From Shannon to Superintelligence: How We Built a Thinking Machine in 70 Years

How language models were created, and how they are changing the world today. This lecture traces a fascinating journey — from Claude Shannon and his information theory, through Fred Jelinek's pioneering work in speech recognition, to Tomáš Mikolov's RNNs. The second part focuses on the current revolution — agentic systems that are already autonomously solving complex tasks today. In closing, we'll reflect on the near future: recursive self-improvement of AI, solving open scientific problems, and the question of whether we stand on the threshold of the greatest technological turning point in human history.

The AI Odyssey

From Shannon to Superintelligence: How We Built a Thinking Machines in 70 Years

Jan Šedivý
ČVUT CIIRC, Praha

17-06-26

From Shannon to Superintelligence

How We Built a Thinking Machine in 70 Years

01 . The Journey

Traces the fascinating evolution of language models—from Claude Shannon's foundational information theory, through Fred Jelinek's speech recognition, to Tomáš Mikolov's RNNs.

02 . The Present

Focuses on the current revolution—agentic systems that are already autonomously solving complex tasks today and changing how we interact with technology.

03 . The Future

Reflects on recursive self-improvement of AI, solving open scientific problems, and the question of whether we stand on the threshold of the greatest technological turning point in human history.

LLM Basic Idea

Sentence probability:

The goal of a language model is to estimate the

$$P(W) = P(w_1, w_2, w_3, w_4, w_5 \ldots w_n)$$

w_1 may be related to w_5 dependency between distant words in a sentence.

Sentence generation:

Compute a **probability of an upcoming word.**

$$P(w_5 \mid w_1, w_2, w_3, w_4)$$

Conditional probability. What is the probability of w_5 given the probability of previous words?

The quick brown fox jumps over the lazy dog.

1913: Andrey Markov



Stochastic Language

Markov analyzed 20,000 letters of Pushkin's *Eugene Onegin* to prove that states can be predicted without memory of the past.

- **Probability:** Foundation for sequence prediction.
- **Memoryless:** Current state determines the next.
- The ancestor of the "Next Token Prediction" used in LLMs today.

1929: Leo Szilard

Information as Physics

Szilard established the link between **Thermodynamics** and **Information** by solving the Maxwell's Demon paradox.

$$S = k \ln(2)$$

- The first recognition that processing information has a physical cost.
- Introduced the fundamental concept of a "bit."



1948: Claude Shannon

Information Theory

Shannon's "A Mathematical Theory of Communication" quantified information for the first time.

$$H = -\sum p_i \log_2 (p_i)$$

The Shannon Game: He challenged humans to predict the next letter in a sentence, measuring the entropy of English—the exact objective function of GPT.



1970s: Fred Jelinek

Speech (ASR)

Moved Automatic Speech Recognition from rules to statistics.

"Fire the Linguists"

Famous quote indicating that data beats human intuition.

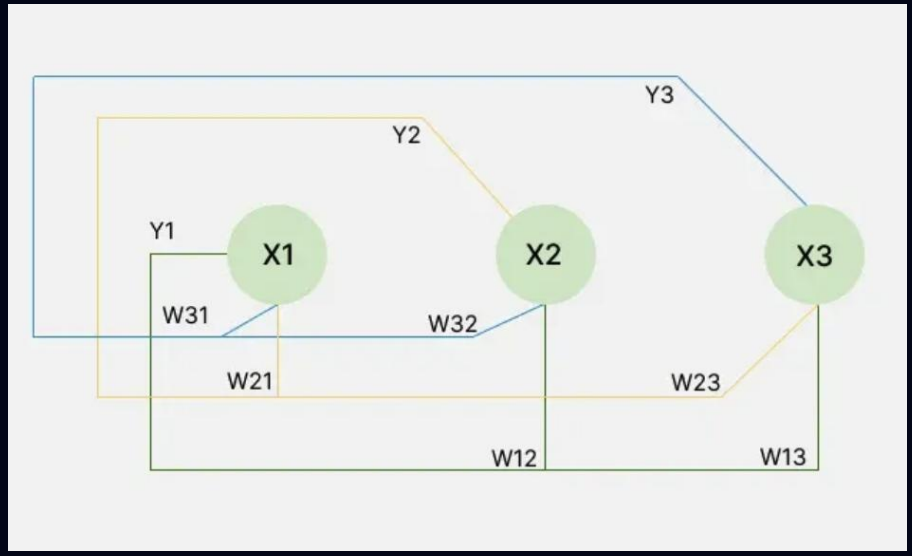


1982: John Hopfield

Associative Memory

Hopfield introduced a network that could store and reconstruct patterns using physics-based energy minima.

- **Content-Addressability:** Recalling a whole image from a fragment.
- **Energy Landscapes:** Computation as a descent into stability.



1986: Geoffrey Hinton



Learning via Error

Hinton popularized **Backpropagation**, allowing multi-layer neural networks to learn internal representations.

- Chain rule applied to neural errors.
- Enabled "Deep Learning" by solving the credit assignment problem.
- Started the connectionist revolution.

2013: Tomas Mikolov

RNNs and Word2Vec

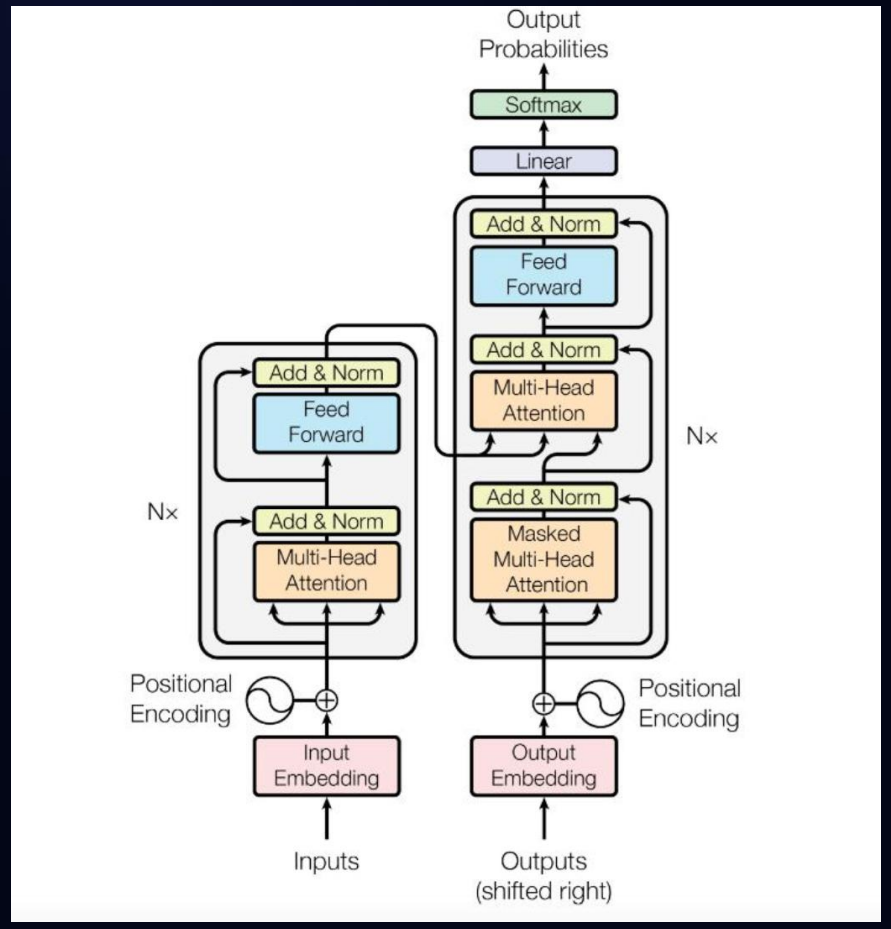
Mikolov proved that RNNs could model long sequences and that words could be represented as vectors.

King - Man + Woman = Queen

- Distributed representations of meaning.
- Efficient training on massive corpora.



2017: Attention Is All You Need



The Transformer

Vaswani et al. introduced the architecture that replaced recurrence with parallel **Self-Attention**.

- Direct links between distant tokens.
- Massive parallelization across GPUs.
- The "Engine" inside GPT, Claude, and Gemini.

2022: OpenAI & ChatGPT



RLHF

Reinforcement Learning from Human Feedback aligned AI with users.



Exponential Growth

Fastest growing consumer application in history.



General Intelligence

Demonstrated zero-shot reasoning across multiple domains.

2024: The Nobel Moment



Validation of a Lifetime

The 2024 Nobel Prize in Physics awarded to **Hopfield** and **Hinton**.

A clear signal that AI is not just a tool, but a fundamental breakthrough in our understanding of information and physical systems.

2025: Agentic Systems

Claude Code

AI moves from *responding* to *doing*. Claude Code represents the shift toward autonomous software engineering.

- Tool Use: Executing terminal commands.
- Long Horizon: Building full applications.
- Agentic Loop: Observe, Plan, Act, Verify.



The Future of AI 1-2 years

Recursive Self-Improvement

AI designing better architectures for the next generation of AI.

Science Accelerated

Discovery of materials and drugs in days, not decades.

Infinite Scalability

Intelligence becoming a universal commodity, like electricity.

"The best way to predict the future is to invent it."

6–10 Years

LOW CONFIDENCE · SPECULATIVE

loop

Recursive Loops at Scale

AI systems begin rewriting their own weights, training procedures, and architectures — with meaningful improvements each cycle. Governance and alignment research become critical bottlenecks.

science

New Physics Discovery

AI moves from accelerating known science to generating genuinely novel hypotheses beyond the Standard Model — or identifying measurement strategies humans haven't conceived.

alt_route

Transformative or Plateau?

Either AGI-adjacent systems emerge, or fundamental limits in reasoning capability demand a new paradigm. The 10-year outcome depends on alignment, compute, and breakthroughs.

WILD CARD: *An unexpected paradigm shift (neuromorphic, quantum ML, or something entirely novel) could accelerate or redirect everything above.*

Q & A