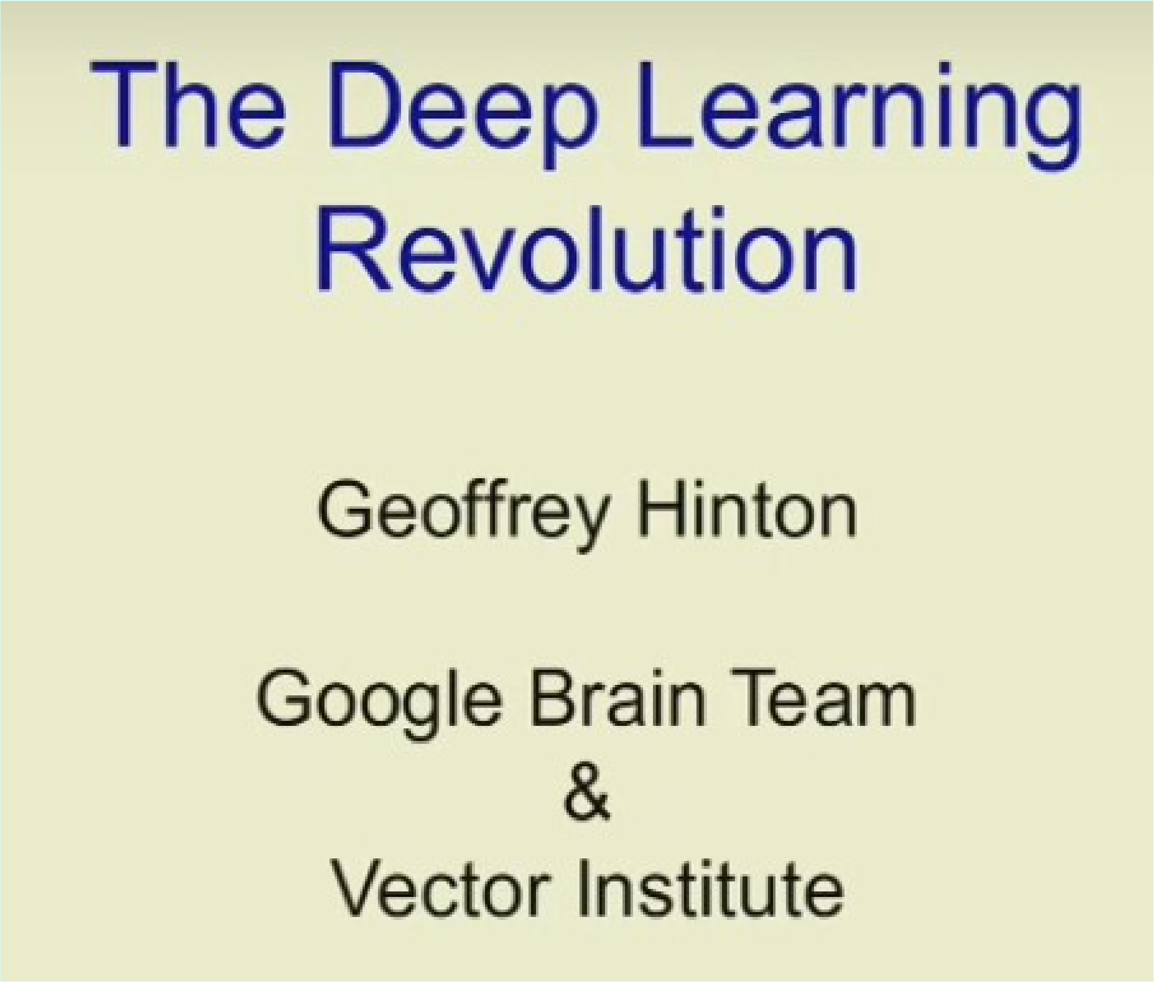


**Source:** https://www.youtube.com/watch?v=VsnQf7exv5I&feature=youtu.be

**To hear Hinton's lecture, move the red dot to the 10 minute mark.**

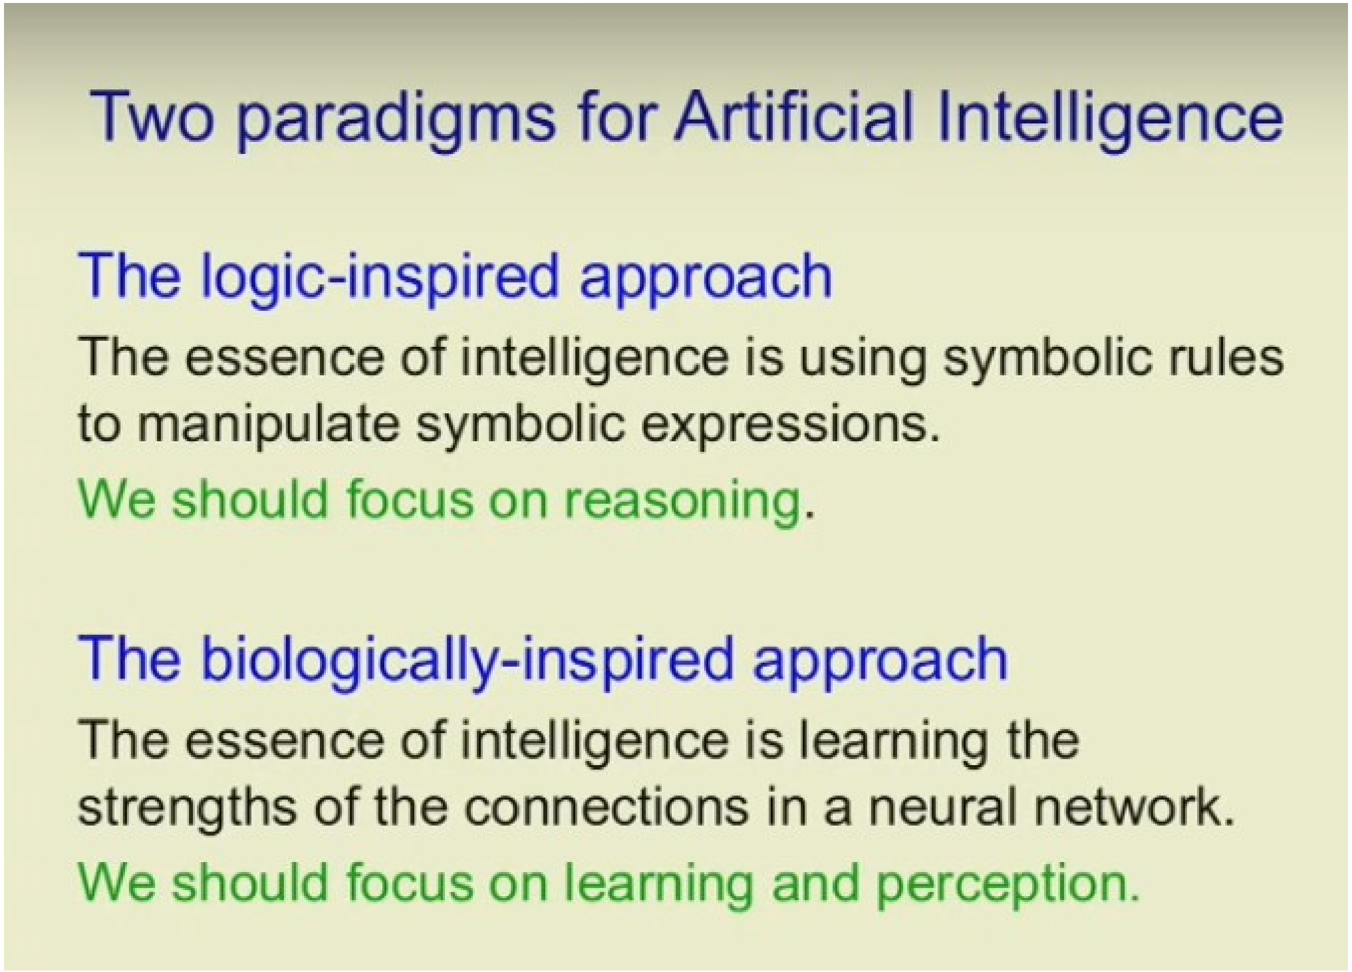


**But there is no limit to the number and sources of AI paradigms.**

- **AI researchers took the lead in organizing cognitive science as a coalition with philosophy, psychology, linguistics, anthropology, and neuroscience.**
- **See the next slide for Marvin Minsky's biologically inspired paradigms.**

**Time: 11 m**

# Hybrid Systems to Support Diversity

## Flexibility and generality are key to intelligence.

- The languages of our stone-age ancestors can be adapted to any subject: science, technology, business, law, finance, and the arts.
- When people invent anything, they find ways to describe it.
- When people in any culture adopt anything from another culture, they borrow or adapt words to describe it in their native language.

## Minsky's theory: A society of heterogeneous modules.

"What magical trick makes us intelligent? The trick is that there is no trick. The power of intelligence stems from our vast diversity, not from any single, perfect principle. Our species has evolved many effective although imperfect methods, and each of us individually develops more on our own. Eventually, very few of our actions and decisions come to depend on any single mechanism. Instead, they emerge from conflicts and negotiations among societies of processes that constantly challenge one another." *

* Marvin Minsky (1986) *The Society of Mind*, New York: Simon & Schuster, §30.8.
See also Push Singh & Marvin Minsky (2005), An architecture for cognitive diversity.

## Two ways to make a computer do what you want

- Intelligent design: Figure out consciously exactly how you would manipulate symbolic representations to perform the task and then tell the computer, in excruciating detail, exactly what to do.

- Learning: Show the computer lots of examples of inputs together with the desired outputs. Let the computer learn how to map inputs to outputs using a general purpose, learning procedure.

**This slide describes learning by perception:**

- **Given a set of input-output pairs (x,y}, derive a function from x to y.**
- **Each x Is an input vector derived by some method of perception.**
- **Each y is an output vector that describes, identifies, or classifies x.**
- **The system learns a function f from any x to the corresponding y: y = f(x).**

**Time: 12 m**

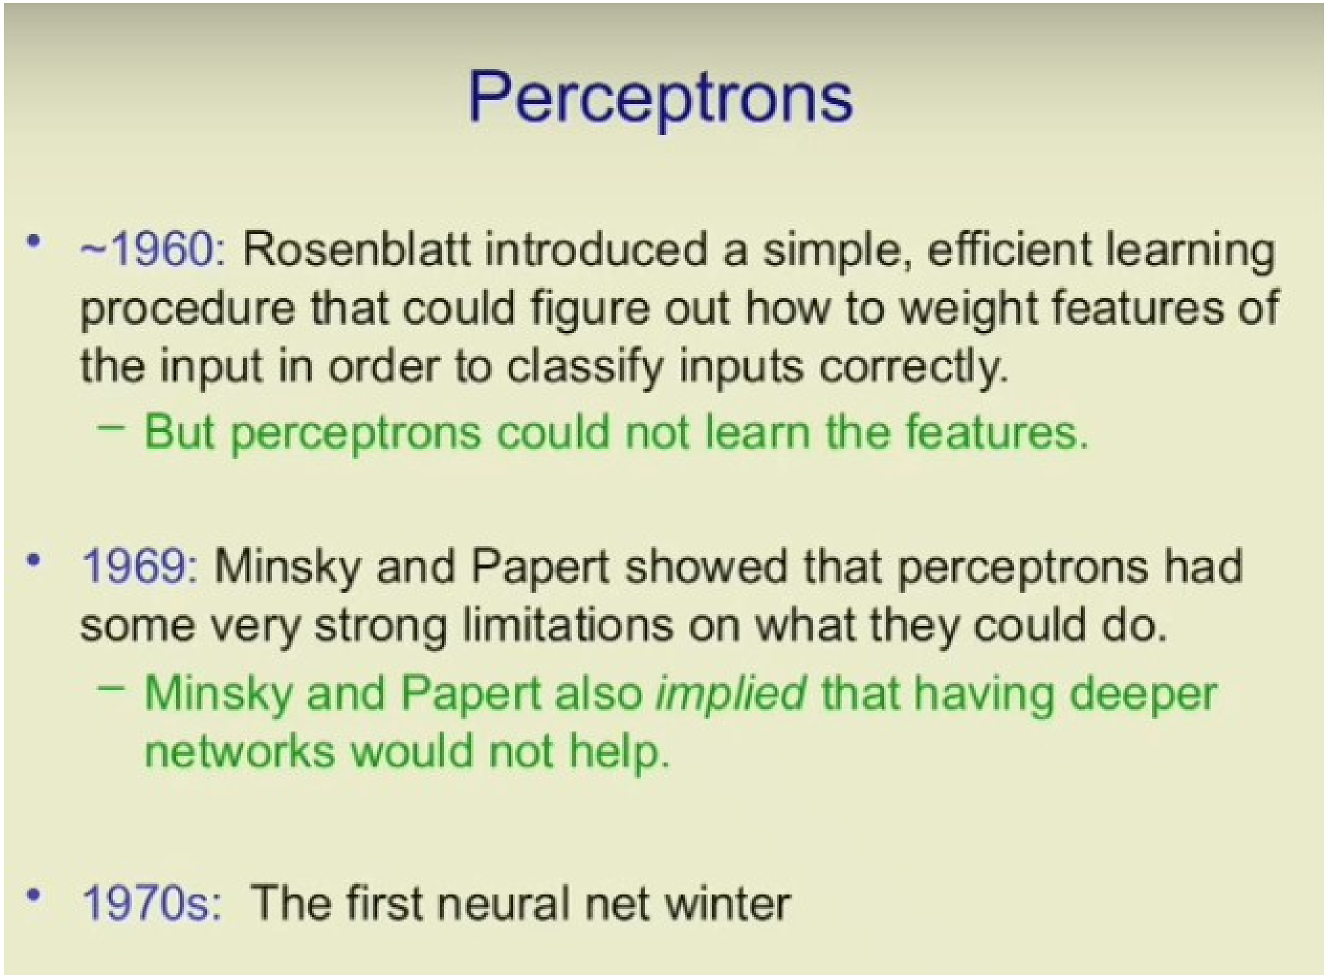


**This slide is correct, but the story is more complex.**

- **The 1970 recession caused funding to dry up for every branch of AI.**
- **Rosenblatt's own hype did as much or more to kill perceptrons as M & P's book. See the next three slides.**

**Time: 14:30 m**

# The Perceptron

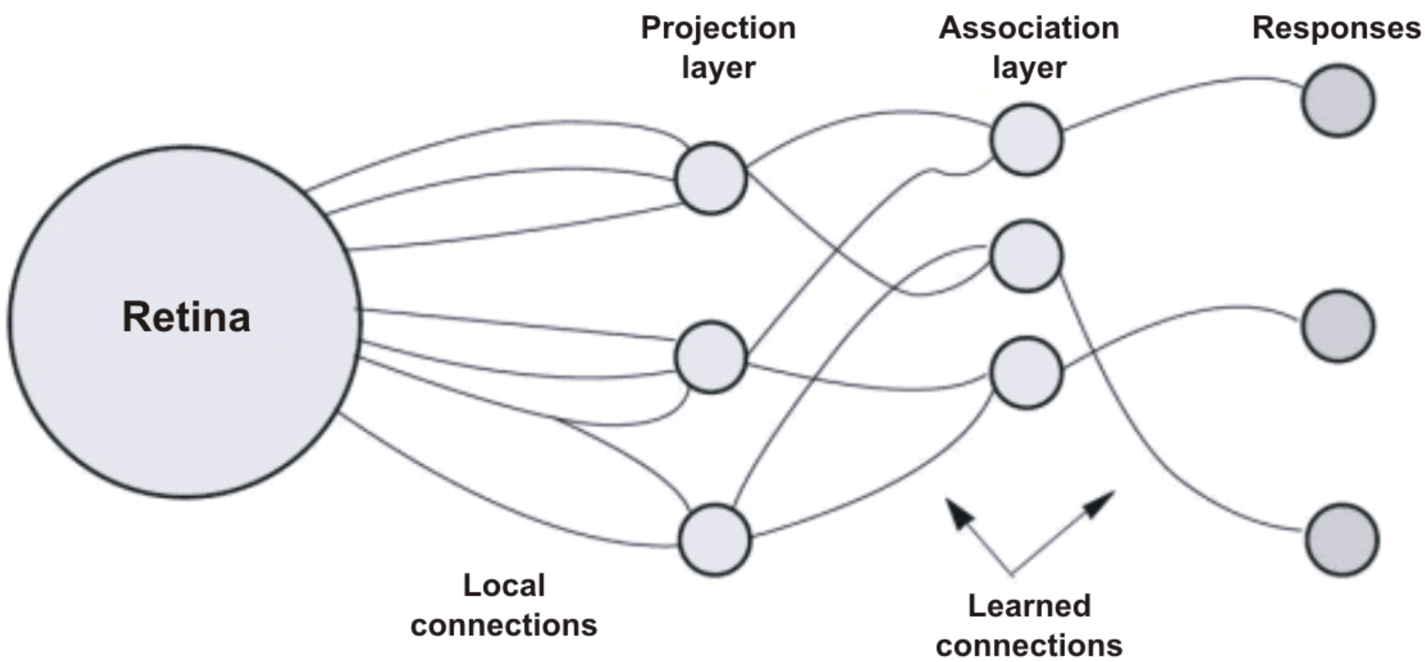


**One-layer neural network invented by Frank Rosenblatt (1957).**

**Mark I: a hardware version funded by the US Navy:**

- **Input: 400 photocells in a 20 x 20 array.**
- **Weights represented by potentiometers updated by electric motors.**

**The New York Times, after a press conference in 1958:**

> ***The perceptron is "the embryo of an electronic computer that [the Navy] expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence." ****



* **http://query.nytimes.com/gst/abstract.html?res=9D03E4D91F3AE73ABC4B52DFB1668383649EDE**

# A Successful Hybrid System

**Program for playing checkers by Art Samuel in 1959:**

- Ran on the IBM 704, later on the IBM 7090.
- The IBM 7090 was comparable in speed to the original IBM PC (1981), and its maximum RAM was only 144K bytes.

**Samuel's program was a hybrid that combined two paradigms:**

- A method for learning a function that evaluates game positions.
- The alpha-beta algorithm for searching game trees.

**Won a game against the Connecticut state checkers champion.**

# Take Advantage of Available Tools

Over sixty years of R & D in AI and computational linguistics.

Tools and resources for a wide variety of paradigms:

- Parsers and translators for natural and artificial languages.
- Grammars, lexicons, ontologies, terminologies, corpora, the Semantic Web, and many ways of linking and indexing.
- Theorem provers and inference engines for formal logic and many kinds of informal and fuzzy reasoning.
- Qualitative, case-based, and analogical reasoning.
- Statistics, data mining, and image data mining,
- NNs, genetic algorithms, and other ML methods.
- Thousands of implementations of all the above.

Like Samuel's checker program, the Alpha Go system, which beat the world Go champion, was a hybrid. It combined an alpha-beta search strategy with a DNN for learning the evaluation function.

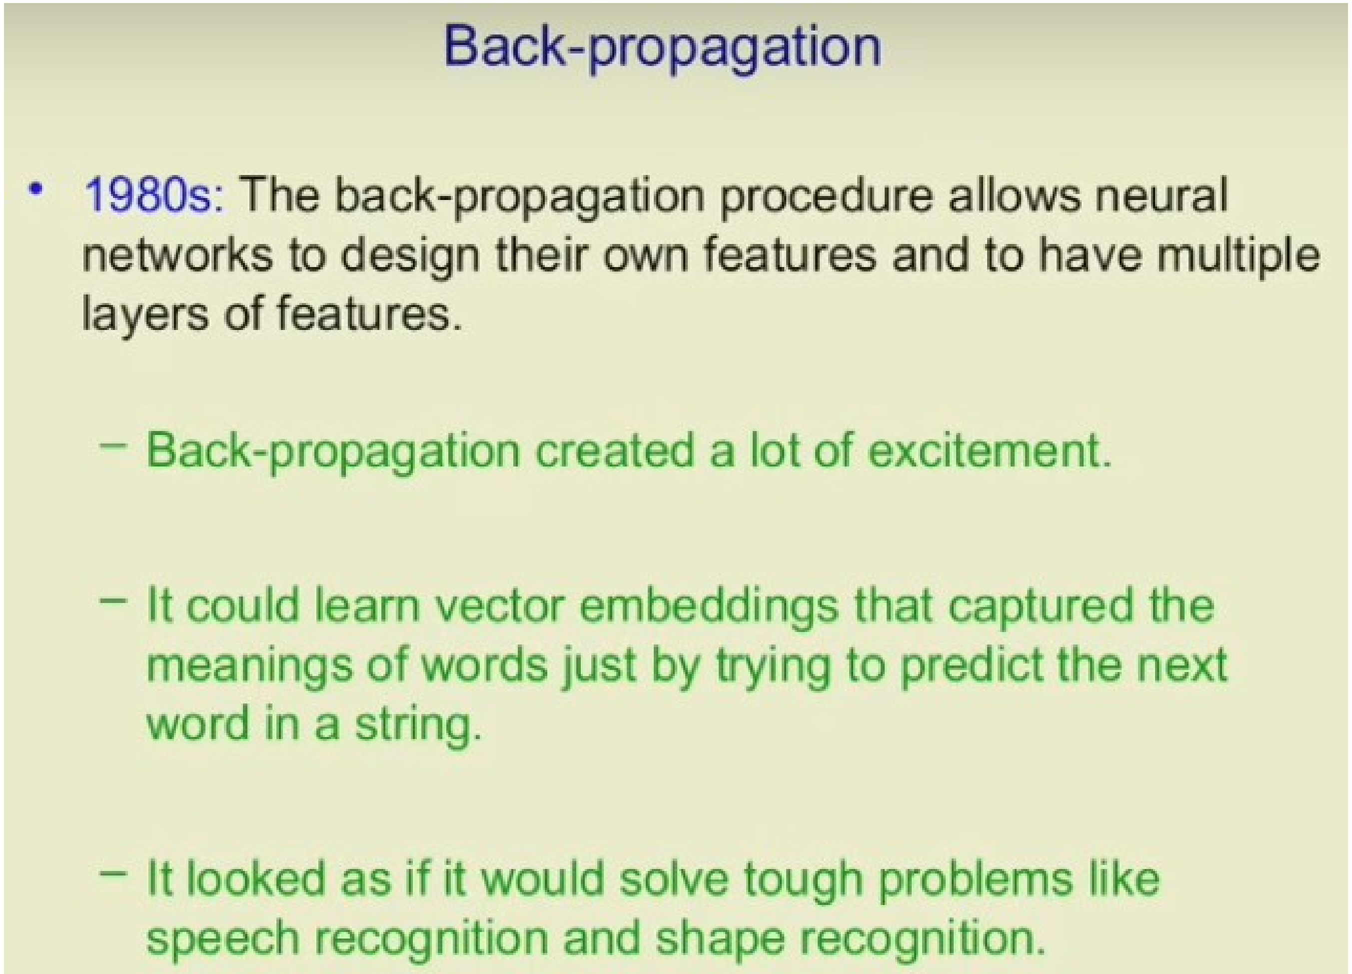


**Paul Werbos invented backpropagation for his PhD thesis (1974).**

- **For his recent papers and slides, see http://www.werbos.com .**
- **For examples, http://www.werbos.com/SPIE_NN_2015_AI_BI_to_NN_CI.pdf**

**Time: 15:30 m**

## How do we train artificial neural networks?

- Supervised training: Show the network an input vector and tell it the correct output.
  - Adjust the weights to reduce the discrepancy between the correct output and the actual output.

- Unsupervised training: Only show the network the input.
  - Adjust the weights to get better at reconstructing the input (or parts of the input) from the activities of the hidden neurons.

**These methods were designed to train animals (AKA rat psychology).**

- **Educational psychologists consider them inadequate for teaching humans.**
- **But various modules in the brain, each consisting of specialized NNs, can supervise the training of other modules.** **See the next slide.**

**Time: 17:30 m**

# Modules Supervising Modules

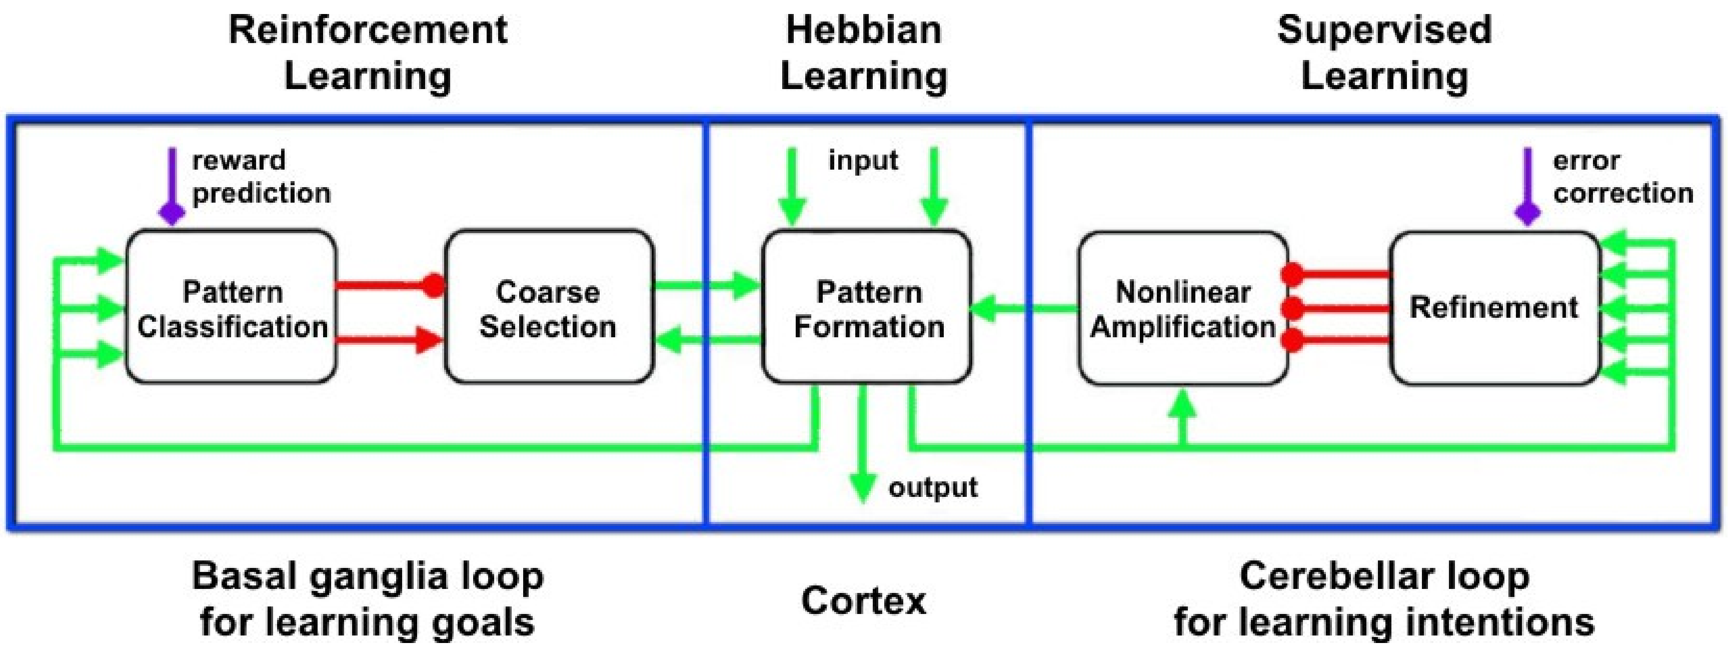


**Evidence that supports both Minsky's theory and NN theories:**

- **Cerebral cortex: Learn patterns by Hebbian-style learning.**
- **Basal ganglia: Learn production rules (if-then style) with dopamine rewards generated by a signal from the cortex to the substantia nigra.**
- **Cerebellum: Supervised learning with the cortex as supervisor.**
- **See "Towards a systems-level view of cerebellar function: the interplay between cerebellum, basal ganglia, and cortex," Caligiore et al. (2016) https://link.springer.com/article/10.1007/s12311-016-0763-3**

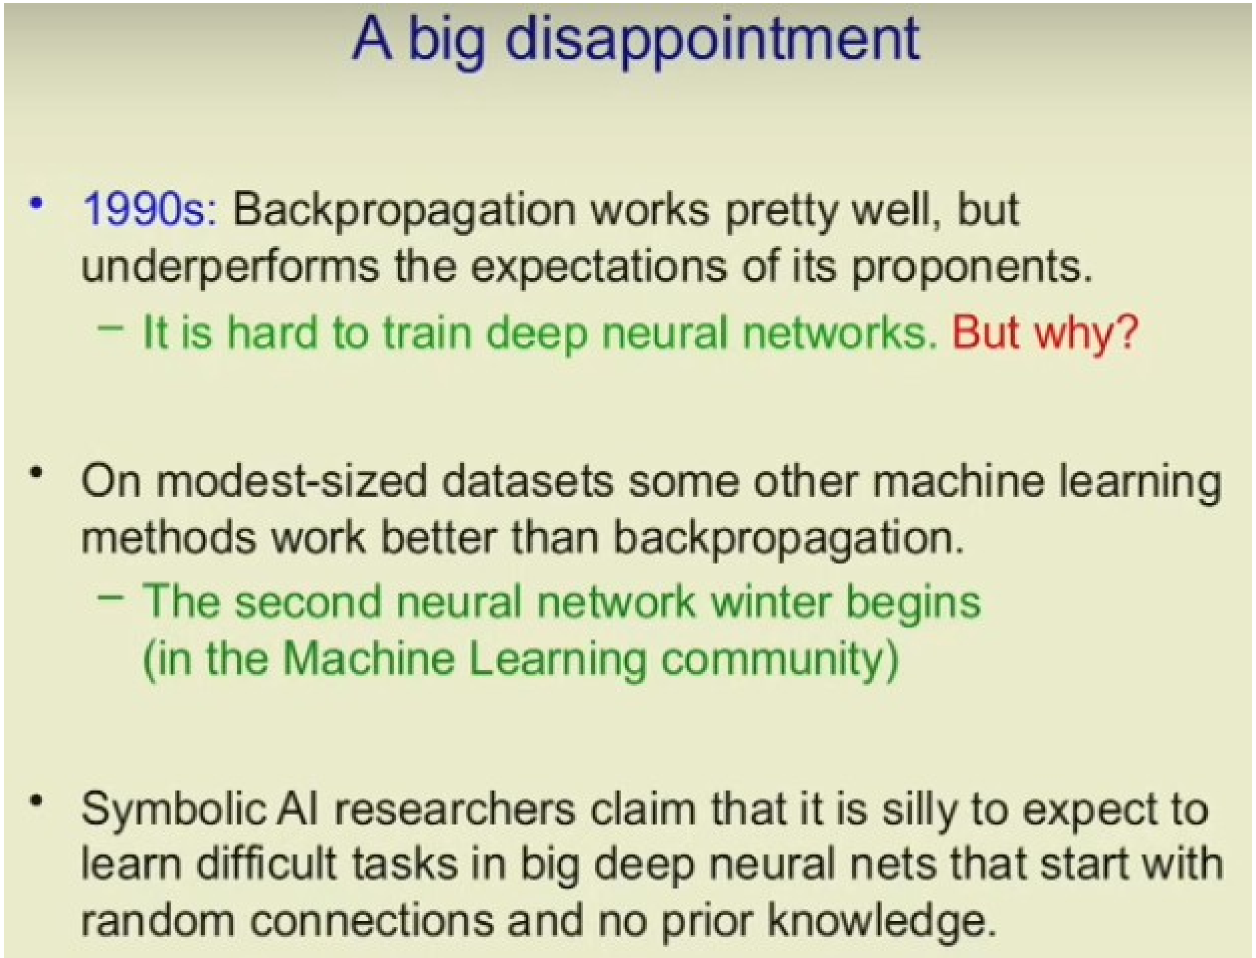


**Variety of ML methods summarized in the 1992 Encyclopedia of AI:**

- **Analytic learning; belief revision; Boltzmann machine; case-based reasoning; cognitive models; concept hierarchies; decision trees; fuzzy concept learning; genetic algorithms; inductive modeling; neural networks or connectionism; rule-guided inference; semantic networks; schema theory."**
- **These and newer learning methods are used today in conjunction with NNs.**



**Time: 21:30 m**

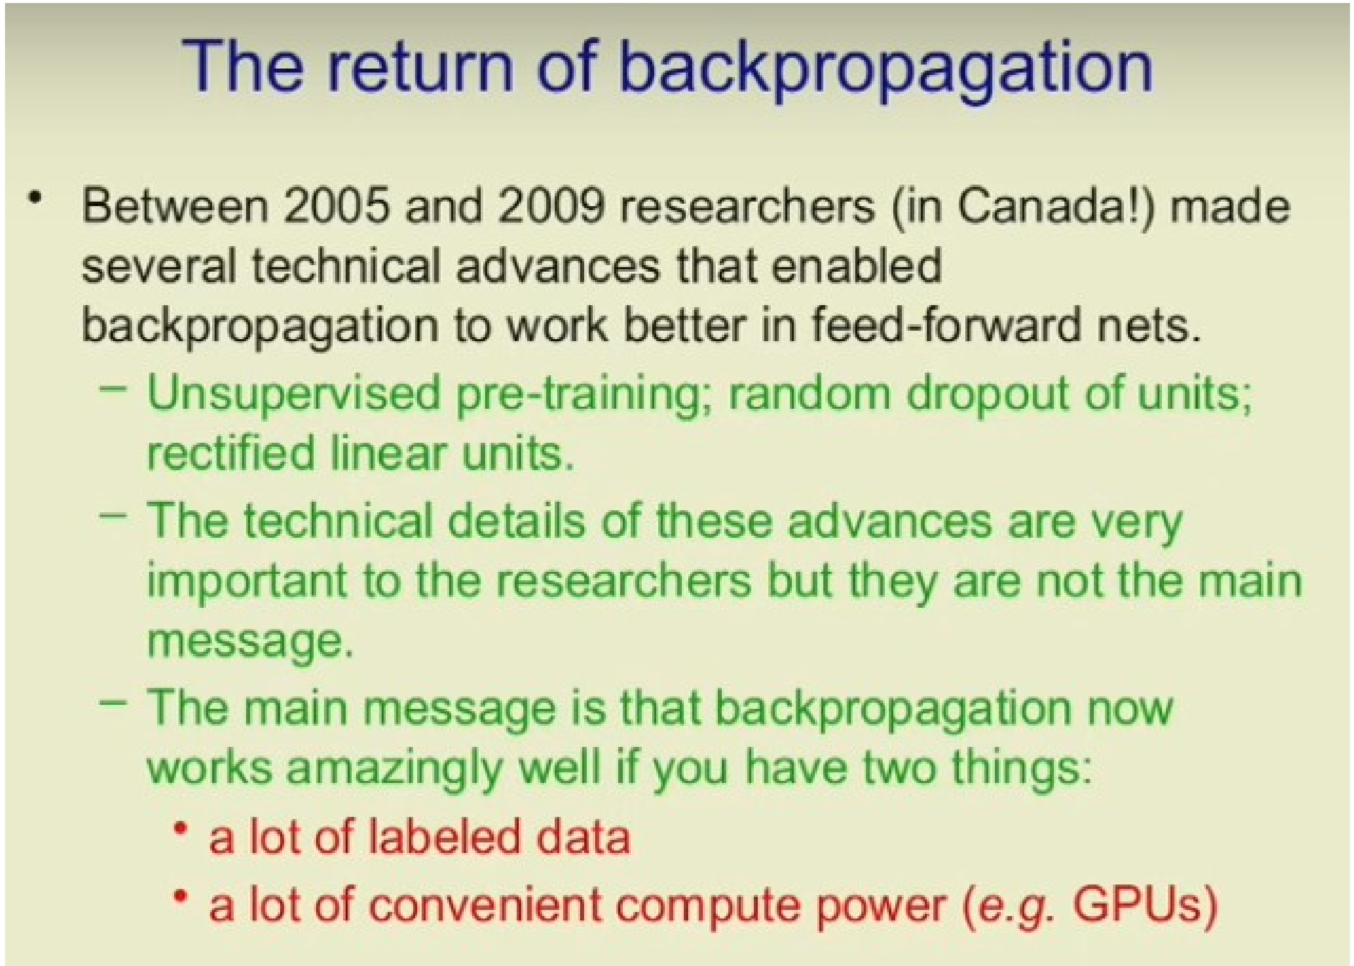


**The success of these methods merits the Turing Award.**

- **They solve a very important problem: perceptual learning and classification.**
- **But they do not solve the issues of cognition. See the next slide.**

**Time: 26 m**

# Applications of DNNs

| Stimulus $x$ | Response $y$ | Application |
|---|---|---|
| Picture | Are there human faces? (0 or 1) | Photo tagging |
| Loan application | Will they repay the loan? (0 or 1) | Loan approvals |
| Ad plus user information | Will user click on ad? (0 or 1) | Targeted online ads |
| Audio clip | Transcript of audio clip | Speech recognition |
| English sentence | French sentence | Language translation |
| Sensors from hard disk, plane engine, etc. | Is it about to fail? | Preventive maintenance |
| Car camera and other sensors | Position of other cars | Self-driving cars |

**Learning a function y=f(x) is only one aspect of intelligence.**

- **Comment by Andrew Ng: Current NN methods automate tasks that take less than a second of mental effort by humans. ***
- **They automate perception, but not cognitive understanding.**
- **Cognition requires the ability to explain what was learned.**

*** Andrew Ng (2016) https://hbr.org/2016/11/what-artificial-intelligence-can-and-cant-do-right-now**

# Thinking Beyond the First Second

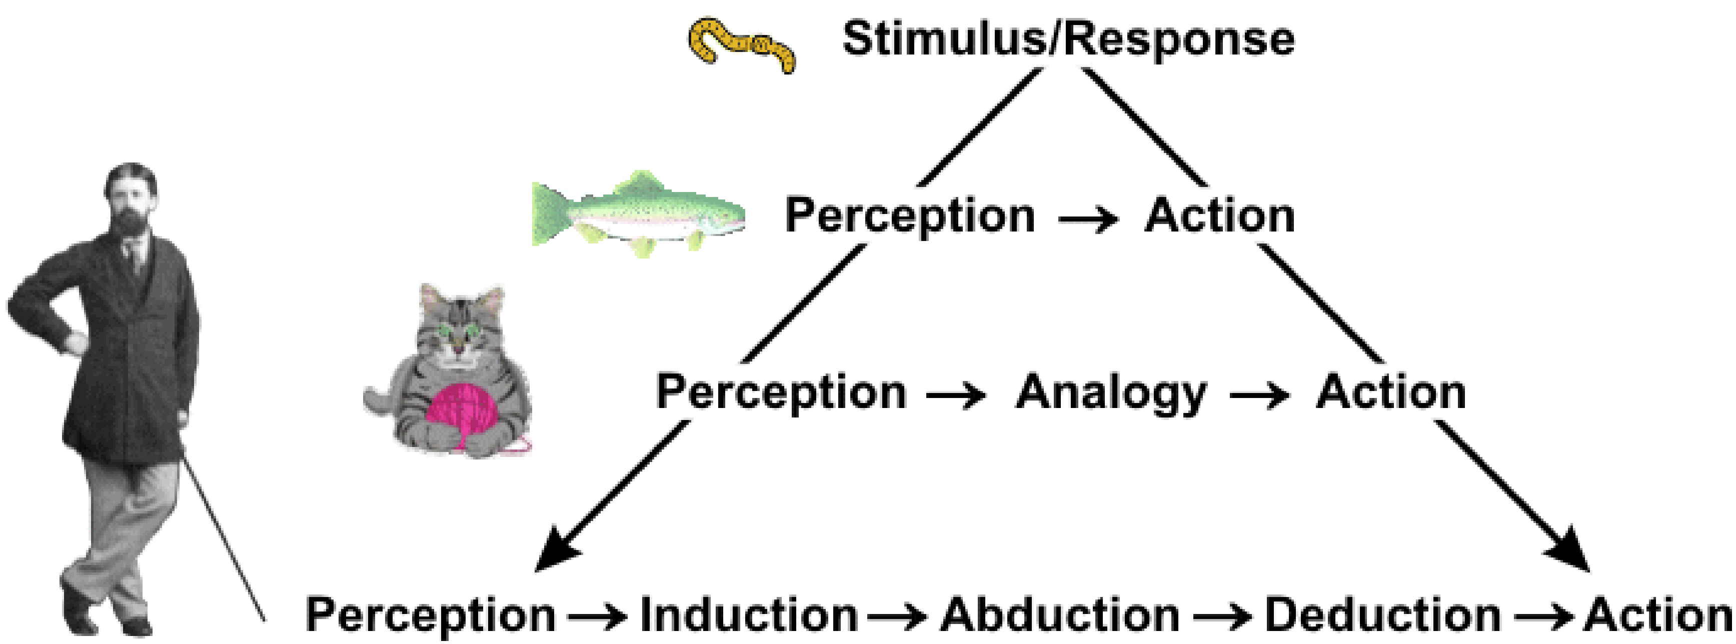


**Perception and classification take one second or less.**

- **Neural nets are valuable for learning and recognizing patterns.**
- **By themselves, NNs support a fish level of intelligence.**
- **With analogies, NNs can support a cat level of intelligence.**

**Analysis, planning, discovery, and innovation take more time.**

- **They require cycles of induction, abduction, deduction...**

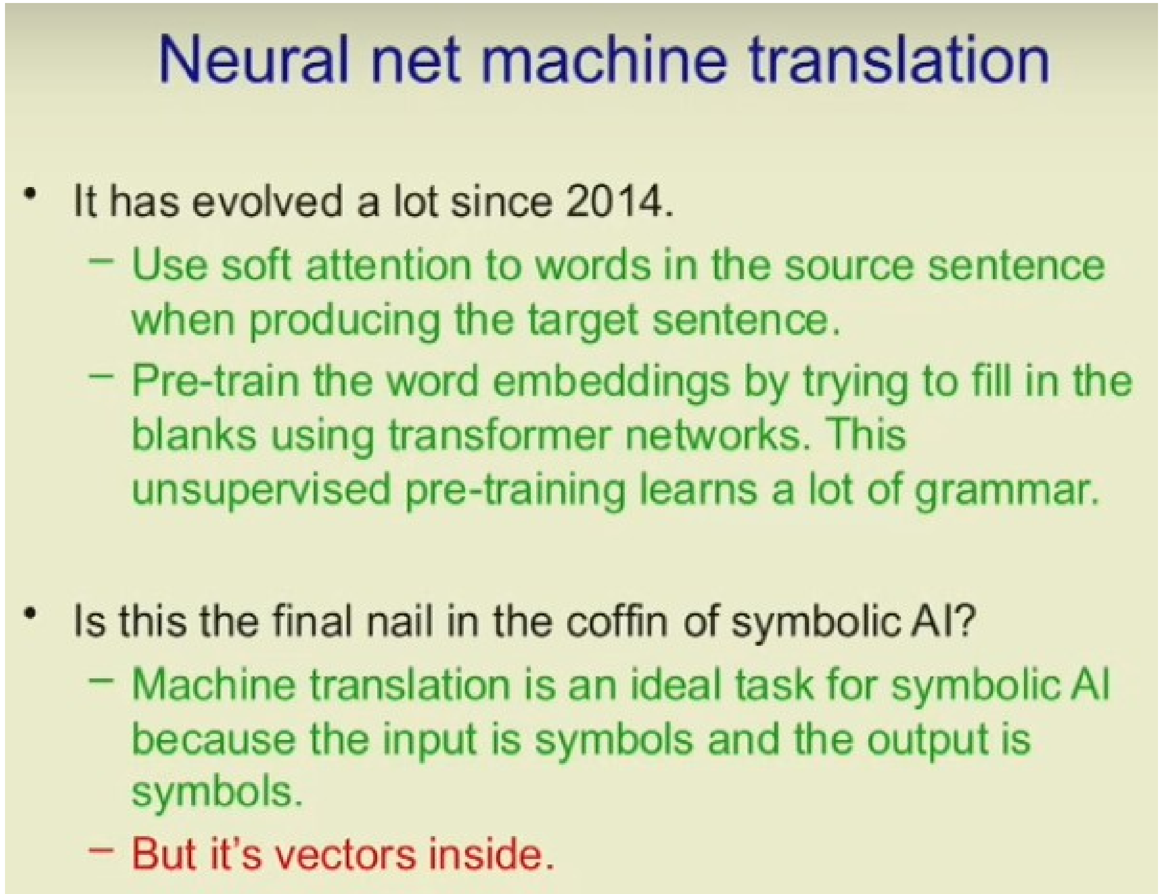


**Current versions of NN-MT learn to map one notation to another.**

- **With NNs, Google translate can generate readable sentences.**
- **They can map a sentence in one language to another (natural or artificial), but they can't explain what they read in either language.** **See the next slide.**

**Time: 32 m**

# Context and Purpose

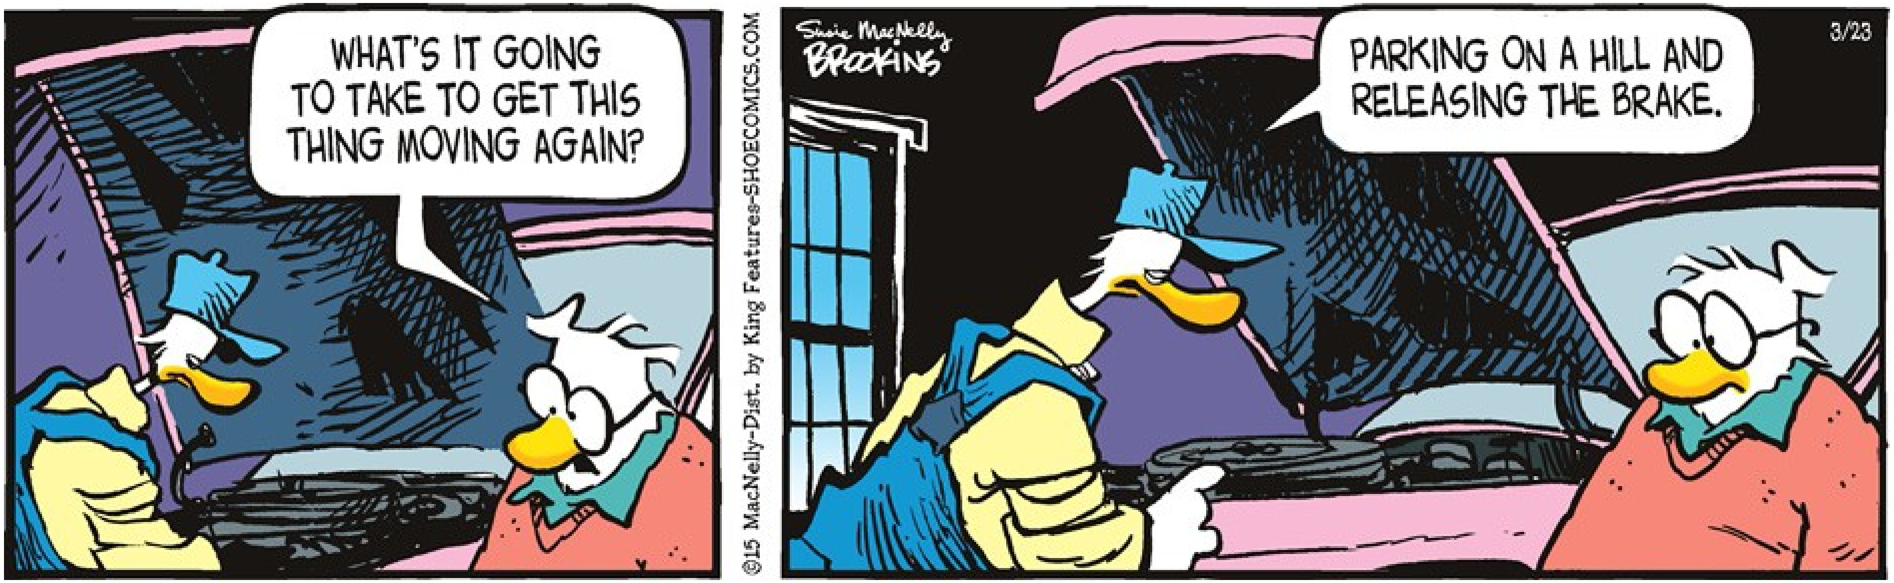


**Syntax is easy: Parse the question and the answer.**

**Semantics depends on context and background knowledge:**

- **Interpret the meaning of *thing*, *take*, and *move* In this situation.**
- **Apply the laws of physics to understand what would happen.**

**Pragmatics depends on the intentions of the participants.**

- **No computer system today can understand that cartoon.**
- **Computers should ask people about purpose or intentions.**

*** Source of cartoon: search for 'moving' at http://www.shoecomics.com/**

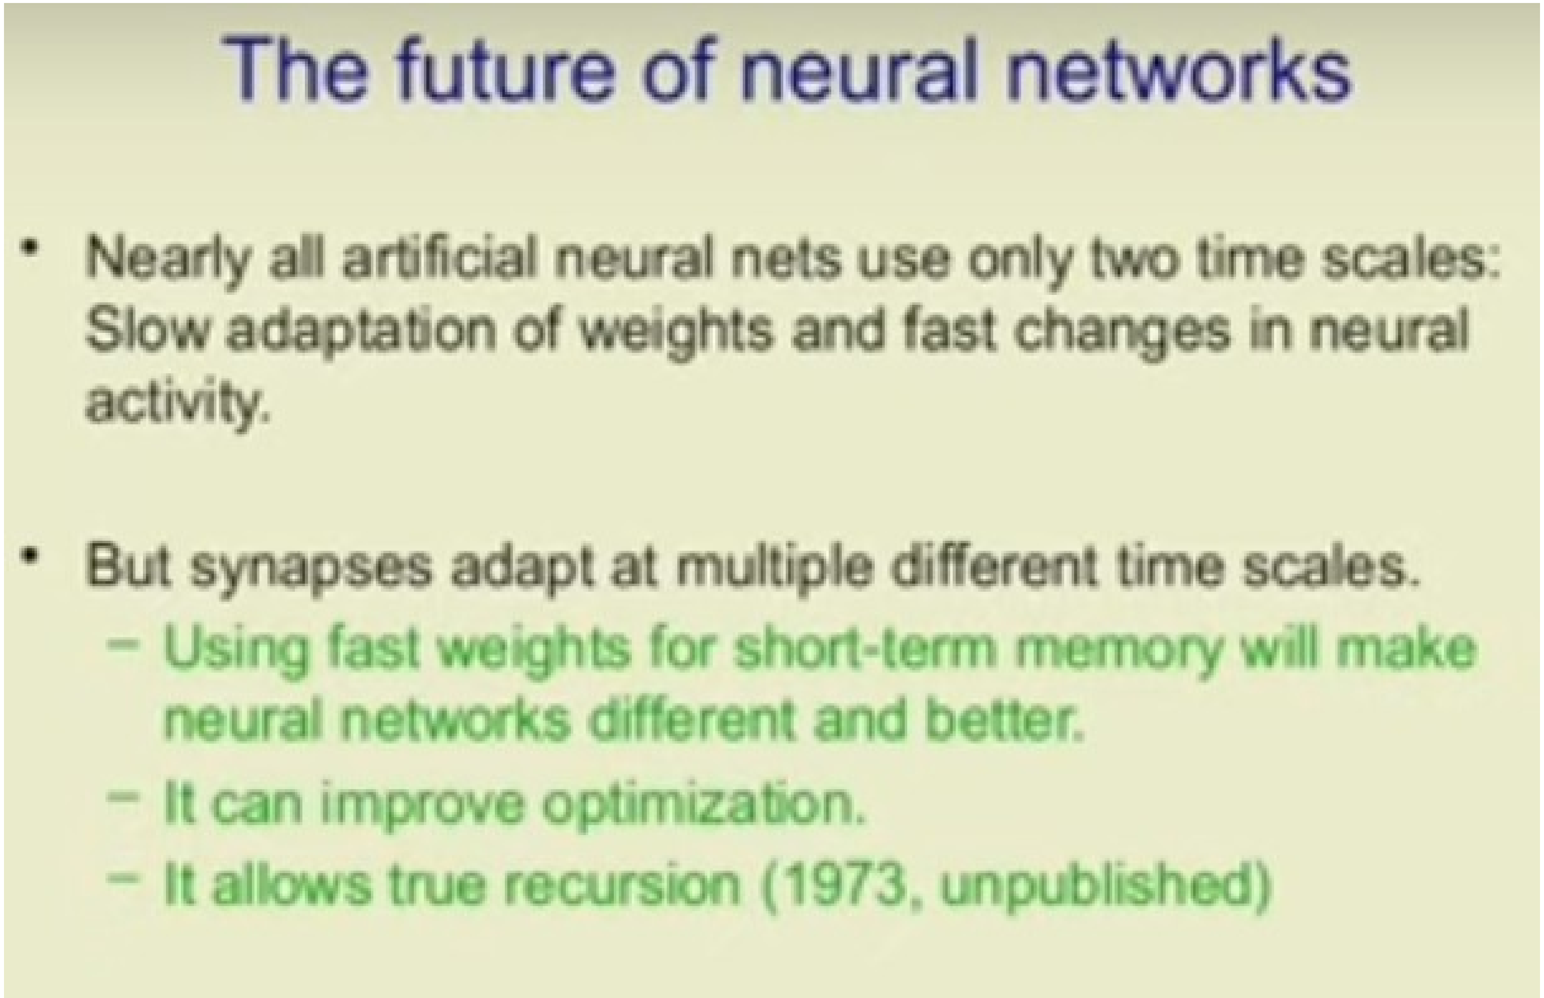


**But those neurons would still be artificial, not biologically realistic.**

- **The brain has many different kinds of neurons for different purposes.**
- **And each neuron is able to encode billions of bits of information.**
- **Neuroscientists admit that they do not know how that information is used.**

**Time: 40 m**

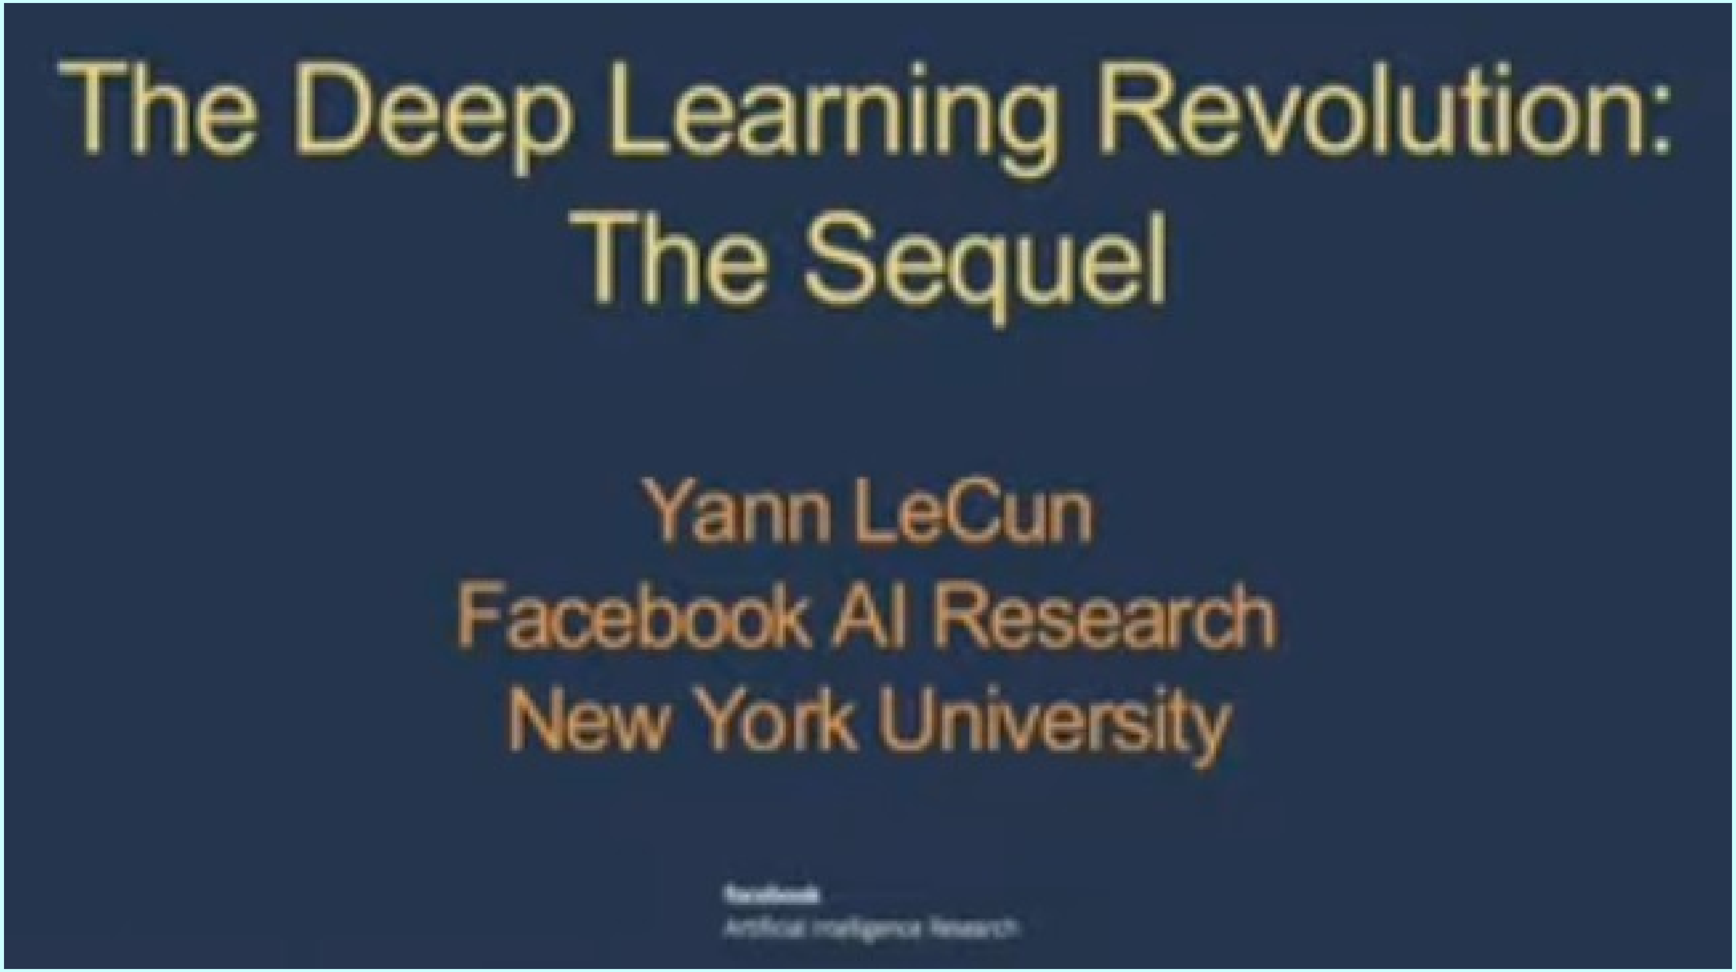
The Deep Learning Revolution:
The Sequel
Yann LeCun
Facebook AI Research
New York University

**Time:  42:45 m**

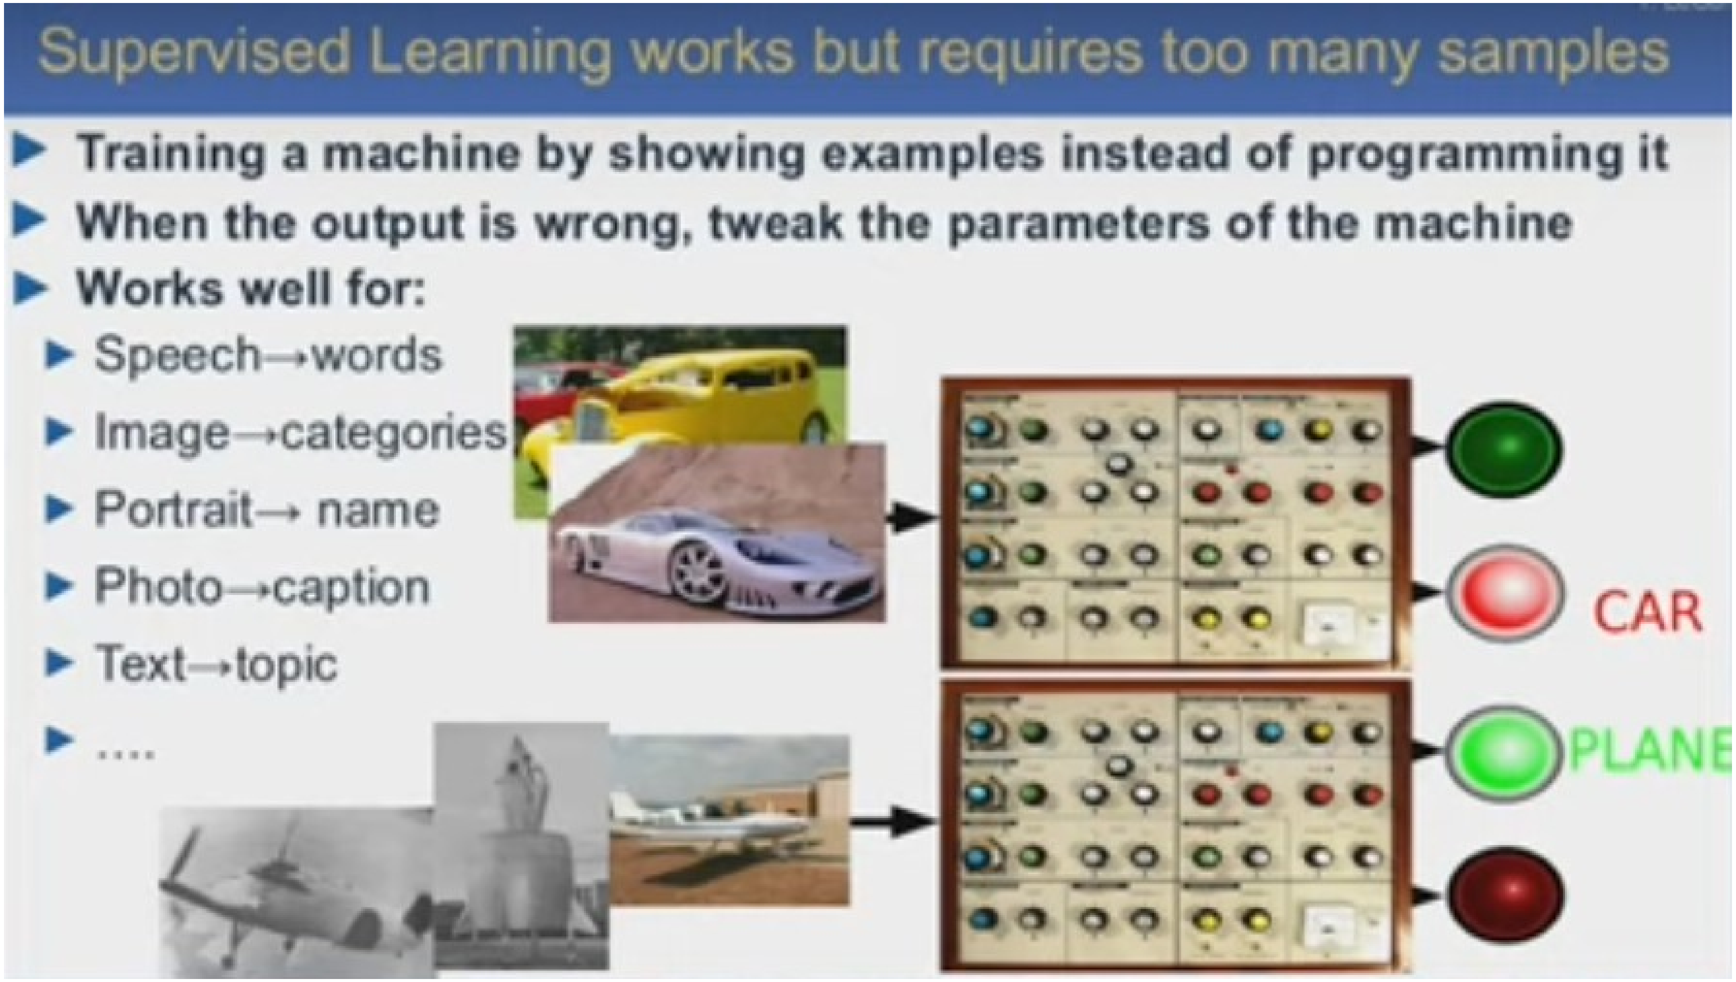


**The word 'works' requires some clarification:**

- **What a NN learns is pattern recognition and classification.**
- **Children have that ability in infancy.  It's a prerequisite for cognitive learning during the rest of their lives..**

**Time:  43 m**

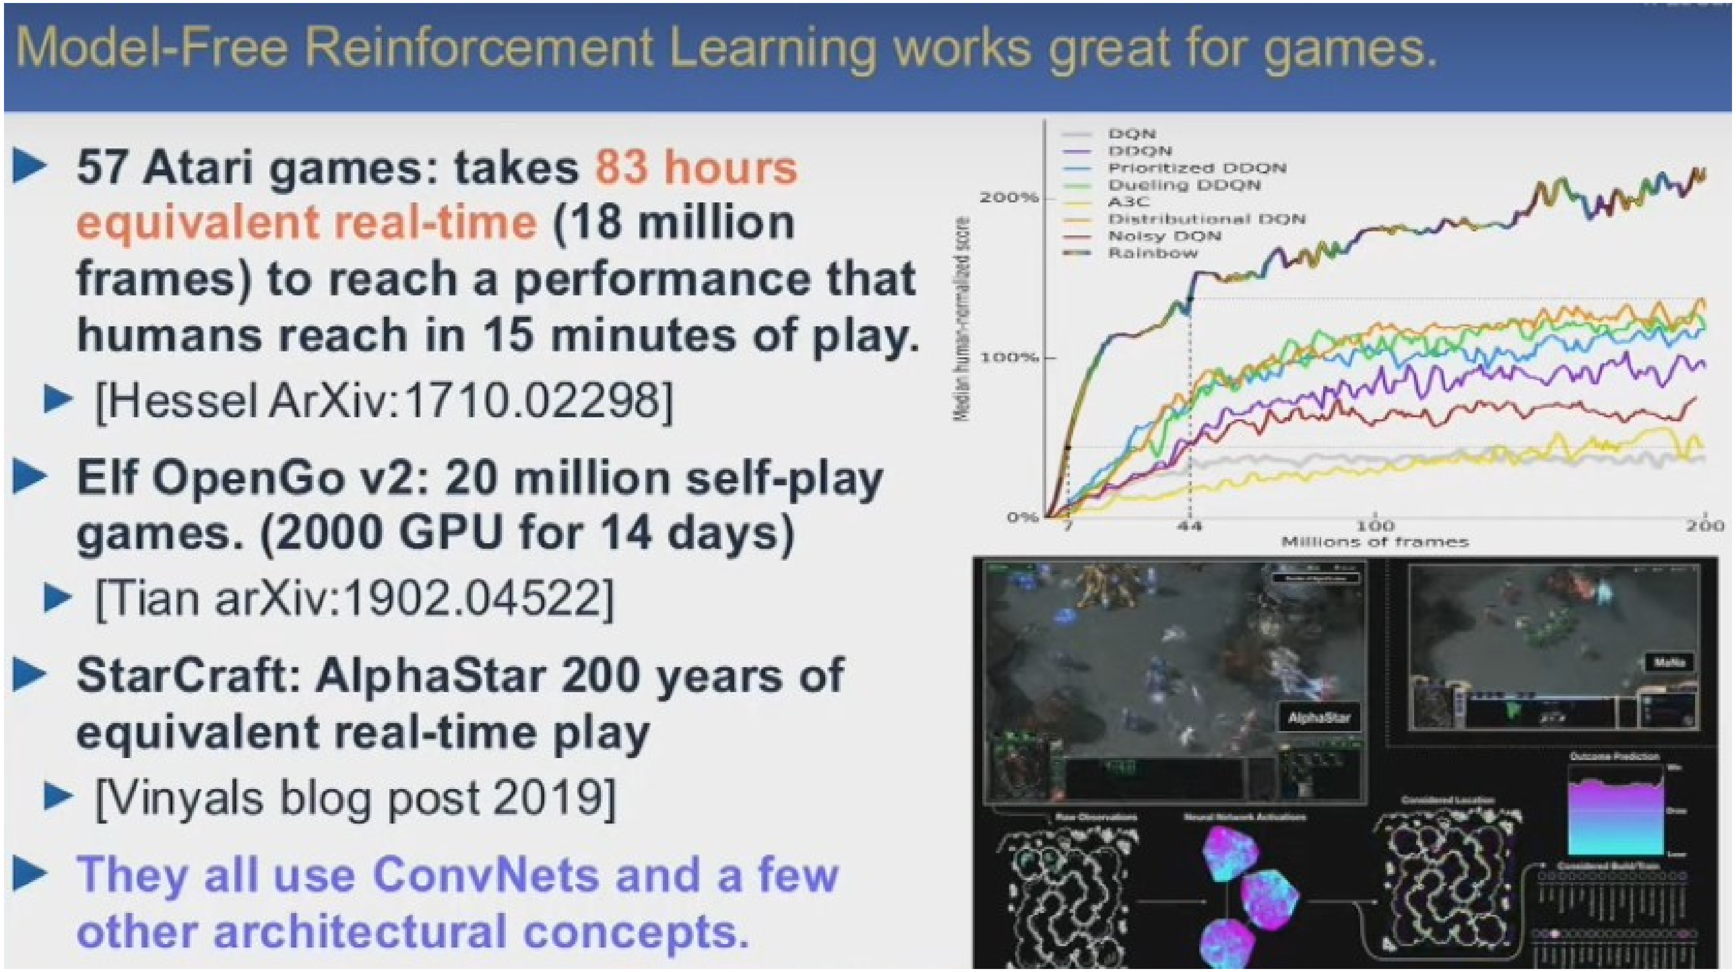


**As this slide shows, humans learn from a tiny fraction of the data.**

- **This criticism does not deny the value of DNNs. They make an outstanding contribution to AI. But much more is required to support cognition.**

**Time: 54 m**


**This slide is critical.**

- **Good drivers can drive for a lifetime without a serious accident.**
- **They learn how to avoid them without ever experiencing them.**
- **And they don't have to simulate millions of them in their minds.**

**Time: 55:30 m**

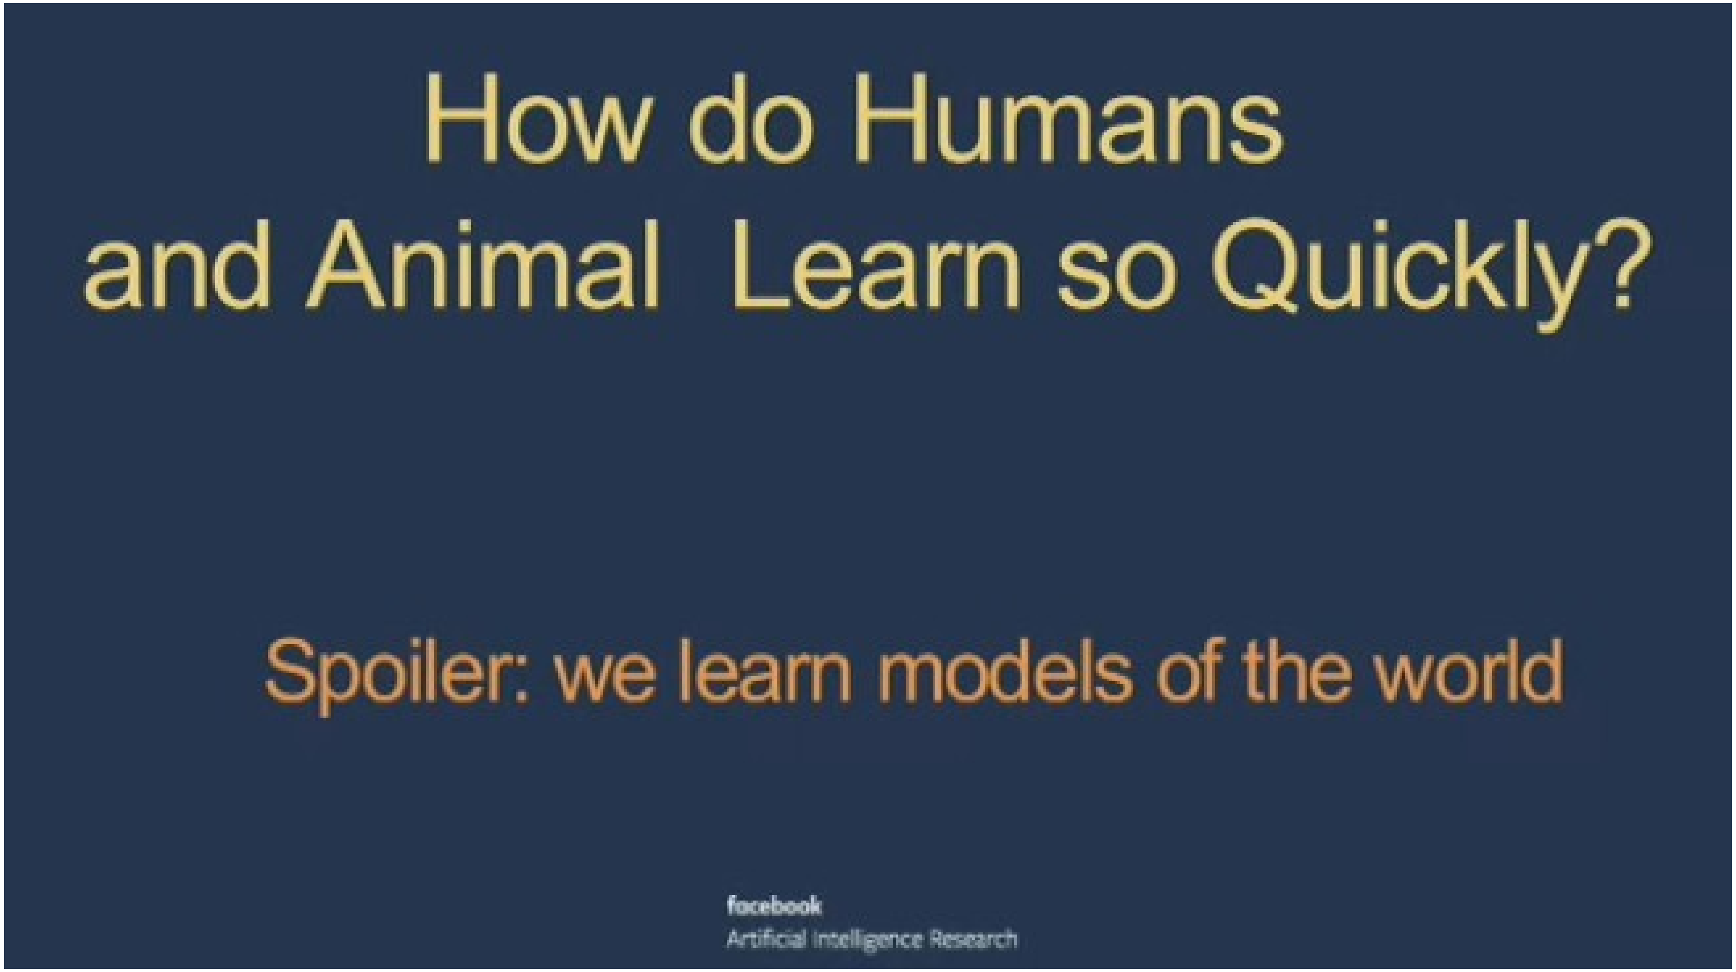


**Yes. That spoiler alert is what Edward Tolman wrote in 1948:**

- **His "Cognitive maps in rats and men" began the cognitive revolution.**
- **Animals "imagine" appropriate maps and models. The next slide shows how birds imagine and build a stable nest from irregular twigs and straw.**

**Time: 56 m**

# Bird Nest Problem

Robots can perform many tasks with great precision.

But they don't have the flexibility to handle unexpected shapes.

They can't wash dishes the way people do — with an open-ended variety of shapes and sizes.

And they can't build a nest in an irregular tree with irregular twigs, straw, and moss.

If a human guides a robot through a complex task with complex material, the robot can repeat the same task in the same way.

But it doesn't have the flexibility of a bird, a beaver, or a human.

# Relating Models to the World

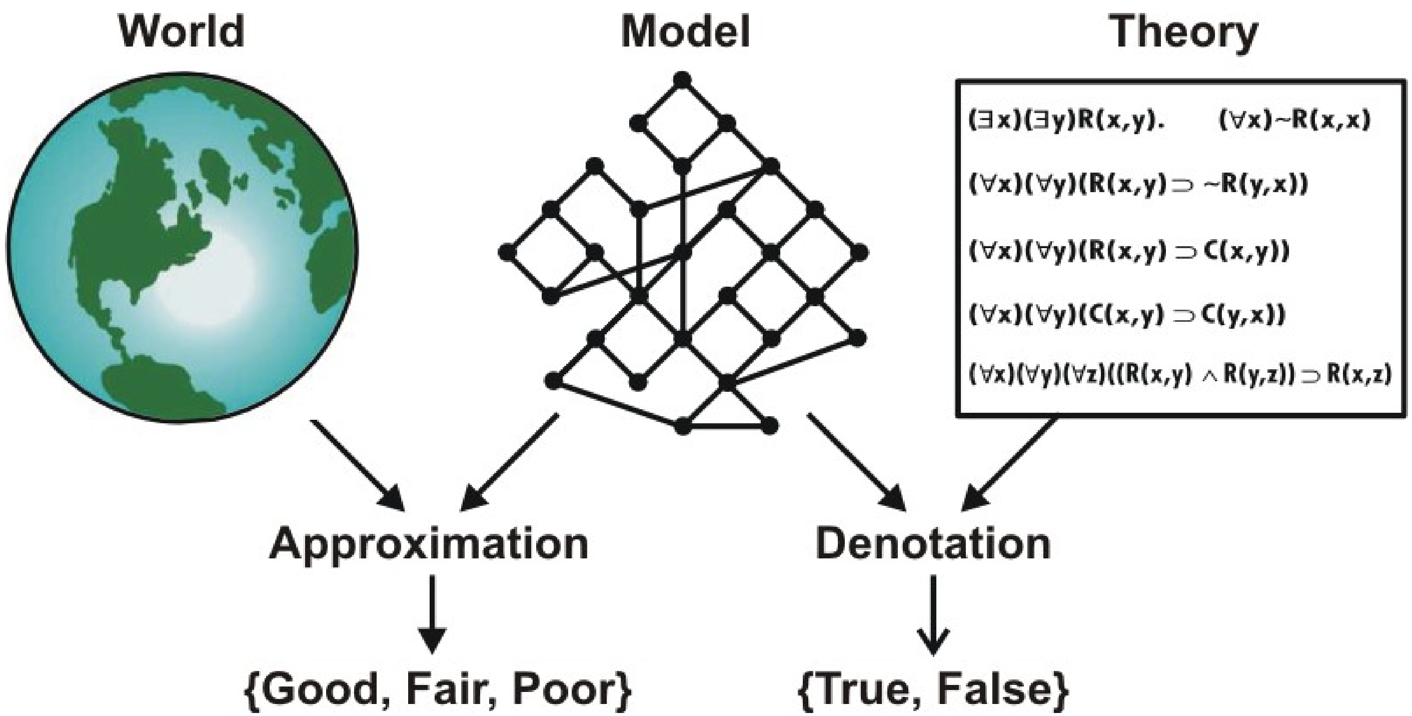


**Engineers: "All models are wrong, but some are useful."**
- **Discrete symbolic models can be clear, sharp, and precise.**
- **But the world is continuous, disordered, and fuzzy.**

**Natural languages are flexible. They can adapt to anything.**
- **They can be as vague or precise as the situation requires.**
- **AI tools should be flexible: Detailed levels must be precise, but the ontology must accommodate anything imaginable.**


**The phrase “remarkably little interaction” is wrong.**

- **Young children and other animals are very “playful”, but that play has a serious purpose: learn what happens when their actions make changes in the world.**
- **Babies of any species learn by doing, not by passively watching television.**

**Time: 57 m**

# The Ultimate Understanding Engine

**Sentences uttered by a child named Laura before the age of 3. ***

> ***Here's a seat.  It must be mine if it's a little one.***
>
> ***I went to the aquarium and saw the fish.***
>
> ***I want this doll because she's big.***
>
> ***When I was a little girl, I could go "geek geek" like that,<br>but now I can go "This is a chair."***

**No computer system today can learn and use language as fast, as accurately, and as flexibly as a child.**

**Preschool children constantly ask "Why?"**

**Those questions get into the pragmatics.  They are the hardest for parents and computer systems to answer.**

*** John Limber, The genesis of complex sentences.**
**http://pubpages.unh.edu/~jel/JLimber/Genesis_complex_sentences.pdf**

# Child Reasoning

**A mother talking with her son, about the same age as Laura: ***

**Mother:** ***Which of your animal friends will come to school today?***

**Son:** ***Big Bunny, because Bear and Platypus are eating.***

**The mother looks in his room, where the stuffed bear and the platypus are sitting in a chair and "eating".**

**The boy built a physical model that reflects his mental model, which is the semantic basis for his language and reasoning:**

- **The bear and the platypus are eating.**
- **Eating and going to school cannot be done at the same time.**
- **Big Bunny isn't doing anything else.**
- **Therefore, Big Bunny is available.**

**This reasoning is more "logical" than anything that Siri says.**

*** Reported by the psychologist Gary Marcus, in an interview with Will Knight (2015)**
**http://www.technologyreview.com/featuredstory/544606/can-this-man-make-ai-more-human/#comments**

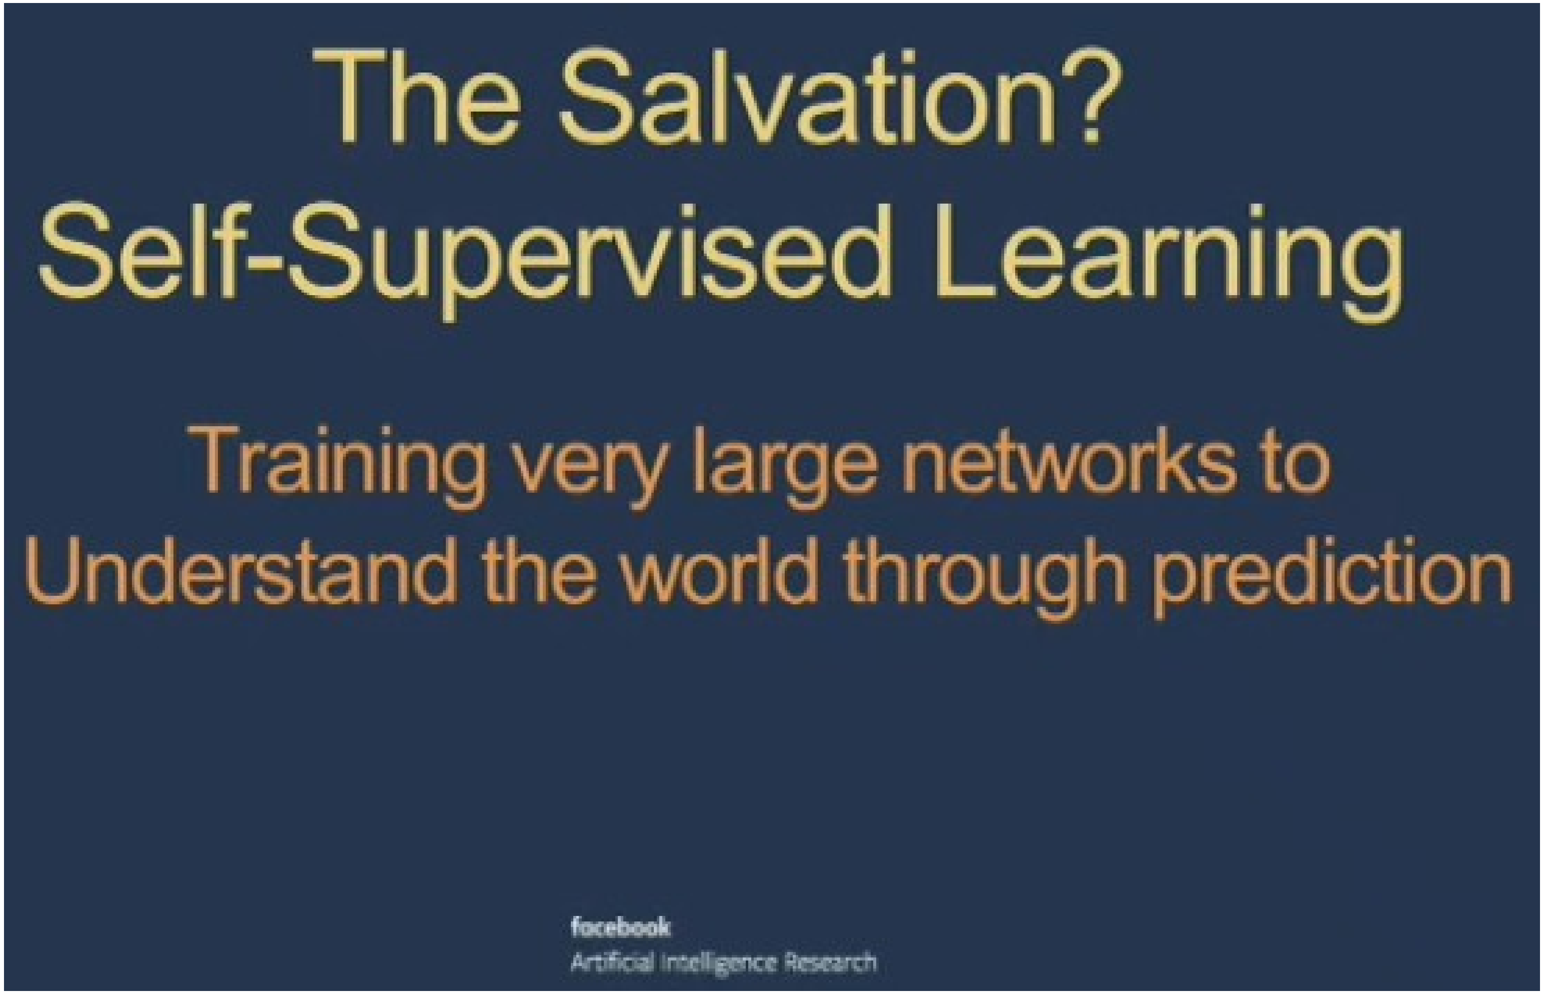


**Yes. But the emphasis on networks is misleading.**

- **The focus should be on the mental models, which are the basis for making predictions. Neurons are the mechanism in the brain, but the current artificial NNs are too primitive to support the complexity of the neurons in any animal..**

**Time: 59:30 m**

## Self-Supervised Learning: Prediction & Reconstruction

- Predict any part of the input from any other part.
- Predict the future from the past.
- Predict the future from the recent past.
- Predict the past from the present.
- Predict the top from the bottom.
- Predict the occluded from the visible
- Pretend there is a part of the input you don't know and predict that.

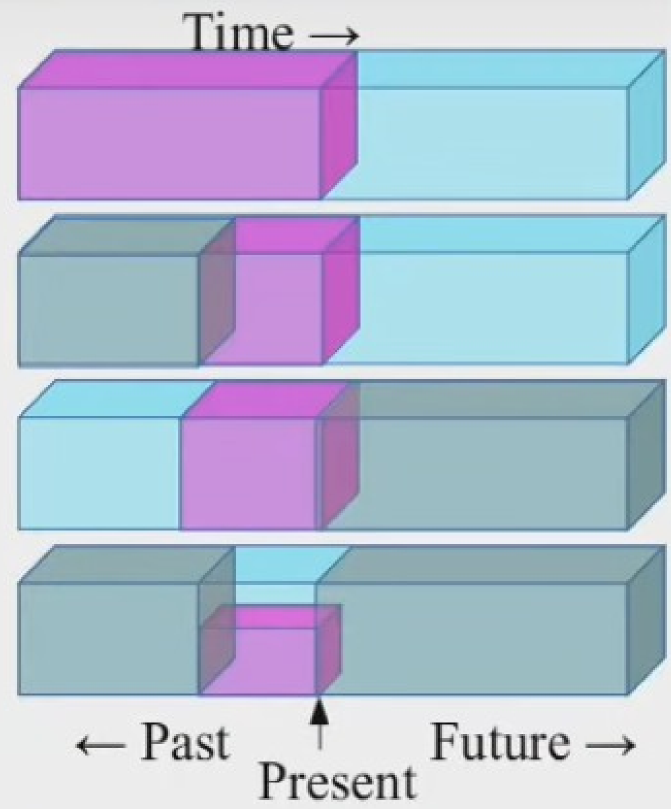


**Yes. Those predictions are routine and essential for everyday life.**

- **Children learn to make such predictions before they begin to talk.**
- **They continue to learn and make more advanced predictions with the mental models they construct during the 20+ years from kindergarten to a PhD.**

**Time: 1 hour**


**Yes. The mental models are key to the coming revolution.**

- **NN-style learning is important for learning to recognize and classify patterns.**
- **But evaluation (self supervision) is based on mental models and methods for operating on them, as constructed from aspects of past experience.**

**Time: 1:04 h**

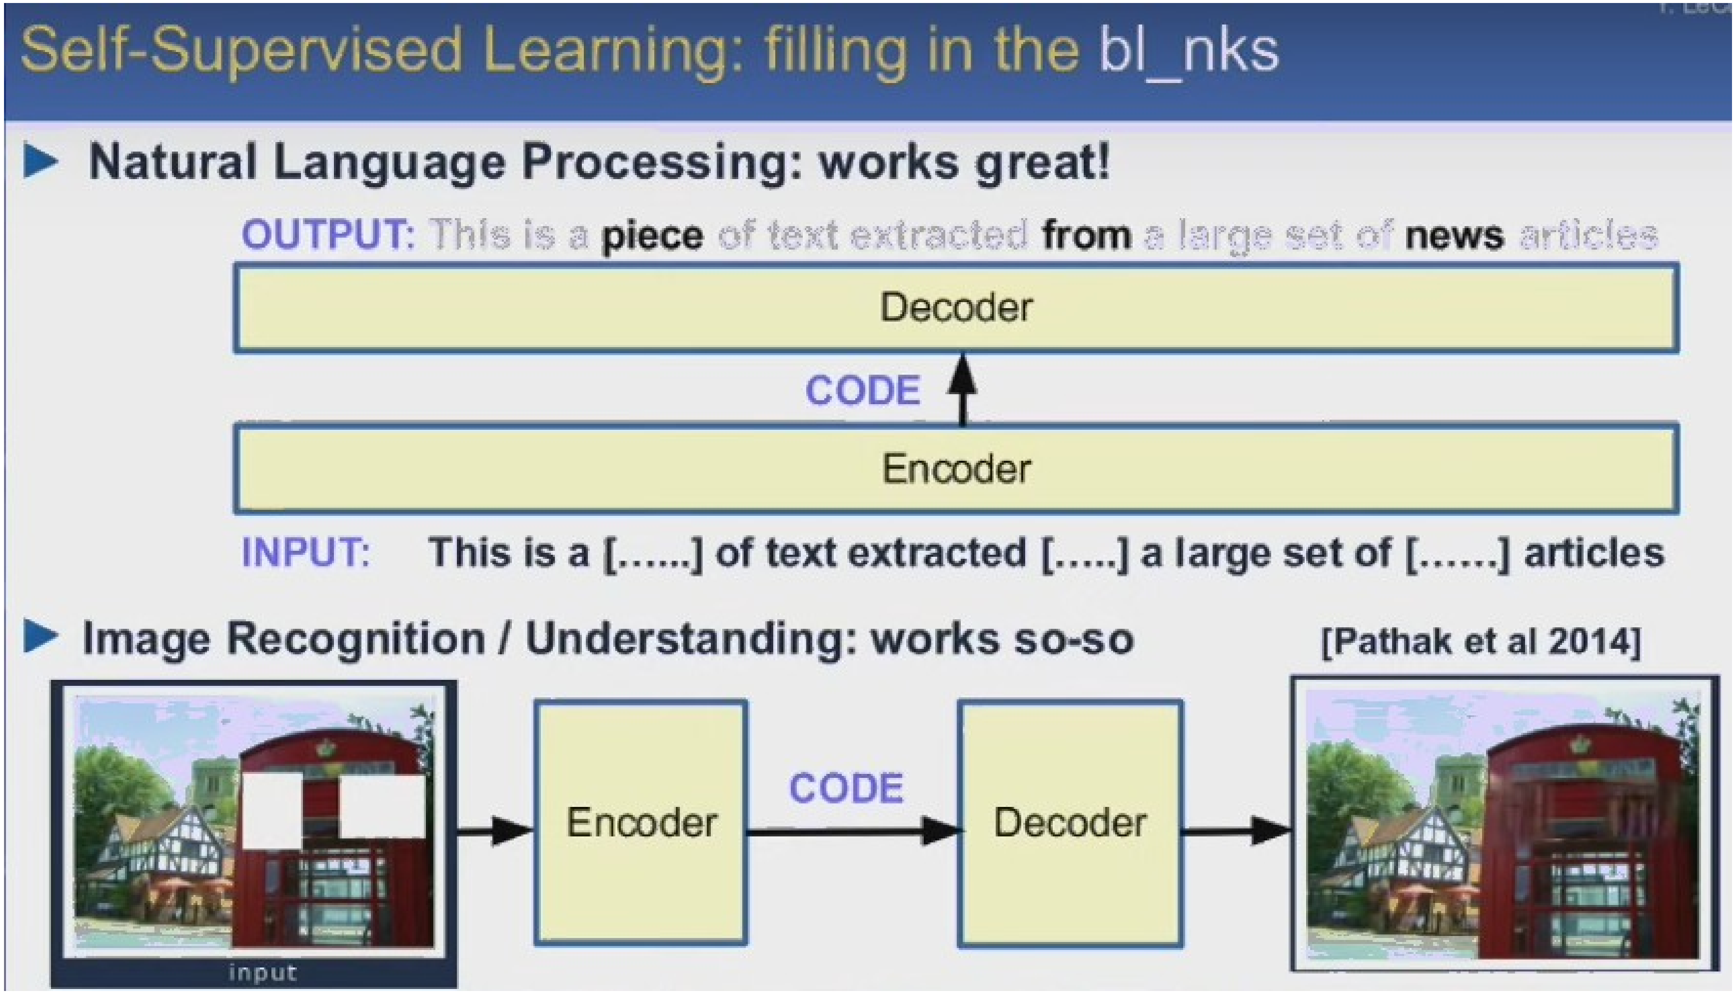


**NLP by NNs works fairly well for syntax, but not for pragmatics.**

- **Pattern-matching enables line-by-line translation from one notation to another.**
- **But language understanding depends on mental models and their relevance to the listener's hopes, fears, and goals. See the remaining slides.**

**Time: 1:04:30 h**

# Paradigms and Hybrids

**Hinton made a claim that AI researchers would endorse:**

> **We used to think of people as rational beings, and what was special about people was that they used reasoning to derive conclusions. Now we understand much better that people are basically massive analogy-making machines. They develop these representations quite slowly, and then the representations they develop determine the kinds of analogies they can make. ***

**C. S. Peirce, a pioneer in logic and semiotic, emphasized the importance of analogies. To support informal reasoning, he invented a graph logic that has direct analogies to the 3-D structures of mental models. But he made a sharp distinction between minds and machines in his 1887 article on logic machines: https://history-computer.com/Library/Peirce.pdf**

**In minds or machines, pattern recognition has always been a hybrid. In the first stage of recognizing a complex image, AI systems search for lines, parts, and boundaries. For the Blocks World of the 1960s, the parts were simple geometrical objects. Today's DNNs can recognize arbitrarily complex parts, such as people and their faces. In mammals, the eyes, the brain stem, the cerebellum, and the cortex collaborate in recognition.**

*** Interview with Geoffrey Hinton, Yann LeCun, and Yoshua Bengio, June 2019,**
**https://cacm.acm.org/magazines/2019/6/236987-reaching-new-heights-with-artificial-neural-networks/fulltext**

# Memory

**During the Q/A session, Hinton made an important comment (time 1:22 h):**

> **Neural nets are pretty good at things you do in parallel in the first few milliseconds. So far, they're not so good at things you do over longer time frames... In order to do things like that, you need some kind of memory.**

**Rote memory, the ability to store and retrieve data before and after any processing is done, is essential for AI systems. For NN applications, memory is an external component, such as a database, a knowledge base, the Semantic Web, or specially designed repositories.**

**In a review of the molecular and systems biology of memory, Eric Kandel and his colleagues analyzed four kinds: short-, intermediate-, and long-term memory in the cerebral cortex and the implicit memory in the cerebellum and other structures beneath the cortex. Much is known, even more is unknown, and the memories that last the longest may be encoded in some molecular form, such as RNA. Only one point is certain: research on artificial NNs has not yet begun to address these issues.**

**In summary, there is no evidence that DNNs have made the 60+ years of research in AI and other branches of cognitive science obsolete. See the argument for symbolic AI by Doug Lenat.**

# Semantics of Natural Languages

Human language is based on the way people think about everything they see, hear, feel, and do.

And thinking is intimately integrated with perception and action.

The semantics and pragmatics of language are

- Situated in time and space,
- Distributed in the brains of every speaker of the language,
- Dynamically generated and interpreted in terms of a constantly developing and changing context,
- Embodied and supported by the sensory and motor organs.

These points summarize current views by psycholinguists.

Philosophers and logicians have debated the issues:

- A continuum between natural languages and formal logics or a sharp dichotomy between NLs and logic.

# Continuum Between Logic and Language

**Richard Montague (1970): A continuum with logic as primary.**

"I reject the contention that an important theoretical difference exists between formal and natural languages."

**Hans Kamp (2001): A qualified continuum.**

"The basic concepts of linguistics — and especially those of semantics — have to be thought through anew... Many more distinctions have to be drawn than are dreamt of in current semantic theory."

**Barbara Partee (2005): More qualifications.**

"The present formalizations of model-theoretic semantics are undoubtedly still rather primitive compared to what is needed to capture many important semantic properties of natural languages..."

**Peirce and Wittgenstein: A continuum with NLs as primary.**

- Every artificial notation is an abstraction from some aspects of NLs.
- No version of logic has all the semantic properties of NLs.
- Any formal logic is just one among many possible language games.

# Relating Language to Logic

**Peirce wrote a succinct summary of the issues:**

> **"It is easy to speak with precision upon a general theme. Only, one must commonly surrender all ambition to be certain. It is equally easy to be certain. One has only to be sufficiently vague. It is not so difficult to be pretty precise and fairly certain at once about a very narrow subject." (CP 4.237)**

**Implications:**

- **A precise formal ontology of everything can be stated in logic, but it's almost certainly false in critical details.**
- **A looser classification, such as WordNet or Roget's *Thesaurus*, can be more flexible for representing lexical patterns.**
- **A specification in logic can be "pretty precise and fairly certain" only for a very narrow subject.**

**A formal logic cannot be vague. But no finite set of symbols can precisely describe every aspect of a continuous world.**

# What is Cognition?

## Lawrence Barsalou's answer:  coordinated non-cognition. *

- Cognition is “embedded in, distributed across, and inseparable from” the “processes of perceiving, acting, and emoting.”
- Visual and motor simulations are essential to language understanding.
- When people view a static object, they anticipate working with it.
- When people view food, they anticipate its taste when eating it.
- Musicians identify their own performances by recognizing the fingering.
- Affect, feelings, rewards, and value judgments are fundamental to all aspects of reasoning and decision making.
- No single aspect is cognition, but all of them together are cognition.
- Social interactions facilitate learning by stimulating more aspects.

## Barsalou’s answer resembles Minsky’s Society of Mind.

- Both of them emphasized open-ended variety of cognitive processes.
- The society must integrate the mind and body at every level.

* Barsalou, Breazeal, & Smith (2007) Cognition as coordinated non-cognition.

# Cognitive Learning

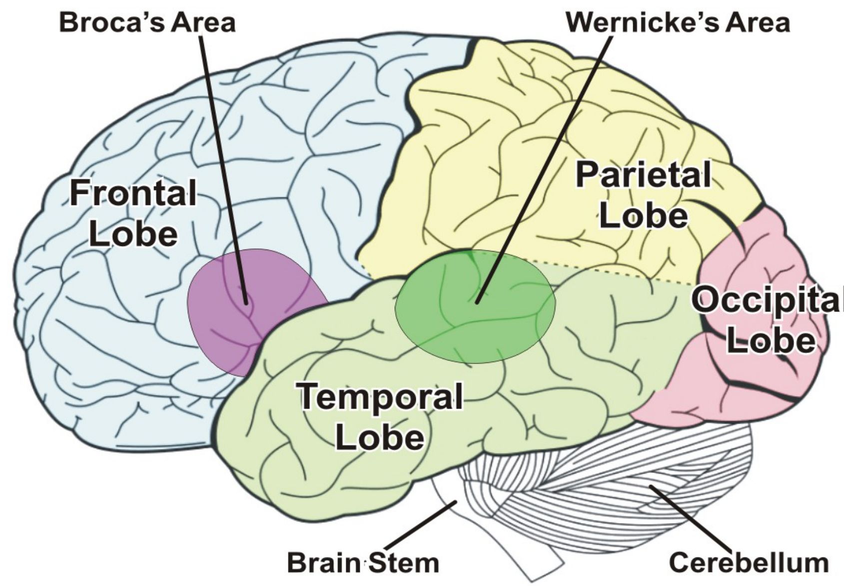


**The areas of the cerebral cortex are highly specialized.**

**A study with fMRI scans showed which areas are active at various stages of learning. ***

**14 participants studied how four devices work: bathroom scale, fire extinguisher, automobile braking system, and trumpet.**

**Cognitive learning is much deeper than deep neural nets:**

1. **Occipital lobes are active in recognizing shapes and patterns.**
2. **Parietal lobes become active in learning mechanical structures.**
3. **All lobes become active as participants are "generating causal hypotheses" about how the system works.**
4. **Finally, the frontal lobes anticipate "how a person (probably oneself) would interact with the system."**

*** R. A. Mason & M. A. Just (2015) http://medicalxpress.com/news/2015-03-science-brain.html**

# Areas Active in Cognition

Most neurons have short links to nearby neurons.

But others make long-distance connections from one lobe to another.

The diagram shows connections among areas of the brain involved in language. *

The colors of the boxes correspond to the colors of the brain areas in the previous slide.

Patterns can be learned and recognized in one area. But cognition links diverse areas across the brain.

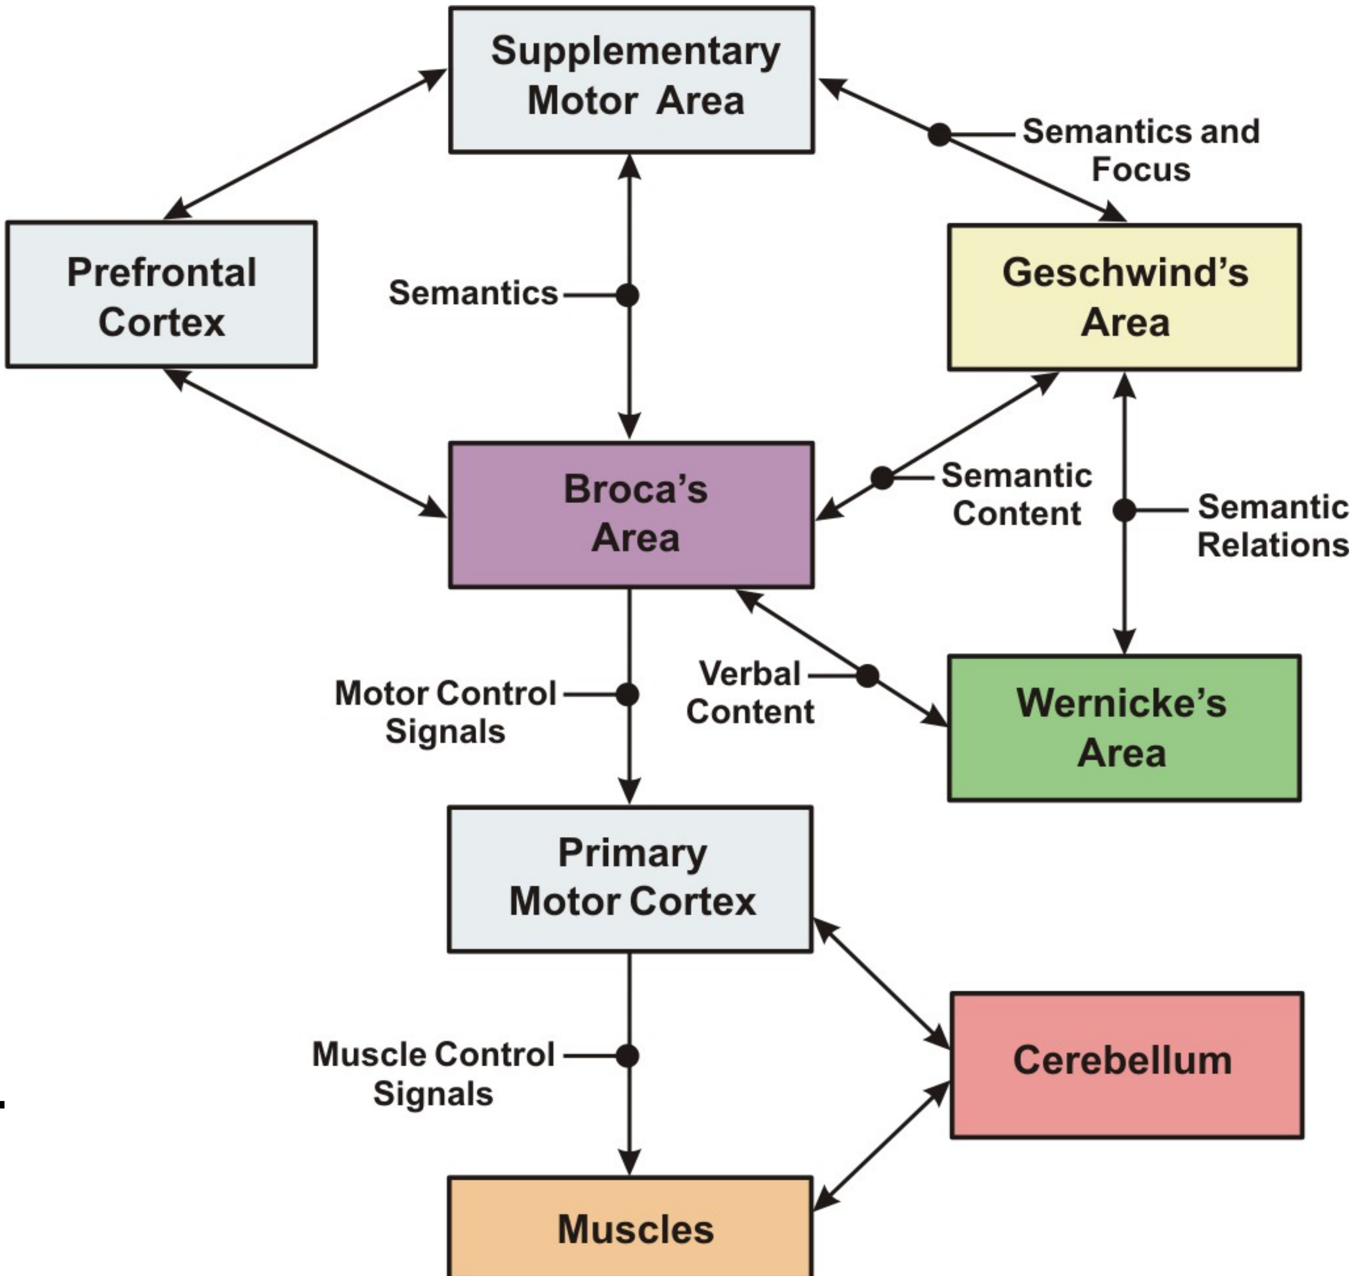


* Diagram adapted from MacNeilage (2008).

# Perceptual and Cognitive Learning

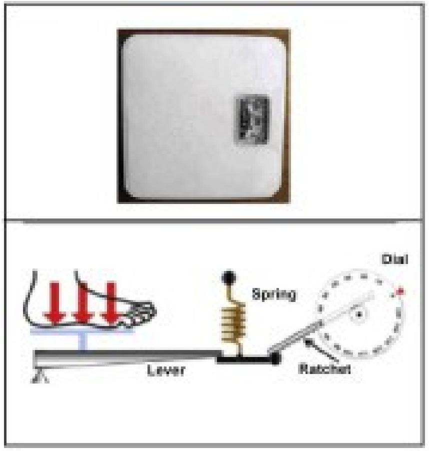


Bathroom Scale


Disc Brake System

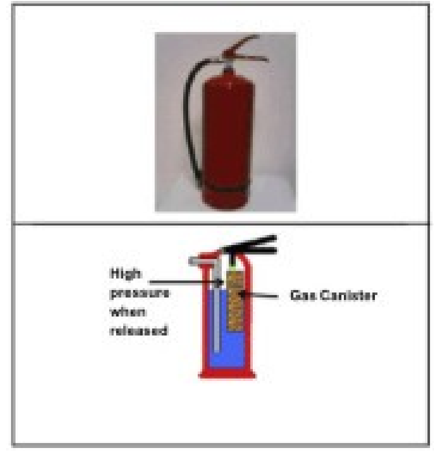


Fire Extinguisher

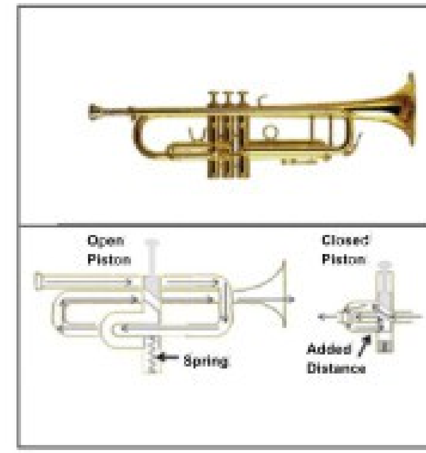


Trumpet

**14 participants studied how four devices work: bathroom scale, fire extinguisher, disc brake system, and trumpet. ***

- **Subjects: college students who were not science or engineering majors.**
- **They had multiple training sessions with each of the four devices.**
- **During test sessions, an fMRI scanner recorded patterns of brain activity.**
- **An early training session just showed pictures and named the parts:** ***A bathroom scale consists of a spring, a lever, a ratchet, and a dial.***
- **Later sessions explained structural and causal relations:** ***The spring pulls a ratchet which rotates a gear attached to a measurement dial.***

*** R. A. Mason & M. A. Just (2015) Physics instruction induces changes in neural representation.**

# Cognitive Learning

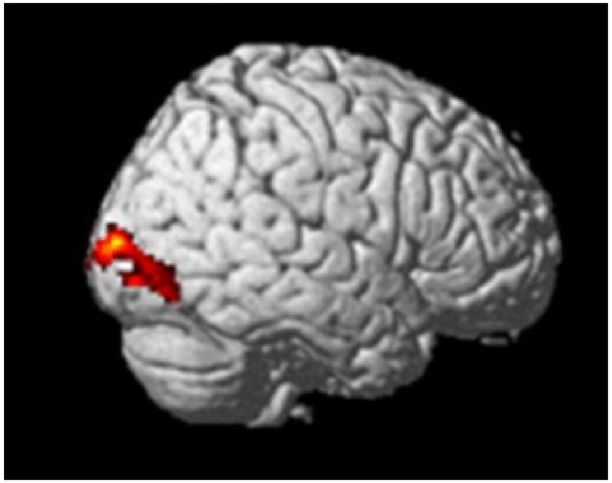

1. Visual perception

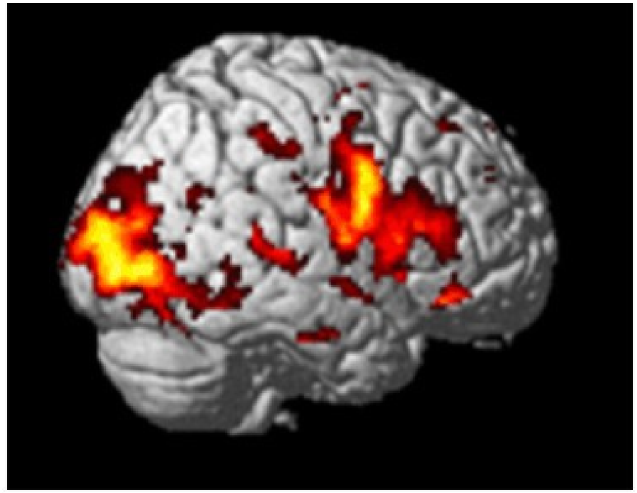

2. Thinking about structure

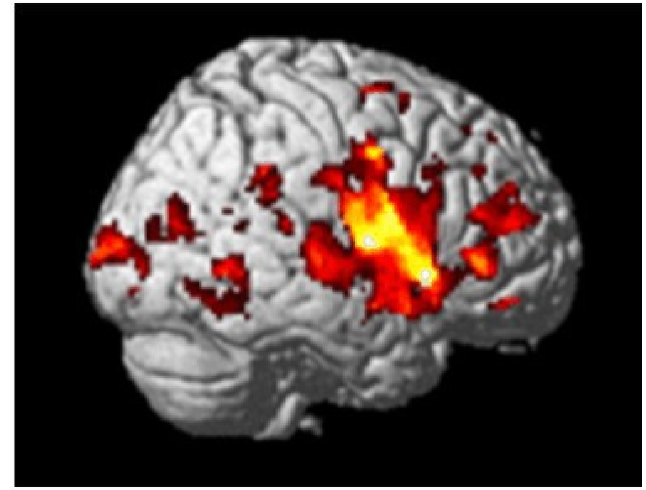

3. Thinking about causality

**Neural activity in the right hemisphere during test sessions:**

- **All 14 students showed similar neural activations.**
- **Questions about the objects and parts activated the visual cortex, the occipital lobes in the back of the brain (fMRI image #1 above).**
- **Questions about structural relations activated the parietal lobes, which link vision to all sensory and motor regions (image #2).**
- **Questions about the causal effects of someone operating the system activated the frontal lobes and connections across the brain (image #3).**
- **Summary: Cognitive learning involves structural and causal relations that link and coordinate perception, action, and reasoning.**

# Intentionality

**Without life, there is no meaning in the universe.**

- **Philosopher Franz Brentano: Intentionality is "*the directedness of thought toward some object, real or imagined.*"**
- **Biologist Lynn Margulis: "*The growth, reproduction, and communication of these moving, alliance-forming bacteria become isomorphic with our thought, with our happiness, our sensitivities and stimulations.*" ***
- **A bacterium swimming upstream in a glucose gradient marks the beginning of goal-directed intentionality.**

**In Peirce's categories, intentionality is a mediating Third.**

- **It's the reason why some mind or quasi-mind directs attention toward some mark, which it interprets as a token of some type.**
- **Some interpretation by some agent makes some mark (an aspect of the universe) meaningful in some way for that agent.**
- **All laws, communications, explanations, value judgments, and social relations depend on the intentions of some agent.**

*** Margulis (1995) Gaia is a tough bitch, http://edge.org/documents/ThirdCulture/n-Ch.7.html**

# Monadic, Dyadic, and Triadic Relations

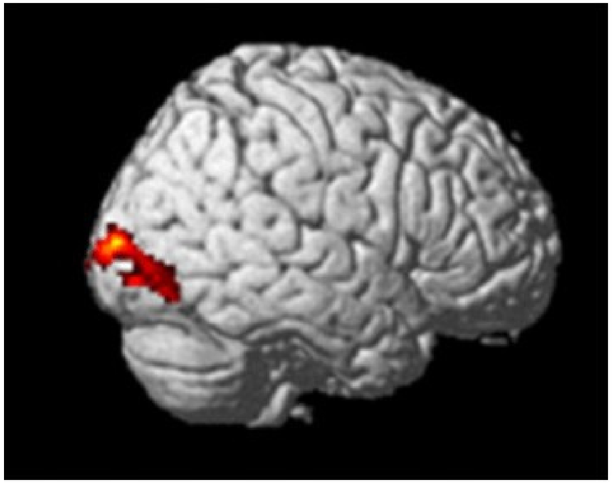

1. Visual perception

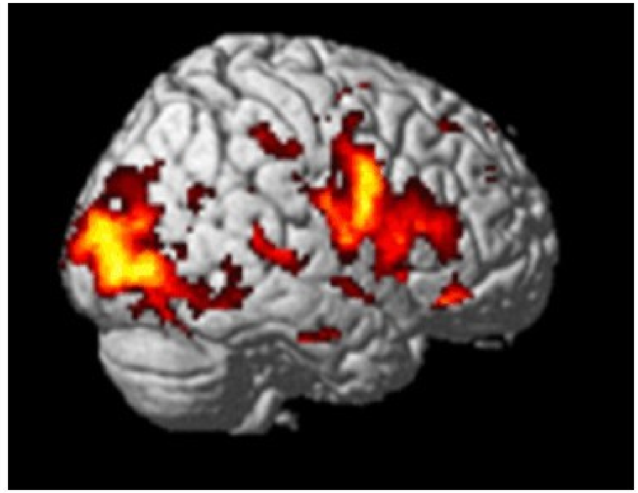

2. Thinking about structure

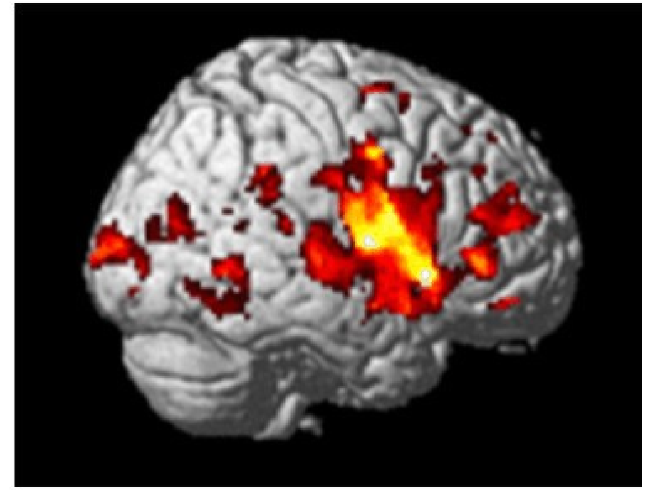

3. Thinking about causality

**Neural correlates of Peirce's First, Second, and Third:**

- **Perception is based on localized percepts or prototypes. It classifies phenomena by monadic predicates (fMRI image #1).**
- **Long-distance connections in the parietal lobes support dyadic relations that connect all sensory and motor modalities (image #2).**
- **The frontal lobes process the mediating triadic relations in reasoning, planning, causality, and intentionality (image #3).**
- **Much more detail must be analyzed and explained, but these examples suggest promising directions for future R & D.**

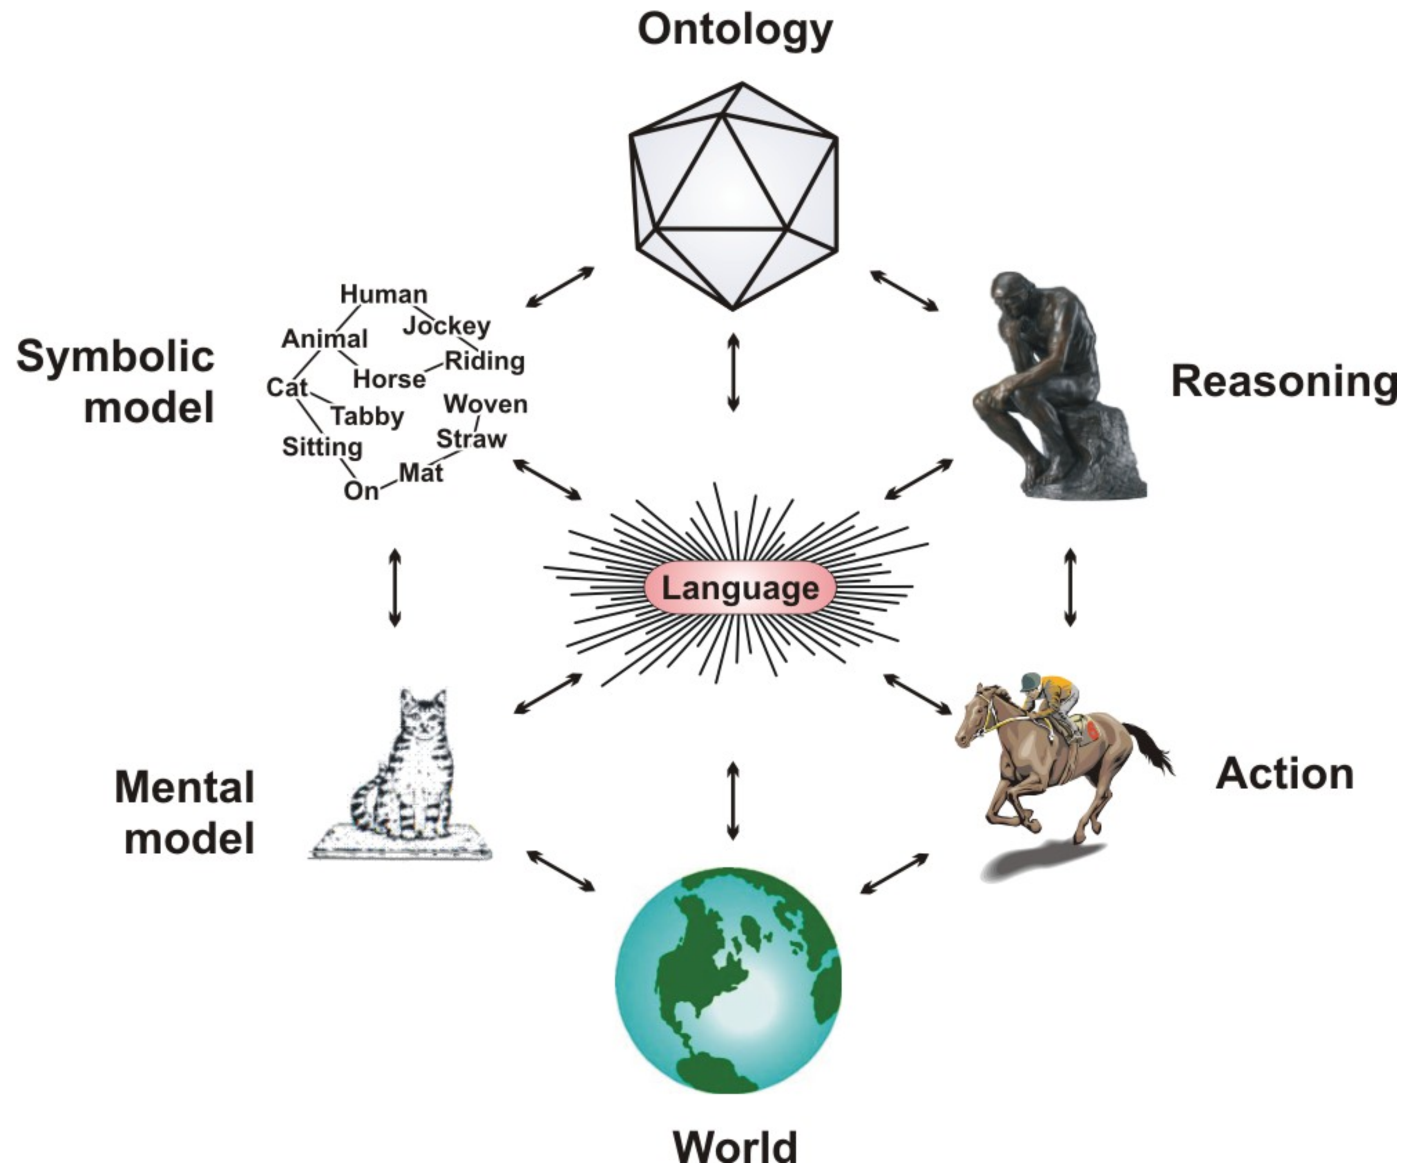


**The hexagon shows how AI relates to language and the world. The world is not an AI component, but the other five corners are. The 12 arrows show mappings between language, the world, and the AI components. Any implementation requires many kinds of tools.**

# The Role of Empathy

Observation by the linguist Zellig Harris:

> "We understand what other people say through empathy – imagining ourselves to be in the situation they were in, including imaging wanting to say what they wanted to say."

That sentence shows a deeper understanding of language than anything that his star pupil, Noam Chomsky, ever wrote.

Subsets of language that map to formal logics are mainly used for mathematical precision in science and engineering.

Baby talk is a dialect that mothers develop through empathy with their babies. People also develop special dialects through empathy with their pets and other animals.

In fact, communication among all higher animals is based on sounds or signs learned through empathy. See slide 26.

Conclusion: Robots will never fully understand human language unless and until they develop empathy with humans.