

Escaping the Local Minimum

Kenneth Friedman

May 9, 2016

Artificial Intelligence performs gradient descent. The AI field discovers a path of success, and then travels that path until progress stops (when a local minimum is reached). Then, the field resets and chooses a new path, thus repeating the process. If this trend continues, AI should soon reach a local minimum, causing the next AI winter. However, recent methods provide an opportunity to escape the local minimum. To continue recent success, it is necessary to compare the current progress to all prior progress in AI.

Contents

<i>Introduction</i>	2
<i>Gradient Descent</i>	2
1960s	3
<i>The First AI Winter</i>	5
1980s	5
<i>The Second AI Winter</i>	6
<i>Modern Day</i>	7
<i>The Next AI Winter</i>	8
<i>Deep Learning Checklist</i>	9
<i>Logic Systems</i>	10
<i>Minsky's Multiplicity</i>	11
<i>Story Understanding</i>	13
<i>Contributions</i>	15

Overview

I begin this paper by pointing out a concerning pattern in the field of AI and describing how it can be useful to model the field's behavior. The paper is then divided into two main sections.

In the first section of this paper, I argue that the field of artificial intelligence, itself, has been performing gradient descent. I catalog a repeating trend in the field: a string of successes, followed by a sudden crash, followed by a change in direction.

In the second section, I describe steps that should be taken to prevent the current trends from falling into a local minimum. I present a

number of examples from the past that deep learning techniques are currently unable to accomplish.

Finally, I summarize my findings and conclude by reiterating the use of the gradient descent model.

Introduction

IMAGINE TRAINING a pattern matching algorithm (say, a neural network) on the following series of data: 10 data points of the value "good", followed by 5 data points of "bad", another 10 data points of "good" and another 5 data points of "bad." Next, you test the system with 10 data points of "good." What will the algorithm predict the next data points to be?

Obviously, and sufficient pattern matching algorithm would predict an incoming set of "bad" points. Unfortunately, this pattern is the trend of the first 60 years of the AI field¹. When the field began, there was been a period of success followed by a period of collapse. Then another period of success emerged, but failure soon followed. Currently, we are in a period of great progress. So should we be predicting failure any day now? Is this pattern doomed to continue, or can we understand why it happens and take steps to prevent it?

I believe we can model the field of AI using a technique from within: gradient descent. By modeling the field's behavior with this well understood concept, we can pinpoint the problem with the technique and take steps to ameliorate its downside.

Gradient Descent

GRADIENT DESCENT IS AN OPTIMIZATION ALGORITHM that attempts to find the minimum of a function. In gradient descent, a series of steps are taken in the direction of the steepest negative slope. These steps repeat along that path until the direction of the slope changes (which proves a local minimum). The problem with gradient descent is clear: a local minimum is not necessarily the global minimum, since a local minimum is only based on a small period of time. In fact, a given local minimum might be worse than a previously discovered minimum (see Figure 1).

The field of Artificial Intelligence can be modeled with this technique. In this case, the function is loosely understood as the distance that "computer intelligence" is from "human intelligence", over time. That is, as the function approaches zero, computer intelligence approaches human level intelligence. Conceptually, a decreasing slope

¹ Of course, this trend is a simplification. A field of academia is not binary in its progress. However, it is a useful simplification to understand periods of significant improvement compared to periods of stagnation.

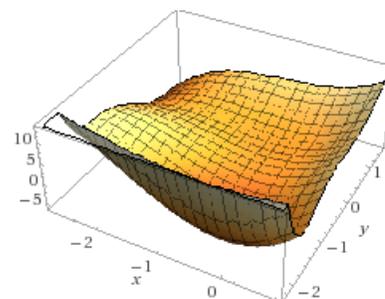


Figure 1: The fallacy of gradient descent can be easily seen here. Starting at (0,0), gradient descent will end at the local minimum near (0, -2), however a far smaller value could be achieved following a path to (-3,1). (Plot of $z = -y^3(y - 2)(y + 2) + x(x + 1)(x + 3) - x^2y$, created with WolframAlpha.)

signifies a progressing state-of-the-art AI, while an increasing slope signifies a lack of progress.

Here, the idea of a local minimum is that the field has chosen a particular path that appeared successful when compared to the recent past. However, the path ran out of areas to improve. Moreover, the path could be worse than a previous minimum — an improvement over the recent past, but a long term dead-end.

In the remainder of this section, I catalog the history of AI falling into local minimum traps.

1960s

EARLY PROGRESS IN ARTIFICIAL INTELLIGENCE was swift and far-reaching. In the years after the Dartmouth Conferences², there were a series of successes that lead to an extremely optimistic timeline.

These successes came largely from attacking a number of challenges using search algorithms. One of the earliest examples of a search based system was SAINT (Symbolic Automatic Integrator)³, a program that could solve integration problems at the level of a college freshman.

SAINT was considered the most advanced problem solving system of its time. In the 108-page paper's concluding chapter, Slagle explains insights that he discovered in both natural and artificial intelligence. While Slagle concedes that the program has a limited domain compared to the human mind, he explains that his integration system could solve problems that were advanced than any prior system.

Another early success was Newell & Simon's General Problem Solver (GPS)⁴. This program was able to solve elementary symbolic logic and elementary algebra. The most impressive demonstration of GPS was its ability to solve the Towers of Hanoi problem (See Figure 2). Instead of attempting problems that were advanced for humans (such as integral calculus), GPS sought to be as general as possible.

The report on GPS discusses the system's focus on two techniques: Means-End Analysis (MEA), and planning (discussed later). MEA was the general term Newell et al. used in order to describe the search heuristics in the program. Specifically, the MAE algorithms in GPS attempted to minimize the difference between the system's *current state and its goal state*.

Other successes include Thomas Evans' ANALOGY program and Daniel Bobrow's STUDENT program. ANALOGY⁵ could solve geometric, IQ-test style analogy problems such as "shape A is to shape B as shape C is to which of the following shapes," as seen in Figure 3.

² The Dartmouth Conferences took place at Dartmouth College in 1956. The complete name of the event was the *Dartmouth Summer Research Project on Artificial Intelligence*. It spanned 2 months and is widely regarded as the founding conference of the field. McCarthy et al. [2006]

³ James Robert Slagle. A heuristic program that solves symbolic integration problems in freshman calculus, symbolic automatic integrator (saint), 1961. URL <http://hdl.handle.net/1721.1/11997>

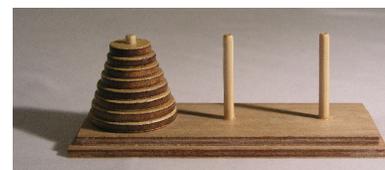


Figure 2: The goal of the Tower of Hanoi is to move the the disks from one ring to another ring, maintaining the smallest-to-largest order of the original. Tower of Hanoi, Wikipedia

⁴ Allen Newell, John C Shaw, and Herbert A Simon. Report on a general problem-solving program, 1959. URL http://129.69.211.95/pdf/rand/ipl/P-1584_Report_On_A_General_Problem-Solving_Program_Feb59.pdf

⁵ T.G Evans. A heuristic program to solve geometric analogy problems, 1962. URL <http://hdl.handle.net/1721.1/6101>

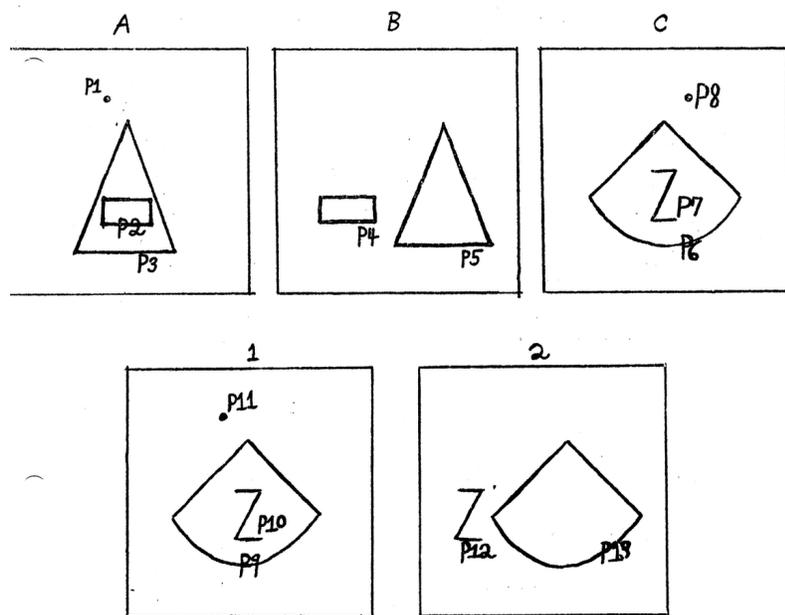


Figure 3: The pictures shown here relate to the type of geometric analogy question that ANALOGY was able to solve.

The question could be asked: Shape A is to shape B as shape C is to which of the following shapes?

ANALOGY was able to determine the answer: shape 2.

Meanwhile, STUDENT⁶ was able to solve word-algebra problems at a high school level. These early successes, many of which were at MIT, along with a large ARPA grant, led Marvin Minsky and John McCarthy to found the MIT AI Lab in the late 1960s.

All of these examples (and more) showed the field of Artificial Intelligence to be progressing well. This progress led to widespread optimism. In 1967, Minsky reportedly said "within a generation ... the problem of creating 'artificial intelligence' will substantially be solved."⁷

It might seem likely that Minsky quickly loosened his prediction. In fact, the opposite occurred. In 1970, three years and many small successes later, Minsky quickened and expanded his prediction, saying:

In from three to eight years we will have a machine with the general intelligence of an average human being. I mean a machine that will be able to read Shakespeare, grease a car, play office politics, tell a joke, have a fight. At that point the machine will begin to educate itself with fantastic speed. In a few months it will be at genius level and a few months after that its powers will be incalculable.⁸

The 1960s marked a period of highly successful gradient descent. Of course, gradient descent enviably leads to a local minimum. In this first golden age of AI, that local minimum hit the field in the mid 1970s and has since been known as the First AI Winter.

⁶ Daniel G Bobrow. *Natural Language Input for a Computer Problem Solving System*. PhD thesis, Massachusetts Institute of Technology, 1964. URL <http://hdl.handle.net/1721.1/5922>

⁷ Daniel Crevier. *AI: The Tumultuous History of the Search for Artificial Intelligence*. Basic Books, Inc., New York, NY, USA, 1993. ISBN 0-465-02997-3

⁸ Brad Darrach. Meet shaky, the first electronic person, 1970. URL <https://books.google.com/books?id=2FMEAAAAMBAJ&lpg=PA57&pg=PA58#v=onepage&q&f=true>

The First AI Winter

BOTH PROGRESS AND FUNDING SLOWED in the mid 1970s. Here, I provide a few examples of the decreased level of success and the reduction of funding that lead this period to be known as the AI Winter. I argue that while there were a number of contributing factors to this winter, lack of diversity in research paths was a leading reason. Simply put, the field hit a local minimum — no further progress could be made without changing course.

The Lighthill report⁹ was the first clear indicator of the AI Winter. The report, released in 1974, conveyed a largely negative outlook on the first quarter century of the field. Lighthill states "Most workers in AI research and in related fields confess to a pronounced feeling of disappointment in what has been achieved in the past 25 years." He goes on to chronicle examples and reasons for this disappointment. A fascinating recording of a debate between Lighthill and John McCarthy concisely summarizes many of these reasons.¹⁰

The overarching theme of Lighthill's pessimism comes from the failure of researchers to account for combinatorial explosion. That is, problems that can be solved within a narrow domain do not scale up to real-world problems due to the branching factor of the search trees (see Figure 4).

Another report, written by the Automatic Language Processing Advisory Committee¹¹, is also known to have caused a dramatic decrease in funding. Early in the report, the recommendations suggest that the National Science Foundation fund four to five centers at \$650,000 each. However readers of the report focused on the conclusion that human language translation was far more cost effective, thus halting AI funding in machine translation.

Fortunately, the winter did not last for an extended period. By the early 1980s, the field had regrouped, and sought a new path towards progress.

1980s

A FOCUS ON EXPERT SYSTEMS and knowledge bases sparked a resurgence of AI in the early 1980s. These systems were commercially viable — a first for AI system — which led to a large influx of AI funding from private corporations.

Programs known as expert systems used specific domain knowledge to answer questions and solve problems. Two of the most well known expert systems include XCON and Cyc.

XCON stands for eXpert CONfigurer, but the system was com-

⁹ Sir James Lighthill. Artificial intelligence: A general survey, 1974. URL http://www.chilton-computing.org.uk/inf/literature/reports/lighthill_report/p001.htm

¹⁰ Sir James Lighthill and John McCarthy. The lighthill debate, 1973. URL <https://www.youtube.com/watch?v=yReDbey7ZMU>

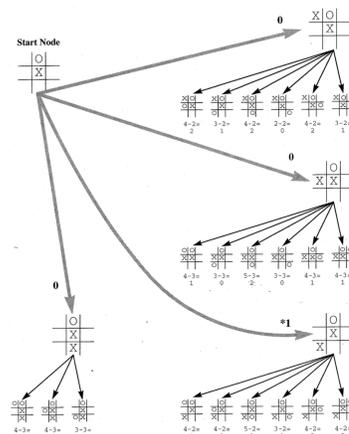


Figure 4: This figure shows the branching factor of tic-tac-toe. It is clear that even though tic-tac-toe is a very simple game, the exponential nature causes massive increases over time. This game has a maximum branching factor of 8. The game Go has a branching factor of 250. The branching factor for real world decisions can be thousands. Clearly, simple search quickly breaks down. Game Play, Charles Dyer

¹¹ Language and machines: Computers in translation and linguistics. URL <http://www.nap.edu/read/9547/chapter/1>

monly known as R1. R1 was built at Carnegie Mellon and used by Digital Equipment Corporation. This AI reportedly saved the company \$40 million annually.

While the system saved millions, R1 was far from optimal. The limitations became increasingly obvious. A 1984 report in AI magazine concluded with a poignant observation: "Even revolutionary ideas and technologies do not solve every problem."¹²

In the same year, Douglas Lenat created Cyc. The goal of Cyc was to have human-like reasoning by maintaining an enormous knowledge base. Lenat's plan involved decades of development, and the project is still ongoing.

Other successes from this time period include Deep Thought¹³ (a chess solving system that won the 1988 North American Computer Chess Championship) and soAR¹⁴ (a cognitive architecture with the goal of performing the "full range of cognitive tasks").

Once again, the successes of the time did not last. Researchers and institutions soon found the fragility of early knowledge based systems. The field ran out of progress on the expert systems path — it reached a local minimum.

The Second AI Winter

¹² Stephen Polit. R1 and beyond: Ai technology transfer at digital equipment corporation. *AI Magazine*, 5(4):76, 1984. URL <http://www.aaai.org/ojs/index.php/aimagazine/article/view/460/396>

¹³ Hans Berliner. Deep thought wins fredkin intermediate prize. *AI Magazine*, 10(2), 1989. URL <http://www.aaai.org/ojs/index.php/aimagazine/article/viewFile/753/671>

¹⁴ John E. Laird, Allen Newell, and Paul S. Rosenbloom. Soar: An architecture for general intelligence. *Artificial Intelligence*, 33, 1987. URL <http://www.sciencedirect.com/science/article/pii/0004370287900506>

FRIDAY, MARCH 4, 1988
Copyright © 1988 The New York Times

Business Day

The New York Times

L+ D1

BUSINESS Digest
FRIDAY, MARCH 4, 1988

Companies
Macy is expected to contest the legality of Campeau's latest tender offer for Federated Department Stores. The Macy strategy apparently will be to contend that the offer represents a new bid, not an intended proposal, and that it violates Federal laws and Delaware anti-takeover statutes. [Page D1.]

A \$1.1 billion bid for Koppers was made by Brazier, a British construction company. Koppers is a Pittsburgh-based road builder and manufacturer of construction materials. The offer was for 145 1/4 shares. Koppers's shares rose to \$51.75, up \$6.625. [D1.]

Chrysler will use self-assembly line, its parts-making unit, but will sell or close four plants, affecting about 2,000 workers. A sale had been previously approved by Chrysler's unions. [D1.]

A new cable entertainment service is planned by Ted Turner, who has gained tentative approval of the plan from a number of cable operators. The eventual goal would be to provide original productions for cable. [D1.]

UKAT entered 50 Boeing 737 jets and took options on 30 more in a deal that could cost \$2.4 billion. [D4.]

UKAT is expected to announce its entry into wholesale clubs through a joint venture with SSI, a Dutch company. [D4.]

Rams sweetened its tender offer for Farmers Group to 95 1/4 shares, up \$1.1 billion. [D1.]

Leading airlines raised fares for the second time in two weeks, reflecting unopposed strength in winter traffic. [D1.]

The Economy
Retailers reported unexpectedly slow sales for February, with a number of large chains posting declines for stores open more than a year. Analysts, while noting that the comparison was with a relatively strong February last year, still expressed concern. [D1.]

Vehicle sales rose 21 percent in late February. For the month, sales of domestic vehicles were 13.6 percent higher, while sales of imported vehicles fell 4.4 percent. Detroit's heavy discounting received most of the credit for the strong results. [D1.]

A Senate panel's compromise on a landmark banking bill is expected to help secure its passage in the fall, despite this strong, but strong opposition is likely in the House. [D1.]

The Fed's chairman addressed new changes in the tax system. Alan Greenspan told Congress that he favors the imposition of consumption taxes and a reduction in the capital gains tax. [D1.]

Northern productivity improved last year by 8.1 percent, just half as much as in 1986, according to revised figures. [D1.]

Markets
The stock market had a mixed day, dominated by takeover activity. The Dow Jones industrial average fell 7.40 points, to 2,895.49, but other indexes had gains. [D1.]

Bond prices were little changed in light activity as many traders waited for fresh data on the economy. [D1.]

The dollar held steady in trading described as the lightest in more than a year. Gold fell to \$428.39 an ounce, down \$2.85. In New York, 1988 1/2 Treasury futures were down 1/4 point.

Setbacks for Artificial Intelligence

Companies Are Hurt By Poor Decisions

Trouble for Many of the Artificial Intelligence Companies

By ANDREW POLLACK
NEW YORK, March 3 — A major retrenchment is occurring in the artificial intelligence industry, that thought they would prosper by providing the technology to make computers "think."

Some of the setbacks stem from the failure of artificial intelligence to quickly live up to its promise of making machines that can understand English, recognize objects or reason like a human expert — to be used for such purposes as diagnosing machinery breakdowns or deciding whether to authorize a loan. Despite this, the technology is making slow but steady progress, and now is being used in a number of ways.

Weak Business Moves
For more of the artificial intelligence industry's problems, credit must be given to poor business decisions by companies that were heavily weighed in with technology rather than business minds. Their main mistake was trying to use special computers for artificial intelligence; the machines were too expensive and did not mesh well with those used by potential customers.

"People believed their own hype," said S. Jerrard Kaplan, co-founder of one leading artificial intelligence company, Teknowledge Inc., and now president of the Co Corporation in San Francisco. "Everyone was planning to do things that were unworkable."

Although computers often appear to be intelligent in their everyday applications, they generally perform repetitive tasks following fixed rules and do not learn or make cognitive decisions, as humans do.

Artificial intelligence aims to make computers do tasks that are said to require intelligence when people do them. One goal has been to make machines that can understand and use other human languages, so that people do not have to use special com-

As the industry realigns, the companies that relied on special purpose machines are languishing.

Company/Headquarters	Description
EXPERT SYSTEM DEVELOPMENT TOOLS	
Teknowledge Folsom, Calif.	Four quarters of losses. 60 workers of 220 laid off. 1987 sales: \$20 million.
Intelliproc Mountain View, Calif.	Six quarters of losses. 50 workers of 200 laid off. 1987 sales: \$29 million.
Carnegie Group* Pittsburgh	Losses. 20-40 workers of 200 laid off. 1987 sales: \$12 million.
Intelligence* Los Angeles	Losses. 20 workers of 130 laid off. 1987 sales: \$12 million.
MACHINE MANIPULATORS	
Synetics Cambridge, Mass.	Continuing losses. Third round of layoffs last year. 1987 sales: \$104 million.
Lisp Machine Andover, Mass.	Filed for bankruptcy last year. 1986 sales: \$12 million.
Xerox Quincy, Conn.	Sluggish sales; recent reorganization of its artificial intelligence business.
Texas Instruments Dallas	Big push in artificial intelligence. Announced a unit consisting of 100 to 150 employees.
EXPERT SYSTEM APPLICATIONS	
Synology Sunnyvale, Calif.	No layoffs. 1987 sales: \$9 million.
Applied Expert Cambridge, Mass.	Layoffs. 1987 sales: \$4 million.
Patterson Cambridge, Mass.	Layoffs. Duast chairman and founder. 1987 sales: \$9 million.
*Private company, unlisted.	

Source: AI Trends

Macy Plans Legal Step In Bidding

Status of Campeau's Last Federated Offer May Face Challenge

By SHADORE BARMAN
In an unrelenting struggle to win Federated Department Stores Inc., R. B. Macy & Company is expected to contest the legality of the Campeau Corporation's latest tender offer for the giant retailer.

Macy is expected to contend that Campeau's latest offer is not an amended proposal but a new one that violates Federal securities laws and Delaware's new anti-takeover statute. With that challenge, Macy would be seeking to delay the current deadline of the Campeau offer in order to improve the chances for its own bid.

According to Securities and Exchange Commission rules, reducing the number of shares or changing the price in a tender offer constitutes an amended offer, allowing shares to be bought in 10 business days. But in the case of a new offer, all business days are required before shares can be purchased.

Eye on Expiration Date
If the proposed bid by Macy's lenders succeeds, the March 15 expiration date of the Campeau bid would be null and void, giving Macy's tender offer an opportunity to raise the stakes.

But if the bid is rejected, it would also mean that Federated shareholders would have to choose between either of the competing tender offers — Macy's for cash and stock or Campeau's for all cash.

Each offer has an estimated value of more than \$1 billion, but even strategists who are said to have accumulated as many as 10 percent of Federated's stock, seem confused by the vague differences between the offers.

Reflecting the confusion, Federated's stock traded independently for a number of days, dropping 2 1/2 cents, to \$66.37 1/2, before trading resumed. Analysts said that would challenge features of the Macy bid.

Figure 5: This article in the March 4, 1988 edition of the New York Times explains the number of commercial failures of expert systems.

The table shows substantial losses and layoffs at Teknowledge, Intelliproc, Carnegie Group, and Intercem, which are all expert-system based companies.

This PDF version is behind a paywall here: <http://timesmachine.nytimes.com/timesmachine/1988/03/04/issue.html>.

It is available in text form publicly, here: <http://www.nytimes.com/1988/03/04/business/setbacks-for-artificial-intelligence.html?pagewanted=all>

THE SECOND AI WINTER FELT VERY SIMILAR to the first winter. Over optimistic views, a reduction in investment, and a lack of new

Continued on Page D5

directions all lead to a period (from approximately 1987 to 1993) known as the second AI winter.

Overoptimism was the first major reason for this winter. One clear example can be seen in the Fifth Generation Computer Systems (FGCS). FGCS¹⁵ was a 1982 initiative developed by Japan's Ministry of International Trade and Industry. The goals of this initiative were far reaching. It sought to act as a supercomputer that could translate languages, interpret pictures, and reason like human beings. After 10 years and \$400 million later, the project failed to meet its goals and was thus canceled.

Both public and private institutions reduced investment after a series of failures. As seen in Figure 5, three different groups of AI focused companies all had dramatic setbacks. First, AI development tool companies suffered many quarters of losses which resulted in laying off 10% to 20% of the workforce. Second, machine manufacturers, such as Symbolics¹⁶ (founded by Danny Hillis) and Xerox, pivoted their goals and reorganized their companies to reduce the focus on Lisp machines. Finally, expert system applications suffered in sales and announced layoffs.

This lack of progress, lack of funding, and commercial failure can easily be explained by comparing the overly optimistic goals of these systems to the implementation details. After domain specific tasks had been mastered, many companies attempted to use expert systems and large knowledge bases to create flexible, general AI systems. However, the fragility and domain specific nature of these methods was not obvious. When more ambitious goals were attempted, they failed again and again.

Eventually, the field exited the winter and reset. The period of time after this winter ended can be categorized as the modern day boom of artificial intelligence.

Modern Day

THERE IS CONSTANT NEWS of AI successes in modern times. In 1997, Deep Blue¹⁷ (a chess AI system) defeated the reigning world chess champion. In 2005, a Stanford robot drove autonomously for 131 miles, winning the DARPA Grand Challenge¹⁸. In 2011, IBM's Watson¹⁹ system competed on the TV game show Jeopardy and defeated both human opponents.

Consumer electronic companies have recently popularized voice based AI assistants. Apple released Siri on iOS in 2011. Google released Google Now on Android in 2012. And Amazon released Echo, voice assistant hardware product which operates hands free, in 2015.

¹⁵ Andrew Pollack. 'fifth generation' became japan's lost generation. *New York Times*, 1992. URL <http://www.nytimes.com/1992/06/05/business/fifth-generation-became-japan-s-lost-generation.html>

¹⁶ John Abell. March 15, 1985: Dot-com revolution starts with a whimper, 1985. URL <http://www.wired.com/2010/03/0315-symbolics-first-dotcom/>

¹⁷ Murray Campbell, A Joseph Hoane, and Feng-hsiung Hsu. Deep blue. *Artificial intelligence*, 134(1):57-83, 2002. URL <http://www.sciencedirect.com/science/article/pii/S0004370201001291>

¹⁸ DARPA Grand Challenge 2005 was a competition hosted by DARPA in order to promote the development of fully autonomous ground vehicles.

¹⁹ AI Magazine. The ai behind watson—the technical article. *AI Magazine*, 2010. URL <http://www.aaai.org/Magazine/Watson/watson.php>

The most recently publicized breakthrough was Google DeepMind's AlphaGo²⁰. AlphaGo is a system that plays the game Go, and it defeated one of the highest ranking human Go players in 2016. AlphaGo mastered Go between 5 to 10 years before many people predicted an AI to be capable of such a high level of play.

This string of success is incredible. However, if the historical pattern continues, an AI winter should be imminent.

The Next AI Winter

IS THE NEXT AI WINTER rapidly approaching? It seems improbable given all of the recent successes. However, it also seemed improbable in 1968 and in 1985, and the winters came nevertheless.

There are many optimistic predictions about the future of the field, as there have always been. Predictions for 2025 include:²¹

- widespread use of self-driving cars (Vint Cerf)
- food will be raised, made, and delivered robotically (Stowe Boyd)
- "AI and robotics will be close to 100% in many areas" (Marc Prensky)

It's possible to imagine the headlines of the next AI winter. Pretend the year is 2019, and the following stories appear on a news site:

- After self-driving cars fail to work in many states, tech companies cancel self driving car plans
- Humans still prove to be much more efficient than robots
- Virtual assistants fall short on understanding complex demands

This would not be a shock although it would be a sign that the field was once again in an AI winter.

However, the next winter is not inevitable. The field can prevent itself from falling into the local minimum trap by ensuring it can do everything that has been possible in the past two AI booms.

In this section, I have shown that a gradient descent & local minimum method is a realistic model to explain the previous 50 years of artificial intelligence. I have cataloged successes of the 1960s and 1980s, as well as the failures of the two AI winters. I have provided examples of recent success, and suggested that a future winter is possible but not inevitable.

In the next section, I suggest steps that should be taken to prevent the field from falling into another local minimum.

²⁰ David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484-489, 2016. URL <http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html>

²¹ Aaron Smith and Janna Anderson. Ai, robotics, and the future of jobs, 2014. URL <http://www.pewinternet.org/files/2014/08/Future-of-AI-Robotics-and-Jobs.pdf>

Deep Learning Checklist

THIS MODERN DAY SUCCESS is largely the result of advances in deep learning. Hinton²² astonished the world with deep neural networks in 2012. Since then, the rise of neural networks (NNs) have taken over as the dominant technique for state-of-the-art AI systems. Unfortunately, there are still many domains for which neural networks have great difficulty.

Deep learning's incredible advantage is its superiority in perception: NNs are the ultimate pattern matching device. In lots of cases, NNs can already find patterns in data at a superhuman level.

In the previous section I explained that the AI field can become trapped in a local minimum if it focuses too heavily on short term progress by only exploring the strength of a given technique.

So, if the AI field is doomed to repeat its mistakes, then the research focus is clear. Deep learning researchers should keep working to make object detection in images nearly perfect. Hinton should work on going from a 15% error rate on ImageNet data to 0.5% error rate. NNs should be developed to find patterns in the stock market to become rich and optimize engagement on social networking websites.

Of course, the field would then be trapped in a local minimum. Fortunately, the statistical, neural network approach²³ can avoid this pitfall. Because NNs are biologically inspired systems, they are architecturally more similar to human thinking than any previous technique. Combining this human-like architecture with the understanding that NNs are Turing Complete, it is possible that NNs can be used as a substrate for all AI systems. That is, all future systems can be built on top of a statistical network base layer.

To prove that the state-of-the-art deep learning techniques are in fact superior to all prior systems, and to show that NNs can be used as a substrate for future programs, it is necessary to apply NNs to previously developed or proposed tasks.

In the remainder of this section, I describe three different categories of AI research that currently do not implement a statistical or NN approach. For each example, I first explain the work that has previously been done on the topic. Then, I describe the value that is to be gained from a successful implementation. Finally, I inventory research that appears to have taken steps — or at least shows promise in taking steps — towards the given topic.

The three categories I describe are: logic systems, Minsky's multiplicity, and story understanding. Each is presented in depth below.

²² Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. URL <https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>

²³ In this section, my use of the term "statistical" is in a neural network sense. Also, my use of the term "neural networks" encompass all subcategories and tangential concepts such as convolutional neural networks, recurrent neural networks, deep belief networks, and the like.

Logic Systems

THE FIRST CATEGORY TO INVESTIGATE is generic symbol manipulation. Symbolic systems are based on high level representations of information with explicit rules. An in-depth look at these early programs, such as SAINT, is discussed on page 3. Figure 6 shows a list gathered by Marvin Minsky of these early successes.

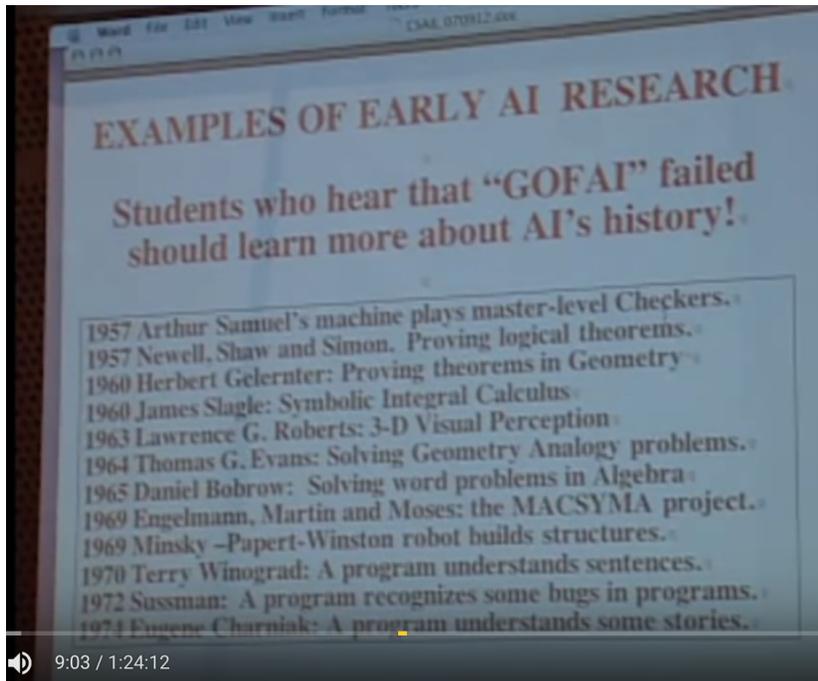


Figure 6: A screenshot of a slide from Marvin Minsky's talk *Artificial Intelligence and the Future of the Human Mind*. The slide shows a list of AI successes from the 1960s and 1970s. Nearly all of the results on the list are rule-based logic systems.

A recording of this 2007 talk can be seen here: <https://www.youtube.com/watch?v=g1DcIJgInF8>.

Fifty years after Slagle's SAINT program, the state-of-the-art calculus systems essentially use the same methodology. The program understands explicit rules²⁴ that were explicitly programmed by the developer. The system then searches for rules it can apply to the current state in an effort to simplify the integral. Because the knowledge representation is at a symbolic level, and the number of rules is limited: there is no need for an integration system to learn rules overtime.

However, it would be useful to have a neural network that is able to perform symbolic algebra. There are two clear reasons for this desire. First, this hypothetical system would demonstrate that neural networks can be used as a substrate for previously-achieved AI systems. Second, a neural network that could perform symbolic algebra would, by definition, be able to manipulate symbols. This would show that high level knowledge representations can be grounded²⁵ in

²⁴ For the example, the power rule, which states $\int x^n dx = \frac{x^{n+1}}{n+1} + C$ when $n \neq -1$.

²⁵ Stevan Harnad. The symbol grounding problem, 1990. URL <http://users.ecs.soton.ac.uk/harnad/Papers/Harnad/harnad90.sgproblem.html>

statistical models.

Two research findings show promise toward this goal. First, a 2010 paper titled "Neural-Symbolic Learning and Reasoning"²⁶ outlines the goals and directions of this space. This paper describes deep learning systems that have been developed to learn first-order logic rules within a NN. This is done by encoding symbols and logic terms as vectors of real numbers.

Second, Alex Graves' "Neural Turing Machine"²⁷ is a neural network system that can use external memory to hold additional information during computation. Early results have shown that the system can learn simple commands such as copying, sorting, and associative recall. These commands can be used as building blocks for future symbol manipulation, which could enable NN systems to be extremely generalizable symbolic manipulators.

These systems show progress, but they are still quite far from the symbolic manipulators of the 1960s. A heavy research focus should be placed on demonstrating that deep learning can succeed in this area.

Minsky's Multiplicity

MARVIN MINSKY EXPLORED a broad range of topics over the course of his career. In two distinct occasions, Minsky provided interesting categorization systems. The first is a list of five areas of heuristic problem solving. The second is a list of the six layers of thinking. I will refer to these two classifications as Minsky's Multiplicity²⁸, though I will address each individually.

The first category to discuss is the list of areas of heuristic problem solving. Minsky first explains these five areas in his groundbreaking 1960 paper, "Steps Towards Artificial Intelligence"²⁹. In this paper, he explains that search, pattern recognition, learning, planning, and induction are five distinct but crucial areas of problem solving.

The strongest advantage of neural networks is in the area of pattern recognition. In the pattern recognition section of Minsky's 1960 paper, he suggests that bayesian nets could be a useful tool in mastering pattern recognition.

Today, it is clear that neural networks are dominant in pattern recognition and learning. NNs are able to find imperceivable patterns in data. Also, NNs do not require hard coded rules, the programs are able to learn based on a vast amount of training data, through back-propogation. While it is obvious that neural networks will continue to become better at pattern recognition and learning, it is worth considering the other 3 areas. In terms of search, the state of NN search-

²⁶ Artur d'Avila Garcez, Tarek R Besold, Luc de Raedt, Peter Földiak, Pascal Hitzler, Thomas Icard, Kai-Uwe Kühnberger, Luis C Lamb, Risto Miikkulainen, and Daniel L Silver. Neural-symbolic learning and reasoning: contributions and challenges. 2015. URL <http://dase-lab.cs.wright.edu/pub/nesy2015.pdf>

²⁷ Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014. URL <http://arxiv.org/abs/1410.5401>

²⁸ Minsky's Multiplicity is not a widely used term for these two topics. In fact, I coined this term specifically for this paper, as no paper that I am aware of has discussed both of these classifications together. The term is a play on words, as Minsky provides multiple multiplicities.

²⁹ Marvin Minsky. Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1):8-30, 1961. URL <http://web.media.mit.edu/~minsky/papers/steps.html>

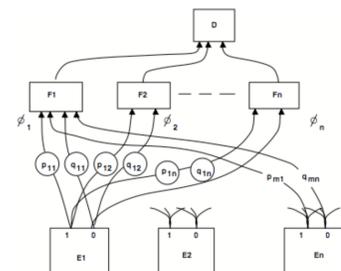


Figure 7: This diagram is an early bayesian network architecture from Minsky's seminal 1960 paper.

ing is dependent on the definition of search. Minsky’s definition was not rigid: "[if] we have a means for checking a proposed solution, then we can solve the problem by testing all possible answers. But this always takes much too long to be of practical interest." There are a number of NN architectures that show promise in the area of search. James Atwood and Don Towsley recently presented "Search-Convolutional Neural Networks"³⁰ (SCNN). This architecture defines features by traversing a local graph by breadth-first search and random walks. Another paper focuses on the more traditional use of the term search: "Using Neural Networks for Evaluation in Heuristic Search Algorithm"³¹ is a NN system that develops its own search heuristics instead of relying on hard-coded search algorithms such as A*.

The areas of planning and induction are often considered to be at a higher level of cognition. Planning breaks down complex problems into subproblems, while induction is the progress of using previous knowledge to come to a conclusion about a related but different scenario. Both areas are often discussed at the symbolic level. Therefore, to progress in these sectors of intelligence, NN must achieve symbolic manipulation abilities discussed in the previous section.

³⁰ James Atwood and Don Towsley. Search-convolutional neural networks. *CoRR*, abs/1511.02136, 2015. URL <http://arxiv.org/abs/1511.02136>

³¹ Hung-Che Chen and Jyh-Da Wei. Using neural networks for evaluation in heuristic search algorithm. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011. URL <http://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/view/3476>



Figure 8: This diagram displays the 6 layers of Minsky’s model. Innate and instinctive urges and drives are realized in the bottom layers. Whereas values, sensors, ideals, and taboos are understood at the upper layers.

The second categorization that Minsky developed included six

layers of thinking, seen in Figure 8. These six layers were initially defined and described in *Emotion Machines*³².

Cognitive architectures, such as The MicroPsi Project³³ and Leabra³⁴ attempt to model all of cognition. These systems strive to reach all six layers. However, neural networks might be able to develop these levels without the need for an entire static architecture.

Both Minsky's layers and neural networks are hierarchical. Therefore it would make logical sense that basic NNs can most easily implement the lowest levels of Minsky's layers. Conversely, NNs would most struggle with the highest level layers, as the system would need a successful implementation of all layers below the goal layer. This logic is accurate.

Modern deep learning systems all "think" at the reactive (base) layer. The neural networks have automatic responses based on perceived information about the world. It is possible to make the argument that more advanced deep learning programs have deliberative thinking. However it would be hard to agree that any NN reaches reflective thinking. No modern system is able to "reflect on and manage deliberative activity".³⁵

This demonstrates that there is universal agreement that neural networks can, so far, only perform at the two lowest levels of thinking. However, reflective thinking is within reach. If neural networks were able to report what features had been detected in the hidden layers of the net, then the NN would be able to explain the reason for its output. This would enable deep learning to reach at least the third layer of cognition.

Minsky's multiplicities provide a straightforward benchmark to compare neural network systems to other techniques. For deep learning to succeed, it must operate on all 6 layers of cognition and be able to perform all five types of problem solving heuristics.

Story Understanding

UNDERSTANDING LANGUAGE IS A CRITICAL task to any sufficiently advanced system. Claims of success in language understanding first started in 1970, when Terry Winograd developed SHRDLU³⁶.

This program was able to carry on a dialog about a blocks-world environment with a user. However, critics of this system argued that the system did not truly understand language — that the system was not actually "thinking." This debate was later formalized in a 1980 paper by John Searle³⁷. In this paper, Searle defined the Chinese Room Argument.

Since that time, two prominent professors have attempted to more

³² Marvin Minsky. *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*. Simon & Schuster, 2006. ISBN 0743276639. URL <http://web.media.mit.edu/~minsky/Introduction.html>

³³ Joscha bach. Micropsi, 2009. URL <http://cognitive-ai.com>

³⁴ Randall C. O'Reilly. The leabra cognitive architecture. URL <http://psych.colorado.edu/~oreilly/papers/oreillyhazyherdip.pdf>

³⁵ There does exist a system that claims to operate on all 6 layers, which I discuss further in the next section.

³⁶ Terry Winograd. Procedures as a representation for data in a computer program for understanding natural language. Technical report, Massachusetts Institute of Technology, 1971. URL <http://hci.stanford.edu/~winograd/shrdlu/AITR-235.pdf>

³⁷ John R Searle. Minds, brains, and programs. *Behavioral and brain sciences*, 3(03):417-424, 1980. URL <http://cogprints.org/7150/1/10.1.1.83.5248.pdf>

formally understand computational language understanding. First, Noam Chomsky, a linguistic and cognitive scientist, has described the "merge" operation³⁸, which he believes is crucial to human language. The merge operation is the ability — which is unique to humans — to take two concepts and form a new concept, indefinitely. In his recent book, *Why Only Us*³⁹, Chomsky explains the evolutionary plausibility of this uniquely-human capability.

In work that pairs well with Chomsky's merge operation, Patrick Winston has spent decades working on computational story understanding. Winston's work largely centers around the *Strong Story Hypothesis*⁴⁰, which states that storytelling and story understanding have a central role in human intelligence.

Winston's implementation of his ideas is a program known as Genesis⁴¹. Genesis is able to understand stories and answer questions based on a very small amount of training information. For example, when given the story of Macbeth (by William Shakespeare), Genesis is able to answer the question "why did Macduff kill Macbeth?" and respond based on perspective. From a personality perspective, Genesis can respond "Macduff kills Macbeth because he is vicious." From a commonsense perspective, Genesis answers "... Macduff kills Macbeth, probably because Macduff wants to kill Macbeth and Macduff is vicious." From a concept level, the response is "Macduff kills Macbeth is part of acts of Revenge, Answered prayer, Mistake because harmed, Pyrrhic victory, and Tragic greed."

With these type of question and answering capabilities, Winston makes the argument that Genesis understands stories at a human level. The argument is also made that Genesis operates on all levels of Minsky's six layer model (described in the pervious section).

No deep learning system has come close to this level of cognition. However, there are introductory steps that show promise. Recent statistical modeling, namely word2vec, is able to take symbolic language and convert words into a vector-space, allowing for computational manipulation. This vectorization of information enables compositionally⁴² of concepts. For example, the vector transformation of *Madrid - Spain + France* results in the vector *Paris*. Therefore, the knowledge representation understands the *Madrid - Spain* to be the concept of a country capitol. Then the capital concept is combined with France to produce the capital of France, which is Paris.

This aligns well with Chomsky's merge operation. Multiple concepts are formed together to produce a new concept. Since the new concept remains a distributed, vector representation, the new concept can be merged with other concepts indefinitely.

Recurrent Neural Network architectures also show promise in an attempt to computationally model language. Bengio⁴³ has described

³⁸ Precisely, Chomsky describes this operation: "an indispensable operation of a recursive system [...] which takes two syntactic objects *A* and *B* and forms the new object $G = \{A, B\}$ "

³⁹ Robert C. Berwick and Noam Chomsky. *Why Only Us*. The MIT Press, 2016

⁴⁰ Patrick Henry Winston. The strong story hypothesis and the directed perception hypothesis. *Proceedings of the AAAI Fall Symposium on Advances in Cognitive Systems*, 2011. URL <http://hdl.handle.net/1721.1/67693>

⁴¹ Patrick Henry Winston. The genesis story understanding and story telling system: A 21st century step toward artificial intelligence. Memo 019, Center for Brains Minds and Machines, MIT, 2014. URL <http://groups.csail.mit.edu/genesis/papers/StoryWhitePaper.pdf>

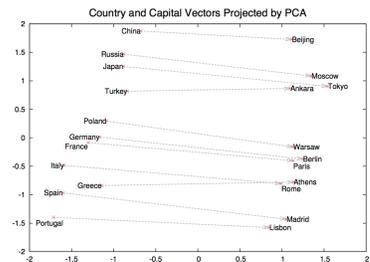


Figure 9: This figure shows a 2 dimensional view of countries and their capitals. The projection necessary to translate from a country to a capital is clear. (This figure is referenced from Mikolov et. al.)

⁴² Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013. URL <https://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf>

⁴³ Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, and Jürgen Schmidhuber. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies, 2001. URL <http://www.bioinf.jku.at/publications/older/ch7.pdf>

the potential for sequential understanding, and recent commercial successes have progressed the field. However, major huddles still lay ahead.

To show that neural networks are a plausible fundamental substrate, two major accomplishments need to be made. First, neural network natural language processing must increase its accuracy when training data is extremely low. Genesis is able to learn stories in under 100 sentences, from a collection of under 25 stories. Today's best systems still use databases of millions of entries. Second, these systems must be able to operate at the upper levels of Minsky's 6 layers. Once the reflective and self-reflective levels are reached, the argument for deep learning models as a base layer will become much stronger. Further work on both of these areas is much needed.

Contributions

IN THIS PAPER, I have argued that the field of Artificial Intelligence preforms a type of gradient descent. I have outlined the periods of success, followed by periods of realization that the current technique has major limitations.

I have explained that deep neural networks are a computationally plausible solution for a substrate that can enable all future AI systems. I explained three areas of research in Artificial Intelligence that deep learning has not solved. For each area, I have described the current state of the art, the reason that area is critical to AI, and what deep learning methods show promise for future success.

After many decades of research, the field has the tools to develop a substrate that can power all AI systems. However, if the field only focuses on the easy and obvious problems, the next AI winter will certainly ensue. To continue progress towards general artificial intelligence — and to escape the local minimum — the deep learning community must focus on the tasks that are far from being reached, but necessary for long term success.

References

Language and machines: Computers in translation and linguistics.

URL <http://www.nap.edu/read/9547/chapter/1>.

Tower of hanoi, 2016. URL https://en.wikipedia.org/wiki/Tower_of_Hanoi.

Tower_of_Hanoi.

John Abell. March 15, 1985: Dot-com revolution starts with a whim-

- per, 1985. URL <http://www.wired.com/2010/03/0315-symbolics-first-dotcom/>.
- James Atwood and Don Towsley. Search-convolutional neural networks. *CoRR*, abs/1511.02136, 2015. URL <http://arxiv.org/abs/1511.02136>.
- Joscha bach. Micropsi, 2009. URL <http://cognitive-ai.com>.
- Hans Berliner. Deep thought wins fredkin intermediate prize. *AI Magazine*, 10(2), 1989. URL <http://www.aaai.org/ojs/index.php/aimagazine/article/viewFile/753/671>.
- Robert C. Berwick and Noam Chomsky. *Why Only Us*. The MIT Press, 2016.
- Daniel G Bobrow. *Natural Language Input for a Computer Problem Solving System*. PhD thesis, Massachusetts Institute of Technology, 1964. URL <http://hdl.handle.net/1721.1/5922>.
- Murray Campbell, A Joseph Hoane, and Feng-hsiung Hsu. Deep blue. *Artificial intelligence*, 134(1):57–83, 2002. URL <http://www.sciencedirect.com/science/article/pii/S0004370201001291>.
- Hung-Che Chen and Jyh-Da Wei. Using neural networks for evaluation in heuristic search algorithm. In *Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011. URL <http://www.aaai.org/ocs/index.php/AAAI/AAAI11/paper/view/3476>.
- Daniel Crevier. *AI: The Tumultuous History of the Search for Artificial Intelligence*. Basic Books, Inc., New York, NY, USA, 1993. ISBN 0-465-02997-3.
- Brad Darrach. Meet shaky, the first electronic person, 1970. URL <https://books.google.com/books?id=2FMEAAAAMBAJ&lpg=PA57&pg=PA58#v=onepage&q&f=true>.
- Charles R. Dyer. Game playing. URL http://pages.cs.wisc.edu/~dyer/cs540/notes/06_gamePlaying.pdf.
- T.G Evans. A heuristic program to solve geometric analogy problems, 1962. URL <http://hdl.handle.net/1721.1/6101>.
- Artur d’Avila Garcez, Tarek R Besold, Luc de Raedt, Peter Földiak, Pascal Hitzler, Thomas Icard, Kai-Uwe Kühnberger, Luis C Lamb, Risto Miikkulainen, and Daniel L Silver. Neural-symbolic learning and reasoning: contributions and challenges. 2015. URL <http://dase-lab.cs.wright.edu/pub/nesy2015.pdf>.

- Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014. URL <http://arxiv.org/abs/1410.5401>.
- Stevan Harnad. The symbol grounding problem, 1990. URL <http://users.ecs.soton.ac.uk/harnad/Papers/Harnad/harnad90.sgproblem.html>.
- Sepp Hochreiter, Yoshua Bengio, Paolo Frasconi, and Jürgen Schmidhuber. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies, 2001. URL <http://www.bioinf.jku.at/publications/older/ch7.pdf>.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. URL <https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf>.
- John E. Laird, Allen Newell, and Paul S. Rosenbloom. Soar: An architecture for general intelligence. *Artificial Intelligence*, 33, 1987. URL <http://www.sciencedirect.com/science/article/pii/0004370287900506>.
- Sir James Lighthill. Artificial intelligence: A general survey, 1974. URL http://www.chilton-computing.org.uk/inf/literature/reports/lighthill_report/p001.htm.
- Sir James Lighthill and John McCarthy. The lighthill debate, 1973. URL <https://www.youtube.com/watch?v=yReDbeY7ZMU>.
- AI Magazine. The ai behind watson—the technical article. *AI Magazine*, 2010. URL <http://www.aaai.org/Magazine/Watson/watson.php>.
- John McCarthy, Marvin L Minsky, Nathaniel Rochester, and Claude E Shannon. A proposal for the dartmouth summer research project on artificial intelligence, 2006. URL <http://www.aaai.org/ojs/index.php/aimagazine/article/view/1904/1802>.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013. URL <https://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf>.

- Marvin Minsky. Steps toward artificial intelligence. *Proceedings of the IRE*, 49(1):8–30, 1961. URL <http://web.media.mit.edu/~minsky/papers/steps.html>.
- Marvin Minsky. *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*. Simon & Schuster, 2006. ISBN 0743276639. URL <http://web.media.mit.edu/~minsky/Introduction.html>.
- Marvin Minsky. Artificial intelligence and the future of the human mind, 2007. URL <https://www.youtube.com/watch?v=g1DcIJgInF8>.
- Allen Newell, John C Shaw, and Herbert A Simon. Report on a general problem-solving program, 1959. URL http://129.69.211.95/pdf/rand/ipl/P-1584_Report_On_A_General_Problem-Solving_Program_Feb59.pdf.
- Randall C. O'Reilly. The leabra cognitive architecture. URL <http://psych.colorado.edu/~oreilly/papers/oreillyhazyherdip.pdf>.
- Stephen Polit. R1 and beyond: Ai technology transfer at digital equipment corporation. *AI Magazine*, 5(4):76, 1984. URL <http://www.aaai.org/ojs/index.php/aimagazine/article/view/460/396>.
- Andrew Pollack. Setbacks for artificial intelligence, 1988. URL <http://www.nytimes.com/1988/03/04/business/setbacks-for-artificial-intelligence.html?pagewanted=all>.
- Andrew Pollack. 'fifth generation' became japan's lost generation. *New York Times*, 1992. URL <http://www.nytimes.com/1992/06/05/business/fifth-generation-became-japan-s-lost-generation.html>.
- John R Searle. Minds, brains, and programs. *Behavioral and brain sciences*, 3(03):417–424, 1980. URL <http://cogprints.org/7150/1/10.1.1.83.5248.pdf>.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016. URL <http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html>.
- Push Singh. Em one: An architecture for reflective commonsense thinking, 2005. URL <http://web.media.mit.edu/%7Epush/push-thesis.html>.

- James Robert Slagle. A heuristic program that solves symbolic integration problems in freshman calculus, symbolic automatic integrator (saint), 1961. URL <http://hdl.handle.net/1721.1/11997>.
- Aaron Smith and Janna Anderson. Ai, robotics, and the future of jobs, 2014. URL <http://www.pewinternet.org/files/2014/08/Future-of-AI-Robotics-and-Jobs.pdf>.
- Terry Winograd. Procedures as a representation for data in a computer program for understanding natural language. Technical report, Massachusetts Institute of Technology, 1971. URL <http://hci.stanford.edu/~winograd/shrdlu/AITR-235.pdf>.
- Patrick Henry Winston. The strong story hypothesis and the directed perception hypothesis. *Proceedings of the AAAI Fall Symposium on Advances in Cognitive Systems*, 2011. URL <http://hdl.handle.net/1721.1/67693>.
- Patrick Henry Winston. The genesis story understanding and story telling system: A 21st century step toward artificial intelligence. Memo 019, Center for Brains Minds and Machines, MIT, 2014. URL <http://groups.csail.mit.edu/genesis/papers/StoryWhitePaper.pdf>.