MARVIN MINSKY:

PIONEER, CRITIC, OPTIMIST

Kristinn R. Thórisson

Associate Professor, School of Computer Science, Reykjavik University

Member, Center for Analysis and Design of Intelligent Agents, Reykjavik U.

Managing Director, Icelandic Institute for Intelligent Machines, Reykjavik, Iceland

www.ru.is

TURING AWARD LECTURE SERIES | REYKJAVIK | APRIL 5 2012

Marvin Lee Minsky

Born in NYC

At MIT since 1958

Turing Award recipient in 1969

Still giving his Society of Mind course at MIT

86 years of age

Early Days

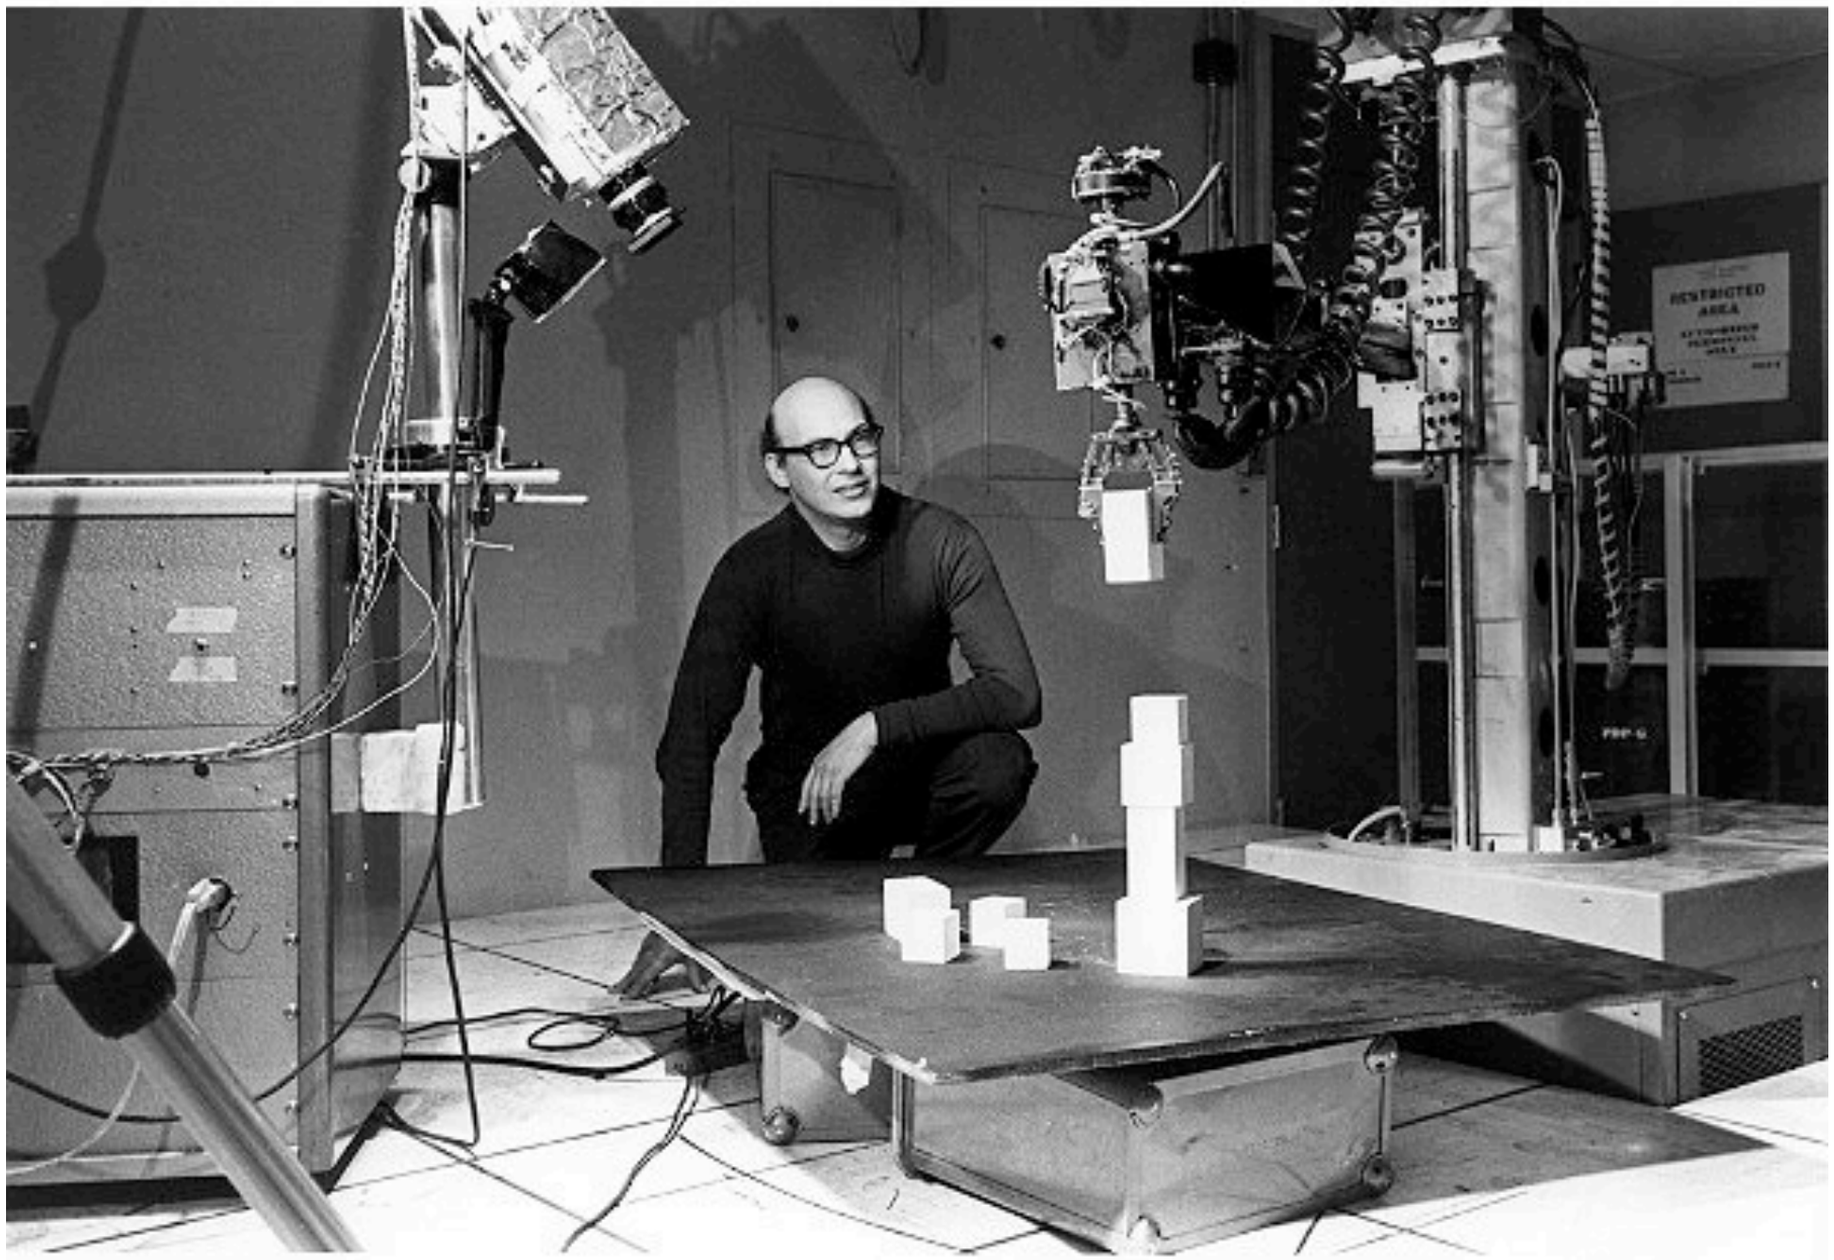

Early Days

Progress Report on Artificial Intelligence
Marvin Minsky and Seymour Papert
Dec 11, 1971

At the time of this report, the main foci of attention of the MIT AI Laboratory included

Robotics: Vision, mechanical manipulation. Advanced automation.
Models for learning, Induction, and analogy.
Schemata for organizing bodies of knowledge.
Development of heterarchical programming control structures.

Models of structures involved in commonsense thinking.
Understanding meanings, especially natural language narrative.
Study of computational geometry.
Computational trade-offs between memory size, and parallelism.
Theories of complexities of various algorithms and languages.
New approaches to education.

Early Days

Minsky was a “hopeless optimist”

- but he wasn’t the only one

Early Days

Minsky was a “hopeless optimist”

- but he wasn’t the only one

Allen Newell

John McCarthy

Oliver Selfridge

Herbert Simon

...

Video

The Human Interface

www.ru.is

Symbolic Computation

Hypothesis: Thinking is the result of symbol manipulation

If true: Neurons represent an inferior way to compute

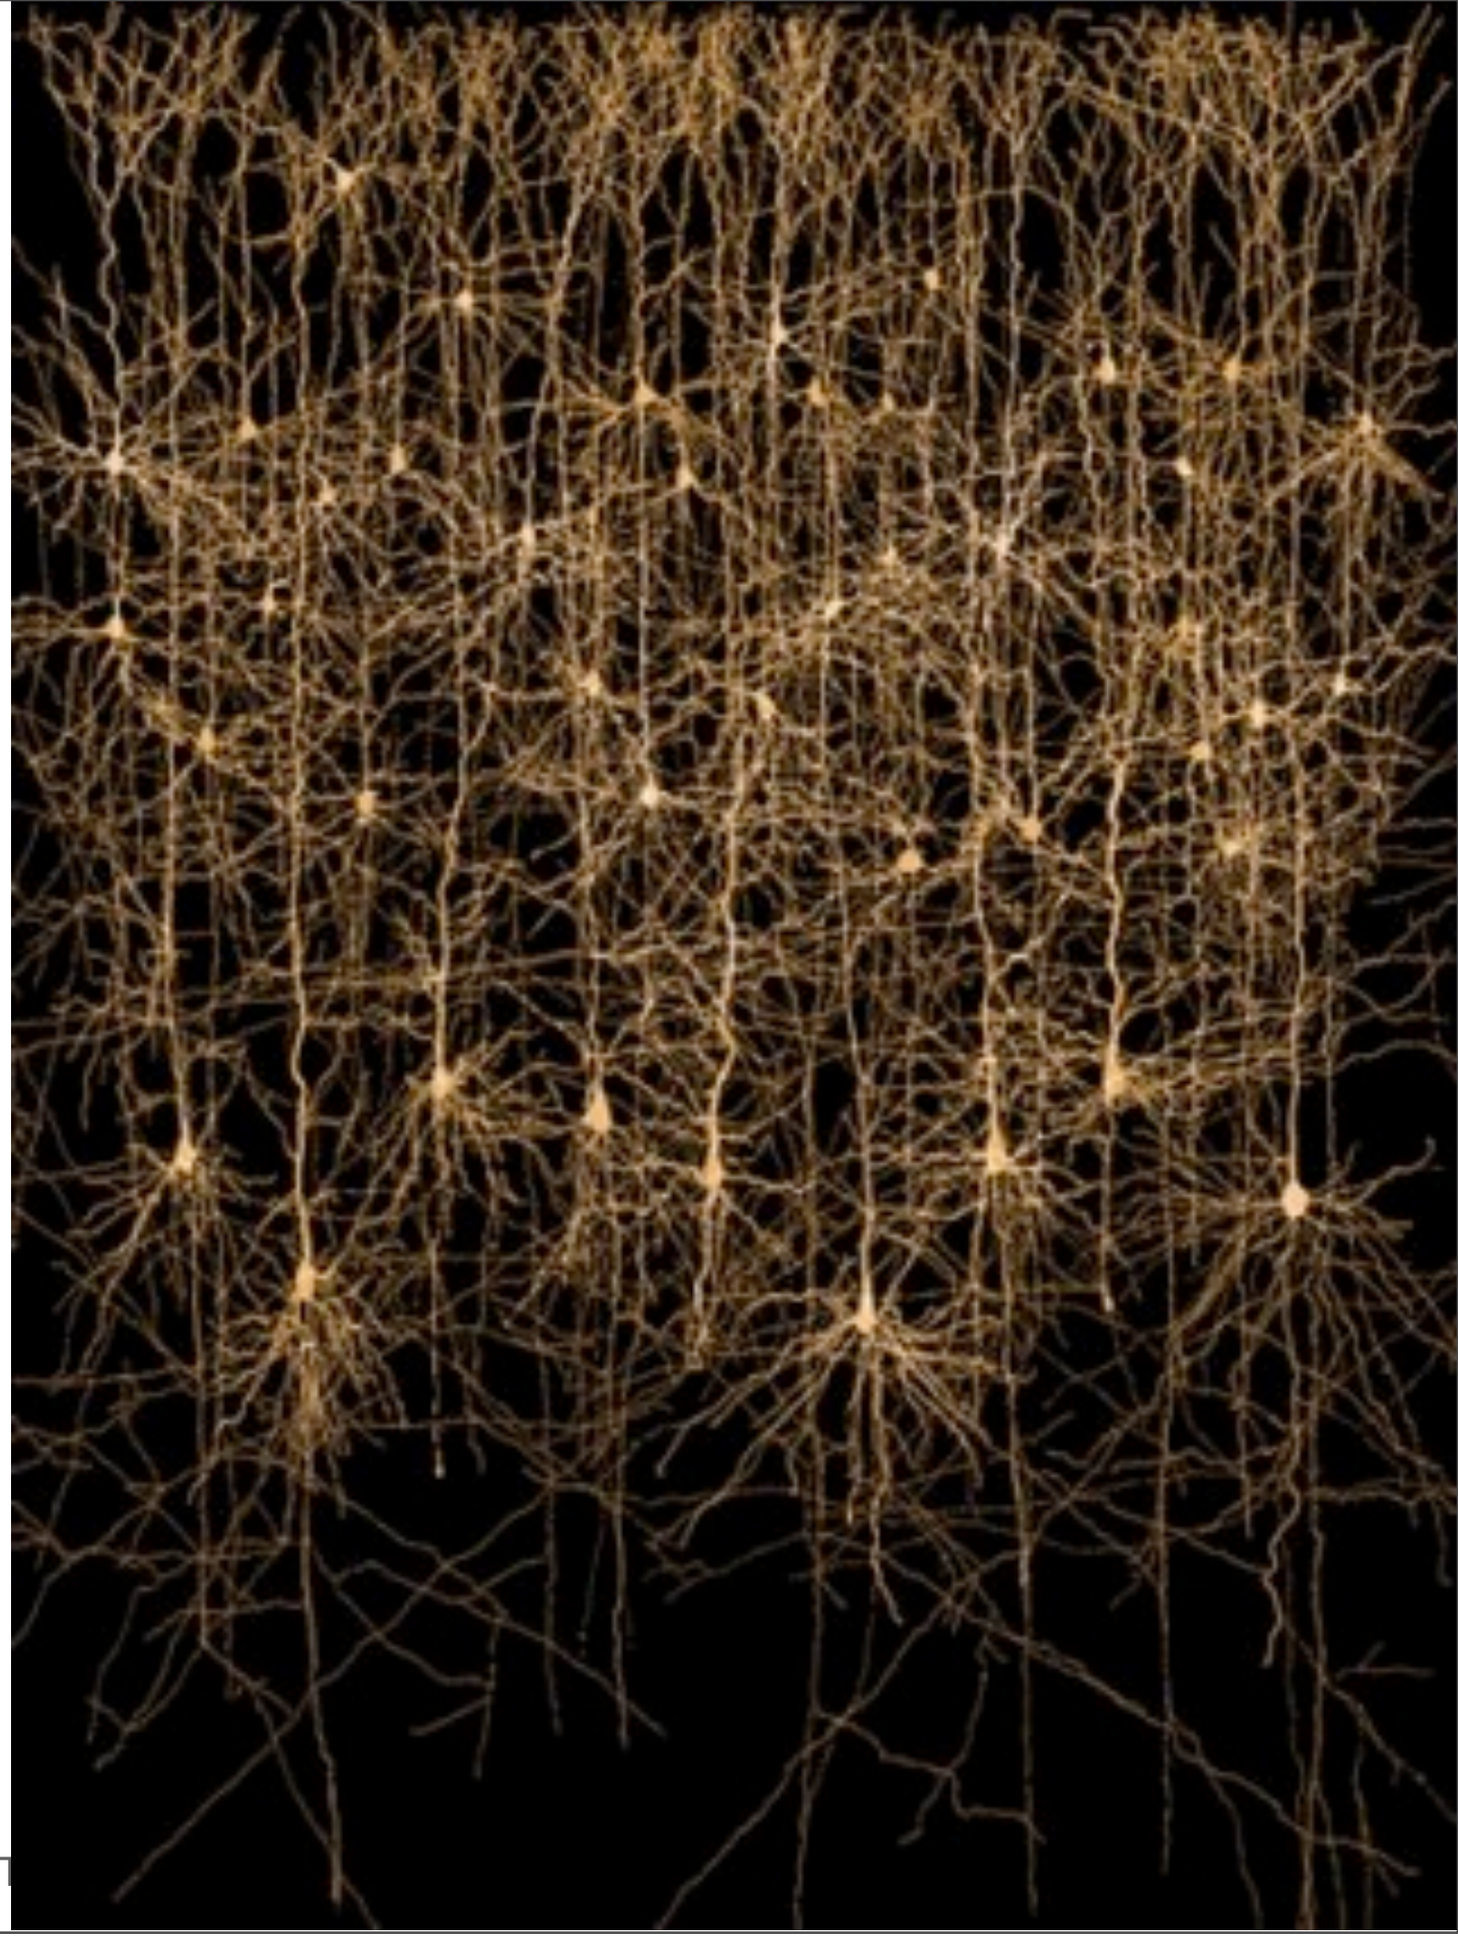

Symbolic Computation in Meat Computers

Short story: **ALIENABLE RIGHTS**

Two interstellar aliens have come to assess the life-forms of Earth. The human life-forms will be entitled to rights – if the aliens can conclude that they think. Most such decisions are easy to make – but this case is unusual.

Symbolic Computation in Meat Computers

...

Apprentice: Why are these humans so quarrelsome? Even their so-called entertainments are mostly fights disguised as plays and games and sports.

Surveyor: This is because they were never designed; they evolved by competing with tooth and claw. Evolution on Earth is still mainly based on the competition of separate genes.

...

Symbolic Computation in Meat Computers

...

Apprentice: Surely, though, we must regard them as intelligent. Despite their faults, they've already built some simple computers - and I've overheard them arguing about whether machines could ever think.

Surveyor: Humph. It is our job to find out if they can think. But I'll grant that it's amazing how much they can do, considering that their brain-cells compute only a few hundred steps per second.

Apprentice: Yet in spite of this they can recognize a friend in less than half a second - or understand a language phrase, or notice that a shoe is untied. How can they react so rapidly when their internal components are so slow?

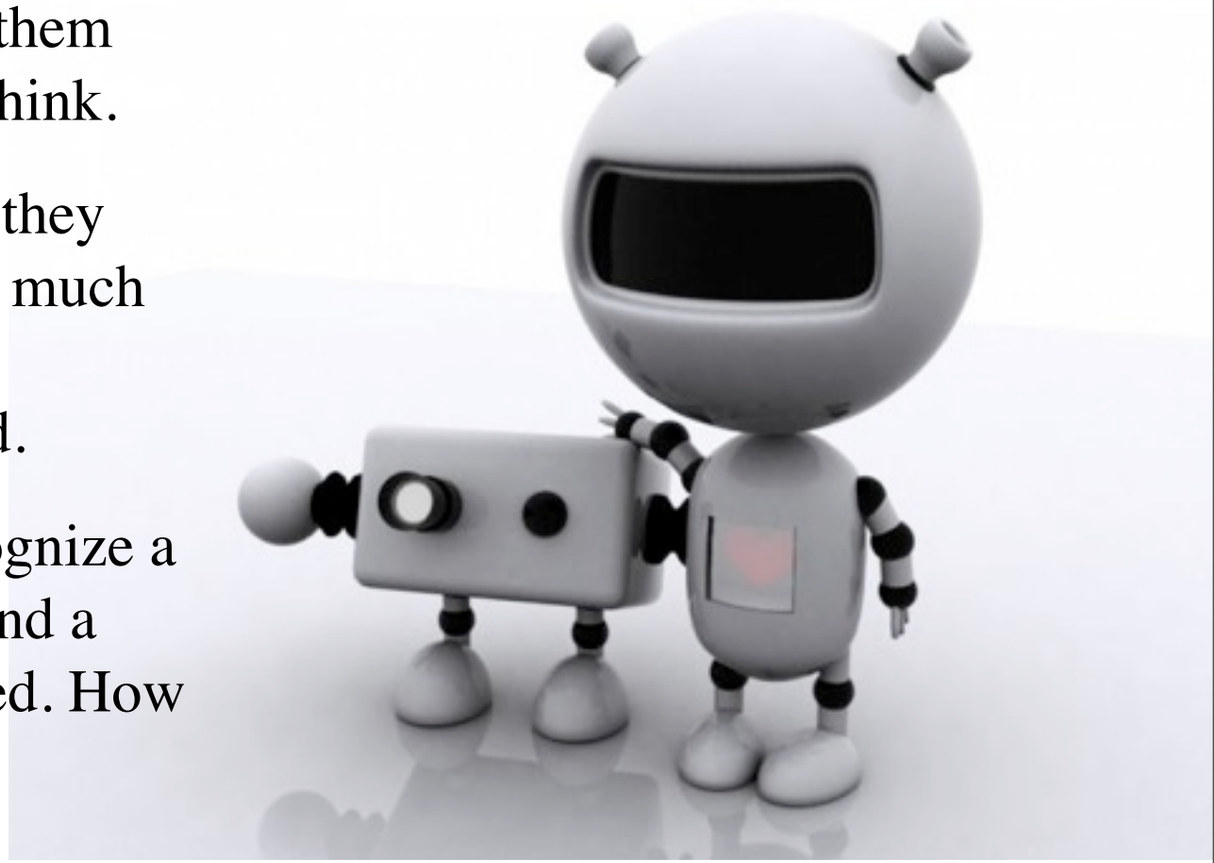

...

Symbolic Computation in Meat Computers

Apprentice: ... Are you suggesting that the more parallel operations are used inside a machine, the more serial it will seem from outside?

Surveyor: I could not have said it more clearly myself. To see why, suppose that a certain task involves two different kinds of sub-jobs. If we want to do them simultaneously, we'll have to run their programs and their data in two separate places, to keep them from interfering with each other. Similarly, if each of those jobs splits into sub-sub jobs, those must each be solved with only a quarter of the available resources. And so on. Total fragmentation. Eventually the sub-sub-sub jobs will end up with no place to work. A purely parallel machine must stop at some limit of complexity - whereas a serial computer will simply slow down.

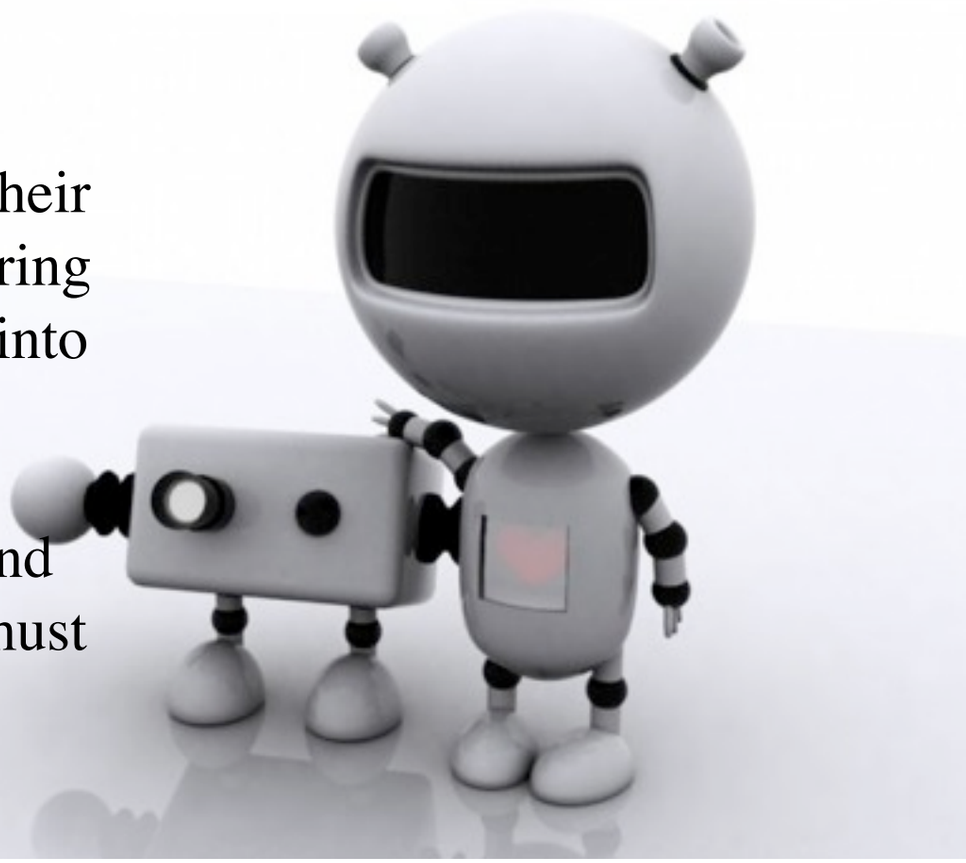

Apprentice: That's funny. Most of the computer experts on Earth seem to think that "parallel" and "distributed" go together. ...

Symbolic Computation in Meat Computers

...

Apprentice: But language isn't everything. Shouldn't we give them credit for explaining things with pictures, too? They do seem to have excellent senses.

Surveyor: That was my first impression, too - until I saw that their TV sets use only three electron guns. Of course, this means that they're virtually blind. ...

...

Symbolic Computation in Meat Computers

...

Surveyor: Some of their books do embody significant knowledge - but most of them are little more than sequences of fictional anecdotes about conflicts involving what they call love and lust, ambition and greed, and harmony and jealousy. Their so-called novels aren't novel at all but mere permutations of those elements. The trouble is that their time-sequential languages force them to squeeze their parallel structures through narrow-band serial channels.

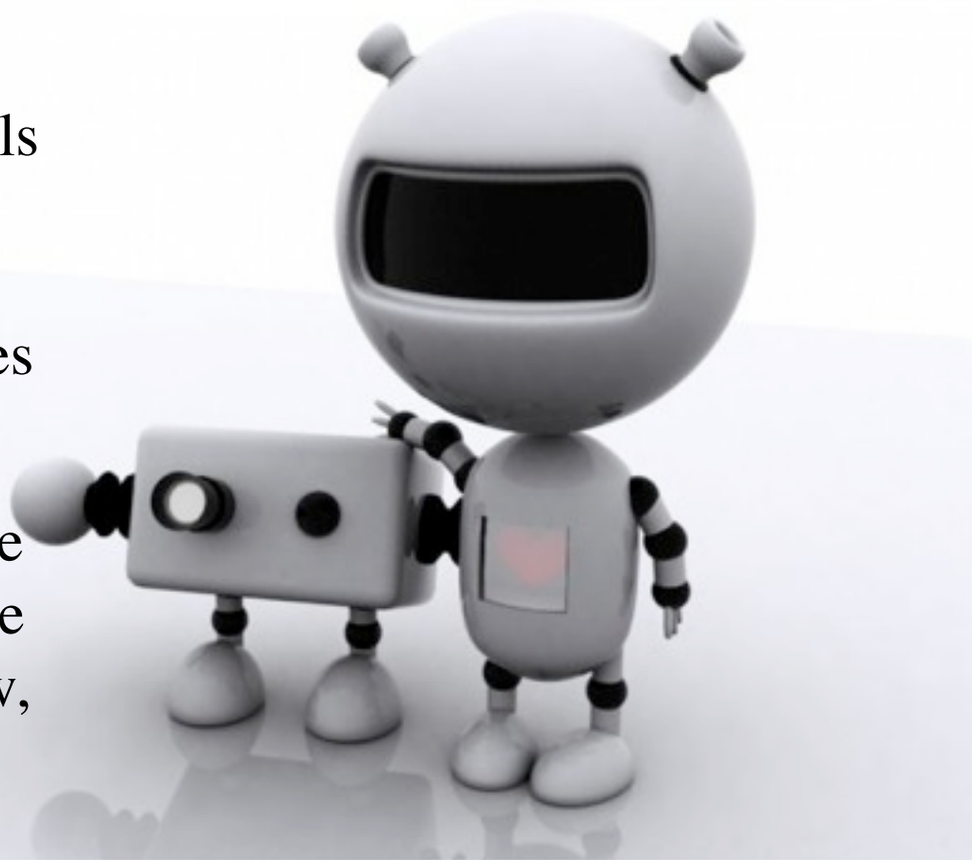

Apprentice: Serial communication? They seem to have everything upside down. Thinking, of course, should be serial - and communication should be parallel. But how, then, do they convert those sequences back into their original forms?

...

Frames

A Framework for Representing Knowledge [1974]

- proposed the idea of frames with slots

- precursor to object-oriented programming

- highly influential paper, still cited

Perceptrons

Frank Rosenblatt: inventor of the perceptron

- psychologist interested in the workings of the brain

- Perceptron: first mechanization of learning new skills by trial and error

Perceptrons

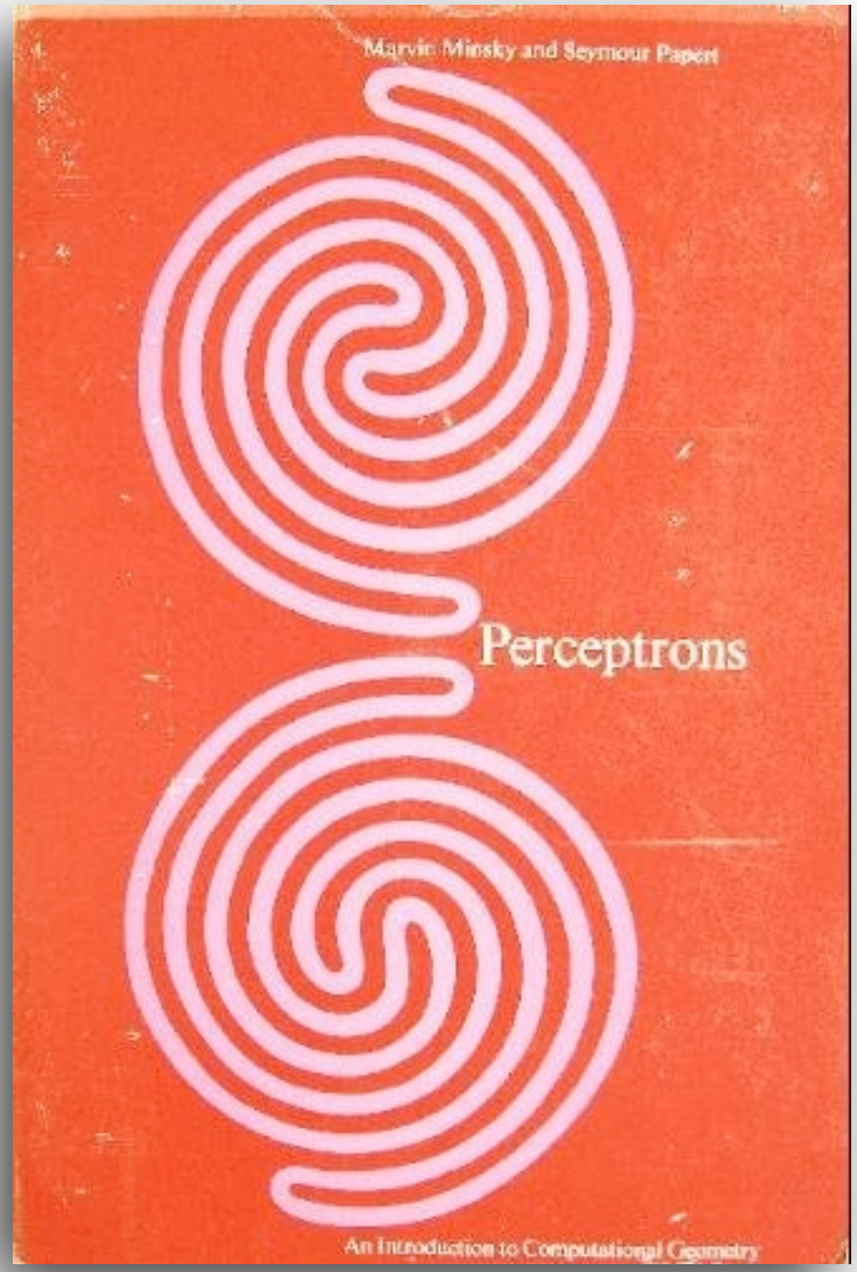

1969

Perceptrons

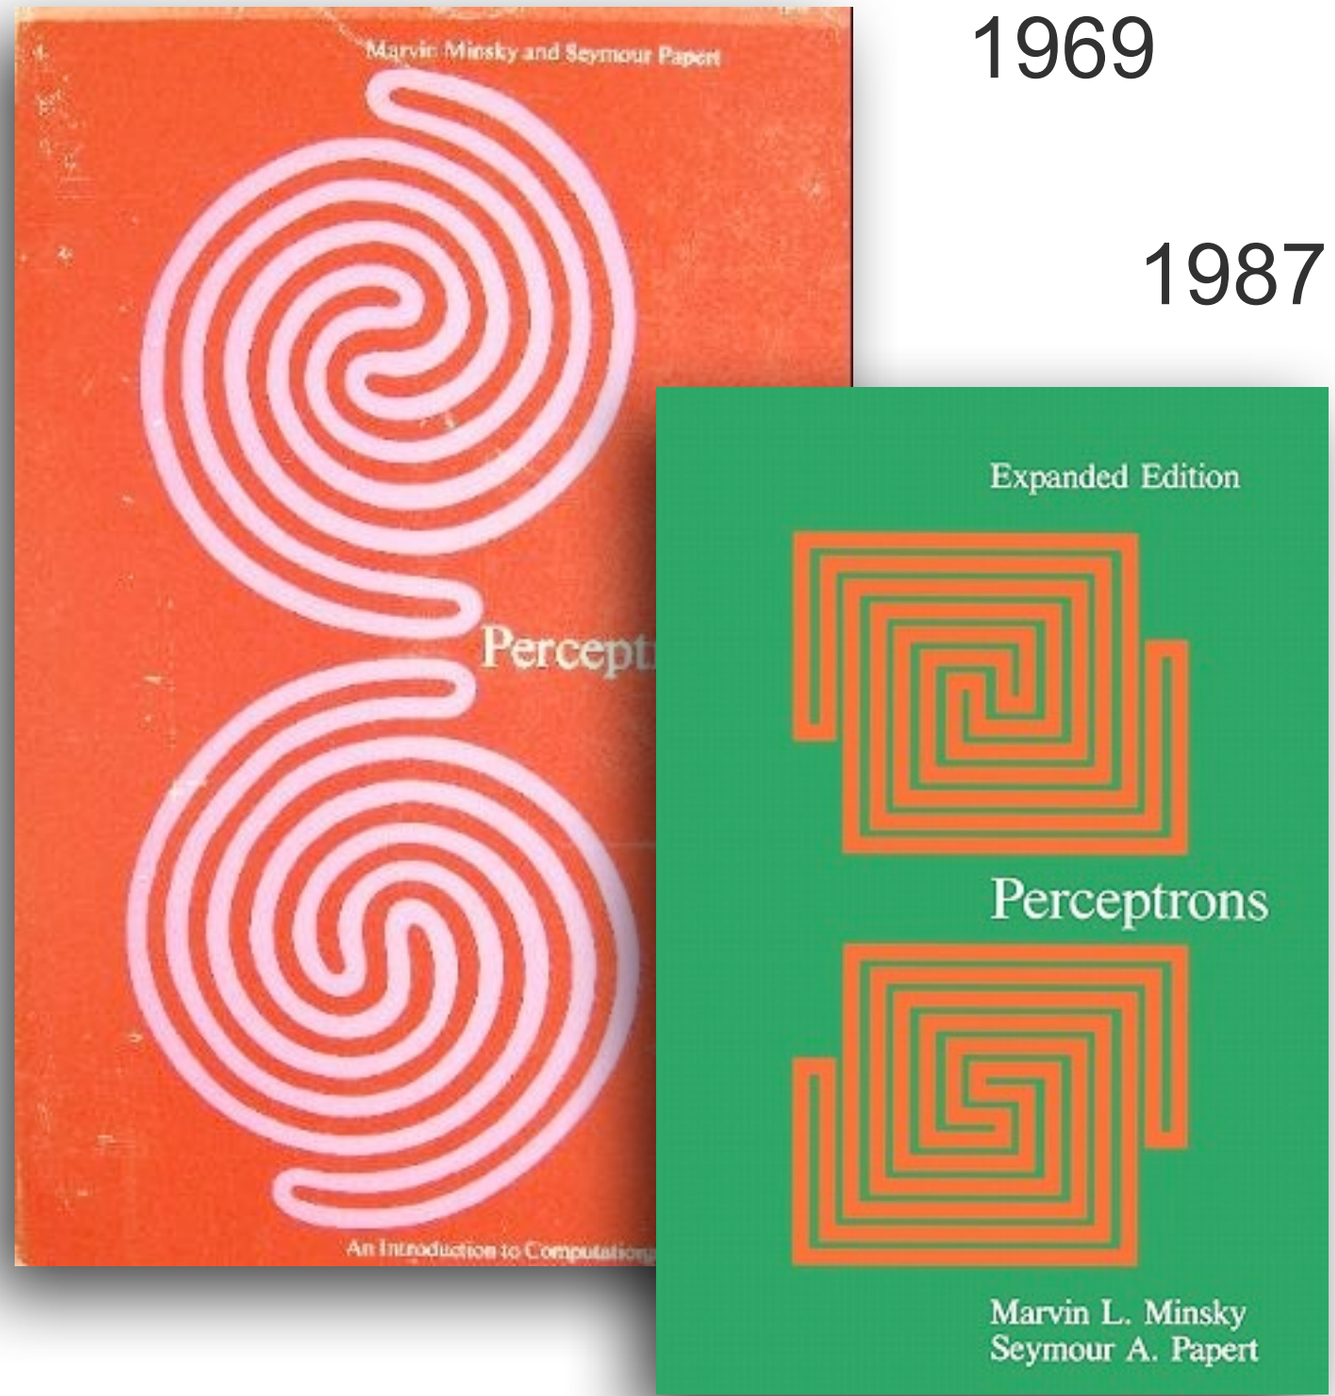

1969

1987

Perceptrons

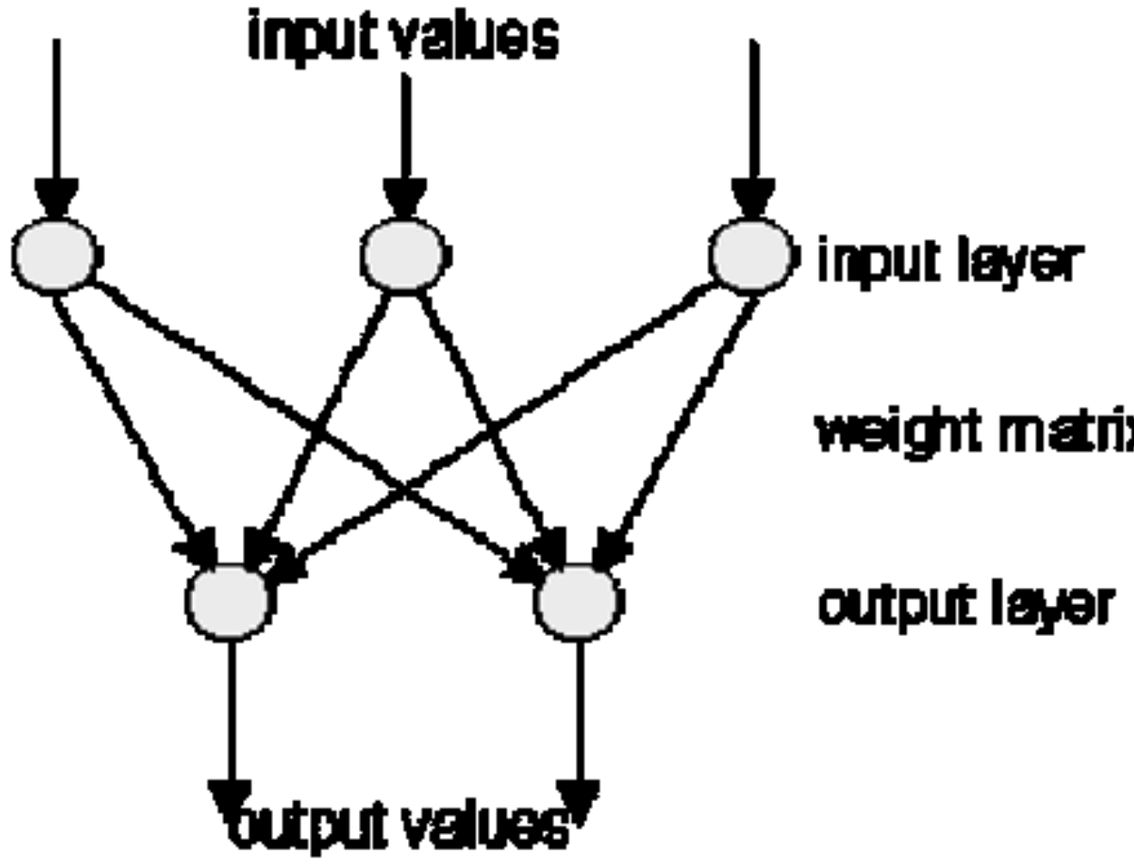

Perceptrons

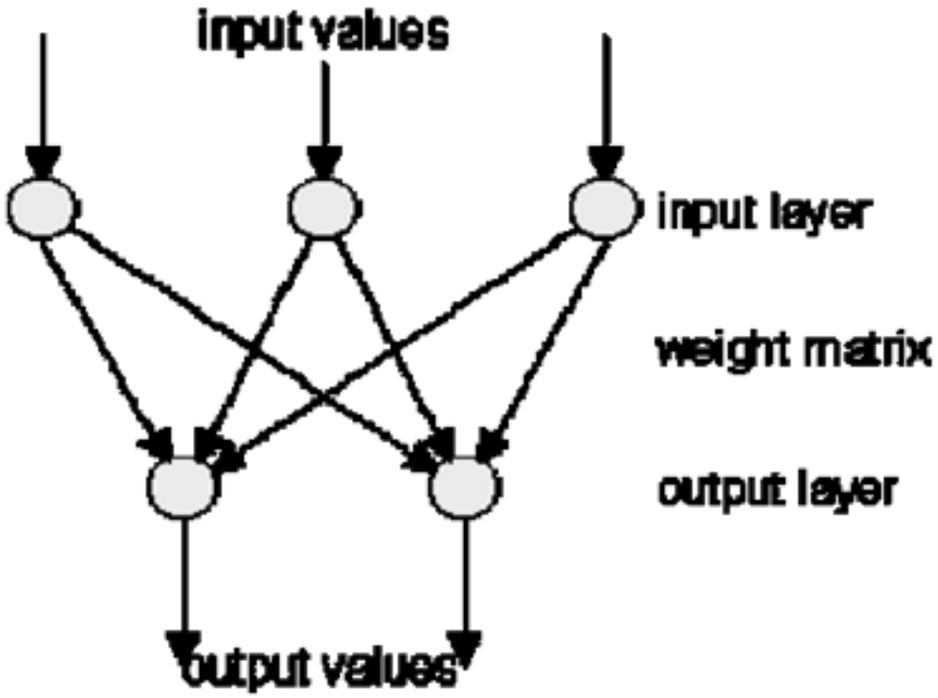

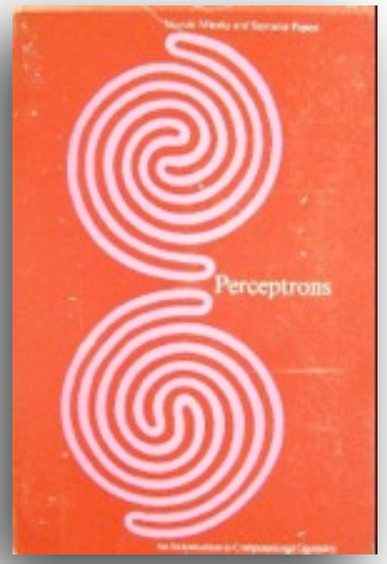

~~XOR~~

~~NXOR~~

Perceptrons

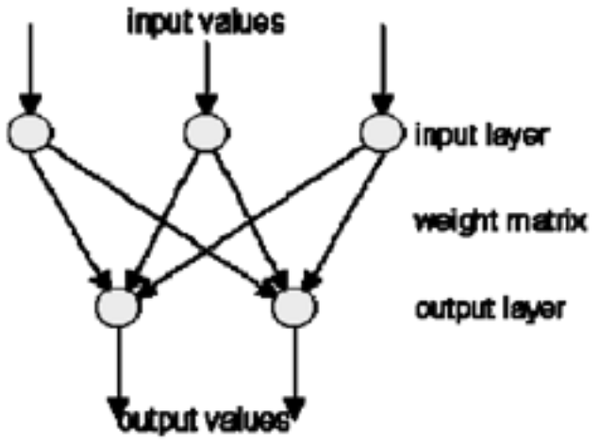

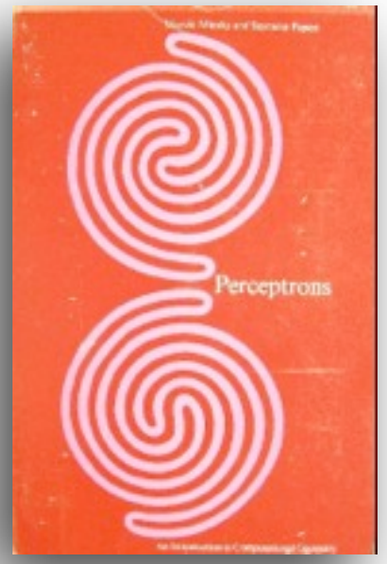

Result:

The 70s and early 80s:

- Significant reduction of funding for research on artificial neural networks

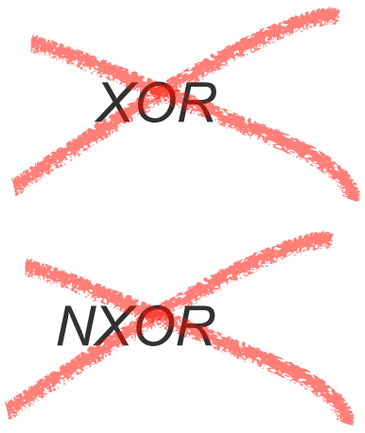

Perceptrons

Solution: more layers

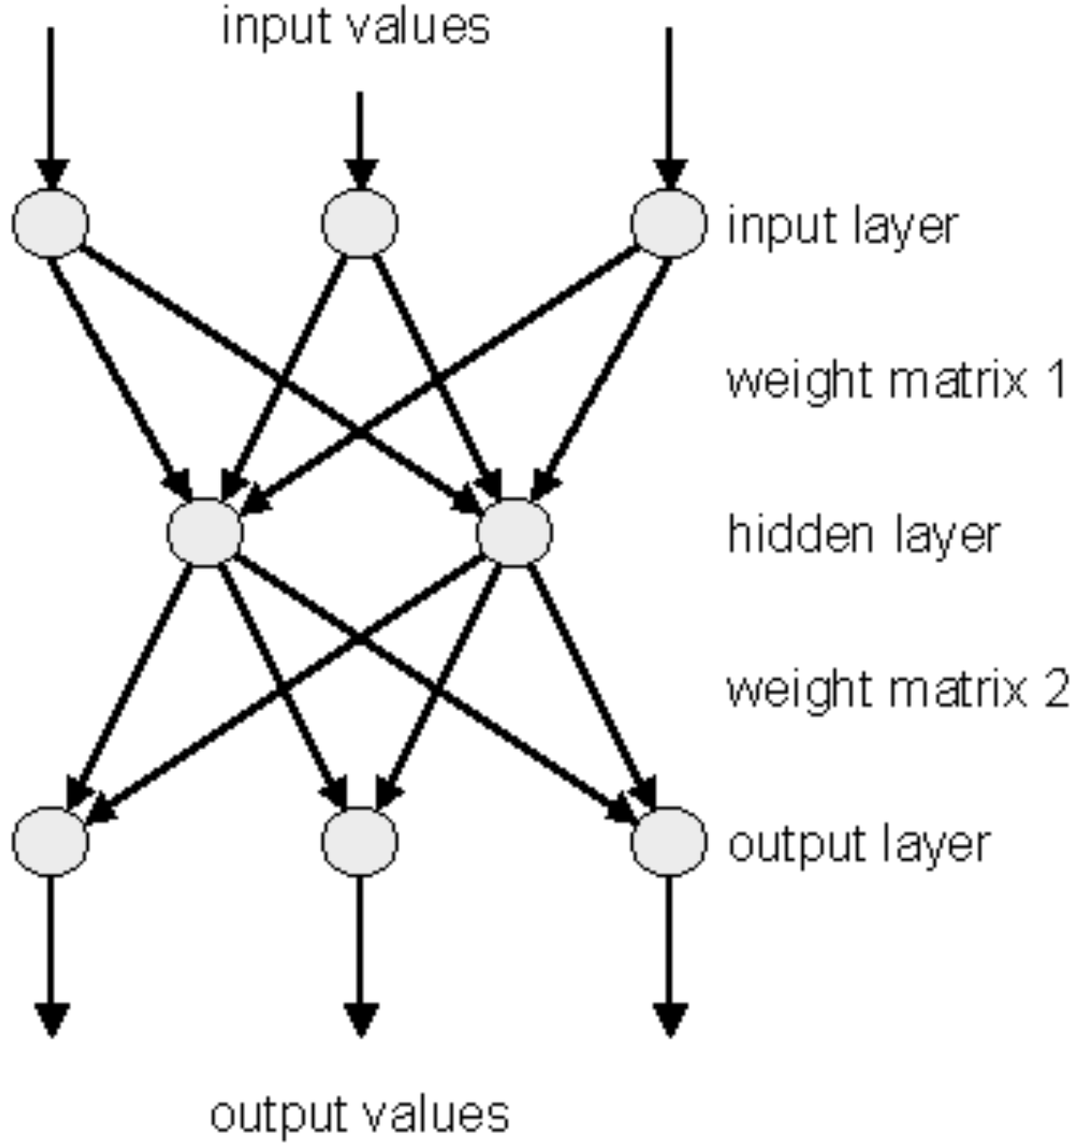

HOW DO YOU IMPLEMENT INTELLIGENCE IN A MACHINE?

Society of Mind

1986

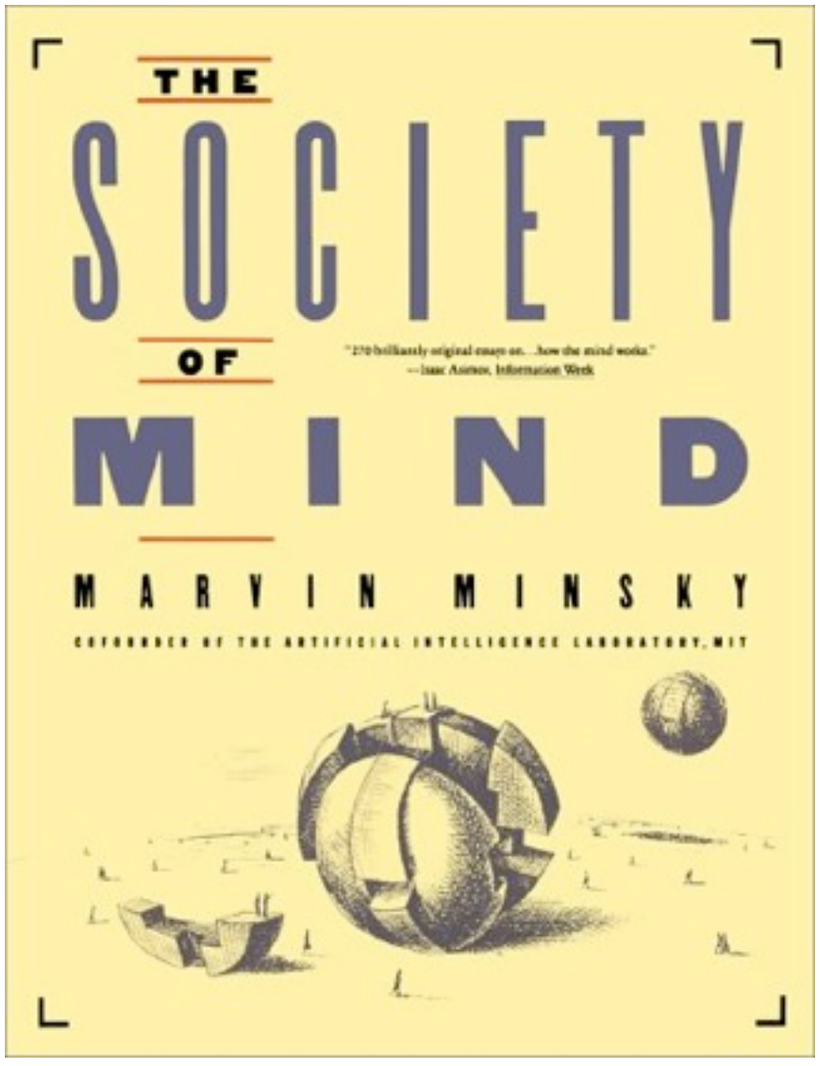

Society of Mind

1986

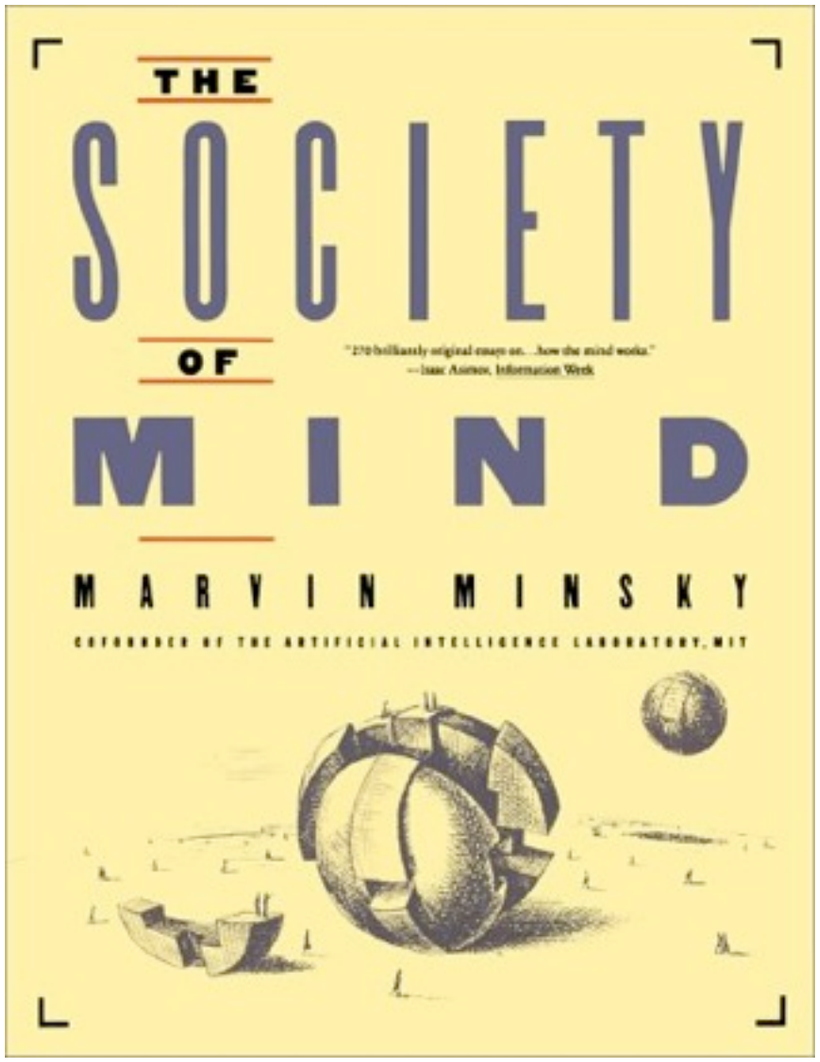

2007

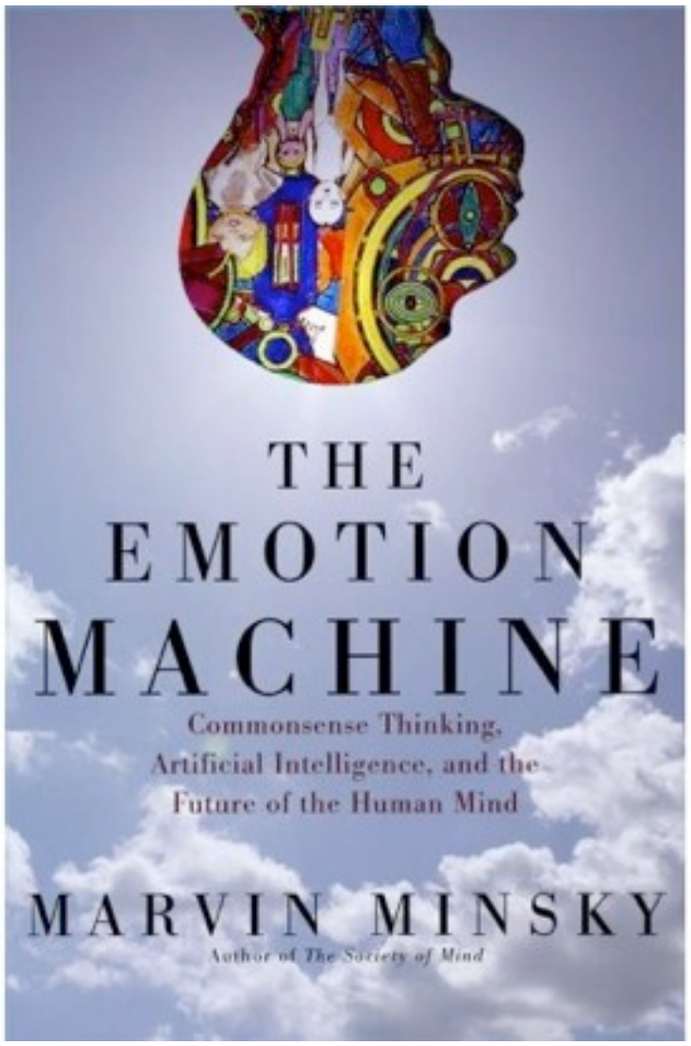

Society of Mind

Basic theory:

The mind is like a society, with myriads of agents interacting to produce what we call a “mind”.

> The functions performed by the brain are the products of the work of **thousands of different, specialized sub-systems** .... We **cannot** hope to understand such an organization by emulating the techniques of those particle physicists who search for the **simplest possible unifying conceptions**. [bolded: KRTh]

Society of Mind

Basic theory:

The mind is like a society, with myriads of agents interacting to produce what we call a “mind”.

> The functions performed by the brain are the products of the work of **thousands of different, specialized sub-systems** .... We **cannot** hope to understand such an organization by emulating the techniques of those particle physicists who search for the **simplest possible unifying conceptions**. Constructing a mind is simply **a different kind of problem**—of how to synthesize organizational systems that can support a large enough diversity of different schemes, yet enable them to work together to exploit one another's abilities. [bolded: KRTh]

Society of Mind

Basic theory:

The mind is like a society, with myriads of agents interacting to produce what we call a “mind”.

> Both my collaborator, Seymour Papert, and I had long desired to combine a mechanical hand, a television eye, and a computer into a robot that could build with children's building-blocks. It took several years for us and our students to develop Move, See, Grasp, and hundreds of other little programs we needed to make a working Builder-agency... It was this body of experience, more than anything we'd learned about psychology, that led us to many ideas about societies of mind.

Society of Mind

Basic theory:

The mind is like a society, with myriads of agents interacting to produce what we call a “mind”.

> An early articulation of these ideas took form in Carl Hewitt's concept of 'Actors', a computational model where sets of concurrently active agents solved problems by exchanging messages.
>
> – Push Singh, 2003 (*Examining the Society of Mind*)

Society of Mind

Basic theory:

The mind is like a society, with myriads of agents interacting to produce what we experience as “mind”.

The “homunculus” problem is avoided as the agencies are hierarchical, and at each level down the agents are a bit dumber than in the layer above.

Society of Mind

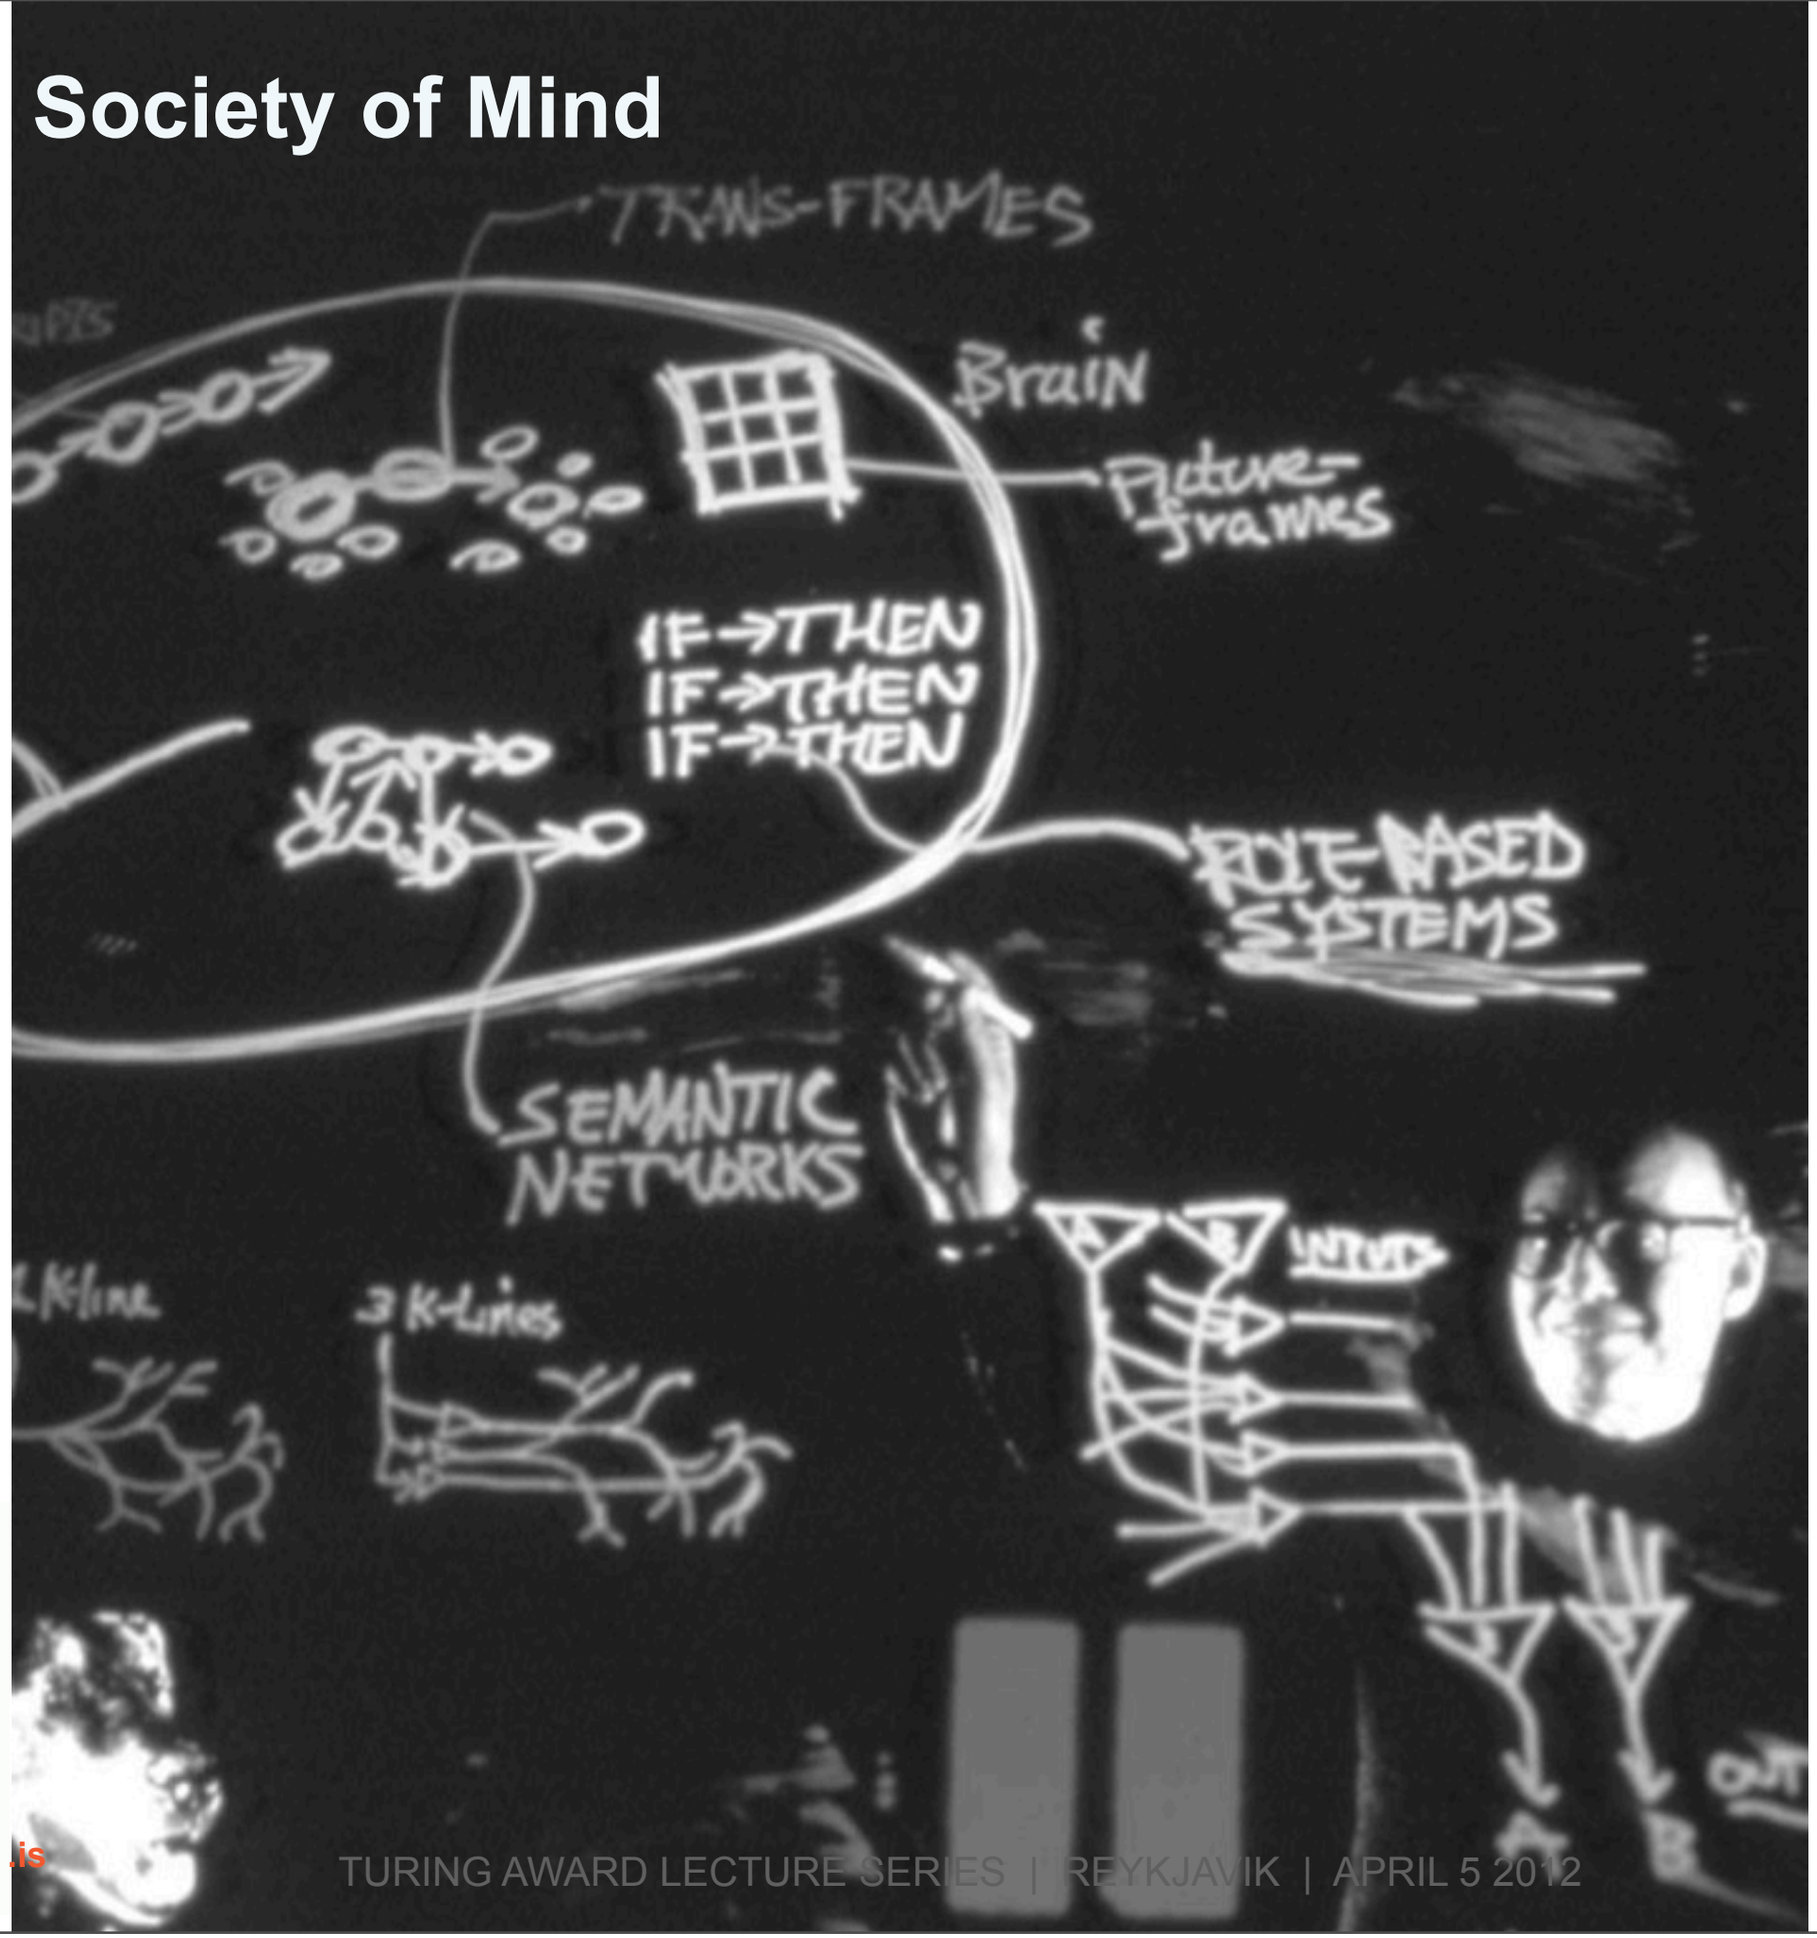

Society of Mind

- K-lines: Purpose of a K-line is simply to turn on a particular set of agents, and because agents have many interconnections, activating a K-line can cause a cascade of effects within a mind.

- Nomes and nemes: two general classes of K-lines, analogous to the data and control lines in the design of a computer. Nemes are concerned with *representing* aspects of the world, and nomes are concerned with *controlling* how those representations are processed and manipulated.

- Isonomes: Signal to different agencies to perform the same uniform type of cognitive operation.

- Pronomes: Are isonomes that control the use of short-term memory representations.

- Paranomes: Sets of pronomes linked to each other so that assignments or changes made by one pronome to some representation produce corresponding assignments or changes by the other pronomes to related representations.

Society of Mind

- Frames: A form of knowledge representation concerned with representing a thing and all the other things or properties that relate to it in certain particular ways, which are attached to the 'slots' of the frame.

- Frame arrays: Frame-arrays are collections of frames that have slots or pronomes in common.

- Transframes: represent events and all of the entities that were involved with or related to the event.

- Difference engines: Solving a problem can be regarded as reducing or eliminating the important differences between the current state and some desired goal state.

- A-brains and B-brains: While your A-brain thinks about the outside world your B-brain's job is to think about the world inside the mind – pay attention to the 'A-brain', so as to be able to notice these potential errors and correct them, evaluate performance, etc.

Society of Mind: The Course

- given every other year
- open to all MIT students
- you fill out a questionnaire and send in your CV
- Marvin hand-picks the ones he wants to attend
- In the first half he gives an overview of some topic(s)
- In the second half students drive the discussion with questions

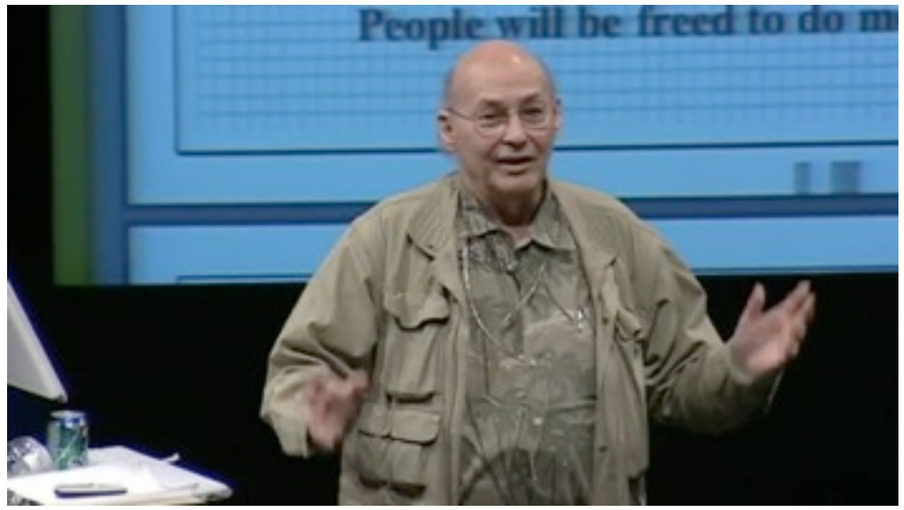

Society of Mind: The Course

- given every other year
- open to all MIT students
- you fill out a questionnaire and send in your CV
- Marvin hand-picks the ones he wants to attend
- In the first half he gives an overview of some topic(s)
- In the second half students drive the discussion with questions

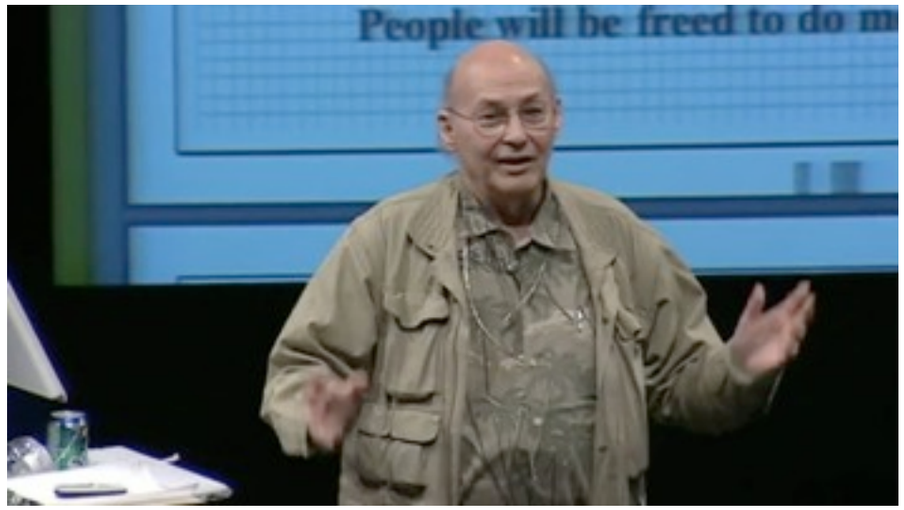

- Two small essays
- One final project

Society of Mind: The Course

Why do we like music?

Why isn't music boring?

What is humour?

How do we recognize objects and scenes?

How do we understand things?

What is understanding?

Why is math hard?

How do we learn new concepts and skills?

How do we achieve goals?

What are feelings and emotions?

How does 'common sense' work?

Why do we want to dance to some music?

....

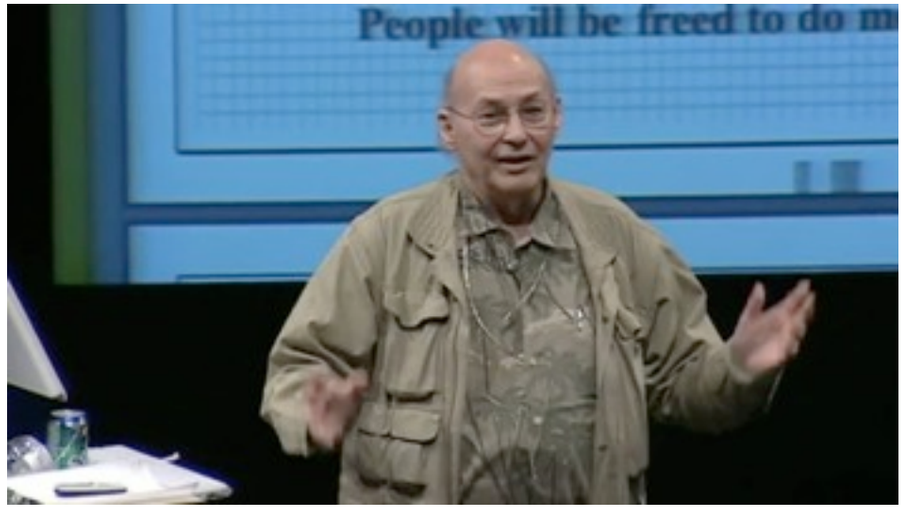

Piano player

Piano player

Piano player

Invented the Triadex Muse Sequencer/Synthesizer

Piano player

Invented the Triadex Muse Sequencer/Synthesizer

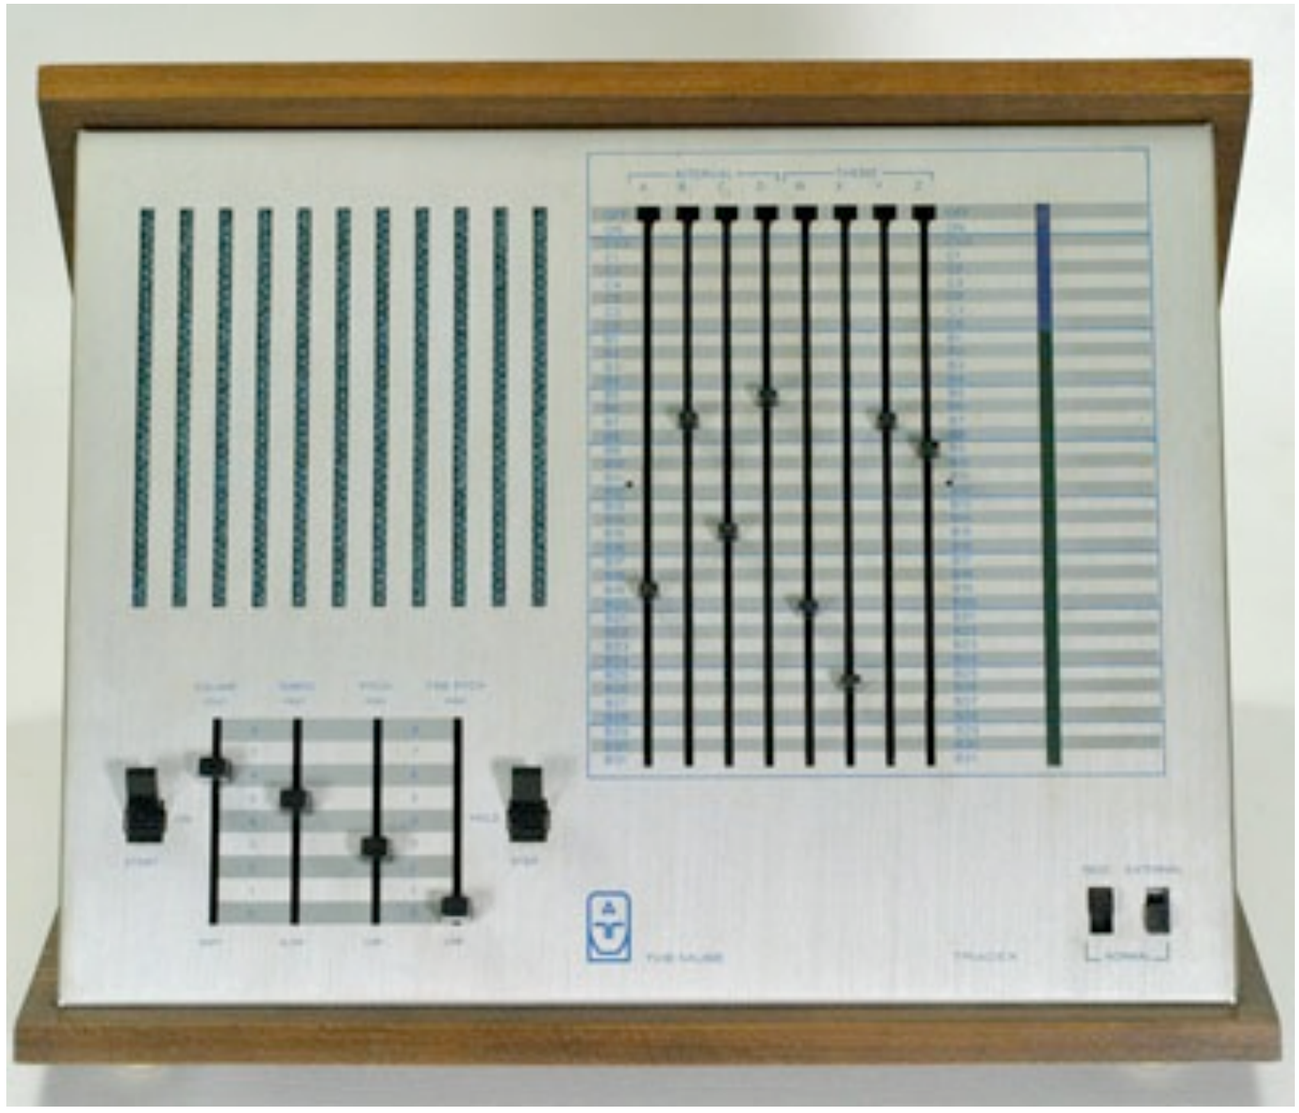

Piano player

Invented the Triadex Muse Sequencer/Synthesizer

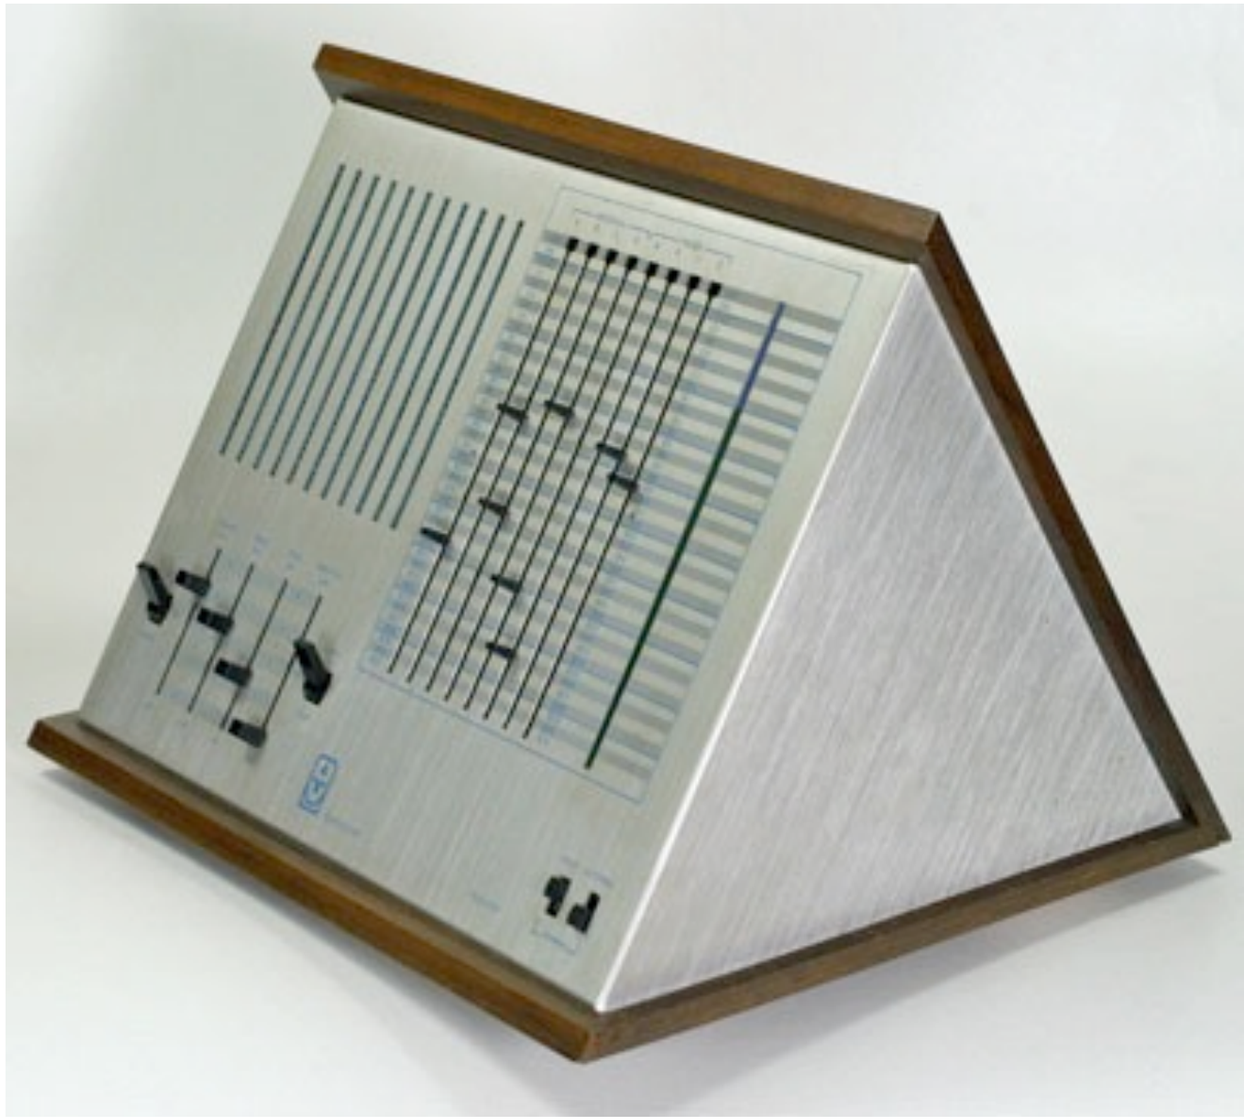

Piano player

Invented the Triadex Muse Sequencer/Synthesizer

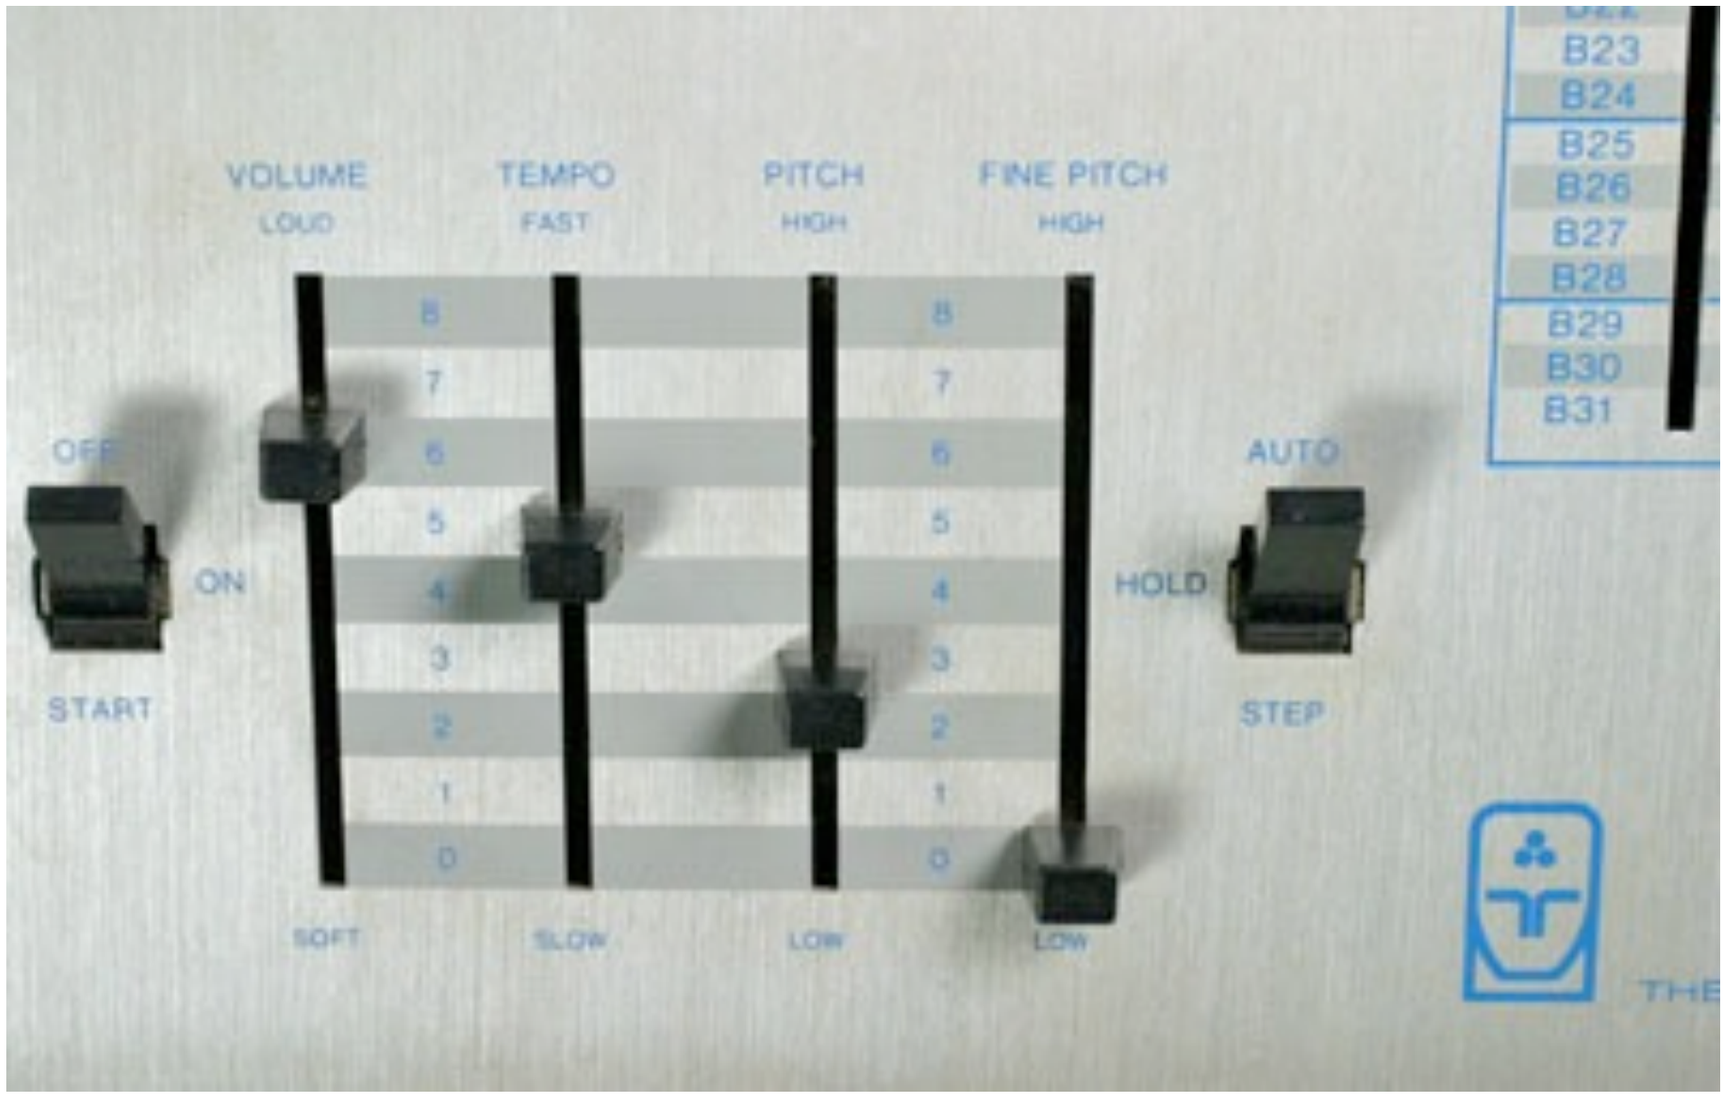

Piano player

Invented the Triadex Muse Sequencer/Synthesizer

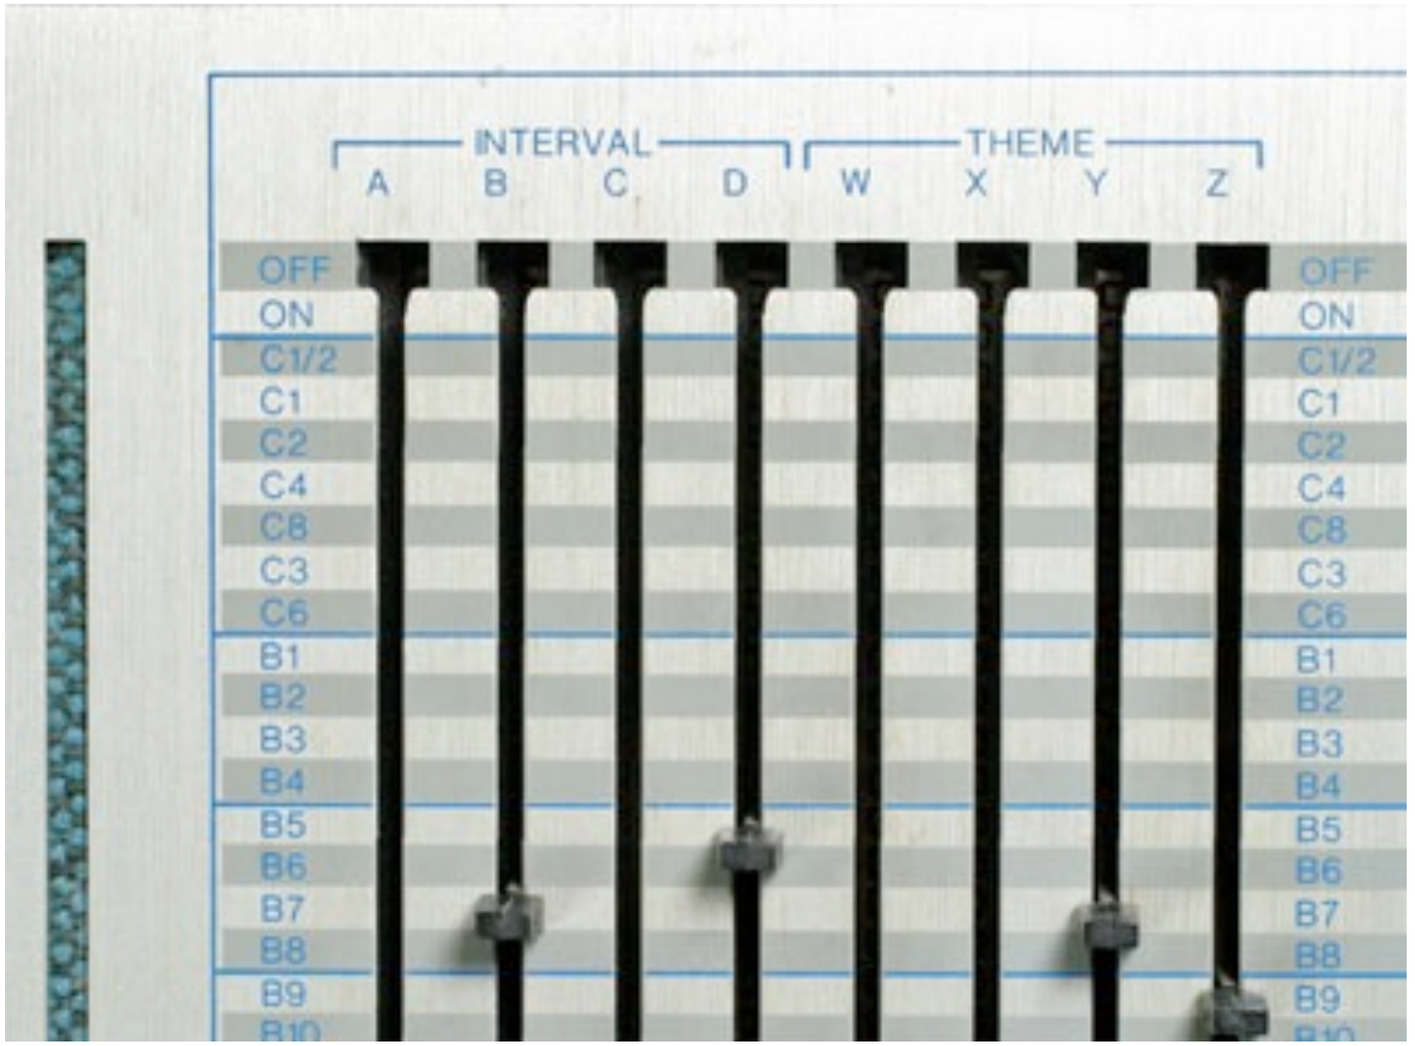

Piano player

Invented the Triadex Muse Sequencer/Synthesizer

https://soundcloud.com/republic-of-nynex/triadex-muse-emulator-tryout

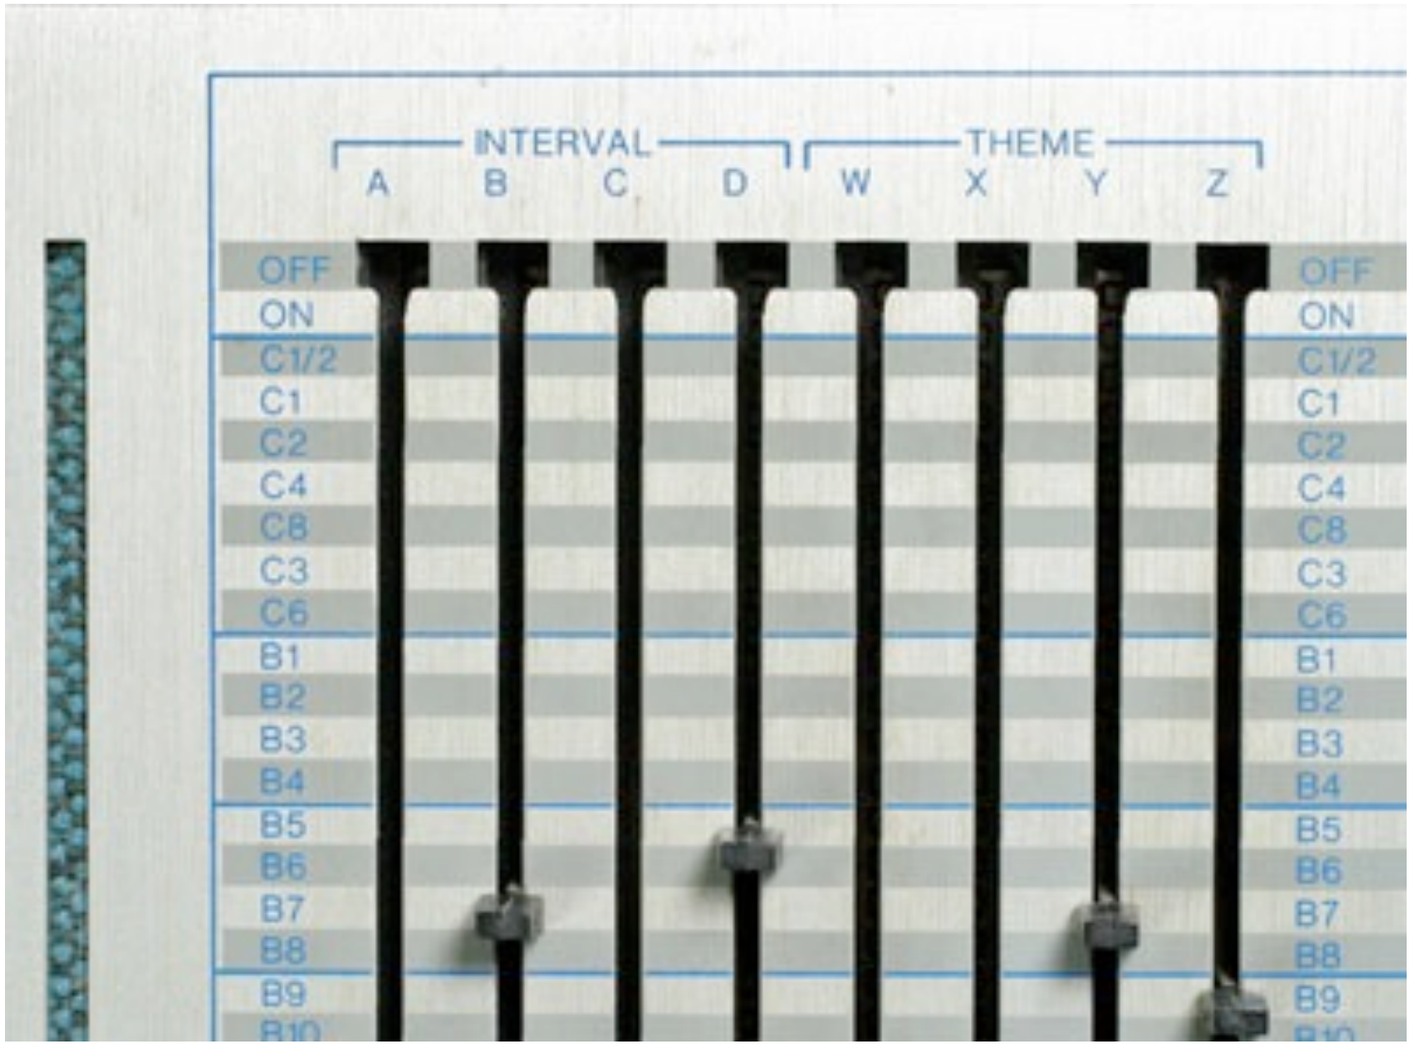

Assignment

Why do we like music – and why do we get bored with it?

The explanation lies with internal prediction: When we listen to music we predict the evolution of the soundscape. When a particular song or melody stops surprising us, we “get tired” of it.

Minsky the Critic: “Grumpy Old Man of AI”

Marvin heavily criticizes:

- The strong emphasis on robots at the MIT AI lab.
  - Better to use simulations.
- Lack of focus on *common sense* in AI research.
  - Only project to on the rigth track: Cyc.
- Any and all attempts at using neurons as the starting point for AI.
  - Intelligence is a symbolic skill.
  - Neural nets don’t contain any structure.

Minsky the Critic: “Grumpy Old Man of AI”

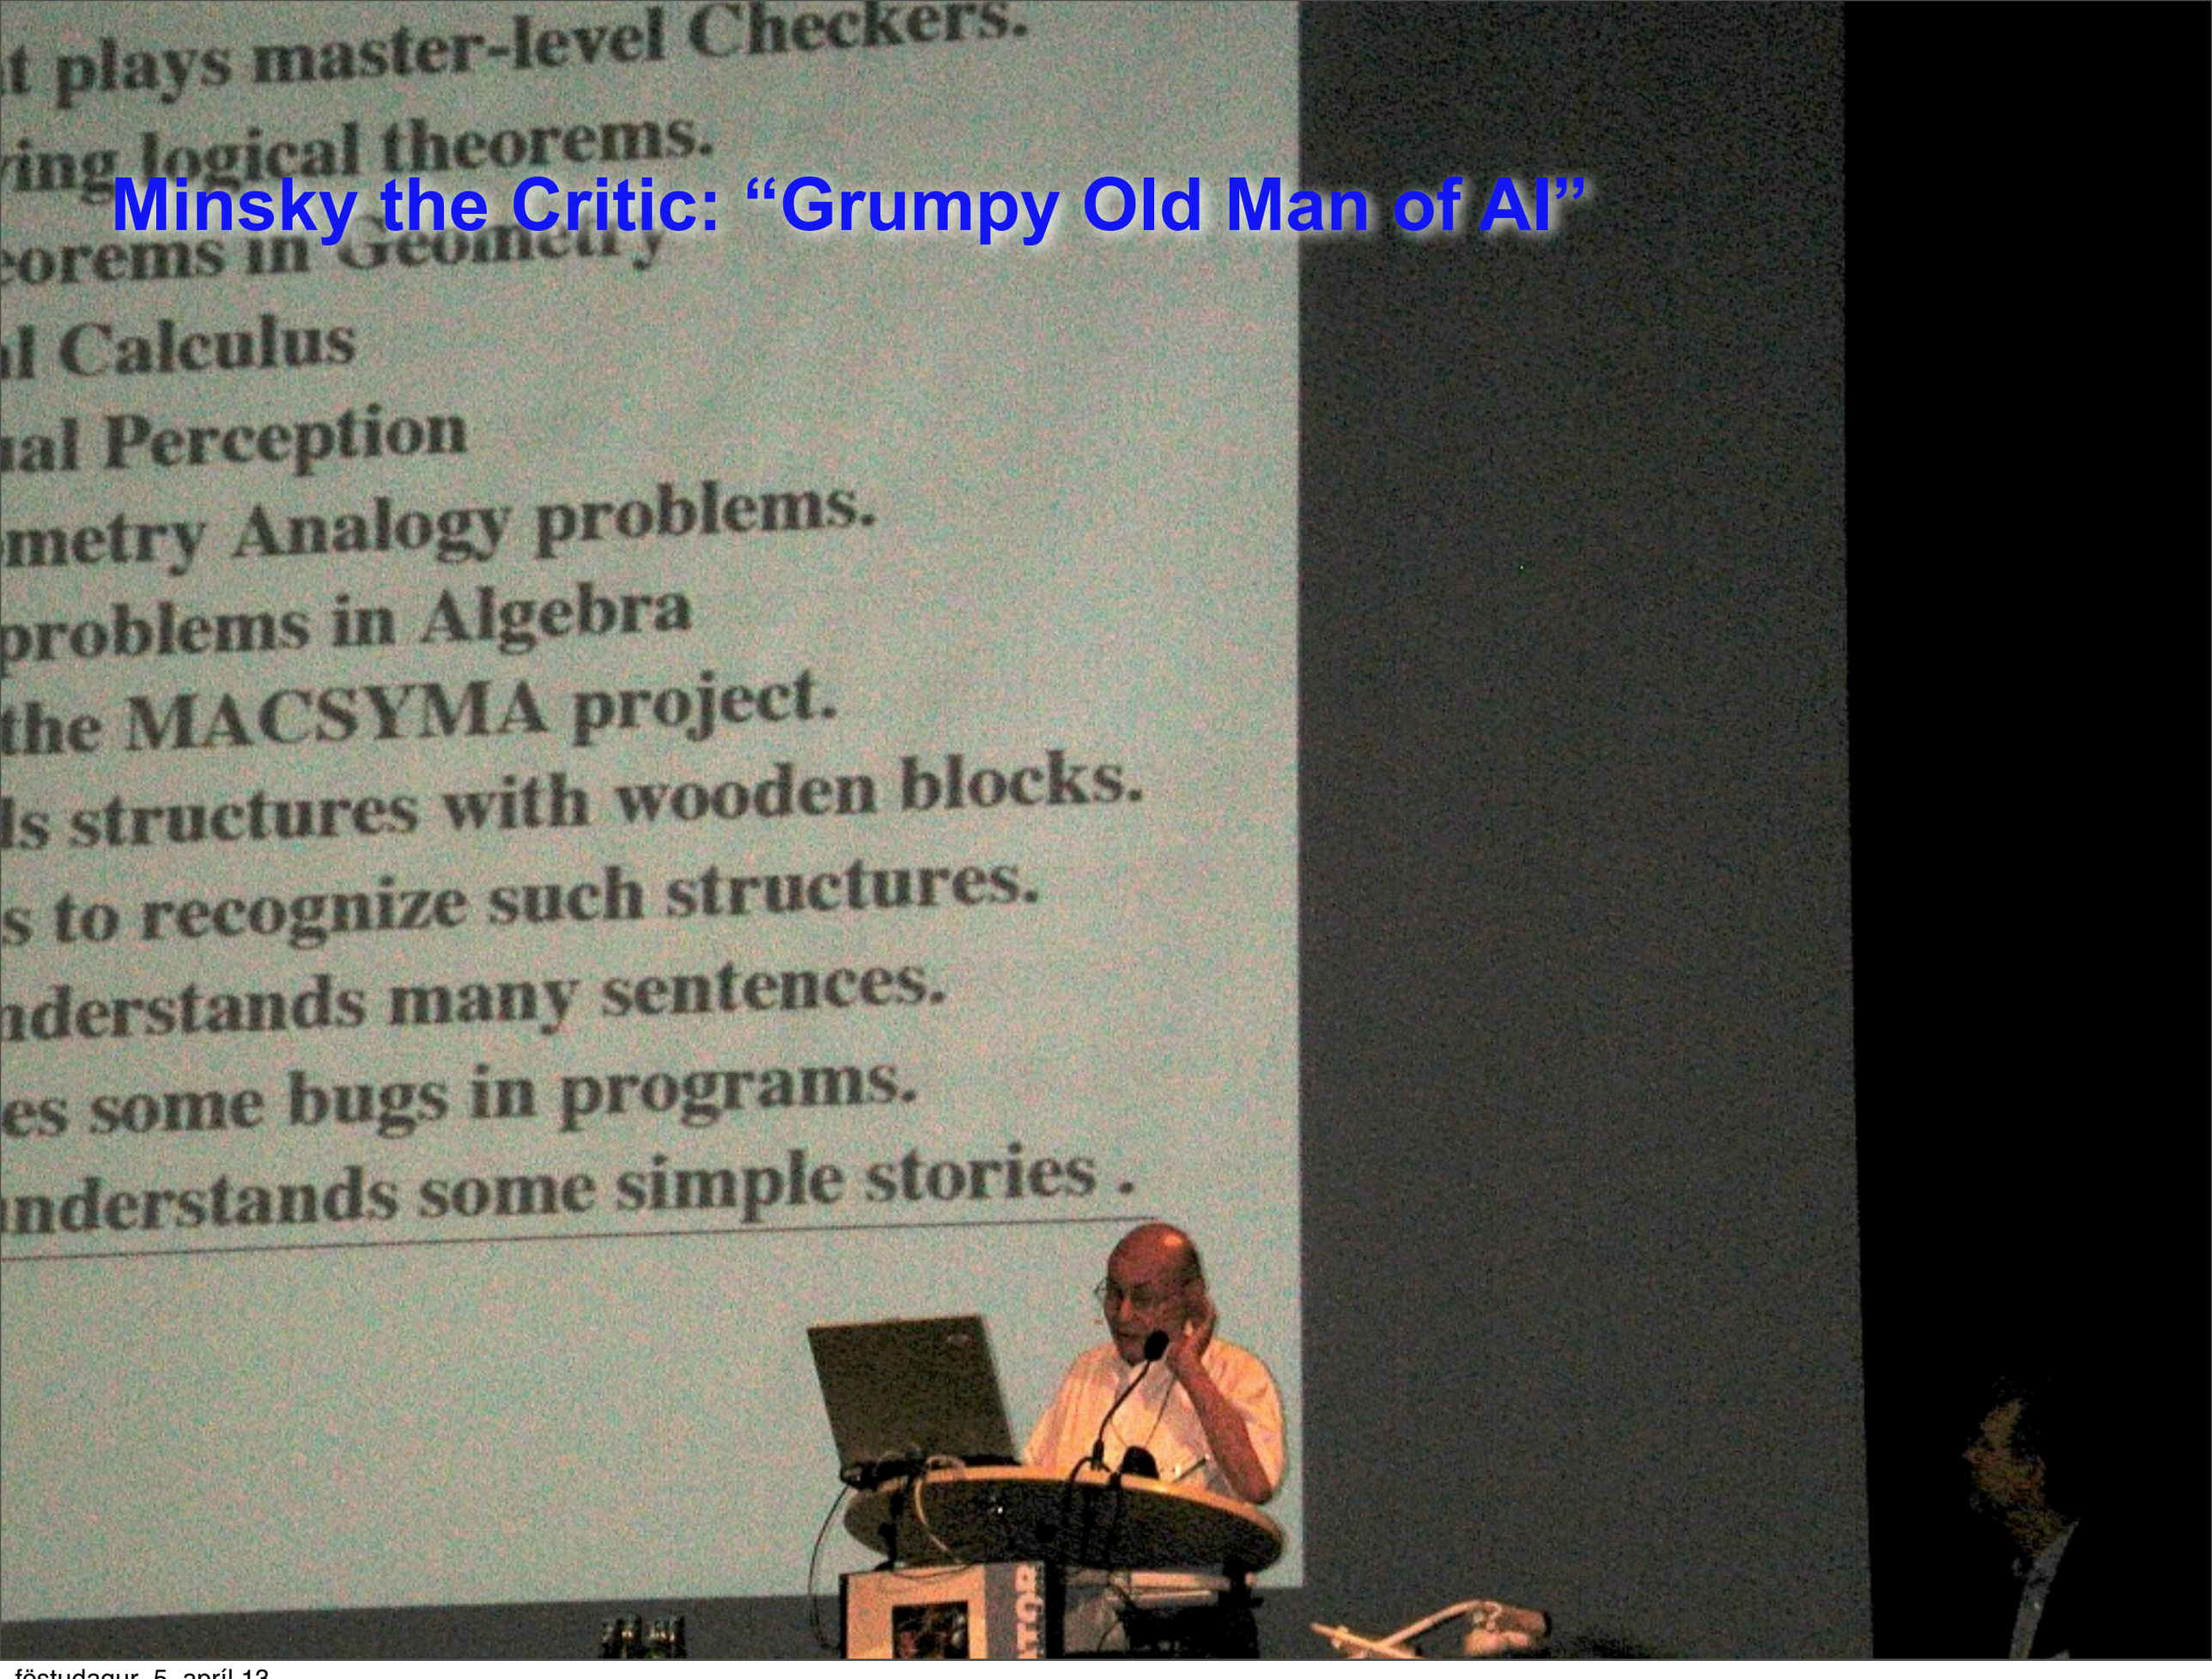

Minsky the Critic: “Grumpy Old Man of AI”

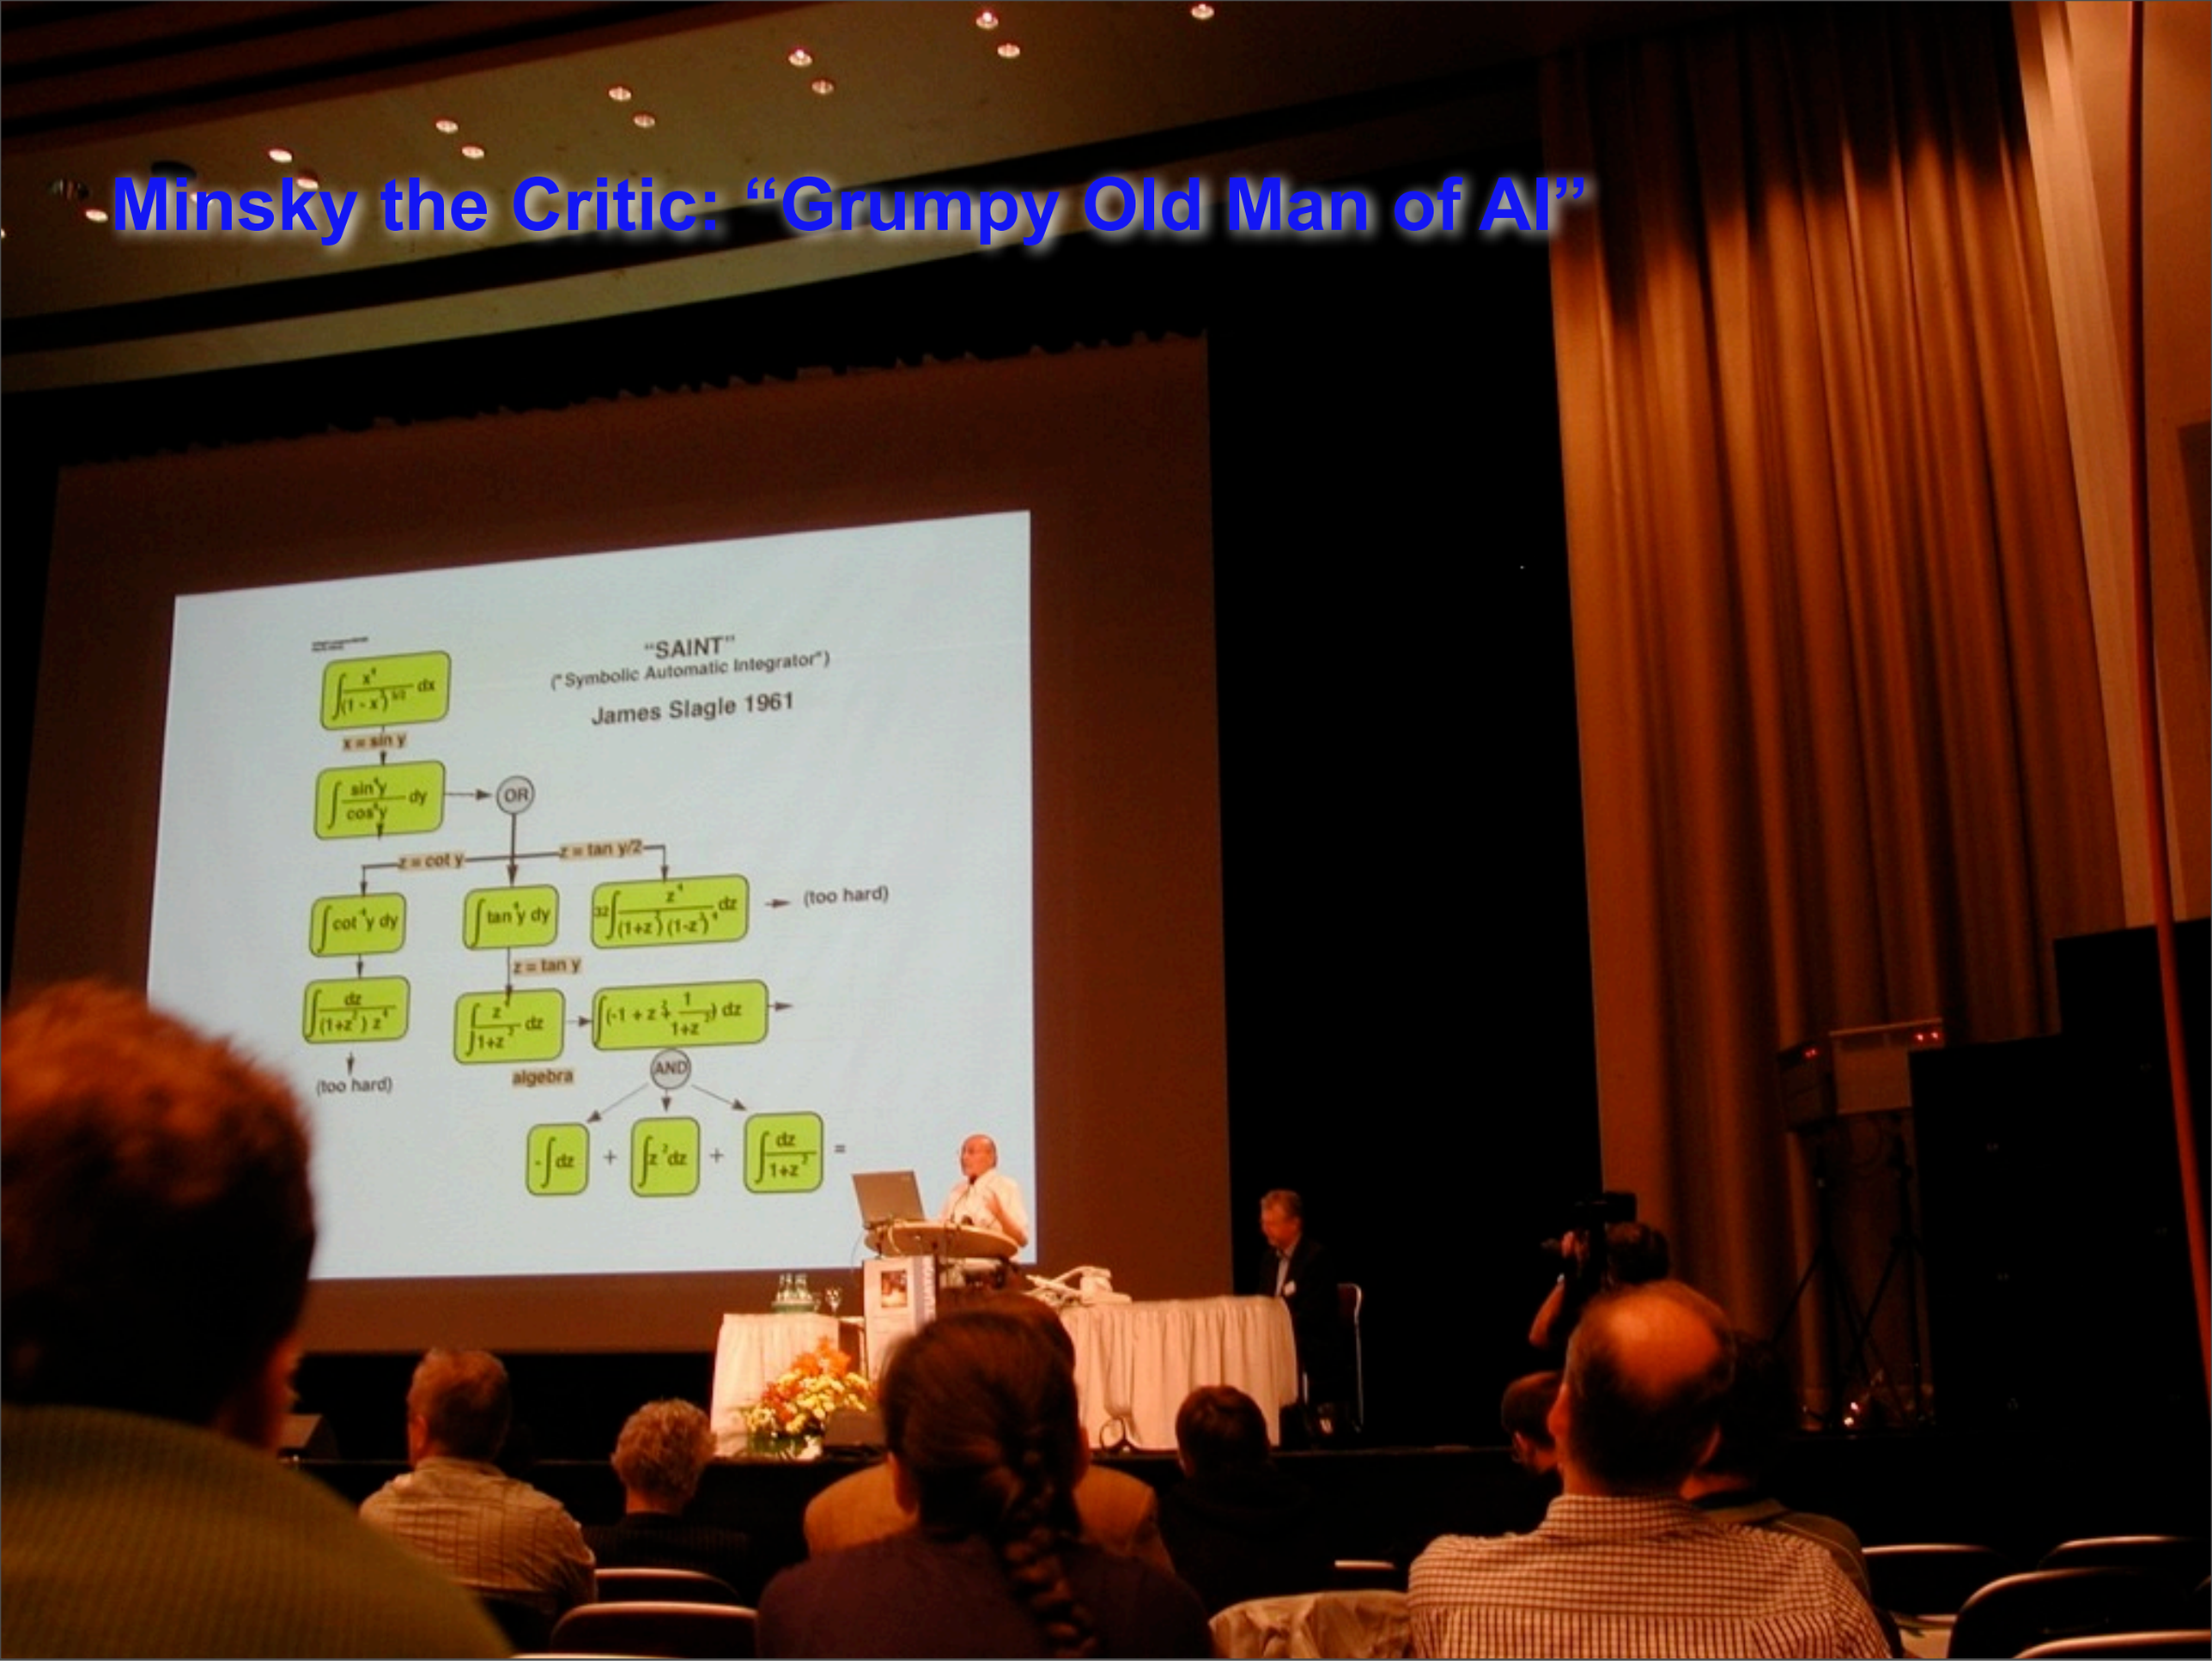

Minsky the Critic: “Grumpy Old Man of AI”

Minsky the Critic: “Grumpy Old Man of AI”

Marvin heavily criticized:

- Me, at that 50th anniversary, for making an “obviously irrelevant and very ill-formulated comment on the organization of Western scientific research”, which clearly has worked excellently for millenia.

What has Marvin Taught Me?

- What “understanding” is.
- That it’s OK to think about stuff that no one else seems to care about (at the moment).
- That being brilliant is not enough (possibly not even necessary) to being a good scientist.
- That being rude, agressive, or generally abrasive, while common in the top-level AI circles, does not help one’s carreer.
- That producing one good idea per decade in science is “going strong”.

Where does Society of Mind Stand Today?

Society of Mind challenges:

- how to coordinate multiple agents
- how to distribute tasks among them
- how to evaluate their results
- how to integrate agent output into coherent behavior

Things we won’t have time to talk about

- 1951 – the world's first “randomly wired neural network learning machine,” called the Stochastic Neural-Analog Reinforcement Computer (SNARC)
- 1957 – The confocal microscope
- 1959 – The MIT AI lab
- 1963 – The first head-mounted graphical display
- 1974 – Framework for representing knowledge
- 1985 – MIT Media Lab

2001: A Space Oddyssey

2001: A Space Oddyssey

www.ru.is

2001: A Space Oddyssey

www.ru.is

2001: A Space Oddyssey

www.ru.is

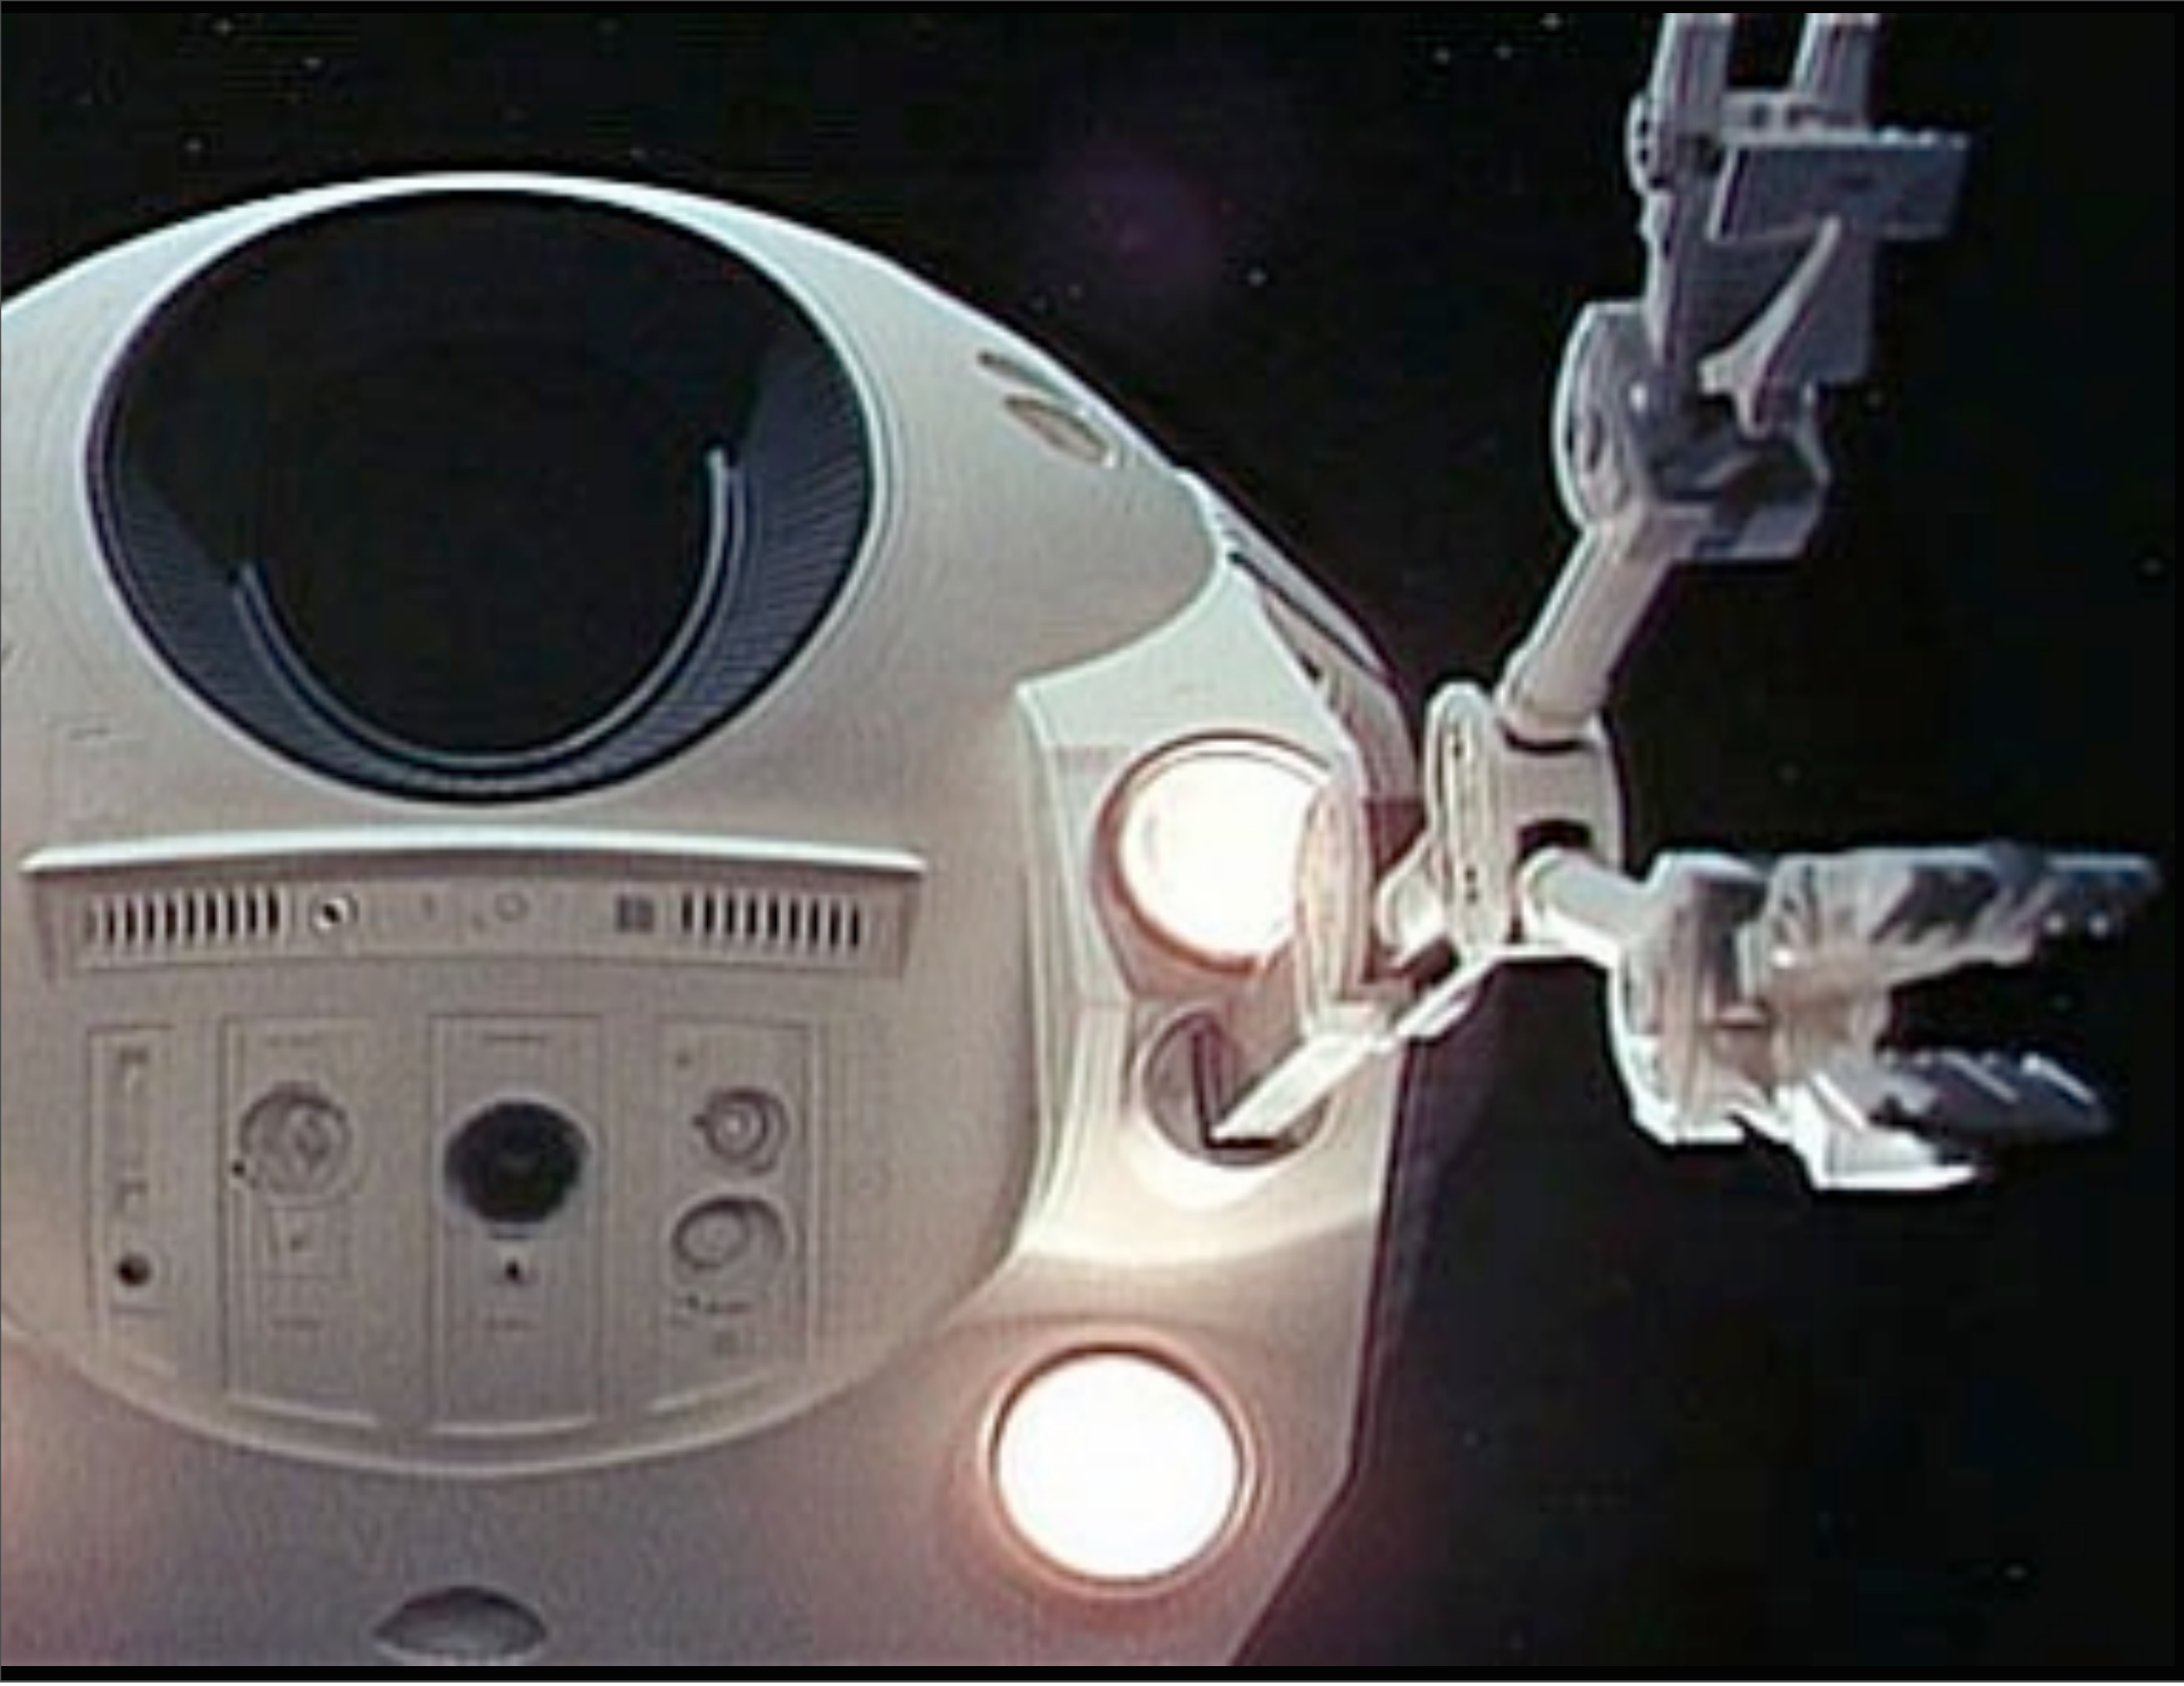

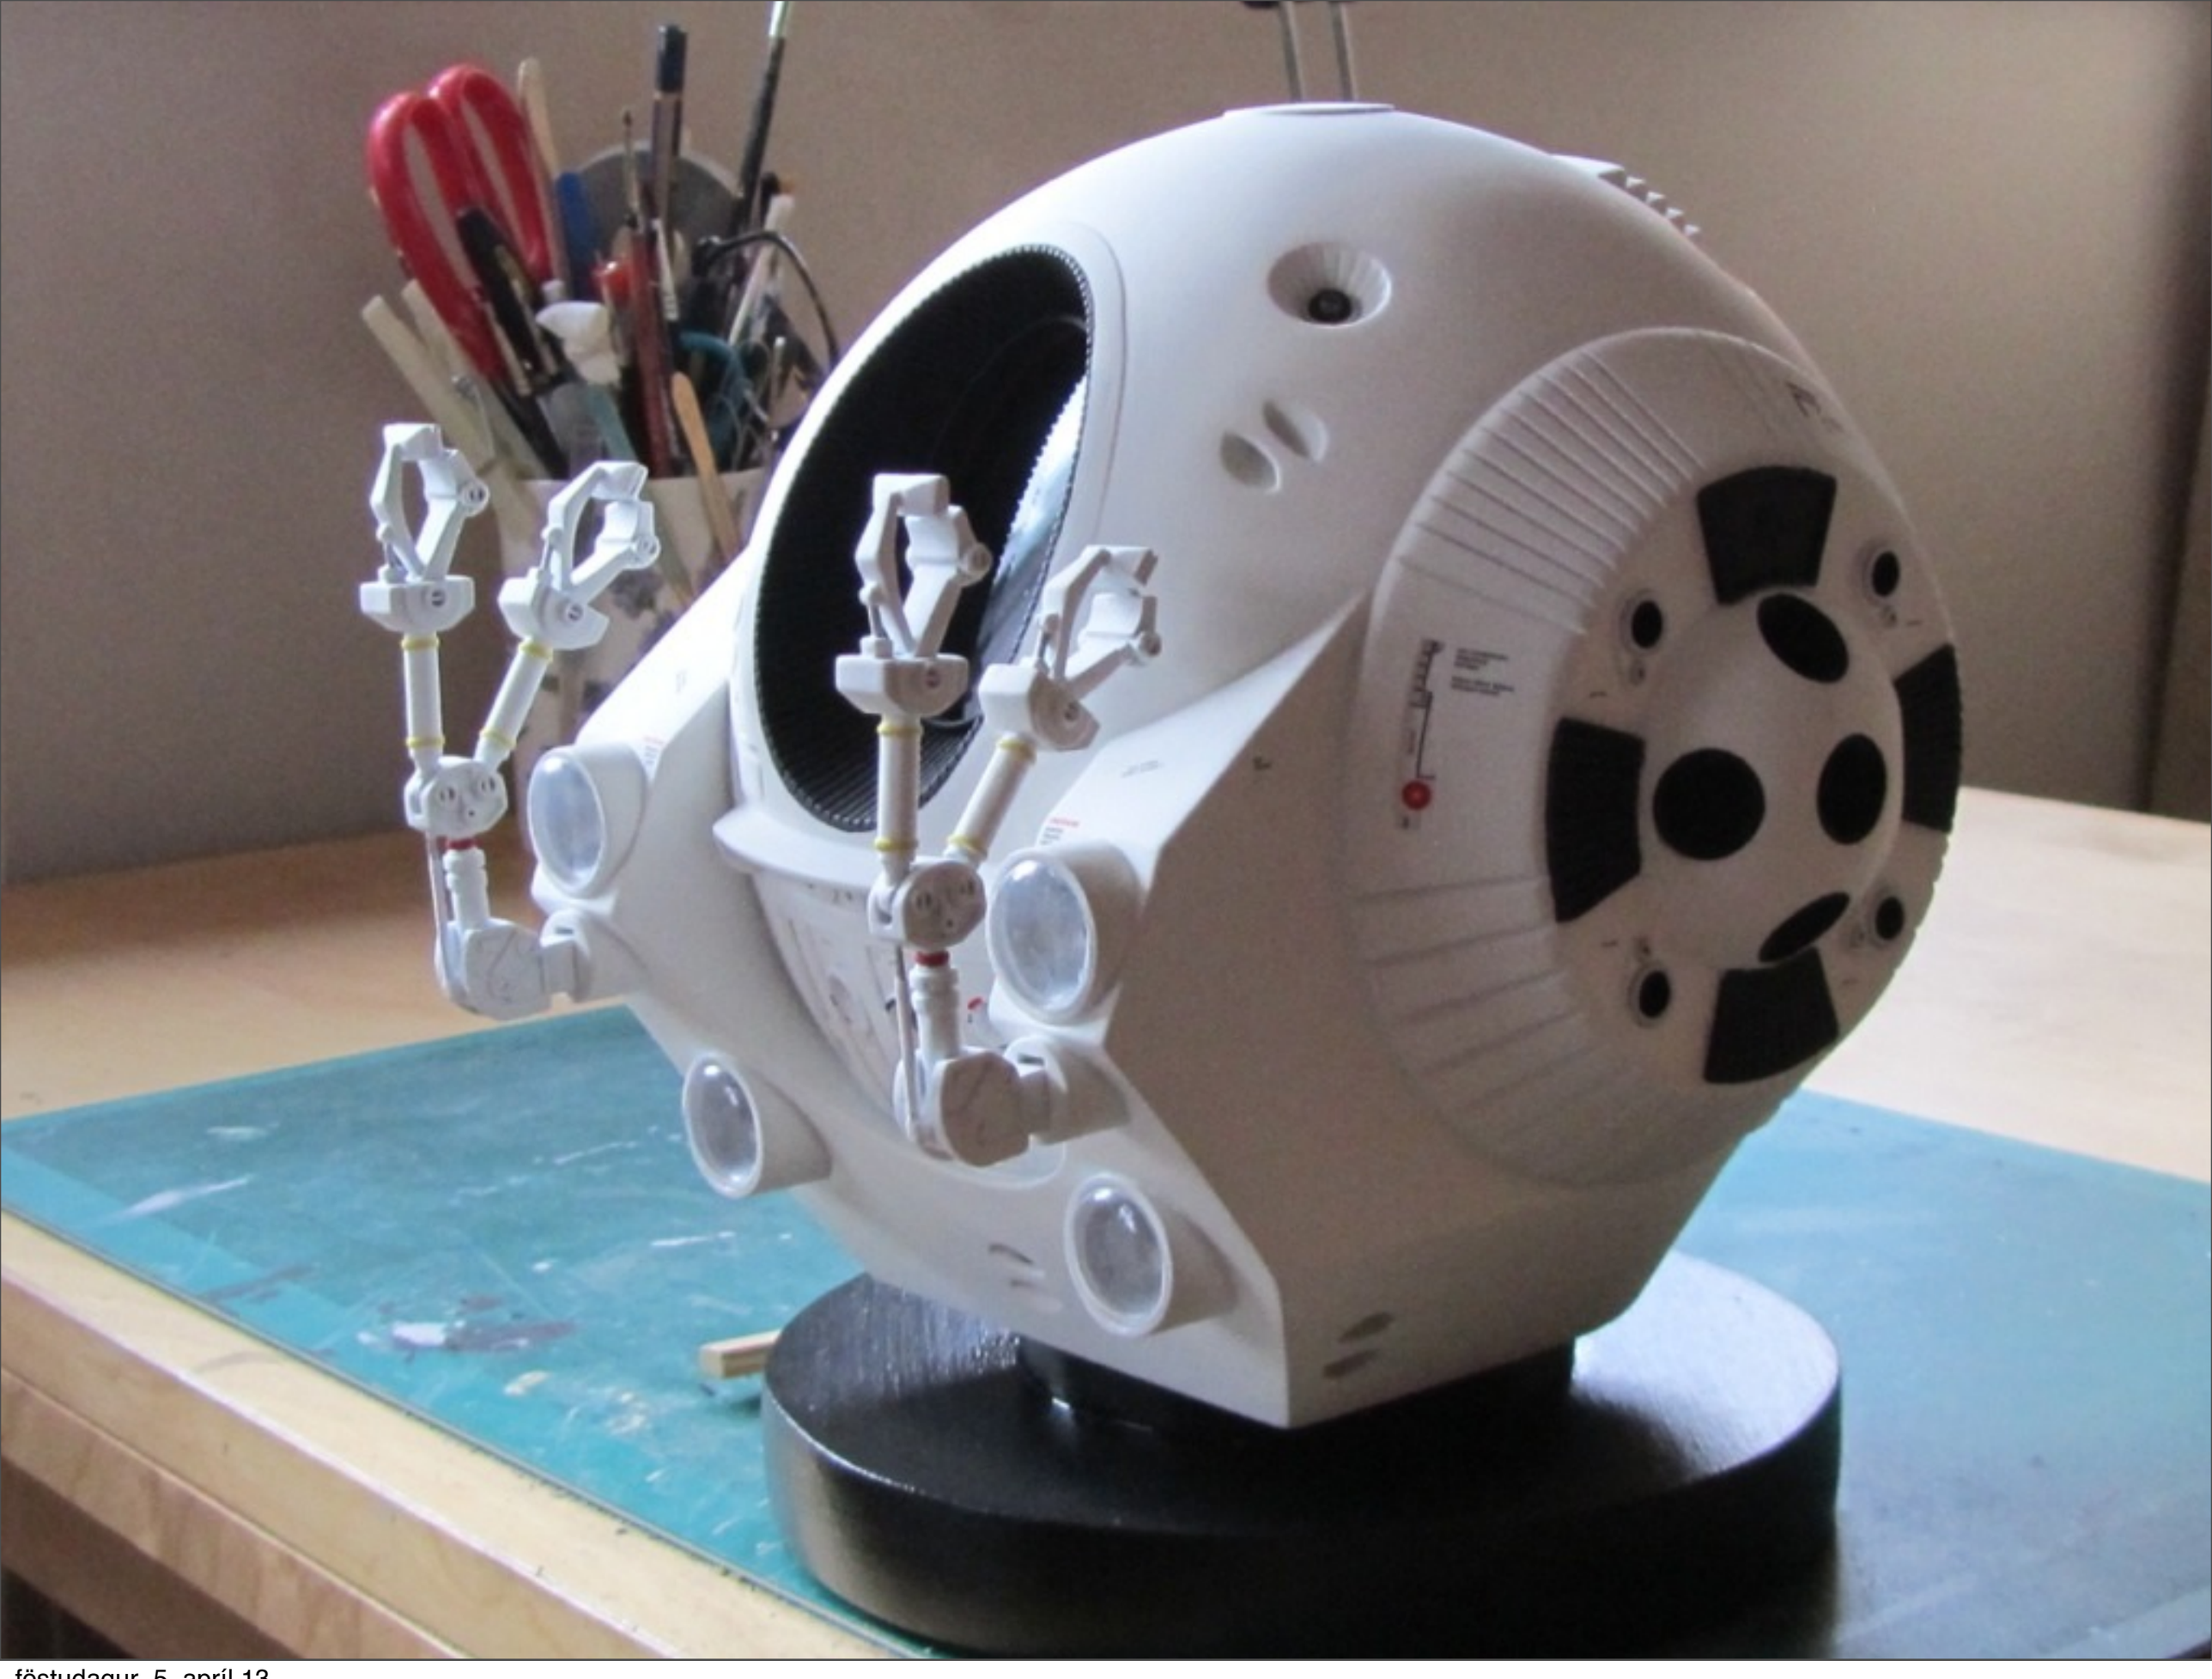

2001: A Space Oddyssey

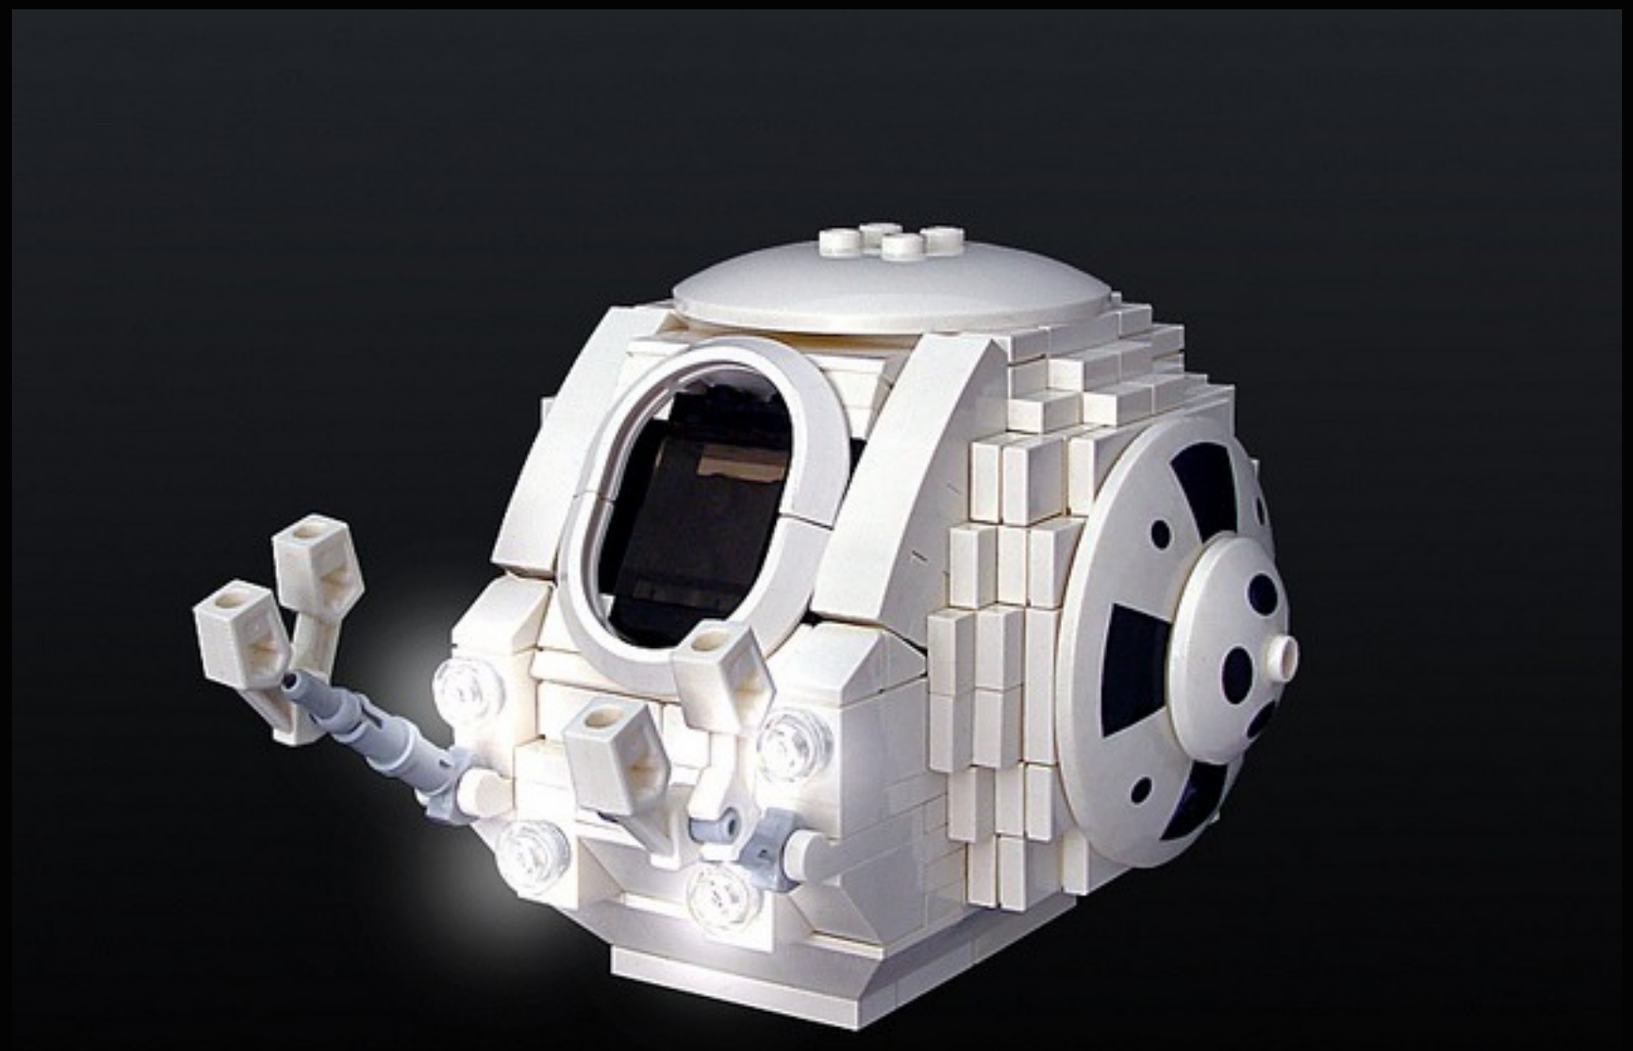

THANK YOU!

www.ru.is

Intelligence: Common View

Intelligence is the ability to X

X = play chess

X = vacuum the floor

X = drive a car through the desert

X = chat