Escaping the local minimum

By Kenny Friedman

MAS.S63
Integrative Theories of
Mind
and Cognition

May 9, 2016

I'M JUST OUTSIDE TOWN, SO I SHOULD BE THERE IN FIFTEEN MINUTES

ACTUALLY, IT'S LOOKING MORE LIKE SIX DAYS

NO, WAIT, THIRTY SECONDS.

THE AUTHOR OF THE WINDOWS FILE COPY DIALOG VISITS SOME FRIENDS

WE MADE A BREAKTHROUGH, SO IT SHOULD BE JUST A FEW YEARS NOW

ACTUALLY, IT'S LOOKING MORE LIKE SIX DAYS

NO, WAIT, THREE HUNDRED YEARS.

A RESEARCHER EXPLAINING WHEN WE'LL HAVE GENERAL ARTIFICIAL INTELLIGENCE

Escaping the local minimum

By Kenny Friedman

MAS.S63
Integrative Theories of
Mind
and Cognition

May 9, 2016

Agenda

Background / The Problem

Vision

3 Examples

Contributions

Agenda

Background / The Problem

Vision

3 Examples

Contributions

For each example:

Define the Problem

Classic Method

Potential Modern Method

Background

Electrical Engineering &
Computer Science

Undergraduate, Junior

Taking 2 classes in AI this semester:

MAS.S63 & 6.803

Background

Electrical Engineering &
Computer Science

Undergraduate, Junior

Taking 2 classes in AI this semester:

MAS.S63 & 6.803

Background

Electrical Engineering &
Computer Science

Undergraduate, Junior

Taking 2 classes in AI this semester:

MAS.S63 & 6.803

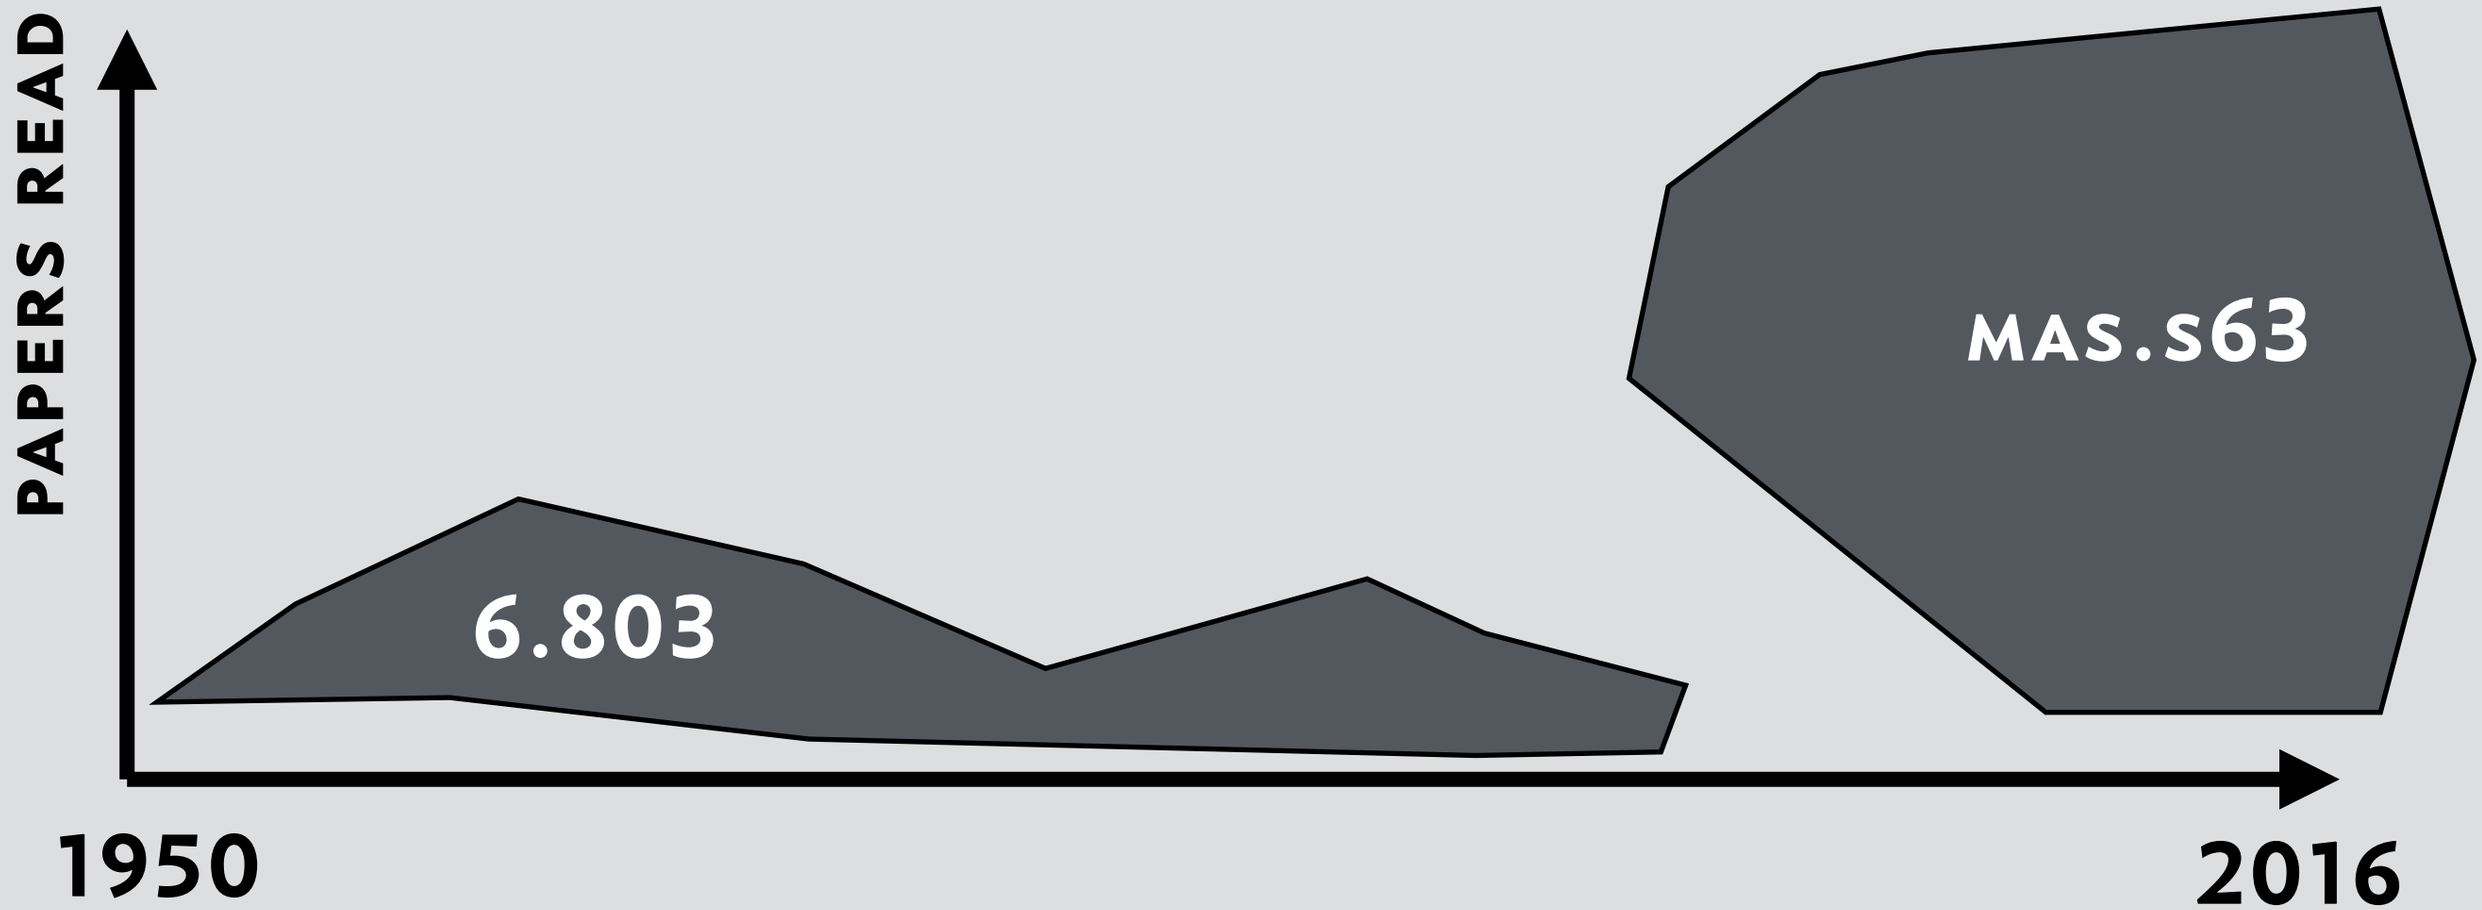

Background

Electrical Engineering &
Computer Science

Undergraduate, Junior

Taking 2 classes in AI this semester:

MAS.S63 & 6.803

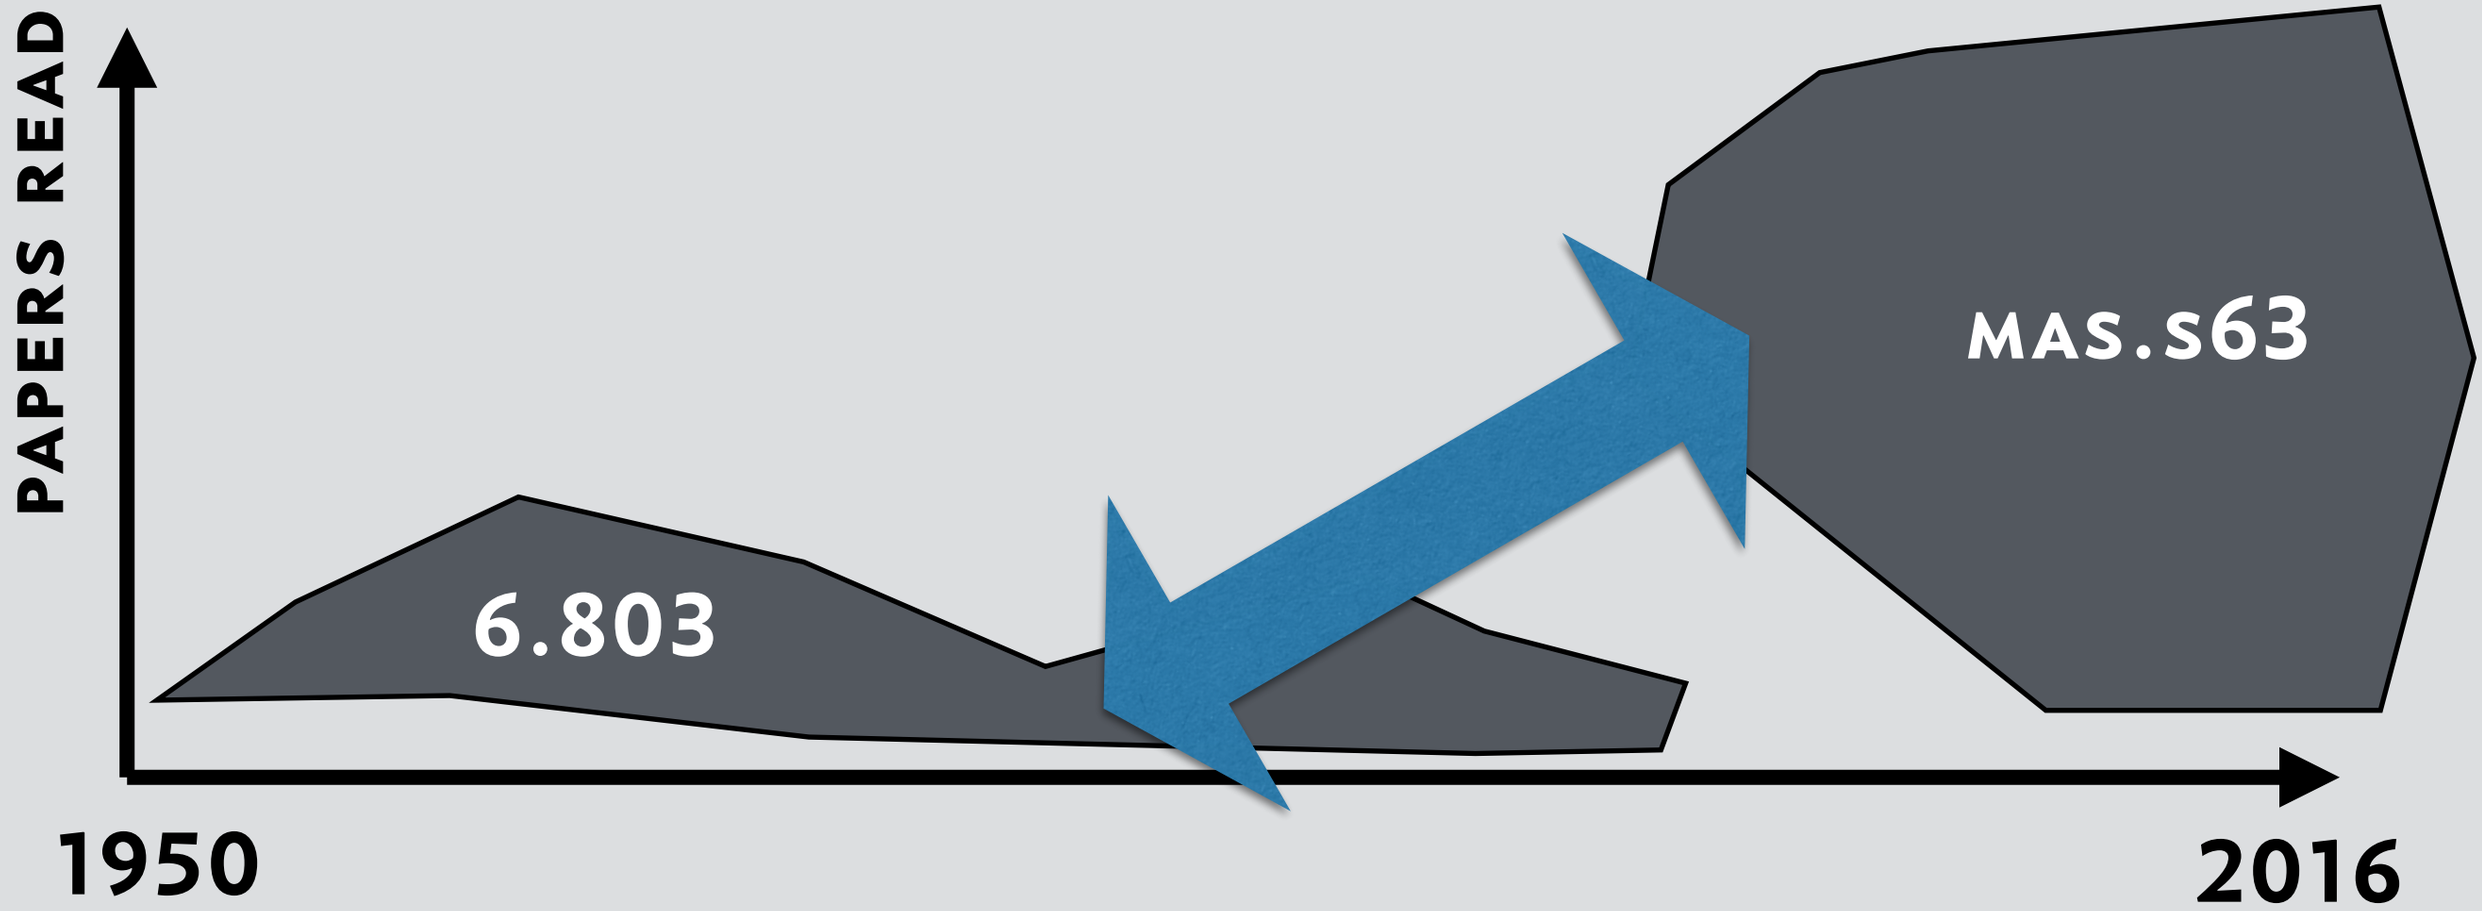

The Vision

AI has been performing gradient descent for decades

Then the field gets caught in a local minimum

Rarely looks around to see if it is stuck

The Vision

AI has been performing gradient descent for decades

Then the field gets caught in a local minimum

Rarely looks around to see if it is stuck

Take methods from recent success

Apply them to classic models & problems

NNs as a substrate for all previous work in the field

Gradient Descent

1960s

SAINT (Symbolic automatic integrator)

ELIZA (chatterbot)

SHRDLU (NPL, Terry Winograd)

Gradient Descent

1960s

SAINT (Symbolic automatic integrator)

ELIZA (chatterbot)

SHRDLU (NPL, Terry Winograd)

"**Within a generation** ... the problem of creating 'artificial intelligence' will substantially be solved."

Minsky, 1967

Gradient Descent

1960s

SAINT (Symbolic automatic integrator)

ELIZA (chatterbot)

SHRDLU (NPL, Terry Winograd)

"**Within a generation** ... the problem of creating 'artificial intelligence' will substantially be solved.”
Minsky, 1967

"In from **three to eight years** we will have a machine with the general intelligence of an average human being.”
Minsky, 1970

AI Winter Strikes

AI Winter Strikes

AI Winter Strikes

AI Winter Strikes

Gradient Descent

1980s

Expert Systems

Knowledge Based Systems (Cyc)

“Intelligent Agents”

AI Magazine Volume 10 Number 2 (1989) (© AAAI)

CHESS REPORT

Deep Thought Wins Fredkin Intermediate Prize

Hans Berliner

Since May 1988, Deep Thought (DT), the creation of a team of students at Carnegie Mellon University, has been attracting a lot of notice. In the Fredkin Masters Open, May 28–30, DT tied for second in a field of over 20 masters and ahead of three other computers, including Hitech and Chiptest (the winner of the 1987 North American Computer Championships). In August at the U.S. Open, DT scored 8.5, 3.5 to tie for eighteenth place with Arnold Denker among others. Its performance was marred by hardware and software bugs. However, DT astounded everyone by beating International Master (IM) Igor Ivanov, the perennial winner of the U.S. Grand Prix circuit prize, who is generally regarded to be as almost all potential contenders. The two first place finishers drew with each other. Hitech led the field at the halfway point but lost to DT in round 3 and threw away a winning position against Fidelity in round four (because of a programming bug). The level of play in this tournament was by far the best ever in a computer event, and the winners clearly deserved their top places.

Ten days later, DT achieved the greatest computer success to date. It tied for first with GM Tony Miles in the prestigious Software Toolworks Open in Los Angeles with a score of 6.5, 1.5. Several GMs played in this tournament, including former World Champion Mikhail Tal of the USSR. In the tournament, DT became the

AI Winter Strikes

AI Winter Strikes

AI Winter Strikes

AI Winter Strikes

Gradient Descent

2010s

Deep Learning

Statistical Models

Voice Recognition Paired with NLP

ROBERT MCMILLAN BUSINESS 01.16.15 6:30 AM

AI HAS ARRIVED, AND THAT REALLY WORRIES THE WORLD'S BRIGHTEST MINDS

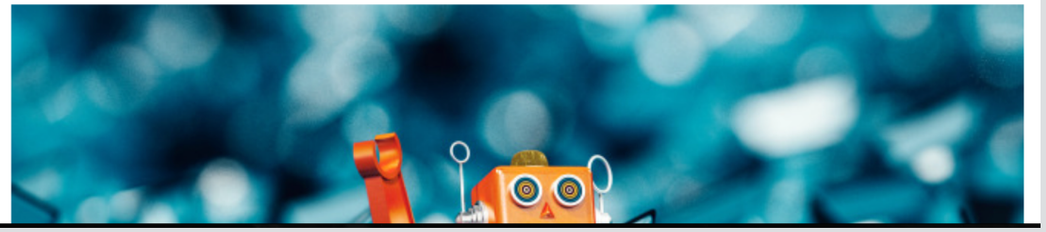

Future?

A Neural Net would predict a third winter any moment

Escape the current local minimum
(of superhuman perception)

Can neural networks become the substrate of all AI systems?
(*What If…* Poggio)

Three Examples

1. Generic Symbol Manipulation / Logic Systems

2. Minsky's Multiplicity

3. Language / Story Understanding

Generic Symbol Manipulation / Logic Systems

Symbolic Integral Calculus (1960)

Recent Examples of Progress:
NTM
Logic as vectors of real numbers

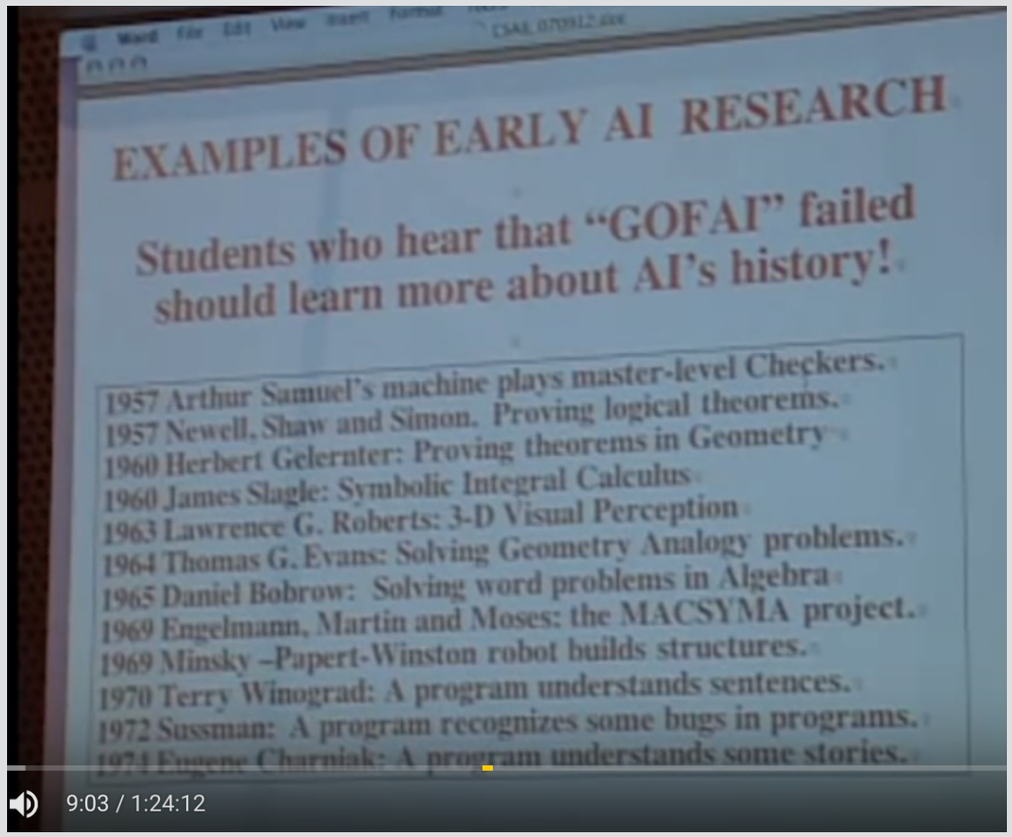

Minsky's Multiplicity

Architectures: Leabra, MicroPsi.

How much can emerge solely from NN?

1960: 5 Areas: Search, Pattern-Recognition, Learning, Planning, and Induction

2006

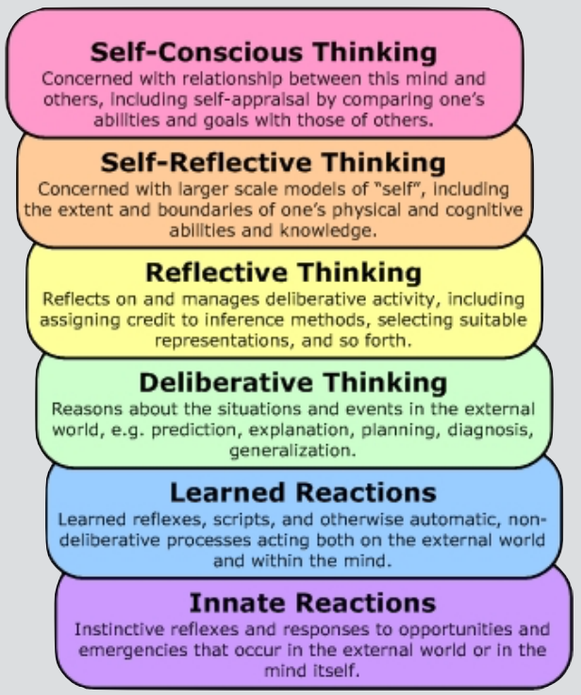

Language / Story Understanding

Chomsky, Winston.

MERGE operator.

Inner Language Hypothesis,
Strong Story Hypothesis

reflection, cultural bias, personality understanding, question answering, onset detection, trouble anticipation, similarity measurement, similarity based retrieval, question driven interpretation, analogical interpretation, reader aware story telling, persuasion, and summary.

Word2Vec, RNNs, show promise but distance

Contributions

Reviewed AI's tendency to fall into a local minimum

Provide three examples of prior systems & models that can not be done with deep learning

Suggest that a research focus should be placed on enabling deep learning to prove neural nets can be the right level of abstraction

Thanks

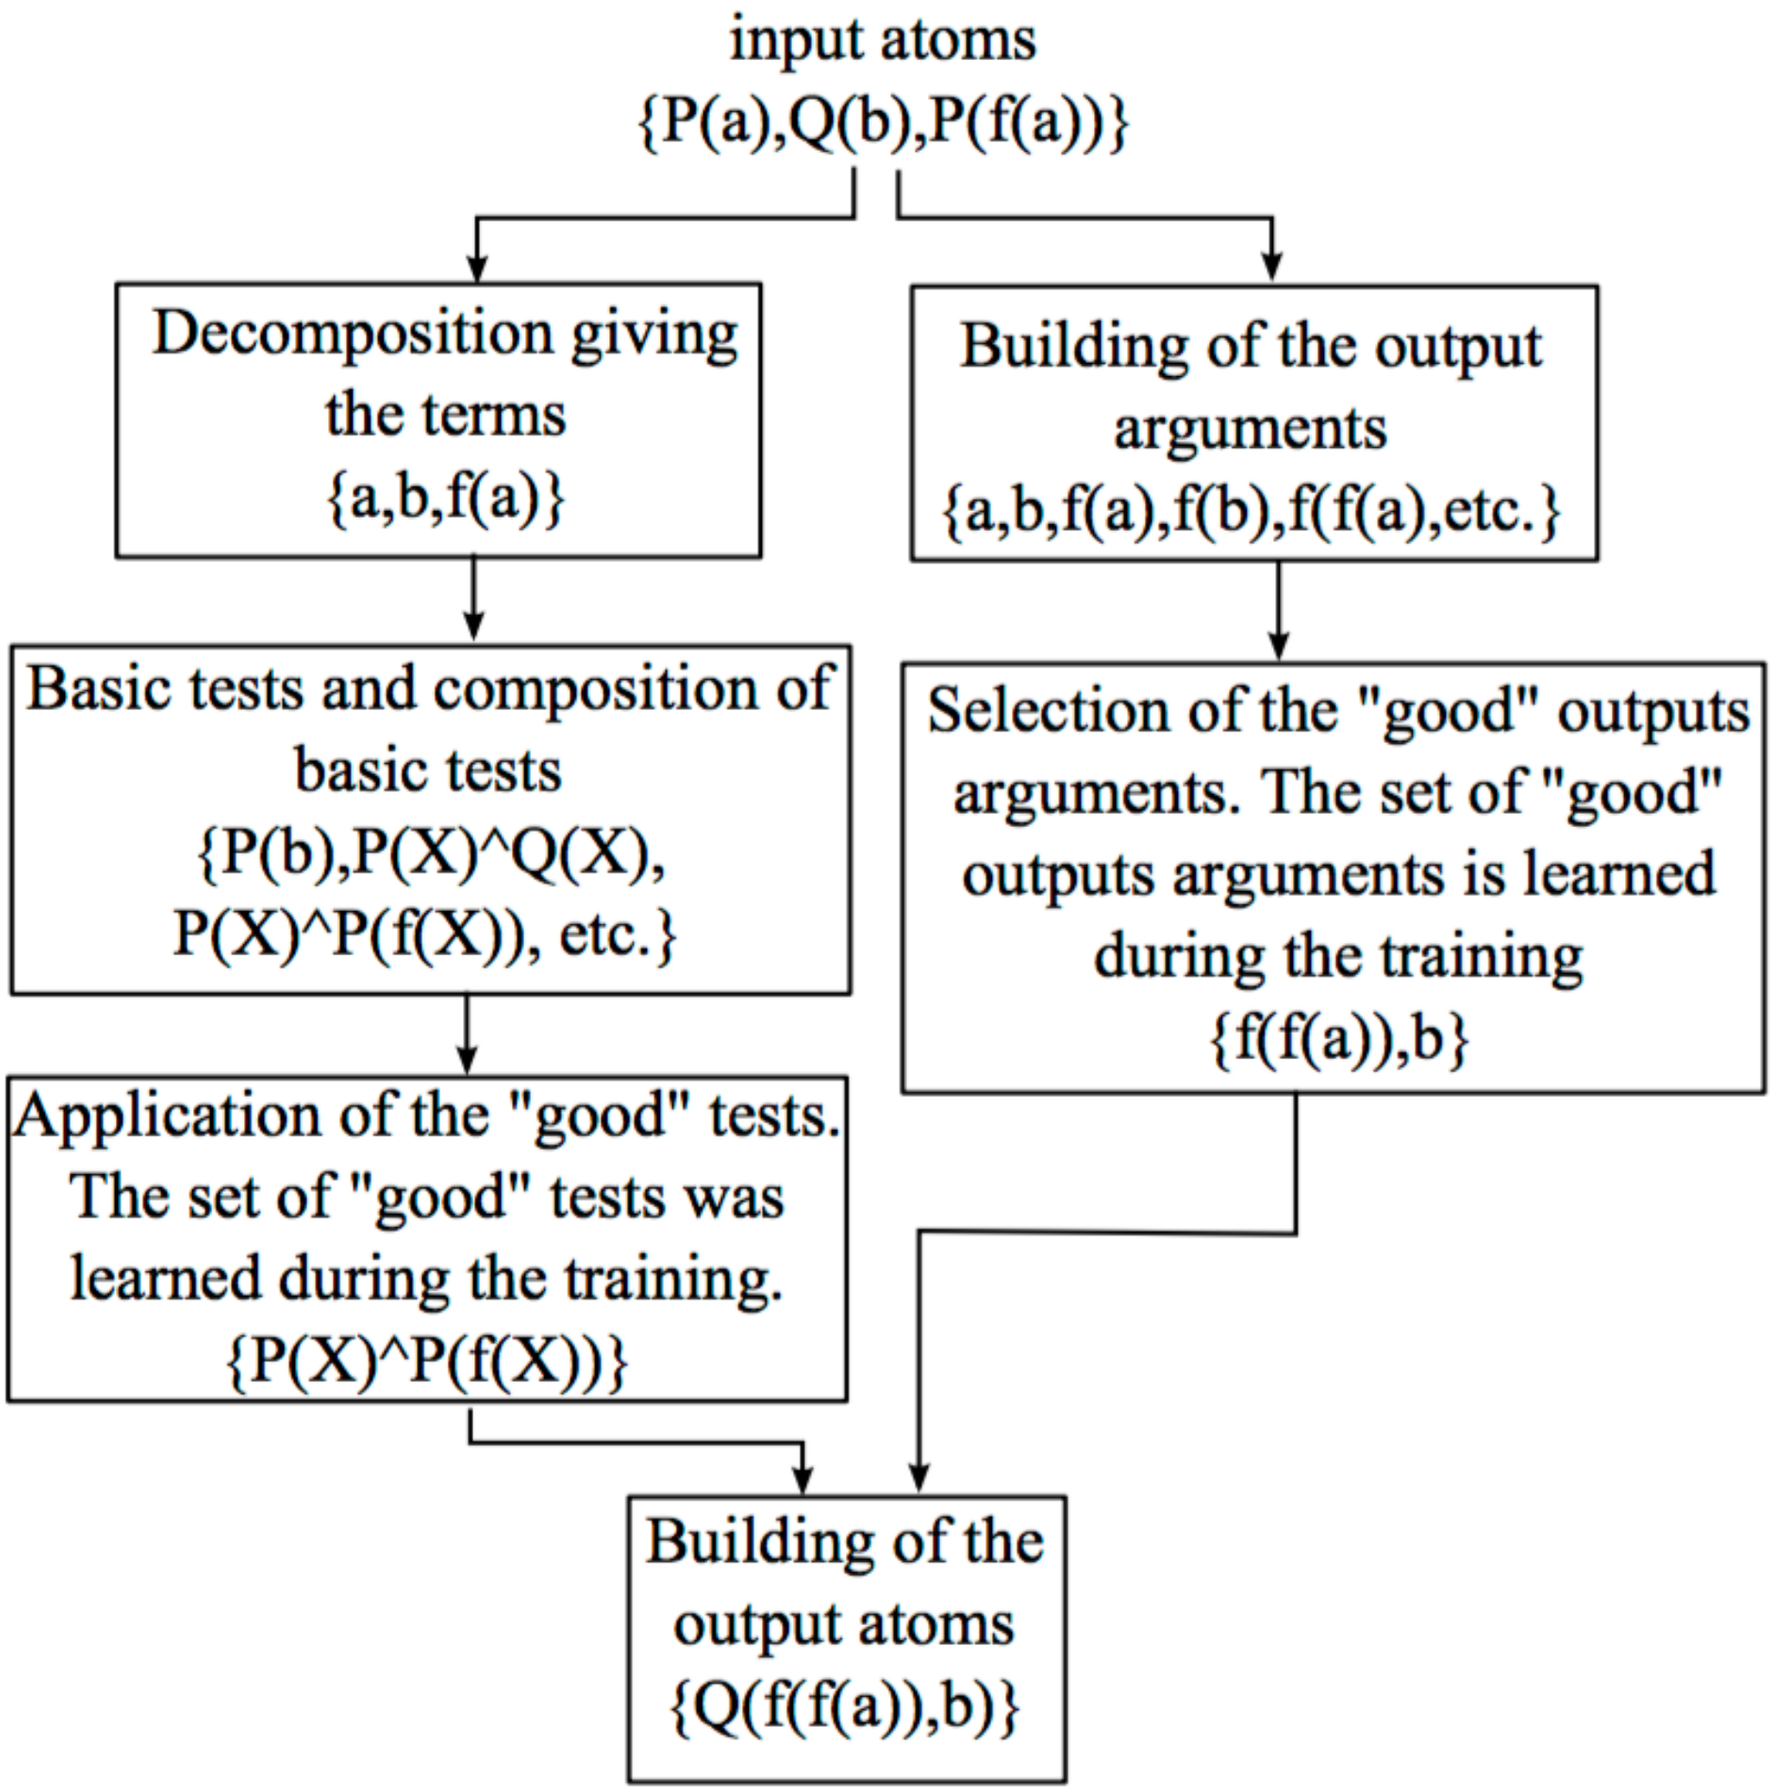
input atoms
{P(a),Q(b),P(f(a))}
Decomposition giving the terms {a,b,f(a)}
Building of the output arguments {a,b,f(a),f(b),f(f(a),etc.}
Basic tests and composition of basic tests {P(b),P(X)^Q(X), P(X)^P(f(X)), etc.}
Selection of the "good" outputs arguments. The set of "good" outputs arguments is learned during the training {f(f(a)),b}
Application of the "good" tests. The set of "good" tests was learned during the training. {P(X)^P(f(X))}
Building of the output atoms {Q(f(f(a)),b)}

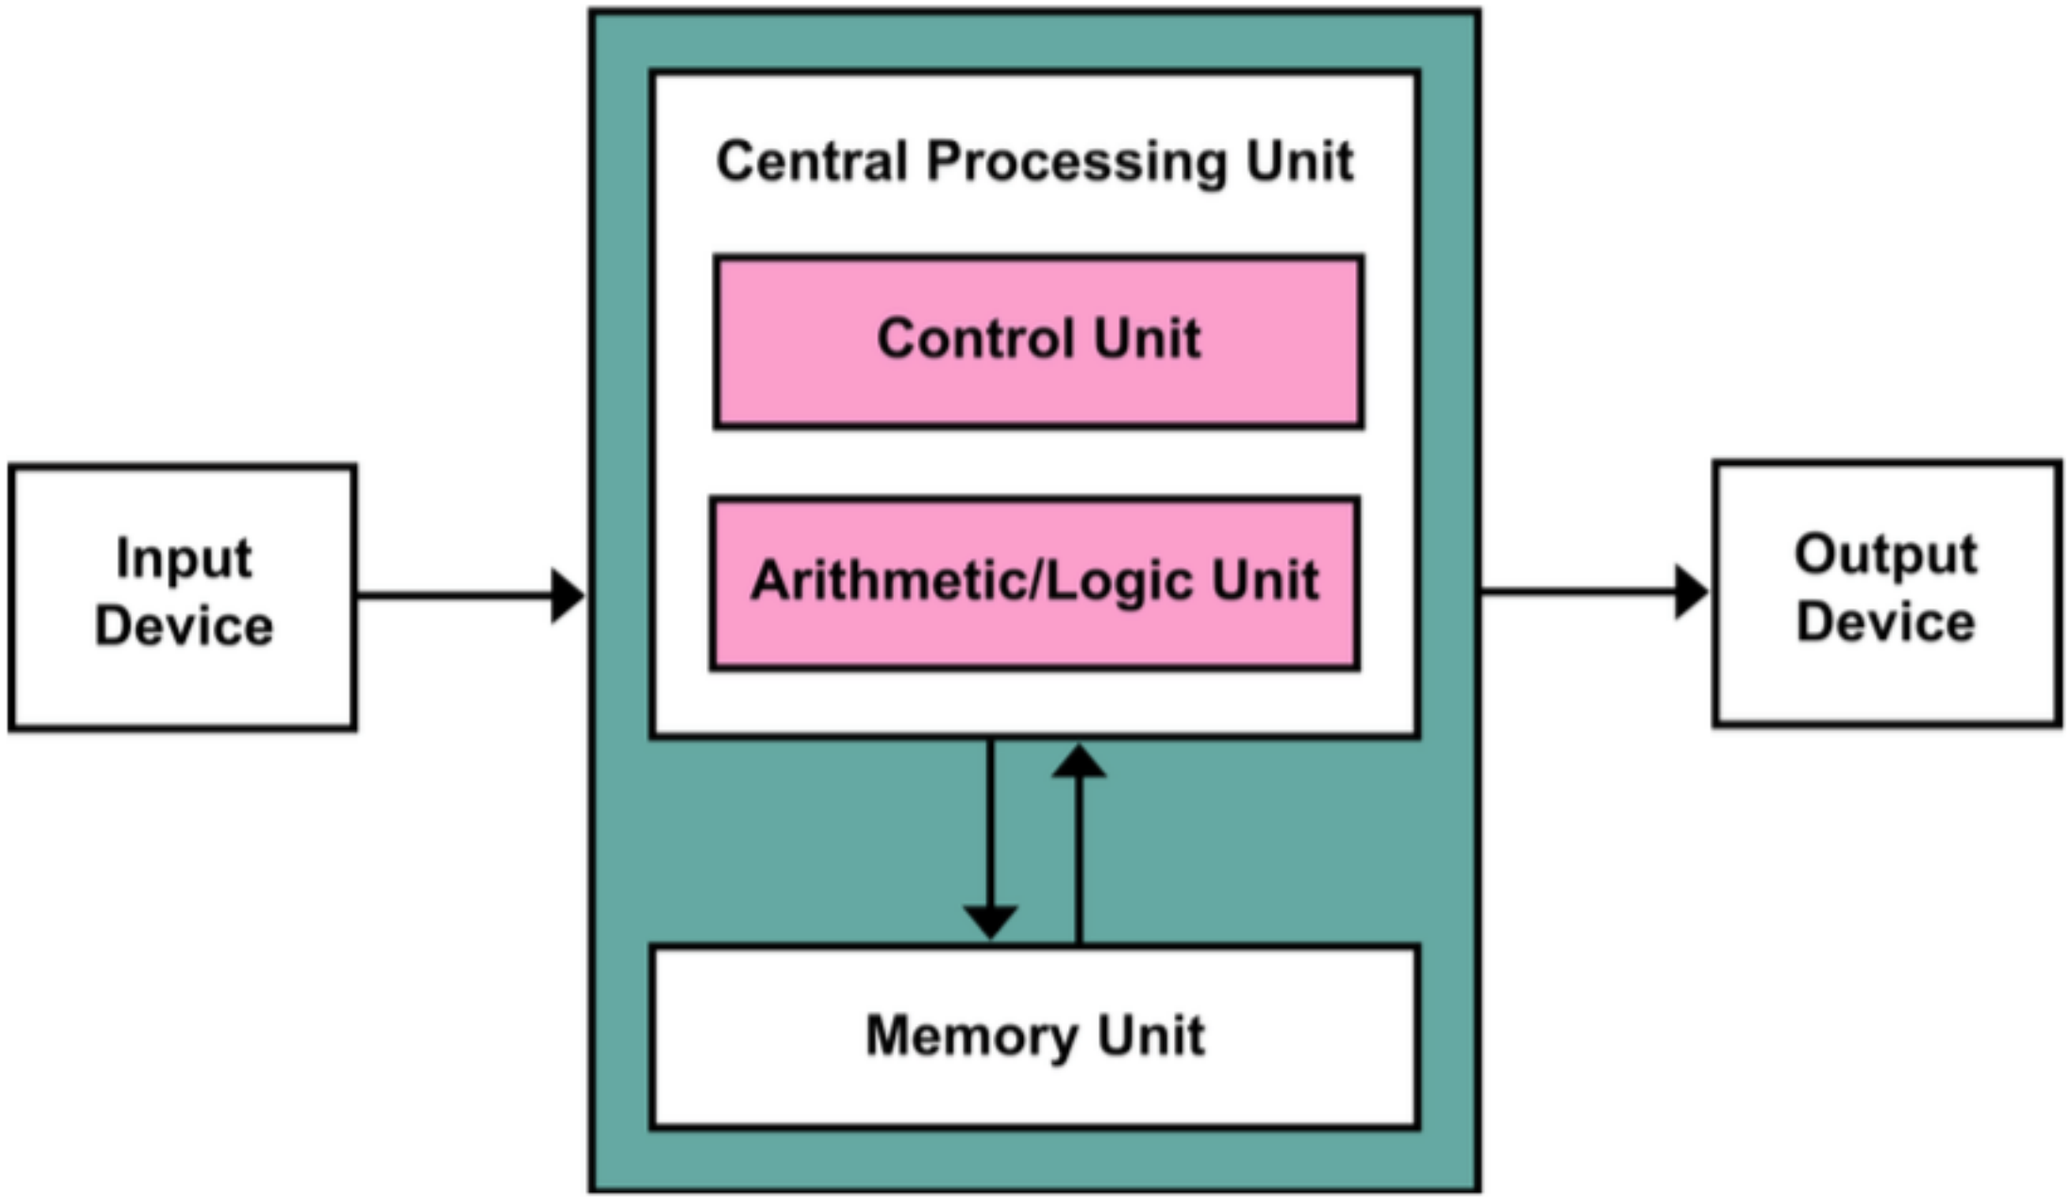
Central Processing Unit
Control Unit
Arithmetic/Logic Unit
Input Device
Output Device
Memory Unit